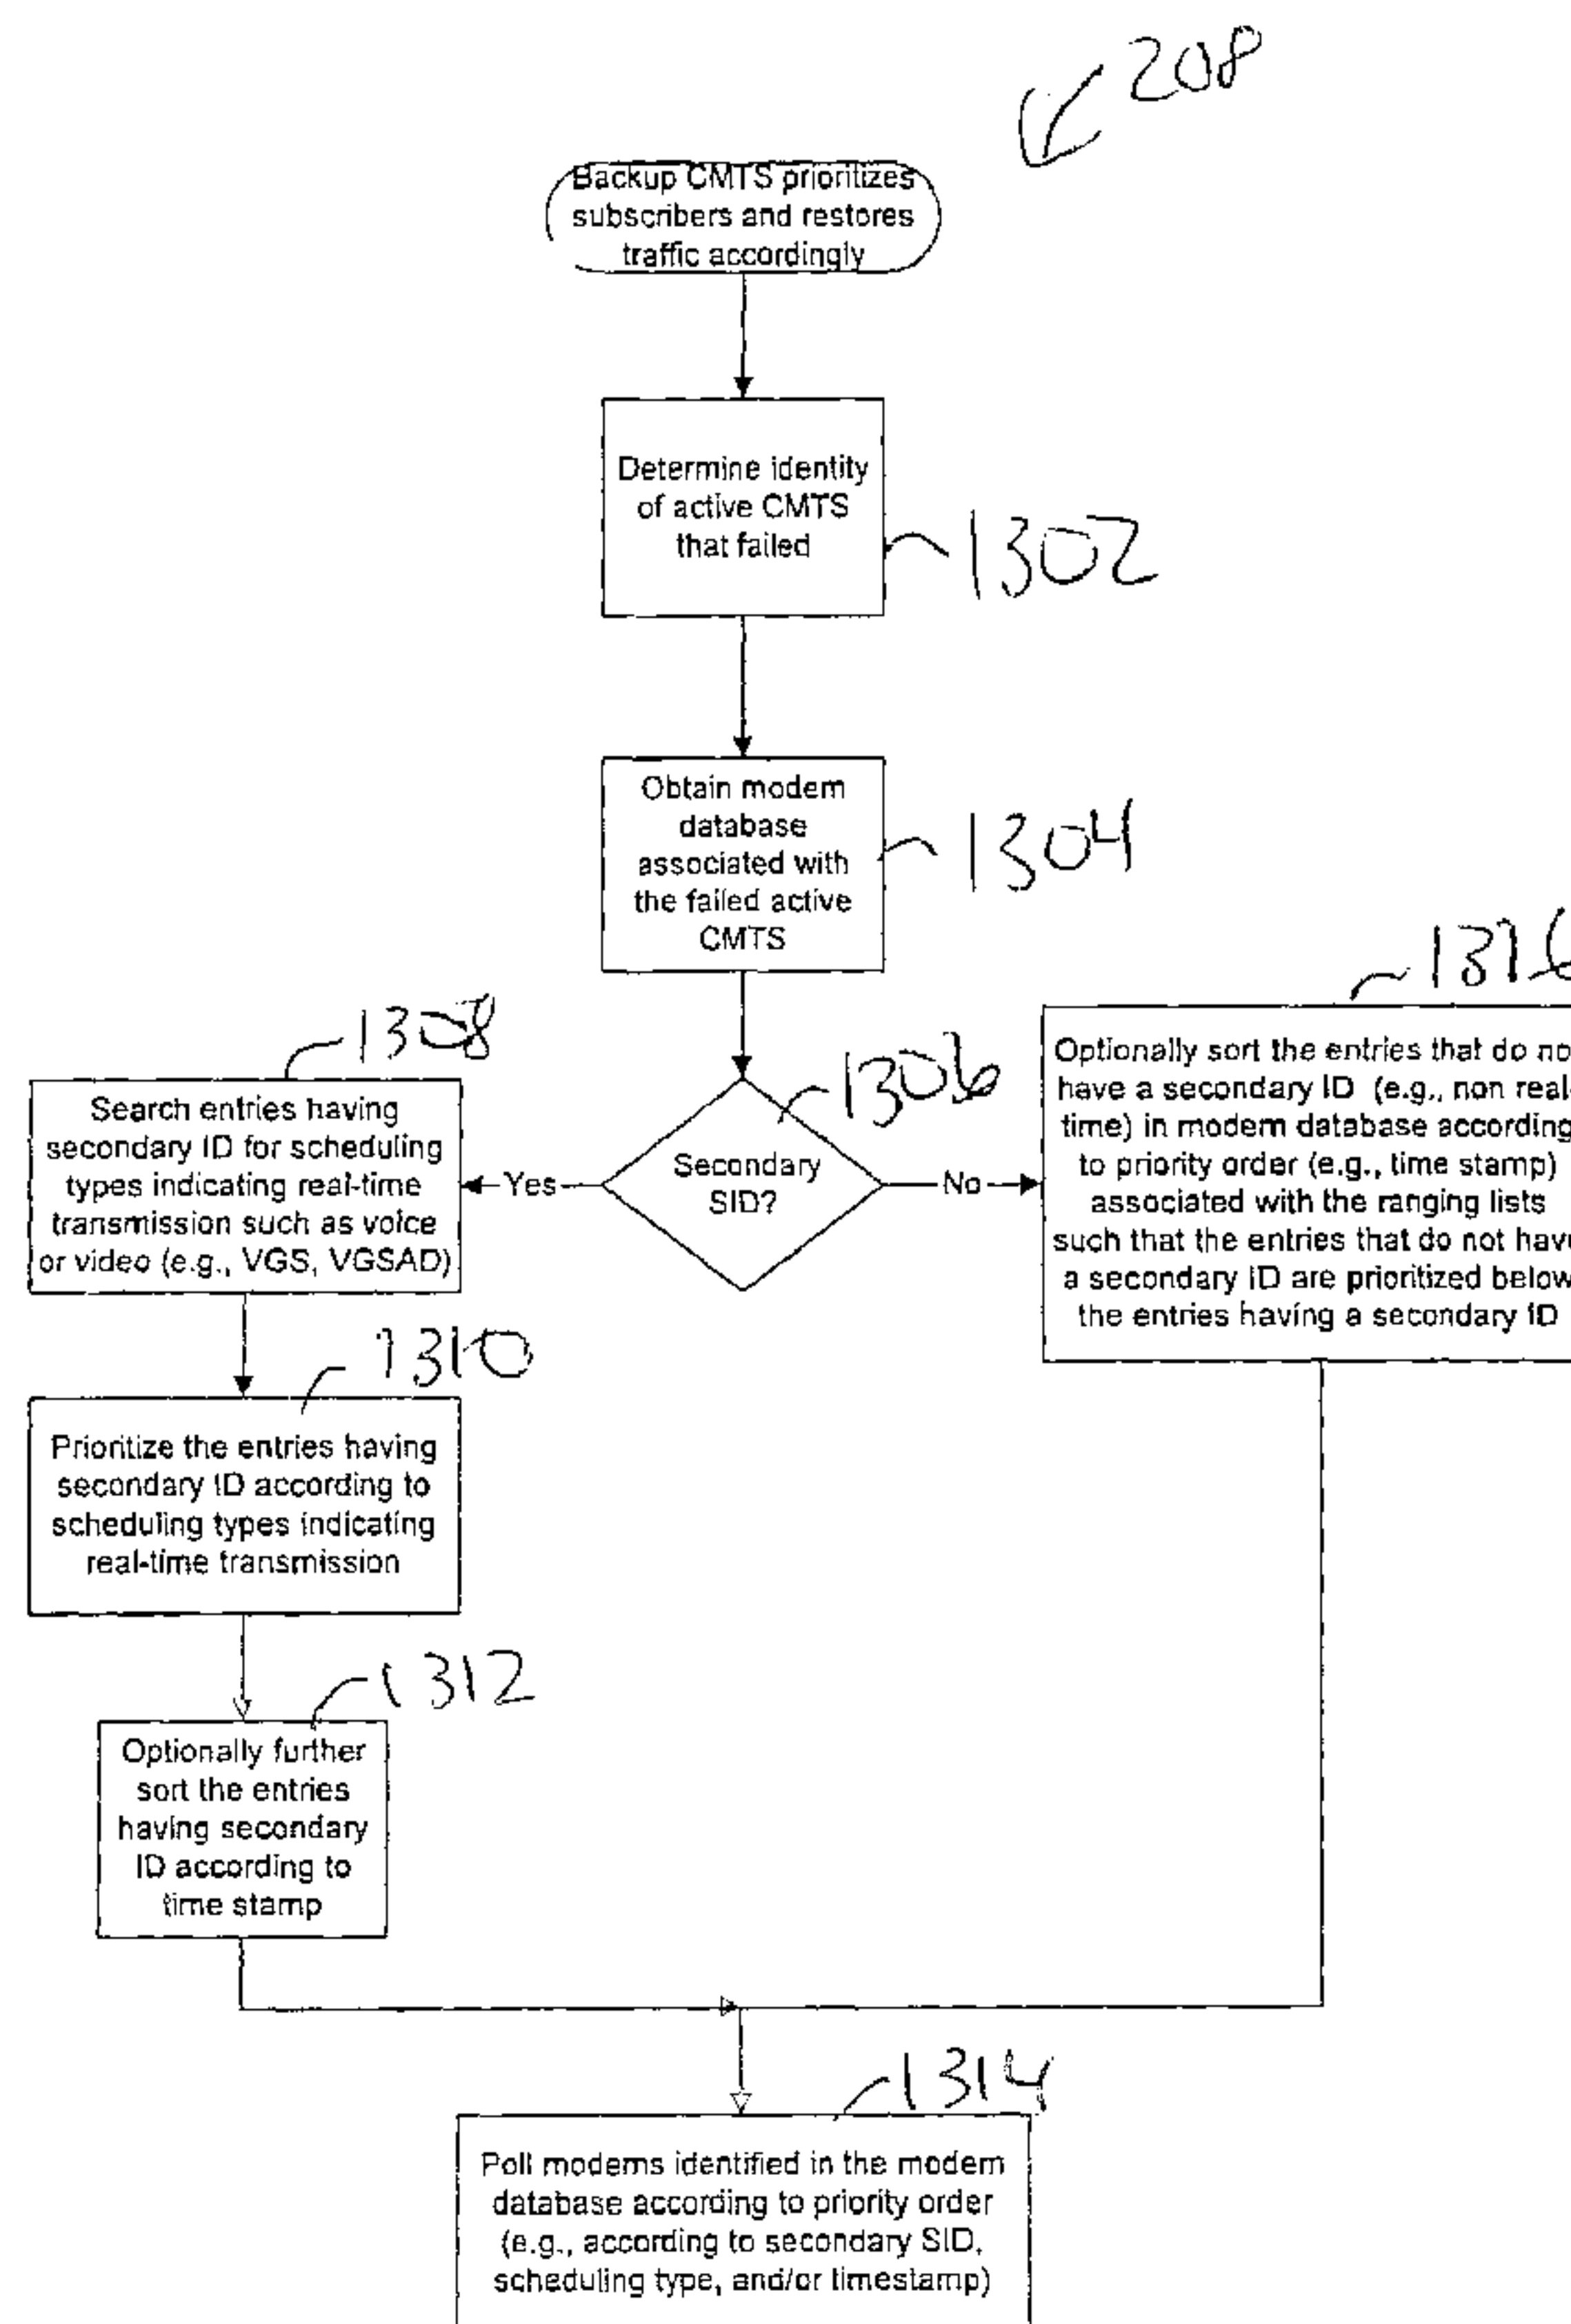




(86) Date de dépôt PCT/PCT Filing Date: 2003/01/24  
 (87) Date publication PCT/PCT Publication Date: 2003/10/16  
 (85) Entrée phase nationale/National Entry: 2004/07/26  
 (86) N° demande PCT/PCT Application No.: US 2003/002225  
 (87) N° publication PCT/PCT Publication No.: 2003/085908  
 (30) Priorité/Priority: 2002/01/28 (10/058,722) US

(51) Cl.Int.<sup>7</sup>/Int.Cl.<sup>7</sup> H04L 12/56  
 (71) Demandeur/Applicant:  
CISCO TECHNOLOGY, INC., US  
 (72) Inventeurs/Inventors:  
LU, YONG, US;  
ZHANG, JIN, US;  
KHAUNTE, SUNIL, US;  
CHANDRAN, KARTIK, US  
 (74) Agent: CASSAN MACLEAN

(54) Titre : APPAREILS ET PROCÉDES DE RETABLISSEMENT DU TRAFIC LORS DE LA REPRISE SUR PANNE  
DANS UN TÊTE DE CABLES  
 (54) Title: METHOD COMPUTER-READABLE MEDIUM AND SYSTEMS FOR RESORTING TRAFFIC DURING  
FAILOVER IN A CABLE HEAD END



(57) Abrégé/Abstract:

Disclosed are methods and systems for guaranteeing restoration of traffic between one or more cable modems and a backup cable modem termination system upon failure of an active cable modem termination system. Subscriber information associated with one

(57) **Abrégé(suite)/Abstract(continued):**

or more cable modems is received from an active cable modem termination system. The subscriber information includes one or more subscriber identifiers. The subscriber information may include a primary subscriber identifier that identifies a particular cable modem as well as a secondary subscriber identifier that is assigned to high priority traffic such as that being transmitted in real-time. In addition, the subscriber information may also include a scheduling type that is further used to categorize the real-time traffic, such as into voice or video traffic. The subscriber information is then prioritized in an order in which the transmission of messages between the one or more cable modems and the backup cable modem are to be restored. For instance, the subscriber information may be prioritized according to time of receipt of the subscriber information, presence of secondary subscriber identifier, and/or scheduling type. The cable modems are scheduled according to this priority order.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
16 October 2003 (16.10.2003)

PCT

(10) International Publication Number  
**WO 2003/085908 A3**

(51) International Patent Classification<sup>7</sup>: **H04L 12/56**

(21) International Application Number:  
PCT/US2003/002225

(22) International Filing Date: 24 January 2003 (24.01.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
10/058,722 28 January 2002 (28.01.2002) US

(71) Applicant (for all designated States except US): **CICSO TECHNOLOGY, INC.** [US/US]; 170 West Tasman Drive, San Jose, CA 95134 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **LU, Yong** [CN/US];

5506 Manderston Drive, San Jose, CA 95138 (US). **ZHANG, Jin** [CN/US]; 10784 Johnson Avenue, Cupertino, CA 95014 (US). **KHAUNTE, Sunil** [IN/US]; 3700 Lillick Drive, Santa Clara, CA 95051 (US). **CHANDRAN, Kartik** [IN/US]; 581 Tahoe Terrace, Mountain View, CA 94041 (US).

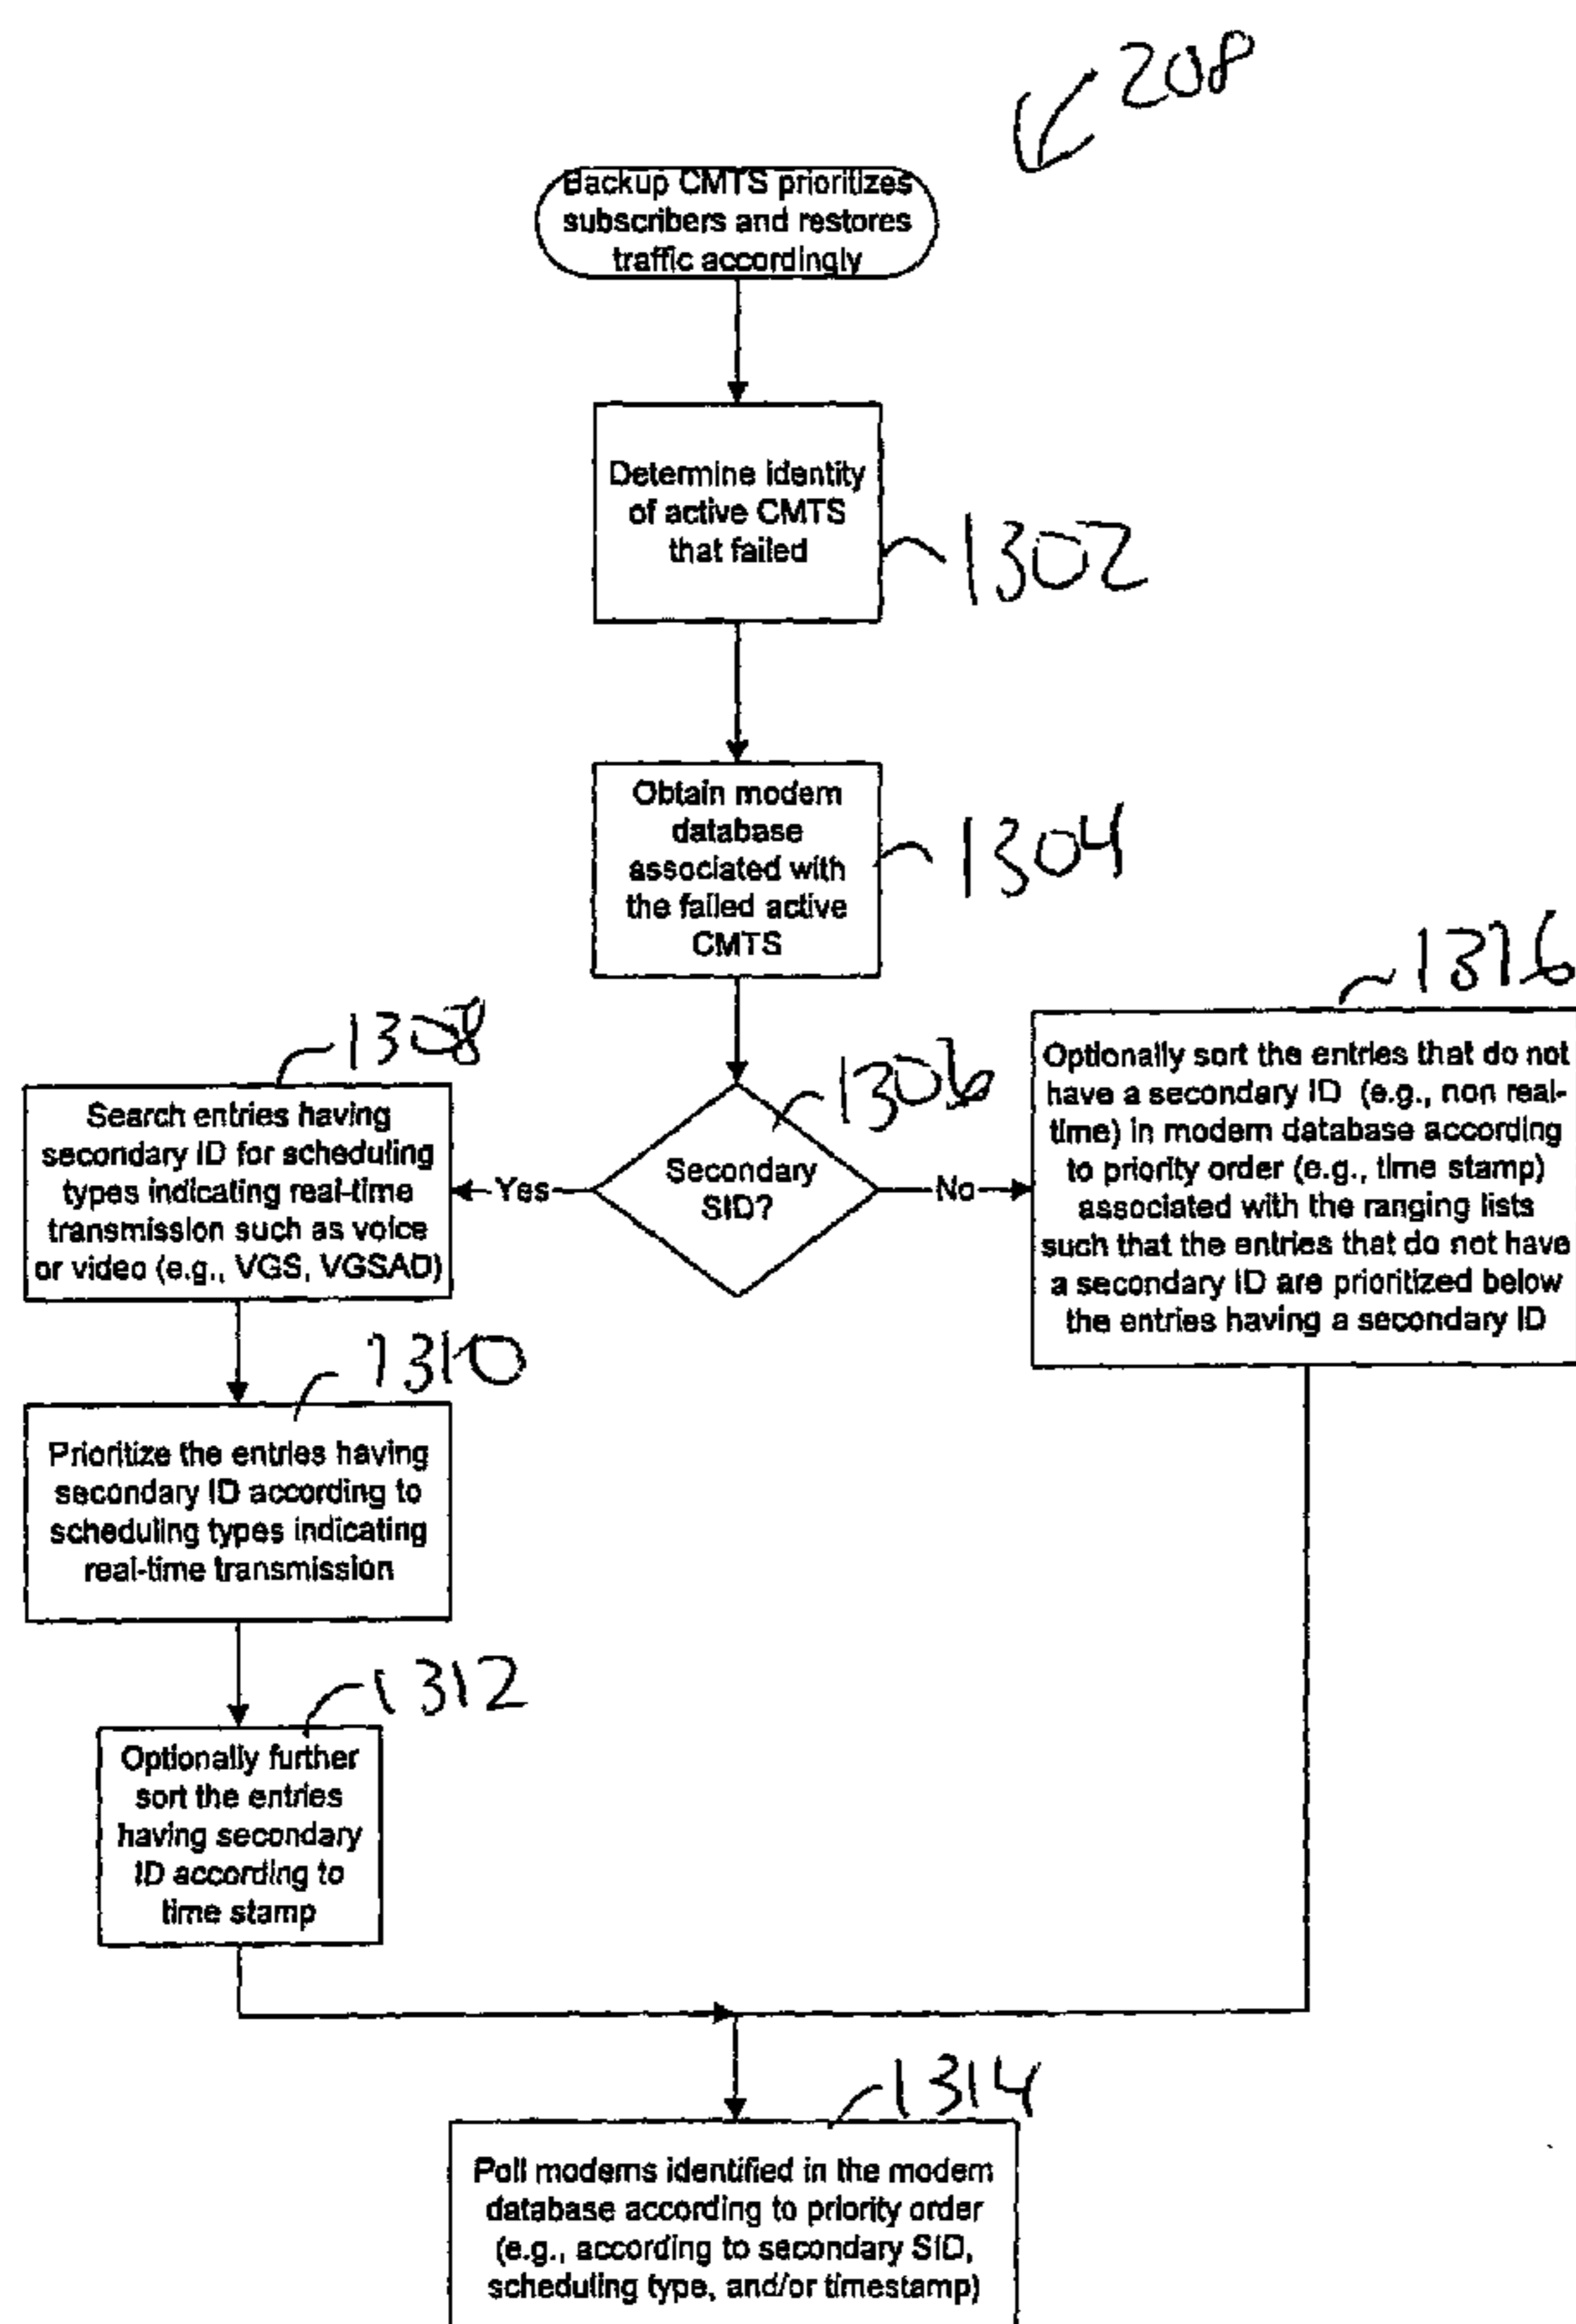
(74) Agent: **HEILBRUNN, Elise, R.**; Beyer Weaver & Thomas LLP, P.O. Box 778, Berkeley, CA 94704 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),

[Continued on next page]

(54) Title: METHOD COMPUTER-READABLE MEDIUM AND SYSTEMS FOR RESORTING TRAFFIC DURING FAILOVER IN A CABLE HEAD END



(57) Abstract: Disclosed are methods and systems for guaranteeing restoration of traffic between one or more cable modems and a backup cable modem termination system upon failure of an active cable modem termination system. Subscriber information associated with one or more cable modems is received from an active cable modem termination system. The subscriber information includes one or more subscriber identifiers. The subscriber information may include a primary subscriber identifier that identifies a particular cable modem as well as a secondary subscriber identifier that is assigned to high priority traffic such as that being transmitted in real-time. In addition, the subscriber information may also include a scheduling type that is further used to categorize the real-time traffic, such as into voice or video traffic. The subscriber information is then prioritized in an order in which the transmission of messages between the one or more cable modems and the backup cable modem are to be restored. For instance, the subscriber information may be prioritized according to time of receipt of the subscriber information, presence of secondary subscriber identifier, and/or scheduling type. The cable modems are polled according to this priority order.

WO 2003/085908 A3

**WO 2003/085908 A3**



Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),  
European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,  
ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, SE, SI,  
SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN,  
GQ, GW, ML, MR, NE, SN, TD, TG).

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**Published:**

— *with international search report*

**(88) Date of publication of the international search report:**

19 February 2004

## APPARATUS AND METHODS FOR RESTORING TRAFFIC DURING FAILOVER IN A CABLE HEAD END

### 5 BACKGROUND OF THE INVENTION

The present invention relates to digital computer network technology. More specifically, it relates to methods and apparatus for restoring traffic within an access network.

10 Broadband access technologies such as cable, fiber optic, and wireless have made rapid progress in recent years. Recently there has been a convergence of voice and data networks which is due in part to United States deregulation of the telecommunications industry. In order to stay competitive, companies offering broadband access technologies need to support voice, video, and other high-bandwidth applications over their local access networks. For networks that use a  
15 shared access medium to communicate between subscribers and the service provider (e.g., cable networks, wireless networks, etc.), providing reliable high-quality voice/video communication over such networks is not an easy task.

One type of broadband access technology relates to cable modem networks. A cable modem network or "cable plant" employs cable modems, which are an  
20 improvement of conventional PC data modems and provide high speed connectivity. Cable modems are therefore instrumental in transforming the cable system into a full service provider of video, voice and data telecommunications services.

A head end complex and cable modem typically communicate. This communication may be one-way or two-way between the cable modem and head end  
25 complex. In one way communication, the cable modem merely receives data from the head end complex, but cannot transmit data to the head end complex. For example, the cable modem cannot request particular types of data from the head end complex. In two-way communication, the cable modem can transmit data (e.g., data requests) to the head end complex.

There are circumstances that may interfere with the communication between a cable Head End and a cable modem. As one example, there may be a hardware failure of the cable Head End. As another example, there may be a software failure within the cable Head End. As yet another example, there may be a power outage  
5 that causes communication between the cable Head End and the cable modem to be disrupted.

In order to ensure that communication between a cable Head End and a cable modem has not been interrupted, a handshake is typically exchanged between the cable Head End and the cable modem. More particularly, this handshake is typically  
10 initiated by the cable Head End. Thus, when the cable modem does not receive a communication from the cable Head End within a period of time, the cable modem declares a time out.

When a time out occurs, the cable modem initiates a new connection with the cable Head End. Unfortunately, this new connection takes time to establish. For  
15 instance, it may take up to 24 seconds to rebuild a connection between a cable modem and its cable Head End.

In the event that a cable Head End fails, it would be undesirable to abruptly disconnect the modem from the cable Head End to re-initiate a connection. This is particularly undesirable when real-time traffic such as voice or video data is being  
20 transmitted. Even if the connection were merely delayed, such a substantial delay during a real-time transmission such as a phone call is unacceptable.

In view of the above, it would be desirable to prevent a time out of a cable modem when a cable Head End fails. Moreover, it would be desirable if the effects of a time out could be minimized.  
25

## SUMMARY OF THE INVENTION

The present invention relates to the prioritization and restoration of communication between a cable Head End and one or more modems upon failure of  
30 a cable modem termination system (CMTS). This is accomplished, through the use of a backup CMTS that supports one or more active CMTSs. The backup CMTS

prioritizes the modems according to associated subscriber information, and polls the modems accordingly. In this manner, the occurrence of a time out of a cable modem upon failure of a CMTS is prevented.

Disclosed are methods and apparatus for guaranteeing restoration of traffic  
5 between one or more cable modems and a backup cable modem termination system upon failure of an active cable modem termination system are disclosed. Subscriber information associated with one or more cable modems are received from an active cable modem termination system. The subscriber information includes one or more  
10 subscriber identifiers. For instance, the subscriber information may include a primary subscriber identifier that identifies a particular cable modem as well as a secondary subscriber identifier that is assigned to high priority traffic such as that being transmitted in real-time. In addition, the subscriber information may also include a scheduling type that is further used to categorize the real-time traffic, such as into voice or video traffic. The subscriber information is then prioritized in an  
15 order in which the transmission of messages between the one or more cable modems and the backup cable modem are to be restored. For instance, the subscriber information may be prioritized according to time of receipt of the subscriber information, presence of secondary subscriber identifier, and/or scheduling type. The cable modems are then polled according to this priority order.

20 Yet another aspect of the invention pertains to computer program products including machine-readable media on which are provided program instructions for implementing the methods and techniques described above, in whole or in part. Any of the methods of this invention may be represented, in whole or in part, as program instructions that can be provided on such machine-readable media. Moreover, the  
25 invention pertains to a cable modem termination system employing the disclosed methods. In addition, the invention pertains to various combinations and arrangements of data generated and/or used as described herein. For example, modem databases having the format described herein and provided on appropriate media are part of this invention.

30

These and other features and advantages of the present invention will be presented in more detail in the following specification of the invention and the accompanying figures which illustrate by way of example the principles of the invention.

5

#### BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a diagram illustrating a specific embodiment of a cable network that may be used with the techniques of the present invention.

10 Figure 2 is a process flow diagram illustrating a method of prioritizing subscribers associated with one or more cable modems in accordance with various embodiments of the present invention.

Figure 3 is a process flow diagram illustrating a method of synchronizing subscriber information for one or more subscribers between an active CMTS and a backup CMTS as shown at block 204 of FIG. 2 in accordance with an embodiment of the invention.

Figure 4 is a block diagram illustrating information stored by a backup CMTS in association with one or more ranging lists received by an active CMTS as shown at block 306 of FIG. 3 in accordance with an embodiment of the invention.

20 Figure 5 is a process flow diagram illustrating a method of providing a failover trigger from an active CMTS to a backup CMTS as shown at block 206 of FIG. 2.

Figure 6 is a process flow diagram illustrating a method of prioritizing subscribers and restoring traffic accordingly as shown at block 208 of FIG. 2 in accordance with one embodiment of the invention.

25

Figure 7 is a diagram illustrating an exemplary ranging list that may be sent by an active CMTS to a backup CMTS as shown at block 304 of FIG. 3.



Figure 8 is a diagram illustrating an exemplary hello message that may be sent by an active CMTS and corresponding exemplary hello acknowledgement message that may be sent by a backup CMTS as shown at blocks 506 and 510 of FIG. 5, respectively.

5 Figure 9 is a diagram illustrating exemplary databases storing prioritized subscriber information for associated active CMTS' that may be maintained by a backup CMTS as shown at block 606 of FIG. 6.

Figure 10 is a process flow diagram illustrating a method of performing a handshake between a backup CMTS and one or more cable modems as shown at  
10 block 210 of FIG. 2 in accordance with various embodiments of the invention.

Figure 11 is a process flow diagram illustrating a method of synchronizing subscriber information for one or more subscribers between an active CMTS and a backup CMTS as shown at block 204 of FIG. 2 in accordance with another embodiment of the invention.

15 Figure 12 is a diagram illustrating an exemplary sync message that may be sent by an active CMTS to a backup CMTS as shown at block 1104 of FIG. 11.

Figure 13 is a process flow diagram illustrating a method of prioritizing subscribers and restoring traffic accordingly as shown at block 208 of FIG. 2 in accordance with another embodiment of the invention.

20 Figure 14 is a diagram illustrating an exemplary tree data structure that may be used to store prioritized subscriber information for an active CMTS in accordance with various embodiments of the invention.

Figure 15 is a process flow diagram illustrating a method of ending a call initiated by a modem in accordance with various embodiments of the invention.

25 Figure 16 is a block diagram illustrating a Cable Modem Termination System (CMTS) which may be used for implementing one or more techniques of the present invention.

Figure 17 is a block diagram illustrating a specific embodiment of a line card which may be used for implementing certain aspects of the present invention.

Figure 18 is a block diagram illustrating a network device suitable for implementing one or more techniques of the present invention.

5

## DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

Reference will now be made in detail to a specific embodiment of the invention. An example of this embodiment is illustrated in the accompanying drawings. While the invention will be described in conjunction with this specific embodiment, it will be understood that it is not intended to limit the invention to one embodiment. On the contrary, it is intended to cover alternatives, modifications, and equivalents as may be included within the spirit and scope of the invention as defined by the appended claims. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. The present invention may be practiced without some or all of these specific details. In other instances, well known process operations have not been described in detail in order not to unnecessarily obscure the present invention.

Figure 1 shows a block diagram of a two-way hybrid fiber-coaxial (HFC) cable network 100 in accordance with one embodiment of the present invention. Of course, the present invention may also be implemented on a one-way cable network. As shown in Figure 1, the cable network 100 includes a Head End complex 102 typically configured to service about 40,000 homes. The Head End complex 102 may include a plurality of components and/or systems (not shown) such as, for example, a Head End, a super Head End, a hub, a primary hub, a second hub, etc. Additionally, as shown in Figure 1, the Head End complex 102 typically includes a Cable Modem Termination System (CMTS) 120. Primary functions of the CMTS include (1) receiving data inputs from external sources 100 and converting the data for transmission over the cable plant; (2) providing appropriate Media Access Control (MAC) level packet headers for data received by the cable system, and (3) modulating and demodulating the data to and from the cable network. Typically, the Head End complex 102 is configured to provide a communication interface between

nodes (e.g. cable modems) in the cable network and external networks such as, for example, the Internet 142. The cable modems typically reside at the subscriber premises 110A-D.

5 The Head End Complex 102 is typically connected to one or more fiber nodes 106 in the cable network. Each fiber node is, in turn, configured to service one or more subscriber groups 110. Each subscriber group typically comprises about 500 to 2000 households. A primary function of the fiber nodes 106 is to provide an optical-electronic signal interface between the Head End Complex 102 and the plurality of cable modems residing at the plurality of subscriber groups 110.

10 In order for data to be able to be transmitted effectively over a computer network such as HFC or other broadband computer networks, a common standard for data transmission is typically adopted by network providers. A commonly used and well known standard for transmission of data or other information over HFC networks is the Data Over Cable System Interface Specification (DOCSIS). The  
15 DOCSIS standard has been publicly presented by Cable Television Laboratories, Inc. (Louisville, Colorado), in a document entitled, DOCSIS 1.1 RF Interface Specification (document control number SP-RFIV1.1-I06-001215, December 15, 2000). This document is incorporated herein by reference for all purposes.

20 Communication between the Head End Complex 102 and fiber node 106a is typically implemented using modulated optical signals which travel over fiber optic cables. More specifically, during the transmission of modulated optical signals, multiple optical frequencies are modulated with data and transmitted over optical fibers such as, for example, optical fiber links 105a and 105b of Figure 1, which are typically referred to as "RF fibers". As shown in Figure 1, the modulated optical  
25 signals transmitted from the Head End Complex 102 eventually terminate at the fiber node 106a. The fiber nodes maintain the signal modulation while converting from the fiber media to the coax media and back.

Each of the fiber nodes 106 is connected by a coaxial cable 107 to a respective group of cable modems residing at subscriber premises 110A-D.  
30 According to the DOCSIS standard, specific frequency ranges are used for transmitting downstream information from the CMTS to the cable modems, and

other specific frequency ranges are used for transmitting upstream information from the cable modems to the CMTS.

In order to allow the cable modems to transmit data to the CMTS, the cable modems share one or more upstream channels within that domain. Access to the upstream channel is controlled using a time division multiplexing (TDM) approach. Such an implementation requires that the CMTS and all cable modems sharing an upstream channel within a particular domain have a common concept of time so that when the CMTS tells a particular cable modem to transmit data at time T, the cable modem understands what to do. "Time" in this context may be tracked using a counter, commonly referred to as a timer counter, which, according to conventional implementations is a 32-bit counter that increments by one every clock pulse.

Typically, digital data on upstream and downstream channels of the cable network is carried over radio frequency ("RF") carrier signals. Cable modems convert digital data to a modulated RF signal for upstream transmission and convert downstream RF signal to digital form. The conversion is done at a subscriber's facility. At a Cable Modem Termination System ("CMTS"), located at a Head End Complex of the cable network, the conversions are reversed. The CMTS converts downstream digital data to a modulated RF signal, which is carried over the fiber and coaxial lines to the subscriber premises. The cable modem then demodulates the RF signal and feeds the digital data to a computer. On the return path, the digital data is fed to the cable modem (from an associated PC for example), which converts it to a modulated RF signal. Once the CMTS receives the upstream RF signal, it demodulates it and transmits the digital data to an external source.

Over time different types of messages are communicated between the head end 102 and one or more cable modem(s). In one embodiment, these messages may include MAP messages, dynamic service messages, and other types of data messages. The head end 102 periodically sends a MAP message that specifies one or more times slots for identified cable modems to the cable modems 110 associated with such head end 102. A cable modem 110 or CMTS may send a dynamic service message for adding, deleting, or changing a particular service type or a parameter of an existing service. A data message is generally defined as any type of message that

is communicated between the head end 102 and one or more cable modems. For instance, the message may include real-time data such as voice or video data.

In accordance with various embodiments of the invention, a backup or “protect” CMTS 150 supports one or more CMTSs 120. These CMTSs 120 will be referred to as “active” CMTSs to distinguish them from the backup CMTS 150 which operates merely in a backup mode. Thus, an active CMTS will operate until it is no longer functional. At that time, the backup CMTS 150 will “replace” or take over those operations of the non-functional active CMTS(s). In other words, upon failure of an active CMTS, the backup CMTS 150 will receive messages from and transmit messages to modems previously supported by the active CMTS. The CMTSs 120 and backup CMTS 150 may be supported in a single head end complex as shown, or in separate head end complexes.

The present invention provides mechanisms for restoring transmission of messages between one or more cable modems and a backup cable modem termination system upon failure of an active cable modem termination system. Subscriber information associated with one or more cable modems is received from an active cable modem termination system. The subscriber information may include one or more subscriber identifiers. For instance, the subscriber information may include one or more primary subscriber identifiers, each of which uniquely identifies a cable modem. As another example, the subscriber information may include one or more secondary subscriber identifiers, which may be used to indicate the type of traffic being transmitted. For instance, the secondary subscriber identifier is typically used to identify real-time data traffic such as voice or video data. The backup cable modem termination system then prioritizes the subscriber information, the prioritized subscriber information indicating an order in which the transmission of messages between the cable modems and the backup cable modem is to be restored. The backup cable modem termination system then poll the cable modems according to the priority order indicated by the prioritized subscriber information.

Messages transmitted between the head end complex 120 and various subscribers may be transmitted or received via any suitable network, such as wide area network 142 (*e.g.*, Internet). In this manner, subscriber devices such as cable modems may transmit and receive messages.

Figure 2 is a process flow diagram illustrating a method of prioritizing subscribers associated with one or more cable modems in accordance with various embodiments of the present invention. As shown, an active or “working” CMTS typically performs a handshake with its supported cable modems at block 202. More specifically, the active CMTS sends a ranging poll to the modems, requesting that the supported modems send its operating parameters to the active CMTS. In response, the modems send a ranging request with its parameters. The active CMTS then sends its ranging response, indicating that the parameters received by the associated modem are accepted.

At block 204 the active CMTS and backup CMTS synchronize subscriber information for one or more subscribers (e.g., modems). In other words, the backup CMTS receives or obtains subscriber information from one or more active CMTSs. Various methods of synchronizing subscriber information will be described in further detail below with reference to FIG. 3 and FIG. 11.

When an active CMTS fails, a failover trigger is provided by the active CMTS to one or more backup CMTSs at block 206. The “failover trigger” indicates to the backup CMTS that the active CMTS is no longer functional. Such a failover trigger is typically implemented through the periodic detection of “hello messages.” In other words, a failover trigger may be implemented through an absence of “hello messages” or failure to detect “hello messages” for a specified period of time. One method of implementing a failover trigger will be described in further detail below with reference to FIG. 5.

Once the backup CMTS determines that an active CMTS has failed, it then prioritizes the subscribers supported by the CMTS and restores traffic accordingly at block 208. The subscribers may be prioritized for each failed CMTS. In addition, various methods of prioritizing the subscriber information for each failed CMTS may be implemented. Two exemplary methods of prioritizing the subscribers will be described in further detail below with reference to FIG. 6 and FIG. 13.

A handshake is then performed between the backup CMTS and cable modems previously supported by the failed CMTS(s) at block 210. For instance, the backup CMTS may simply perform a handshake such as that described above with

reference to block 202. One method of performing a handshake by a backup CMTS will be described in further detail below with reference to FIG. 10.

Figure 3 is a process flow diagram illustrating a method of synchronizing subscriber information for one or more subscribers between an active CMTS and a backup CMTS as shown at block 204 of FIG. 2 in accordance with an embodiment of the invention. Any event may trigger the sending of subscriber information from an active CMTS to a backup CMTS. In this example, a lapse of a specified period of time at block 302 triggers the sending of the subscriber information as shown at block 304. In other words, the subscriber information is sent periodically to the backup CMTS. The subscriber information preferably identifies those modems that have ranged successfully, which may be referred to as a "ranging list." An exemplary ranging list will be described in further detail below with reference to FIG. 7.

When the backup CMTS receives the subscriber information (e.g., ranging list), it preferably stores the time of receipt of the subscriber information by the backup CMTS such that the subscriber information is associated with the time of receipt as shown at block 306. Various data structures may be used to store the time-stamped subscriber information. Such exemplary data structures are described below with reference to FIG. 4. The backup CMTS may then send an acknowledgement of the subscriber information to the active CMTS at block 308.

Figure 4 is a block diagram illustrating information stored by a backup CMTS in association with one or more ranging lists received by an active CMTS as shown at block 306 of FIG. 3 in accordance with an embodiment of the invention. The information stored by a backup CMTS for a single active CMTS may include subscriber information received at multiple points in time. For instance, as shown, a first ranging list 402 includes multiple entries, each of which includes a primary subscriber identifier 404 that uniquely identifies a cable modem. Each entry preferably further includes a media access control (MAC) address 406 associated with the cable modem. A set of subscriber information such as a ranging list 402 is associated with a time of receipt 408 of the ranging list 402 by the backup CMTS. Similarly, a second ranging list 410 includes multiple entries, each of which include a primary subscriber identifier 412 and a MAC address 414. The second ranging list

416 is similarly associated with a time of receipt 416 of the ranging list 416 by the backup CMTS.

Figure 5 is a process flow diagram illustrating a method of providing a failover trigger from an active CMTS to a backup CMTS as shown at block 206 of FIG. 2. As shown, heartbeat timing is configured in the active CMTS at block 502, which indicates the frequency with which hello messages are to be sent from the active CMTS. Similarly, the backup CMTS is also configured with failure detection timing at block 504 so that the backup CMTS is configured with a number of consecutive heartbeats to be missed (undetected) to signal a failure of the active CMTS. Thus, a heartbeat detection counter is preferably initialized to zero to indicate the number of heartbeats that have been missed. It is then determined at block 506 whether the backup CMTS has received a hello message from the active CMTS. If so, the heartbeat detection counter is reset to zero at block 508 and the backup CMTS sends a hello acknowledgement message to the active CMTS at block 510. An exemplary hello and hello acknowledgement message will be described in further detail below with reference to FIG. 8. If a hello message is not received from the active CMTS, the heartbeat detection counter is incremented at block 512. If the heartbeat detection counter is equal to the number of consecutive heartbeats to miss at block 514, the backup CMTS declares the active CMTS dead at block 516.

Figure 6 is a process flow diagram illustrating a method of prioritizing subscribers and restoring traffic accordingly as shown at block 208 of FIG. 2 in accordance with one embodiment of the invention. As shown, the backup CMTS determines the identity of the failed CMTS at block 602. Subscriber information associated with one or more cable modems previously supported by the active CMTS is then obtained at block 604. For instance, a database maintained by the backup CMTS storing subscriber information associated with the active CMTS may be obtained. The subscriber information may then be sorted according to a priority order (e.g., specified by time stamp of each ranging list) at block 606, as will be described in further detail below with reference to FIG. 9. However, such sorting may be unnecessary. In other words, the subscriber information may be automatically prioritized according to time stamp when the subscriber information is stored, since the subscriber information (e.g., ranging list) and associated time of



receipt may be stored consecutively (in order of the time of receipt) upon receipt in a file or database. The modems identified in this prioritized subscriber information are then polled according to the implicit or explicit priority order at block 608. In this manner, the modems that have been most recently polled are polled last and the  
5 modems not polled recently are polled first.

As described above, the subscriber information may include one or more ranging lists. Figure 7 is a diagram illustrating an exemplary ranging list that may be sent by an active CMTS to a backup CMTS as shown at block 304 of FIG. 3. As shown, the ranging list 700 includes a header 702, type 704, length 706, and value  
10 708 field. For instance, the length may indicate the number of primary subscriber identifiers present in the value field 708.

Identification of a failed CMTS may occur when a specified number of consecutive hello messages are not detected by a backup CMTS. Figure 8 is a diagram illustrating an exemplary hello message that may be sent by an active CMTS  
15 and corresponding exemplary hello acknowledgement message that may be sent by a backup CMTS as shown at blocks 506 and 510 of FIG. 5, respectively. As shown, a hello or hello acknowledgement message 800 includes a header 802, type 804, length 806 and value 808 field. As shown, the value field 808 may indicate whether the message is a hello or hello acknowledgement message. The hello and hello  
20 acknowledgement messages may be implemented in a variety of protocols such as the Hot Chassis to Chassis Protocol (HCCP).

As described above, subscriber information associated with one or more cable modems may be received periodically from one or more active cable modem termination systems, which may then be separately prioritized and stored in memory  
25 (e.g., in a file or database). Figure 9 is a diagram illustrating exemplary databases storing prioritized subscriber information for associated active CMTS' that may be maintained by a backup CMTS as shown at block 606 of FIG. 6. As shown, a first database 902 stores a plurality of ranging lists obtained from a first CMTS that are prioritized according to time stamp. Similarly, a second database 904 stores a  
30 plurality of ranging lists obtained from a second CMTS that are prioritized according to time stamp.

Figure 10 is a process flow diagram illustrating a method of performing a handshake between a backup CMTS and one or more cable modems as shown at block 210 of FIG. 2 in accordance with various embodiments of the invention. As shown, the backup CMTS sends a ranging poll message at block 1002 to one or more modems. The modems then send a ranging request at block 1004. The CMTS receives the ranging request at block 1006 and sends a ranging response. When the CMTS receives a DOCSIS message at block 1008 including a Dynamic Service Address (DSA) request, the CMTS sends a DSA response including a secondary subscriber identifier 1010. A secondary subscriber identifier is typically assigned to a particular modem to indicate the type of data being transmitted. For instance, the secondary subscriber identifier may indicate that the data is being transmitted in real-time (e.g., voice or video data). The modem then sends a DSA acknowledgement at block 1012.

As described above, subscriber information is preferably prioritized separately for each CMTS rather than for all CMTSs combined. The subscriber information may be prioritized according to time stamp. However, the subscriber information may also be prioritized according to other criteria, instead of, or in combination with the time stamps. For instance, the presence of a secondary subscriber identifier may also be used to prioritize various subscribers.

When information such as a secondary subscriber identifier is used to prioritize a particular call, it may be desirable to trigger the sending of the associated subscriber information upon receipt of the call by the active CMTS from a cable modem (rather than periodically sending the subscriber information after a specified period of time). Figure 11 is a process flow diagram illustrating a method of synchronizing subscriber information for one or more subscribers between an active CMTS and a backup CMTS as shown at block 204 of FIG. 2 in accordance with another embodiment of the invention. The event that triggers the sending of subscriber information may be the receipt of a new call as shown at block 1102. Once received, the active CMTS sends subscriber information in the form of a sync message to the backup CMTS at block 1104. An exemplary sync message will be described in further detail below with reference to FIG. 12.

When the backup CMTS receives the sync message, the backup CMTS stores the subscriber information obtained from the sync message at block 1106. The backup CMTS may also store the time of receipt of the sync message at block 1108. The backup CMTS may also send an acknowledgement to the active CMTS at block 5 1110. Thus, upon subsequent failure of an active CMTS, the backup CMTS may prioritize the stored subscriber information.

Figure 12 is a diagram illustrating an exemplary sync message that may be sent by an active CMTS to a backup CMTS as shown at block 1104 of FIG. 11. As shown, a sync message 1200 and therefore subscriber information may include a 10 primary subscriber identifier 1202 as well as a secondary subscriber identifier 1204. The sync message 1200 may also include a service template including quality of service requirements 1206. The quality of service requirements 1206 may define attributes of the traffic, which may include a variety of attributes, such as packet size, maximum traffic rate, For instance, the quality of service requirements 1206 may 15 include attributes of voice traffic. The quality of service requirements 1206 also typically include a scheduling type. For instance, various possible scheduling types include best effort, variable bit rate, Unsolicited Grant Service (UGS) or Unsolicited Grant with Activity Detection (UGSAD). In addition, the sync message 1200 may also include a packet classifier 1208. A packet classifier typically includes a source 20 and destination address, which are used to determine the treatment to be applied to packets as specified in the service template according to the specified source and destination address. For instance, calls initiated by and/or received from various IP addresses may be treated differently, and therefore may be prioritized accordingly.

As described above, prioritization of subscribers and restoration of traffic 25 may be performed solely based upon priority order as described above with reference to FIG. 6. However, as described above with reference to FIG. 12, subscriber information (e.g., sent in the form of one or more sync messages) may include various information such as a secondary subscriber identifier. Figure 13 is a process flow diagram illustrating a method of prioritizing subscribers and restoring traffic 30 accordingly as shown at block 208 of FIG. 2 in accordance with another embodiment of the invention. As shown, the backup CMTS determines the identity of the previously active CMTS that has failed at block 1302. Subscriber information

associated with the failed active CMTS are then obtained at block 1304. For instance, one or more modem databases storing subscriber information for one or more modems supported by the active CMTS may be obtained from memory by the backup CMTS.

5           Once the subscriber information for the failed active CMTS is obtained by the backup CMTS, it is then prioritized according to the traffic type (instead of or in addition to time stamp). More specifically, as described above, the secondary subscriber identifier indicates that the messages to be transmitted between the backup CMTS and the associated cable modem are to be transmitted in real-time. For  
10 instance, the messages to be transmitted may include voice or video data. Thus, the traffic having an assigned secondary subscriber identifier is prioritized above that having solely a primary subscriber identifier. As shown at block 1306, the subscriber information (e.g., each entry) is searched for a secondary subscriber identifier. The subscriber information (and therefore the associated modems) having  
15 a secondary subscriber identifier is then prioritized such that the subscriber information has a higher priority than the subscriber information that does not have a secondary subscriber identifier.

In addition, the scheduling type further indicates the type of traffic to be transmitted for traffic having an associated secondary subscriber identifier. More  
20 specifically, the scheduling type may indicate whether the messages to be transmitted between the backup CMTS and the associated cable modem include voice data or video data. Possible scheduling types may include, for example, Unsolicited Grant Service (UGS) and Unsolicited Grant Service with Activity Detection (UGS-AD). Thus, within the entries having a secondary subscriber identifier, the entries may be  
25 further searched and prioritized according to scheduling type at blocks 1308 and 1310. It is also possible to further prioritize the subscriber information having a secondary identifier according to time of receipt (e.g., time stamp) of the subscriber information by the backup CMTS from the active CMTS at block 1312. Of course, as described above with reference to FIG. 6, it is possible to prioritize the entries that  
30 do not have a secondary identifier (e.g., non real-time traffic) according to time of receipt of the ranging lists as shown at block 1316. These sorted entries that do not have a secondary subscriber identifier associated therewith will therefore be

prioritized among themselves, but will be prioritized below those having a secondary subscriber identifier. Alternatively, those entries not having a secondary subscriber identifier may simply be prioritized as a group below those having a secondary subscriber identifier. Once prioritized, the modems identified in the modem database  
5 are polled according to priority order at block 1314. As described above, this priority order may be performed according to secondary subscriber identifier, scheduling type, and/or time stamp.

Figure 14 is a diagram illustrating an exemplary tree data structure 1400 that may be used to store prioritized subscriber information for an active CMTS in  
10 accordance with various embodiments of the invention. As shown, those entries 1402 having both a primary subscriber identifier and a secondary subscriber identifier are prioritized above those entries 1404 having only a primary subscriber identifier. Within the entries 1402 having both a primary subscriber identifier and a  
15 secondary subscriber identifier, the entries are prioritized according to scheduling type. Similarly, within the entries 1404 having solely a primary subscriber identifier, the entries are prioritized according to time stamp.

Once a modem call is completed, it is desirable to modify the stored subscriber information maintained by the backup CMTS. At this time, the active CMTS is non-functional, and therefore the switchover to the backup CMTS has  
20 already been performed. Thus, although the telephone call was initiated when the active CMTS was operational, the telephone survived the crash of the active CMTS that subsequently occurred. Figure 15 is a process flow diagram illustrating a method of ending a call initiated by a modem in accordance with various  
25 embodiments of the invention. As shown, the backup CMTS receives a dynamic service completion request from a modem at block 1502 that indicates that the call initiated by the cable modem has been terminated. The backup CMTS sends a dynamic service completion response to the modem at block 1504. The backup CMTS then removes the subscriber information (e.g., entry) associated with the  
30 cable modem from memory associated with a previously failed active CMTS at block 1506. For instance, as described above, a separate CMTS database may be maintained by the backup CMTS for each active CMTS it supports.

Generally, the techniques of the present invention may be implemented on software and/or hardware. For example, it can be implemented in an operating system kernel, in a separate user process, in a library package bound into network applications, on a specially constructed machine, or on a network interface card. In a  
5 specific embodiment of this invention, the technique of the present invention may be implemented in software such as an operating system or in an application running on an operating system.

A software or software/hardware hybrid system of this invention is preferably implemented on a general-purpose programmable machine selectively activated or reconfigured by a computer program stored in memory. Such a programmable  
10 machine may be a network device designed to handle network traffic. Such network devices typically have multiple network interfaces. One important class of device that may be used to implement the present invention is the Cable Modem Termination System. Preferably, the CMTS is a "routing" CMTS, which handles at  
15 least some routing functions.

Figure 16 shows a block diagram of a specific embodiment of a Cable Modem Termination System (CMTS) 1600 which may be used to implement certain aspects of the present invention. As shown in Figure 16, the CMTS 1600 may comprise a plurality of routing engines (e.g. 1601a, 1601b). In a specific  
20 implementation, Routing Engine A 1601a may be configured as a primary or working routing engine, while Routing Engine B 1601b may be configured as a backup or standby routing engine which provides redundancy functionality.

As shown in the embodiment of Figure 16, each of the routing engines may include a variety of similar modules and/or components. In order to avoid confusion,  
25 the various components and/or modules relating to Routing Engine A 1601a will now be described in greater detail with the understanding that such descriptions may also be applied to the corresponding components and modules of Routing Engine B 1601b.

According to a specific embodiment, Routing Engine A may be configured or  
30 designed to include a plurality of functionally different modules or components, including, for example, a Forwarding Processor (FP) Module 1611a adapted to

provide packet forwarding functionality; a Route Processor (RP) Module 1603a adapted to implement routing or forwarding operations; a utility component 1602a adapted to provide system clock and timestamp functionality; etc. The routing engine components may be configured to provide layer one, layer two, layer three  
5 and layer four functionality as well as quality of service (QoS) functionality.

According to a specific implementation, the RP Module 1603a may be configured as a processor-based routing system comprising functionality incorporated within a typical router, such as, for example, specially configured router models 1600, 2500, 2600, 3600, 4500, 4700, 7200, 7500, 10012, and 12000 available  
10 from Cisco Systems, Inc. of San Jose, California. For example, as shown in the embodiment of Figure 16, the RP Module 1603a comprises a general-purpose processor 1605a (e.g., a MIPS route processor) coupled to a system controller 1609a and memory 1607a. It should be noted that components have been described in singular form for clarity. One skilled in the art would appreciate that multiple  
15 processors, a variety of memory formats, or multiple system controllers, for example, can be used in this context as well as in other contexts while falling within the scope of the present invention. The memory 1607a may comprise synchronous dynamic random access memory (SDRAM) storage locations addressable by the processor 1605a for storing software programs and data structures accessed by the components.  
20 A network routing operating system, portions of which may reside in memory and executed by the route processor, functionally organizes the router by invoking network operations in support of software processes executing on the router.

The RP processor 1605a may be configured to construct and load routing tables used by the FP Module 1611a. The processor 1605a may also be configured  
25 or designed to perform configuration management functions of the routing engine 1601a, and to communicate with neighboring peer, standby, and/or backup routers to exchange protocol data units used to construct the routing tables in accordance with conventional routing algorithms. It will be apparent to those skilled in the art that other memory types, including various computer readable media, may be used for  
30 storing and executing program instructions pertaining to the operation of the routing engine.

Interface circuitry 1627a may be coupled to the respective interface circuitry 1633a, 1633b of line cards 1631a, 1631b. According to a specific implementation, interface circuitry 1627a may be configured to reside on a backplane logic circuit 1623a of the routing engine. In one example, the backplane logic circuit 1623a is embodied as a high performance, application specific integrated circuit (ASIC). An example of a backplane logic circuit that may be advantageously used with the present invention is disclosed in co-pending and commonly owned U.S. Patent Application Serial No. 09/791,063, filed on February 22, 2001, the entirety of which is hereby incorporated by reference for all purposes.

According to a specific embodiment, the backplane logic circuit (which, according to a specific implementation, may be configured as an ASIC), may be configured to further interface the line cards to a packet buffer 1625a and a forwarding engine 1621a of the FP Module 1611a. The packet buffer 1625a may include memory which is configured to store packets as the forwarding engine 1621a performs its packet forwarding functions. For example, the packet buffer may be used to store low priority data packets while high priority, low latency voice packets are forwarded by the forwarding engine to a data network interface 1635a. According to various embodiments, the FP Module 1611 may comprise a processor 1613a and memory 1615a for handling transport layer 1617 and network layer 1619 functionality. In one implementation, the processor 1613a may be configured to track accounting, port, and billing information for various users on a cable modem network 1651. The processor 1613a may also be configured to maintain desired service flow or session state information in memory 1615a such as, for example, for voice calls initiated over the cable modem network. The FP Module 1611a may also be configured to provide transaction compacting functionality, data parcel tunneling functionality, switching functionality, filtering option selection and filtering functionality, etc.

According to a specific implementation, Routing Engine A 1601a may be connected to Routing Engine B 1601b via at least one link 1646, such as, for example, a backplane line or system bus. Routing engine redundancy may be provided by designating one of the routing engines as the working or primary routing engine and designating the other routing engine(s) as the redundant or standby



routing engine(s). When configured as a working routing engine, the Routing Engine A may perform all appropriate forwarding and routing functions. When a failure occurs at the working routing engine, the redundant routing engine (e.g. Routing Engine B) may then take over the operations of the working routing engine.

5 Thereafter, when Routing Engine A recovers, it may assume the functions of the redundant routing engine, or it may take over the functions of the working routing engine.

According to different embodiments of the present invention, one or more of the routing engines may be configured to communicate with a plurality of line cards (e.g. 1631, 1635) via point-to-point links. For example, as shown in Figure 16, each of the plurality of line cards 1631 and 1635 are connected to each of the routing engines 1601a, 1601b via point-to-point links 1641 and 1643. One advantage of the point-to-point link configuration is that it provides additional reliability in that the failure of one or more line cards will not interfere with communications between

10 other line cards and the routing engine(s). For example, if Line Card A 1631a suddenly failed, each of the routing engines would still be able to communicate with the other line cards.

15

According to a specific embodiment, the plurality of line cards may include different types of line cards which have been specifically configured to perform specific functions. For example, line cards 1631 may correspond to radio-frequency (RF) line cards which have been configured or designed for use in a cable network. Additionally, line cards 1635 may correspond to network interface cards which have been configured or designed to interface with different types of external networks (e.g. WANs, LANs,) utilizing different types of communication protocols (e.g. Ethernet, Frame Relay, ATM, TCP/IP, etc). For example, the data network interface 1635a functions as an interface component between external data sources and the cable system. The external data sources transmit data to the data network interface 1635a via, for example, optical fiber, microwave link, satellite link, or through various media. A data network interface may include hardware and software for

20

25

30 interfacing to various networks. According to various embodiments, a data network interface may be implemented on a line card as part of a conventional router for a packet-switched network. Using this type of configuration, the CMTS is able to send

and/or receive IP packets to and from the data network interface using, for example, network layer software 1619a.

According to a specific implementation, the operations associated with obtaining an IP address for cable modems may be implemented by the network layer software. This may involve the CMTS communicating with a DHCP server (not shown) via a data network interface, for example.

As shown in Figure 16, at least a portion of the line cards includes interface circuitry for providing an appropriate interface between the host line card, other line cards, and/or the routing engine(s). For example, interface circuitry 1633a may include interconnect ports coupled to one or more of the point-to-point links 1641, 1643. According to a specific implementation, the interface circuitry functions as a translator that converts conventional formats of data received at the line cards to a suitable protocol format for transmission from the line card to the appropriate routing engine. In one implementation, the interface circuitry 1633a may also include circuitry to perform cyclic redundancy code (CRC) generation and checking on packets, along with interconnect format checking.

According to a specific embodiment, the point-to-point links 1641, 1643 may be configured as clock forwarded links such that each point-to-point link comprises at least one data wire for transporting data signals and at least one clock wire for carrying clock signals. However, it will be understood to those skilled in the art that the clock forwarding technique may be scaled to accommodate other clock forwarding arrangements such as, for example, connections comprising a plurality of data signals and/or clock signals. Additionally, according to a specific embodiment, each line card may be configured to provide at least one communication interface between the routing engines (1601a, 1601b) and a portion of the cable network. The data network interface 1635a may couple the routing engine 1601a to an external data network 1655 such as, for example, the Internet.

According to one embodiment, all or selected lines cards, routing engines and/or data network interfaces may be configured to use at least one common dedicated line or backplane (e.g. 1645). According to other embodiments, the routing engines 1601a, 1601b may have an additional dedicated connection(s) for

supporting redundancy. In a specific implementation, the backplane may be configured as an Ethernet medium that is shared by the CMTS. When the line cards are inserted into the backplane, they communicate with the routing engines over the lines 1645 in accordance with a “capabilities” exchange that identifies the types of  
5 line cards and their various characteristics/parameters.

According to a specific implementation, during initialization of the CMTS, the routing engines 1601a and 1601b negotiate for working routing engine status over the backplane. Assertion of working status causes the line cards 1631 to configure their respective interface circuitry to communicate with the designated  
10 working routing engine (e.g. Routing Engine A 1601a). The Routing Engine A 1601a then configures the CMTS and line cards, establishes routing relationships, and initiates traffic forwarding operations. The redundant routing engine 1601b may complete a self-test and perform initialization of its various functions. The two routing engine assemblies may then exchange conventional negotiation messages  
15 (which may include, for example, health and status messages) via the backplane lines 1645. According to a specific implementation, the exchanged messages are defined by an Enhanced High System Availability (EHSA) negotiation algorithm available from Cisco Systems, Inc. of San Jose, California. The redundant routing engine may also request transaction information from the working routing engine.

20 When the redundant routing engine 1601b detects that the primary routing engine has failed, the redundant routing engine may take over as the new working routing engine, and initiate a “cutover” operation to thereby cause the line card interface circuitry (e.g. 1633a, 1633b) to identify and communicate with the new working routing engine 1601b. The new working routing engine 1601b may then  
25 access and retrieve state information (such as, for example, telephone call state information, service flow state information, etc.) stored on selected line cards in order to maintain existing service flows.

Prior to a failure situation, the redundant routing engine 1601b may be configured to monitor the status of the working routing engine 1601a, and may  
30 further be configured or designed to receive updated configuration, transaction and/or state information, which may then be stored in an appropriate location in the redundant routing engine 1601b.

The line cards may further comprise circuitry for “looping” packets back onto the redundant routing engine 1601b over the point-to-point links. This allows the redundant routing engine 1601b to send and receive test packets to evaluate its own operation in addition to the operation of the dedicated lines prior to the occurrence of  
5 a system failure.

The techniques of the present invention may be implemented on various general purpose Cable Modem Termination Systems. In a specific embodiment, the systems of this invention may be specially configured CMTSs such as, for example, specially configured models in the uBR-7200 and uBR-10012 series of CMTSs  
10 available from Cisco Systems, Inc. of San Jose, California. In an alternative embodiment, the methods of this invention may be implemented on a general-purpose network host machine such as a personal computer or workstation. Further, the invention may be at least partially implemented on a card (e.g., an interface card) for a network device or a general-purpose computing device.

15 Although the system shown in Figure 16 represents one specific CMTS architecture of the present invention, it is by no means the only CMTS architecture on which the present invention can be implemented. For example, other types of interfaces and media could also be used with the CMTS.

Regardless of network device’s configuration (for cable plants or otherwise),  
20 it may employ one or more memories or memory modules (e.g., memory 1607a, 1615a, etc.) configured to store program instructions for the network operations and other functions of the present invention described herein. The program instructions may specify an operating system and one or more applications, for example. Such memory or memories may also be configured to store data structures, selected  
25 filtering criteria, or other specific non-program information described herein.

Because such information and program instructions may be employed to implement the systems/methods described herein, the present invention relates to machine-readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable  
30 media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical

media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). The invention may also be embodied in a carrier wave travelling over an appropriate medium such as airwaves, optical lines, electric lines, etc. Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher level code that may be executed by the computer using an interpreter.

Figure 17 shows a specific embodiment of a line card 1700 which may be used for implementing certain aspects of the present invention. According to a specific embodiment, the line card 1700 may be configured or designed to implement selected aspects of the DOCSIS functionality which were conventionally implemented by the CMTS, such as, for example, DOCSIS MAC functionality.

In the specific embodiment as shown in Figure 17, line card 1700 provides functions on several network layers, including a physical layer 1732, and a Media Access Control (MAC) layer 1730. Generally, the physical layer is responsible for receiving and transmitting RF signals on the cable plant. Hardware portions of the physical layer include at least one downstream modulator and transmitter 1706 and/or at least one upstream demodulator and receiver 1714. The physical layer also includes software 1786 for driving the hardware components of the physical layer.

Upstream optical data signals (packets) arriving via an optical fiber node are converted to electrical signals, and then demodulated by the demodulator/receiver 1714. The demodulated information is then passed to MAC layer block 1730.

A primary purpose of MAC layer 1730 is to encapsulate, with MAC headers, downstream packets and decapsulate, of MAC headers, upstream packets. In one embodiment, the encapsulation and decapsulation proceed as dictated by the above-mentioned DOCSIS standard for transmission of data or other information. The MAC headers include addresses to specific modems (if sent downstream), or to the CMTS (if sent upstream). Note that the cable modems also include MAC addressing components. In the cable modems, these components encapsulate upstream data with a header containing the MAC address of the CMTS.

MAC layer 1730 includes a MAC hardware portion 1734 and a MAC software portion 1784. The MAC layer software portion may include software relating to DOCSIS MAC functionality, filtering option selection and filtering functionality, etc. The MAC layer hardware and software portions operate together to provide the above-described DOCSIS MAC functionality. In a preferred embodiment, MAC controller 1734 is dedicated to performing some MAC layer functions, and is distinct from processor 1755.

After MAC layer block 1730 has processed the upstream information, it is then passed to interface circuitry 1702. As described previously, interface circuitry 1702 includes the appropriate hardware and/or software for converting data formats received at the line cards to a suitable protocol format for transmission from the line card to an appropriate routing engine.

When a packet is received from the routing engine at the interface circuitry 1702, the packet is then passed to MAC layer 1730. The MAC layer 1730 also transmits information via a one-way communication medium to downstream modulator and transmitter 1706. Downstream modulator and transmitter 1706 takes the data (or other information) in a packet structure and converts it to modulated downstream frames, such as MPEG or ATM frames, on the downstream carrier using, for example, QAM64 modulation. Other methods of modulation may also be used such as, for example, QAM256 modulation, CDMA (Code Division Multiple Access), OFDM (Orthogonal Frequency Division Multiplexing), FSK (FREQ Shift Keying), etc. The return data is likewise modulated using, for example, QAM16 or QSPK. According to a specific embodiment, the modulated data is converted from IF electrical signals to RF electrical signals (or vice-versa) using one or more electrical signal converters (not shown).

As shown in Figure 17, line card 1700 includes a central hardware block 1750 including one or more processors 1755 and memory 1757. These hardware components interact with software and other hardware portions of the various layers within the line card. They provide general purpose computing power for much of the software. Memory 1757 may include, for example, I/O memory (e.g. buffers), program memory, shared memory, etc. One or more data structures used for implementing the technique of the present invention may reside in such memory. In

one embodiment, the software entities 1782, 1784, and 1786 are implemented as part of a network operating system running on hardware 1750. Preferably, at least a part of the filtering functionality of this invention is implemented in software as part of the operating system. In Figure 17, such software may be part of MAC layer software 1784, or may be closely associated therewith. Of course, the filtering logic of the present invention could reside in hardware, software, or some combination of the two.

According to a specific implementation, the procedures typically employed by the CMTS during registration and pre-registration may be performed at the MAC layer of the line card 1700. In such an embodiment, most of the registration operations may be performed by the hardware and software provided for MAC layer logic 1730.

It will be appreciated that, according to a specific embodiments, at least a portion of functions described herein which are performed by the CMTS (e.g. Figure 7), line cards (e.g. Figure 17), or selected components thereof, may be implemented in a centralized CMTS system (e.g. residing at the Head End Complex of the cable network and/or may be implemented at one or more distributed CMTS (DCMTS) systems (e.g. residing at one or more fiber nodes).

Further, it will be appreciated by one having ordinary skill in the art that the technique of the present invention may be implemented in any computer network having a standardized protocol for utilizing a central termination system (e.g. Head End) to schedule timeslots for remote stations or nodes on a return (or upstream) channel. In wireless networks, the central termination system may be referred to as a Head End or wireless base station. In satellite networks, the central termination system may be referred to as a master controlling station.

Generally, the techniques (e.g., performed by the backup CMTS) of the present invention may be implemented on software and/or hardware. That is, the subscriber information may be obtained and prioritized by any suitable device. The prioritizing techniques may also be implemented within the same device on which the prioritized data is stored. By way of other examples, the prioritizing techniques can be implemented in an operating system kernel, in a separate user

process, in a library package bound into network applications, on a specially constructed machine, or on a network interface card. In a specific embodiment of this invention, the prioritizing techniques of the present invention is implemented in software such as an operating system or in an application running on an operating  
5 system.

A software or software/hardware hybrid implementation of the monitoring technique of this invention may be implemented on a general-purpose programmable machine selectively activated or reconfigured by a computer program stored in memory. Such programmable machine may be a network device designed to handle  
10 network traffic, such as, for example, a router or a switch. Such network devices may have multiple network interfaces including frame relay and ISDN interfaces, for example. Specific examples of such network devices include routers and switches. For example, the monitoring devices of this invention may be specially configured routers or servers such as specially configured router models 1600, 2500, 2600,  
15 3600, 4500, 4700, 7200, 7500, and 12000 available from Cisco Systems, Inc. of San Jose, California. A general architecture for some of these machines will appear from the description given below. In an alternative embodiment, the monitoring techniques of this invention may be implemented on a general-purpose network host machine such as a personal computer or workstation. Further, the invention may be  
20 at least partially implemented on a card (e.g., an interface card) for a network device or a general-purpose computing device.

Referring now to Figure 18, a network device 1860 suitable for implementing various techniques of the present invention includes a master central processing unit (CPU) 1862, interfaces 1868, and a bus 1867 (e.g., a PCI bus). When acting under  
25 the control of appropriate software or firmware, the CPU 1862 may be responsible for implementing specific functions associated with the functions of a desired network device. For example, when configured as a prioritizing device, the CPU 1862 may be responsible for analyzing packets, encapsulating packets, forwarding packets to appropriate network devices, etc. The CPU 1862 preferably accomplishes  
30 all these functions under the control of software including an operating system (e.g. Windows NT), and any appropriate applications software.



CPU 1862 may include one or more processors 1863 such as a processor from the Motorola family of microprocessors or the MIPS family of microprocessors. In an alternative embodiment, processor 1863 is specially designed hardware for controlling the operations of network device 1860. In a specific  
5 embodiment, a memory 1861 (such as non-volatile RAM and/or ROM) also forms part of CPU 1862. However, there are many different ways in which memory could be coupled to the system. Memory block 1861 may be used for a variety of purposes such as, for example, caching and/or storing data, programming instructions, etc.

The interfaces 1868 are typically provided as interface cards (sometimes  
10 referred to as "line cards"). Generally, they control the sending and receiving of data packets over the network and sometimes support other peripherals used with the network device 1860. Among the interfaces that may be provided are Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces, token ring interfaces, and the like. In addition, various very high-speed interfaces may be  
15 provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, HSSI interfaces, POS interfaces, FDDI interfaces and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM. The independent processors may control such  
20 communications intensive tasks as packet switching, media control and management. By providing separate processors for the communications intensive tasks, these interfaces allow the master microprocessor 1862 to efficiently perform routing computations, network diagnostics, security functions, etc.

Although the system shown in Figure 18 illustrates one specific network  
25 device of the present invention, it is by no means the only network device architecture on which the present invention can be implemented. For example, an architecture having a single processor that handles communications as well as routing computations, etc. is often used. Further, other types of interfaces and media could also be used with the network device.

30 Regardless of network device's configuration, it may employ one or more memories or memory modules (such as, for example, memory block 1865) configured to store data, program instructions for the general-purpose network

operations and/or other information relating to the functionality of the monitoring techniques described herein. The program instructions may control the operation of an operating system and/or one or more applications, for example. The memory or memories may also be configured to include the subscriber information, etc.

5           Because such information and program instructions may be employed to implement the systems/methods described herein, the present invention relates to machine readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable media are listed above.

10           It will be appreciated that the techniques of the present invention are not limited to cable networks, and may be applied to any access data network which uses at least one shared access communication channel to communicate between a plurality of nodes in the network and a Head End of the network.

15           Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. Therefore, the described embodiments should be taken as illustrative and not restrictive, and the invention should not be limited to the details given herein but should be defined by the following claims and their full scope of equivalents.

WHAT IS CLAIMED IS:

1. In a backup cable modem termination system, a method of restoring transmission of messages between one or more cable modems and the backup cable modem termination system upon failure of an active cable modem termination system, the method comprising:

receiving subscriber information associated with one or more cable modems from an active cable modem termination system, the subscriber information including one or more subscriber identifiers;

prioritizing the subscriber information, the prioritized subscriber information indicating an order in which the transmission of messages between the one or more cable modems and the backup cable modem is to be restored; and

polling the cable modems in the order indicated by the prioritized subscriber information.

2. The method as recited in claim 1, wherein prioritizing and polling are performed in response to a failover trigger from the active cable modem termination system.

3. The method as recited in claim 1, further comprising:  
determining that the active cable modem termination system has failed;  
wherein prioritizing and polling are performed after determining that the active cable modem termination system has failed.

4. The method as recited in claim 1, wherein the subscriber information identifies those modems that have ranged successfully.

5. The method as recited in claim 1, wherein receiving the subscriber information occurs after a specified period of time or after a call is received by the active cable modem termination system from one or more of the cable modems.

6. The method as recited in claim 1, further comprising:

storing the subscriber information after receiving the subscriber information;

wherein prioritizing the subscriber information comprises prioritizing the stored subscriber information.

5

7. The method as recited in claim 1, wherein the subscriber information associated with each of the cable modems comprises a primary subscriber identifier that identifies the associated cable modem.

10

8. The method as recited in claim 7, wherein the subscriber information further comprises a MAC address associated with the cable modem.

15

9. The method as recited in claim 7, wherein at least a portion of the subscriber information further comprises a secondary subscriber identifier.

10. The method as recited in claim 9, wherein the secondary subscriber identifier indicates that the messages to be transmitted between the backup cable modem termination system and the associated cable modem are to be transmitted in real-time.

20

11. The method as recited in claim 9, wherein the secondary subscriber identifier indicates whether the messages to be transmitted between the backup cable modem termination system and the associated cable modem include voice data or video data.

25

12. The method as recited in claim 9, wherein the subscriber information further comprises quality of service requirements.

30

13. The method as recited in claim 9, wherein at least a portion of the subscriber information further comprises a scheduling type.

14. The method as recited in claim 13, wherein the scheduling type indicates a type of real-time traffic to be transmitted.

5 15. The method as recited in claim 13, wherein the secondary subscriber identifier indicates that the messages to be transmitted between the backup cable modem termination system and the associated cable modem are to be transmitted in real-time.

10 16. The method as recited in claim 13, wherein the scheduling type indicates whether the messages to be transmitted between the backup cable modem termination system and the associated cable modem include voice data or video data.

15 17. The method as recited in claim 13, wherein the scheduling type is UGS or UGS-AD.

18. The method as recited in claim 13, wherein prioritizing the subscriber information comprises:

searching the subscriber information for a secondary subscriber identifier;  
and

20 prioritizing the subscriber information having a secondary subscriber identifier such that the subscriber information has a higher priority than the subscriber information not having a secondary subscriber identifier;

25 19. The method as recited in claim 18, wherein prioritizing the subscriber information having a secondary subscriber identifier further comprises:

prioritizing the subscriber information having a secondary subscriber identifier according to the scheduling type.

30 20. The method as recited in claim 19, wherein prioritizing the subscriber information having a secondary identifier further comprises:

prioritizing the subscriber information having a secondary identifier according to time of receipt of the subscriber information from the active cable modem termination system.

5           21.     The method as recited in claim 18, wherein prioritizing the subscriber information not having a secondary identifier comprises:

prioritizing the subscriber information not having a secondary identifier according to time of receipt of the subscriber information from the active cable modem termination system.

10

22.     The method as recited in claim 1, wherein prioritizing the subscriber information comprises:

storing the subscriber information and a time of receipt of the subscriber information by the backup cable modem termination system such that the subscriber information is associated with the time of receipt.

15

23.     The method as recited in claim 22, wherein the stored subscriber information is stored in order of the time of receipt.

20           24.     The method as recited in claim 1, further comprising:

storing the subscriber information and a time of receipt of the subscriber information by the backup cable modem termination system such that the subscriber information is associated with the time of receipt.

25           25.     The method as recited in claim 24, further comprising:

prioritizing the subscriber information according to the time of receipt.

26.     The method as recited in claim 1, further comprising:

after receiving the subscriber information, sending an acknowledgement of the subscriber information to the active cable modem termination system.

27. The method as recited in claim 1, further comprising:

5 repeatedly receiving subscriber information associated with one or more cable modems from an active cable modem termination system prior to prioritizing the subscriber information.

28. The method as recited in claim 27, further comprising:

10 wherein receiving subscriber information comprises receiving subscriber information associated with one or more cable modems from a first active cable modem termination system and receiving subscriber information associated with one or more cable modems from a second active cable modem termination system; and

15 wherein prioritizing the subscriber information comprises prioritizing the subscriber information associated with the first active cable modem termination system is performed separately from prioritizing the subscriber information associated with the second active cable modem termination system.

29. The method as recited in claim 28, further comprising:

20 storing the prioritized subscriber information associated with the first active cable modem termination system separately from the prioritized subscriber information associated with the second active cable modem termination system.

30. The method as recited in claim 1, further comprising:

25 receiving an indication that an active cable modem termination system has failed;

determining an identity of the failed active cable modem termination system; and

wherein receiving subscriber information associated with one or more cable modems from the active cable modem termination system comprises obtaining the subscriber information associated with the failed active cable modem termination system.

5

31. The method as recited in claim 1, further comprising:

receiving an indication that a call initiated by one of the cable modems has been terminated; and

removing subscriber information associated with the one of the cable modems from memory associated with a previously failed active cable modem termination system.

10

32. A computer-readable medium storing thereon computer-readable instructions for performing a method in a backup cable modem termination system of restoring transmission of messages between one or more cable modems and the backup cable modem termination system upon failure of an active cable modem termination system, comprising:

15

instructions for receiving subscriber information associated with one or more cable modems from an active cable modem termination system, the subscriber information including one or more subscriber identifiers;

20

instructions for prioritizing the subscriber information, the prioritized subscriber information indicating an order in which the transmission of messages between the one or more cable modems and the backup cable modem is to be restored; and

25

instructions for polling the cable modems in the order indicated by the prioritized subscriber information.

33. A backup cable modem termination system adapted for restoring transmission of messages between one or more cable modems and the backup cable



modem termination system upon failure of an active cable modem termination system, comprising:

means for receiving subscriber information associated with one or more cable modems from an active cable modem termination system, the subscriber information  
5 including one or more subscriber identifiers;

means for prioritizing the subscriber information, the prioritized subscriber information indicating an order in which the transmission of messages between the one or more cable modems and the backup cable modem is to be restored; and

means for polling the cable modems in the order indicated by the prioritized  
10 subscriber information.

34. A backup cable modem termination system adapted for restoring transmission of messages between one or more cable modems and the backup cable modem termination system upon failure of an active cable modem termination  
15 system, comprising:

a processor; and

a memory, at least one of the processor and the memory being adapted for:

receiving subscriber information associated with one or more cable modems from an active cable modem termination system, the subscriber information  
20 including one or more subscriber identifiers;

prioritizing the subscriber information, the prioritized subscriber information indicating an order in which the transmission of messages between the one or more cable modems and the backup cable modem is to be restored; and

polling the cable modems in the order indicated by the prioritized subscriber  
25 information.




Fig. 1

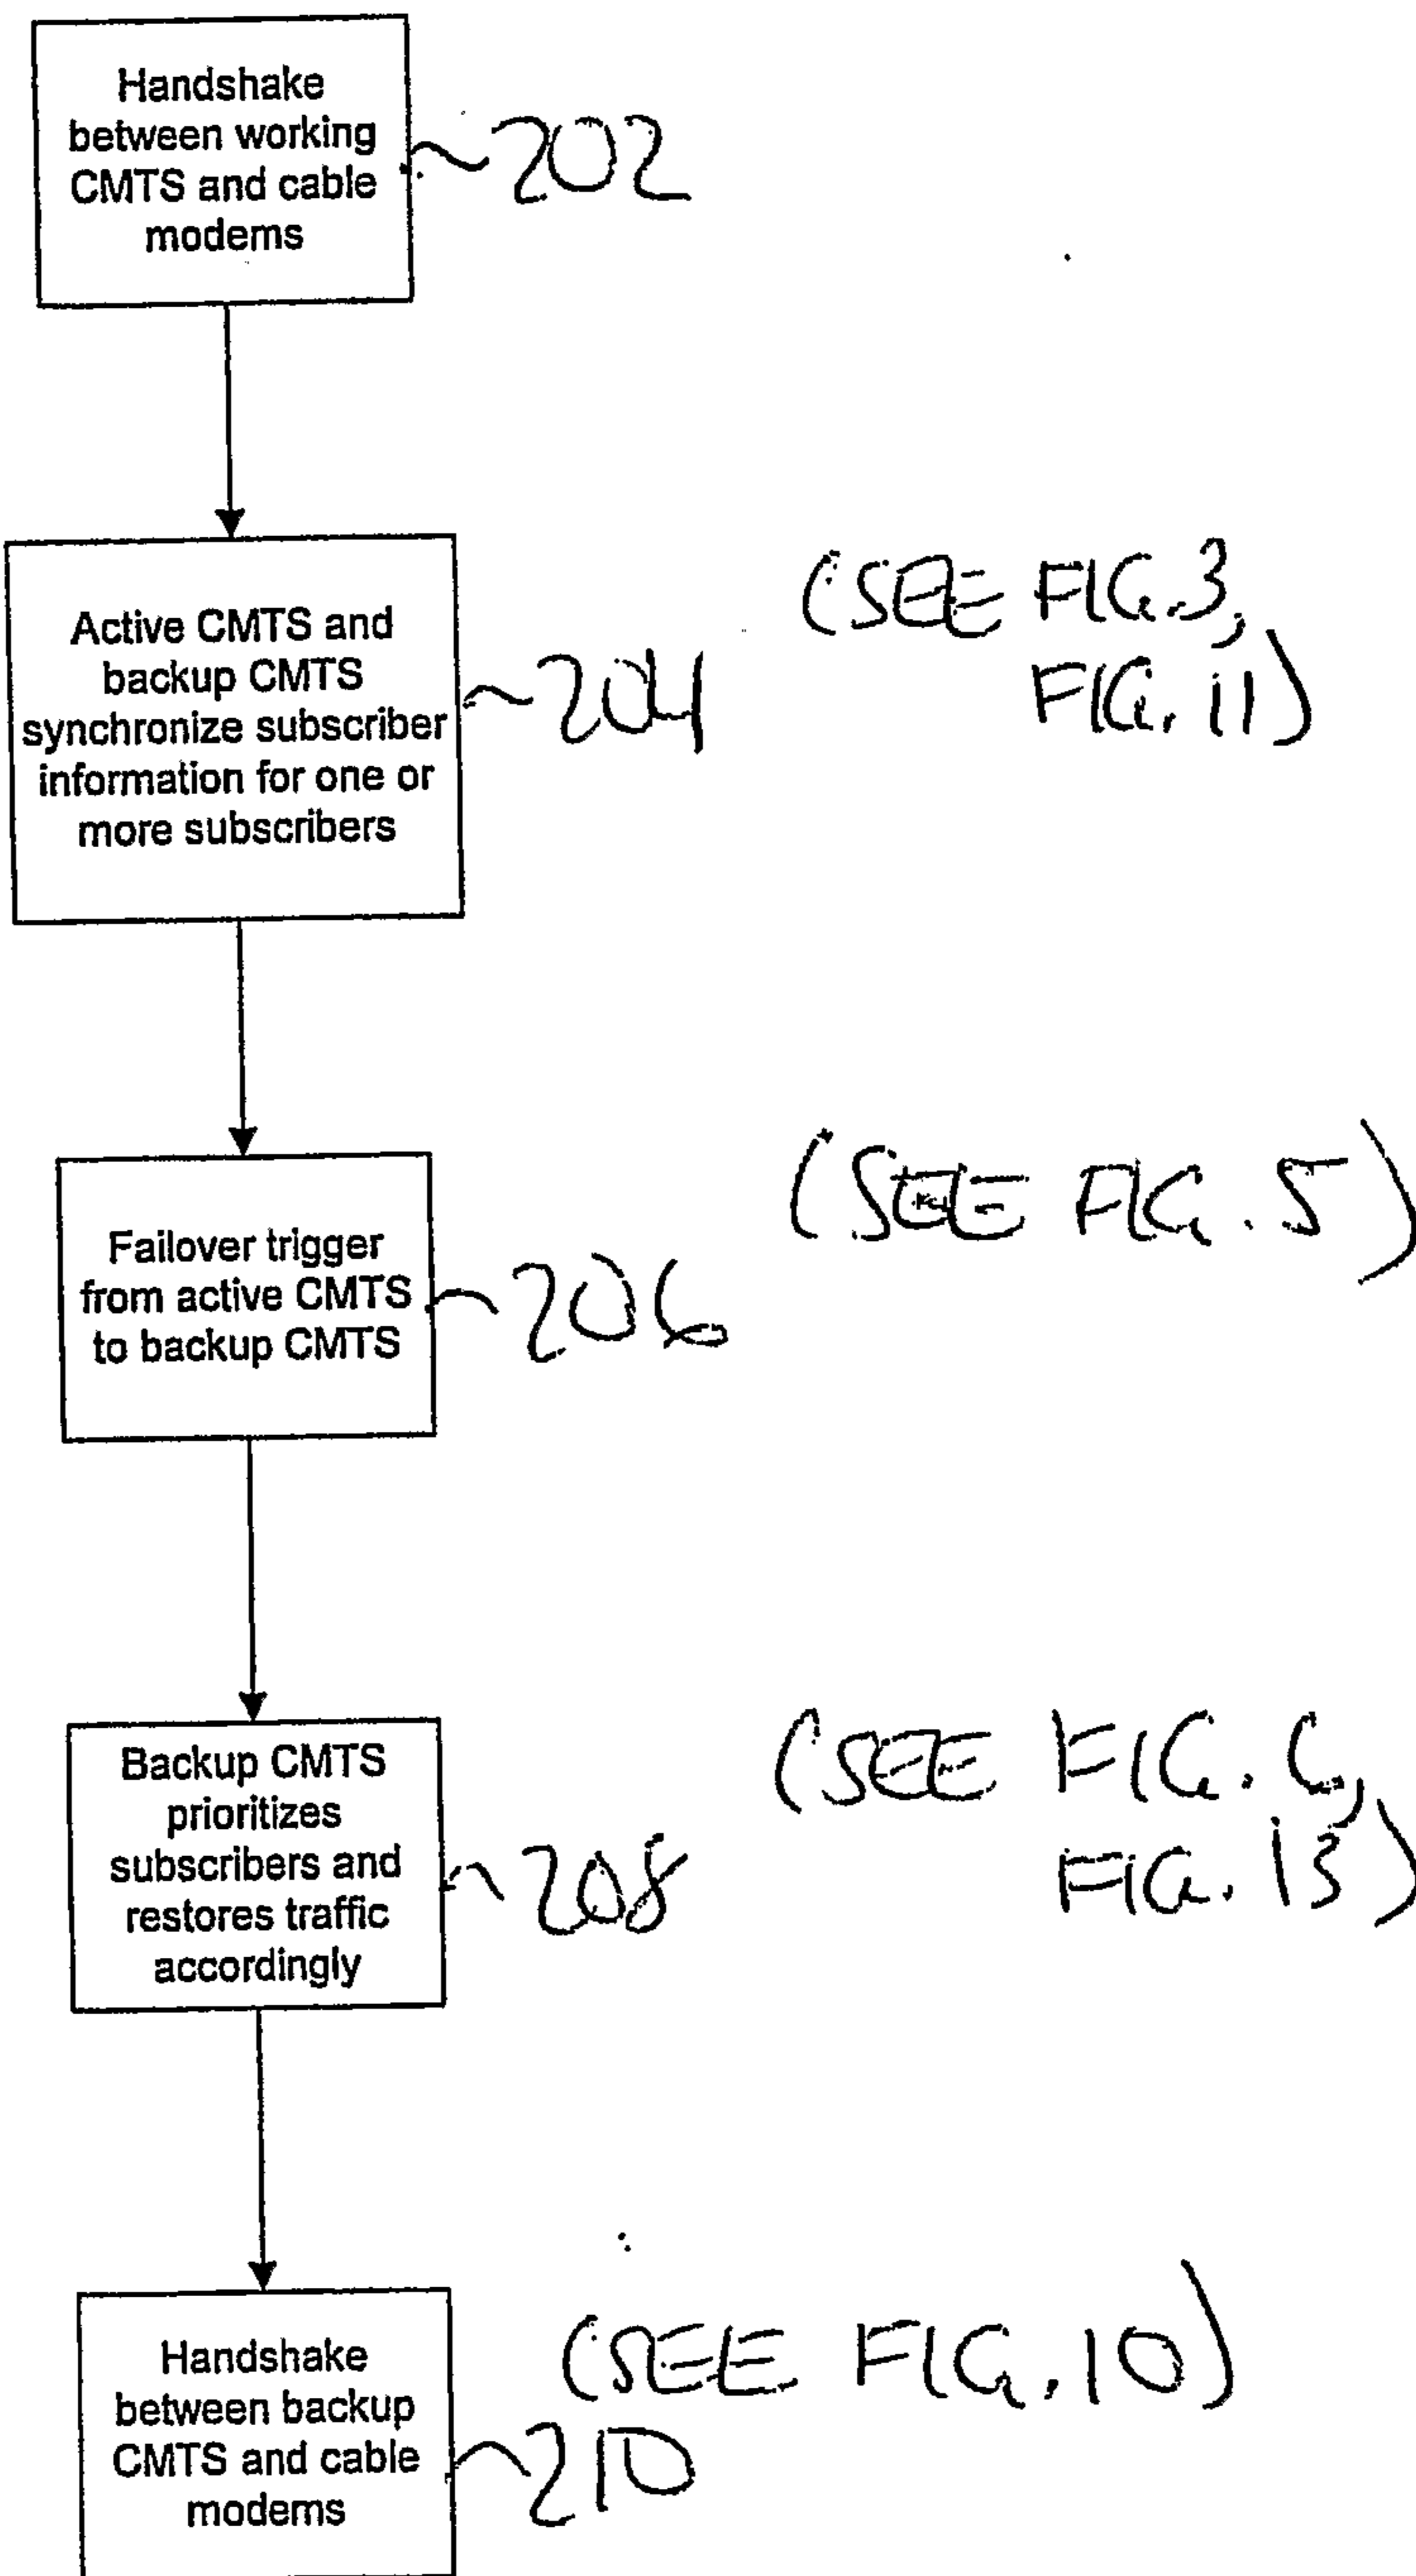


FIG. 2

204

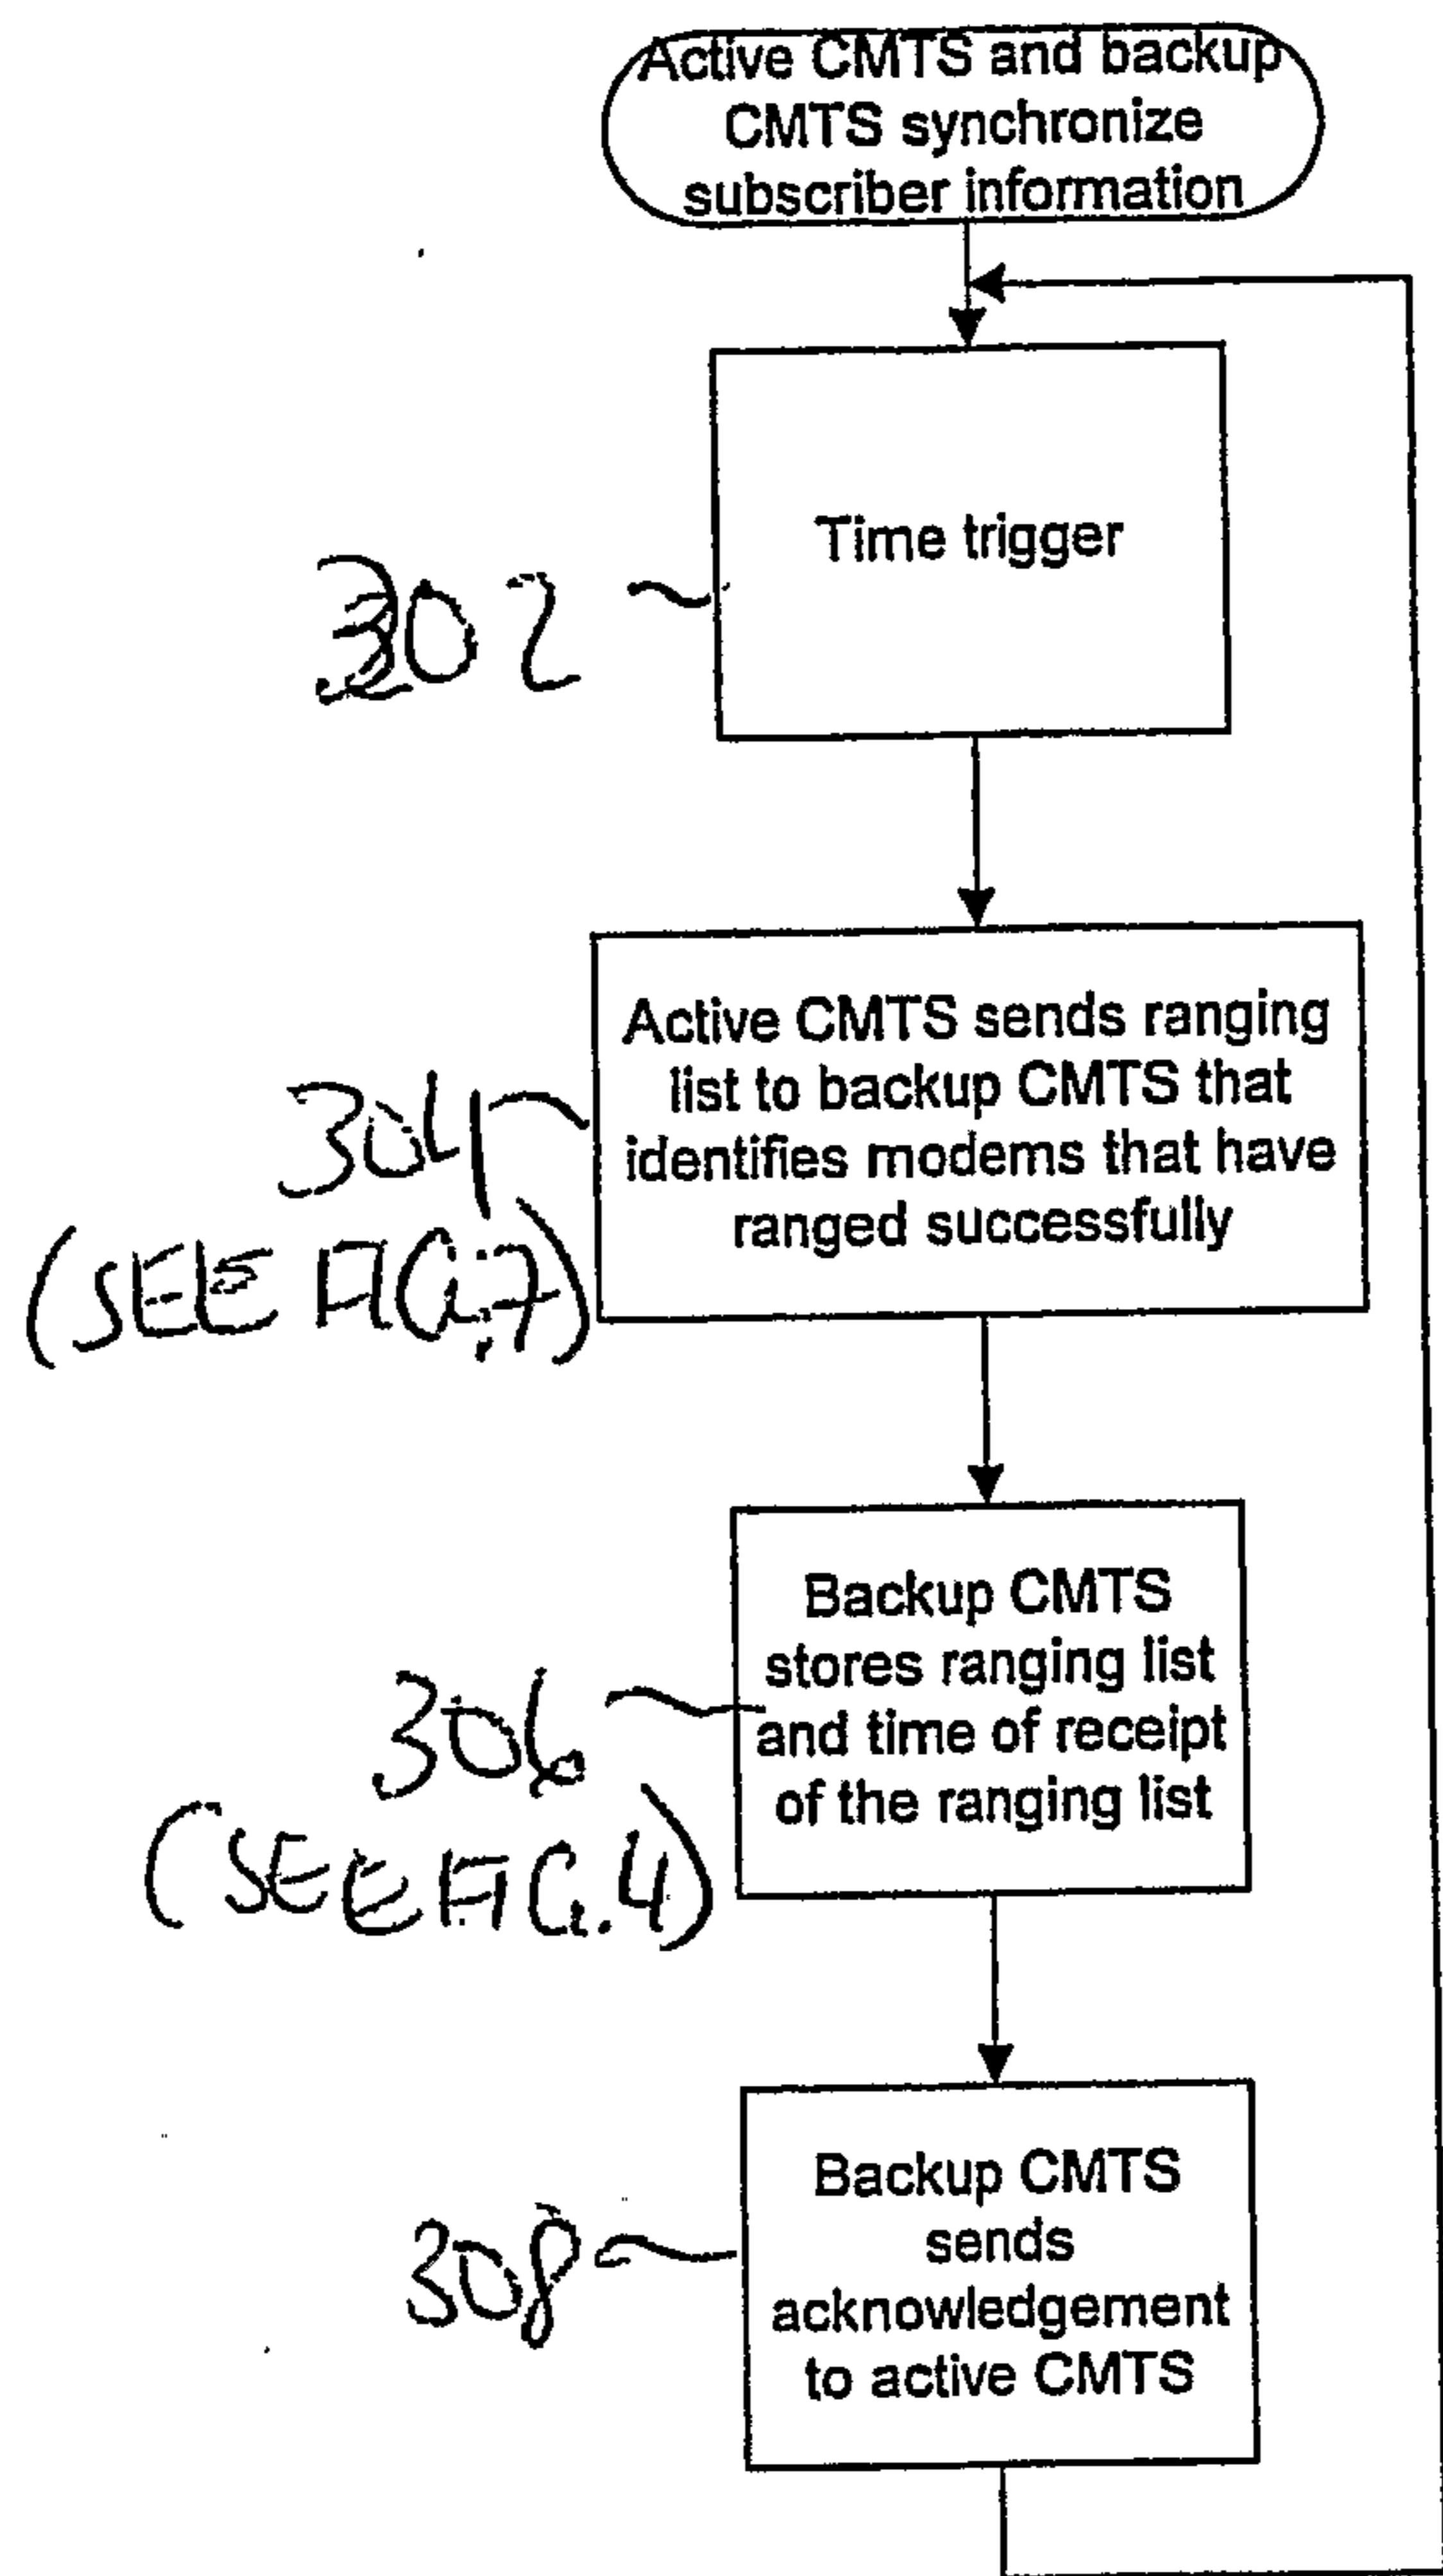


FIG. 3

4/17

404      406

Modem ID	MAC address
ID1	MAC1
ID2	MAC2
.	.
.	.
IDN	MACN
Time of arrival t1	

402

408

412      414

Modem ID	MAC address
ID1	MAC1
ID2	MAC2
.	.
.	.
IDN	MACN
Time of arrival t2	

410

416

FIG. 4

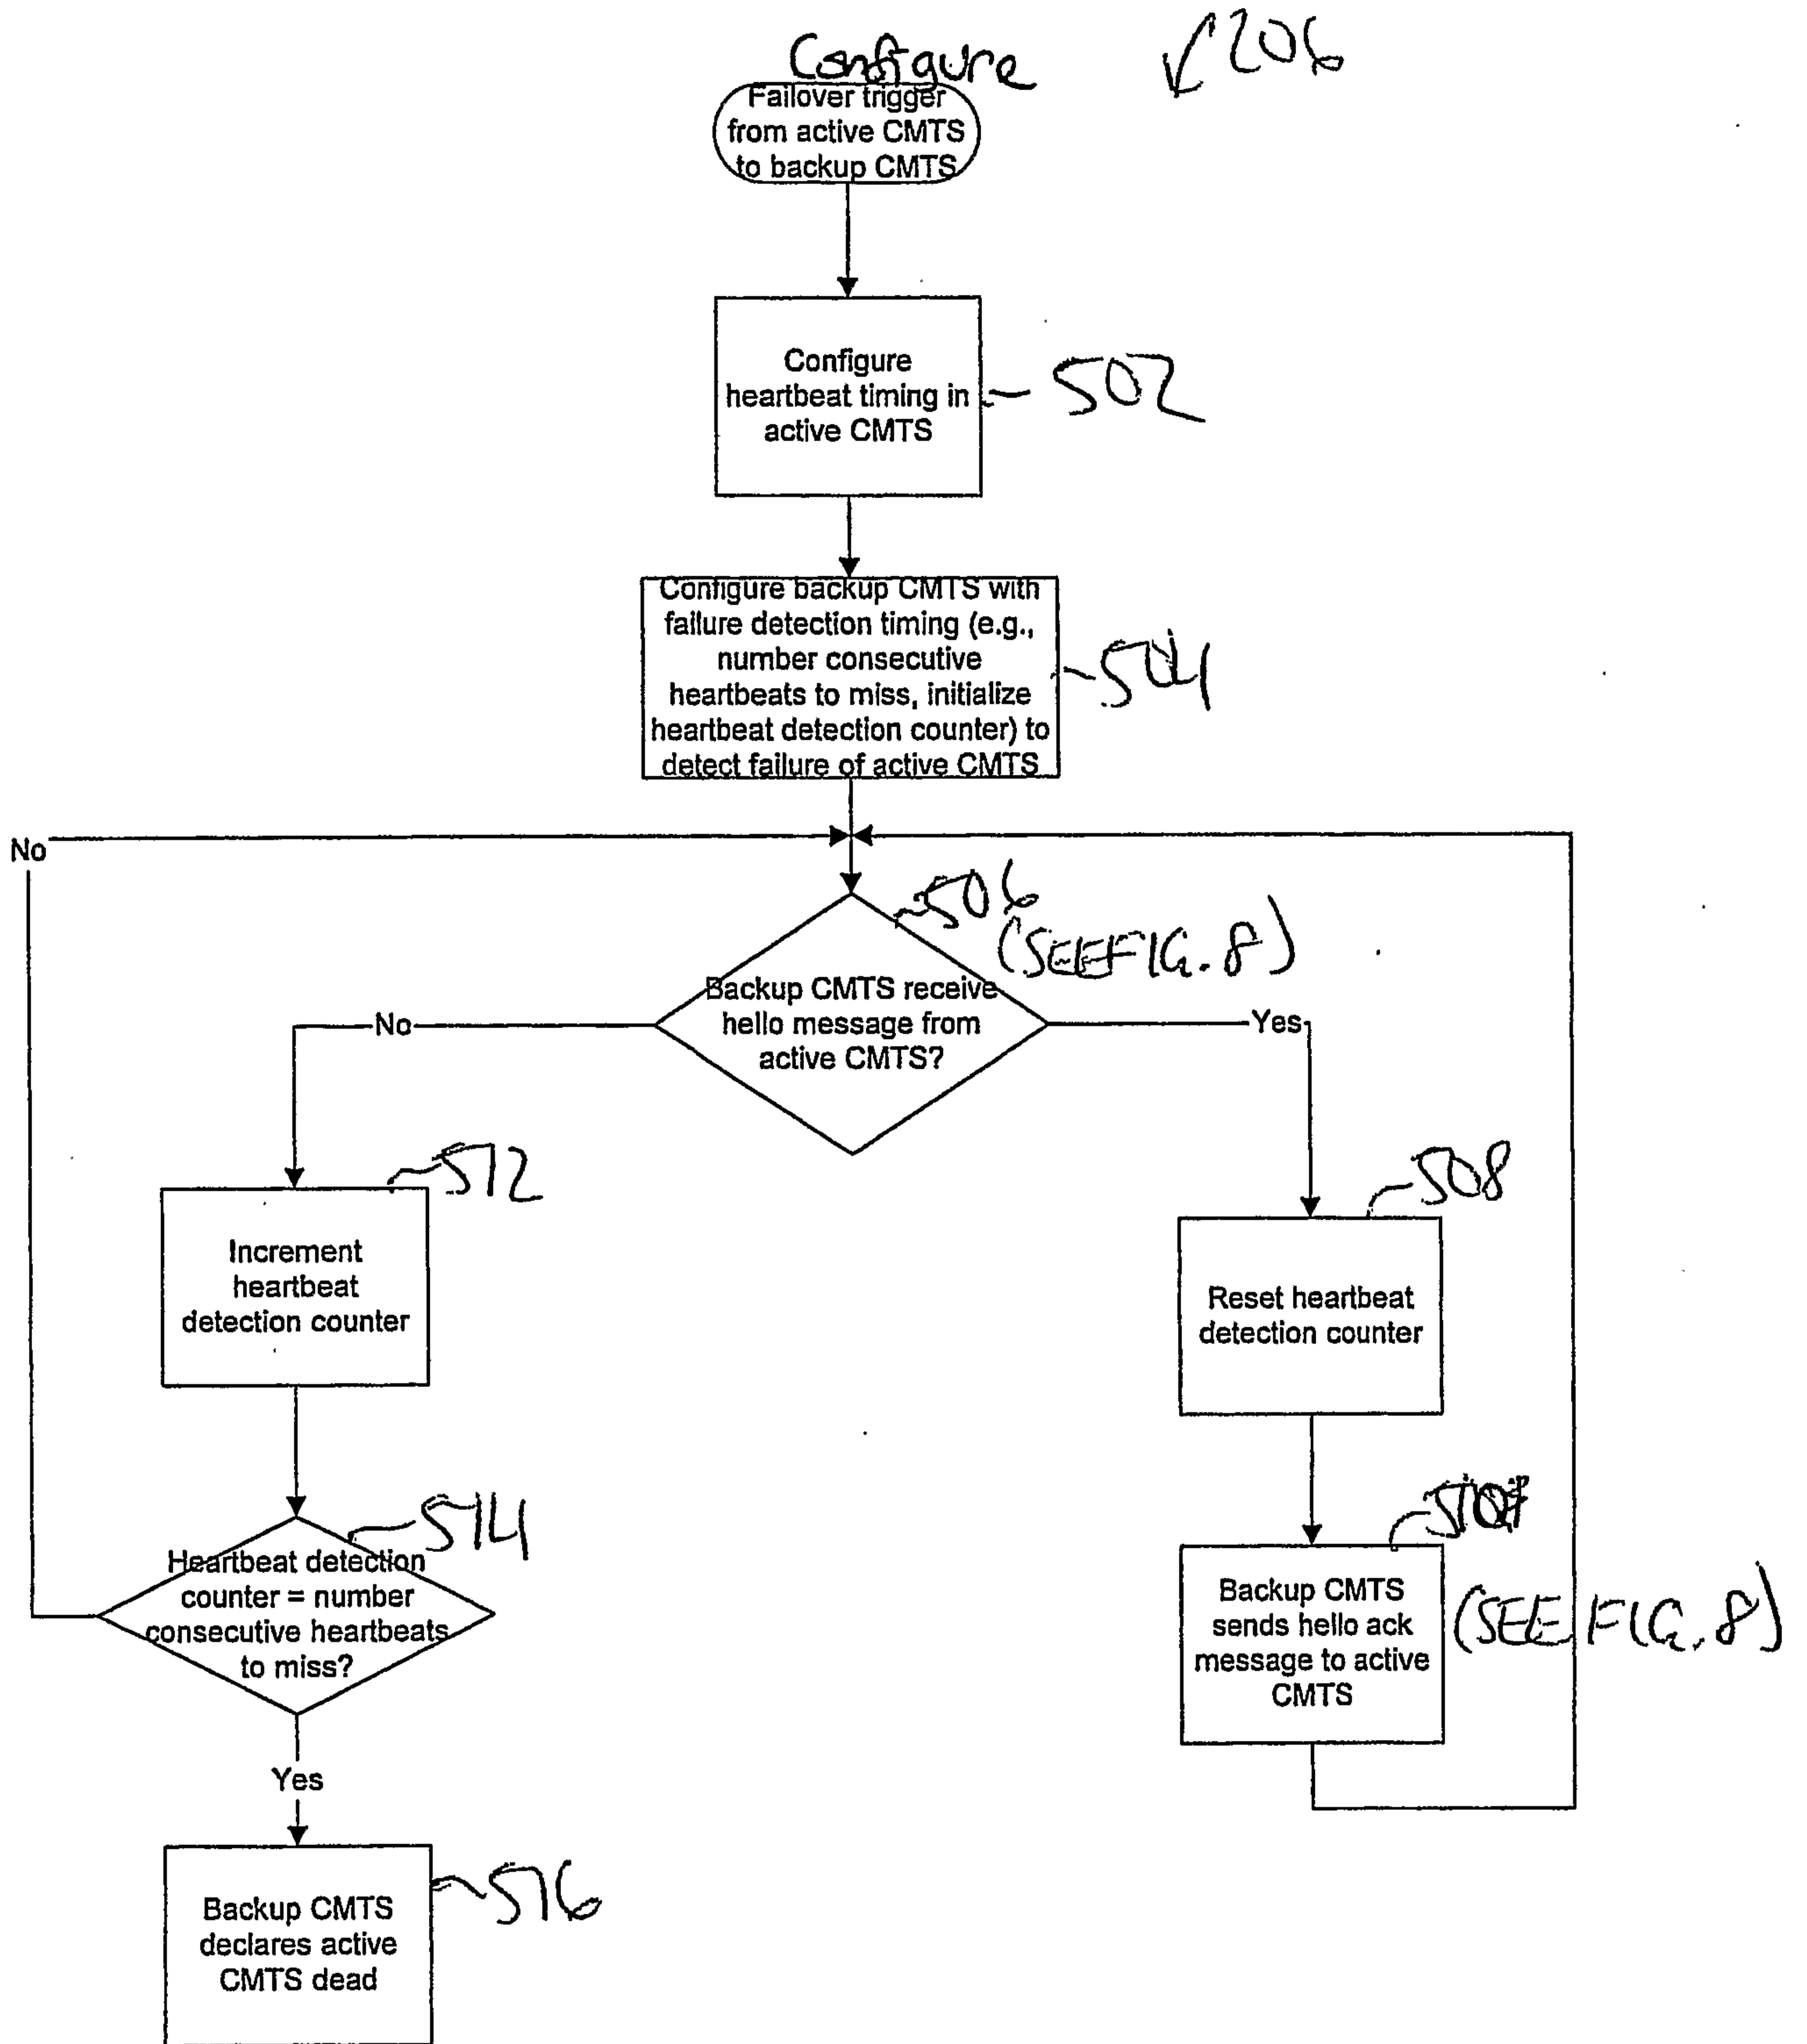


FIG. 5

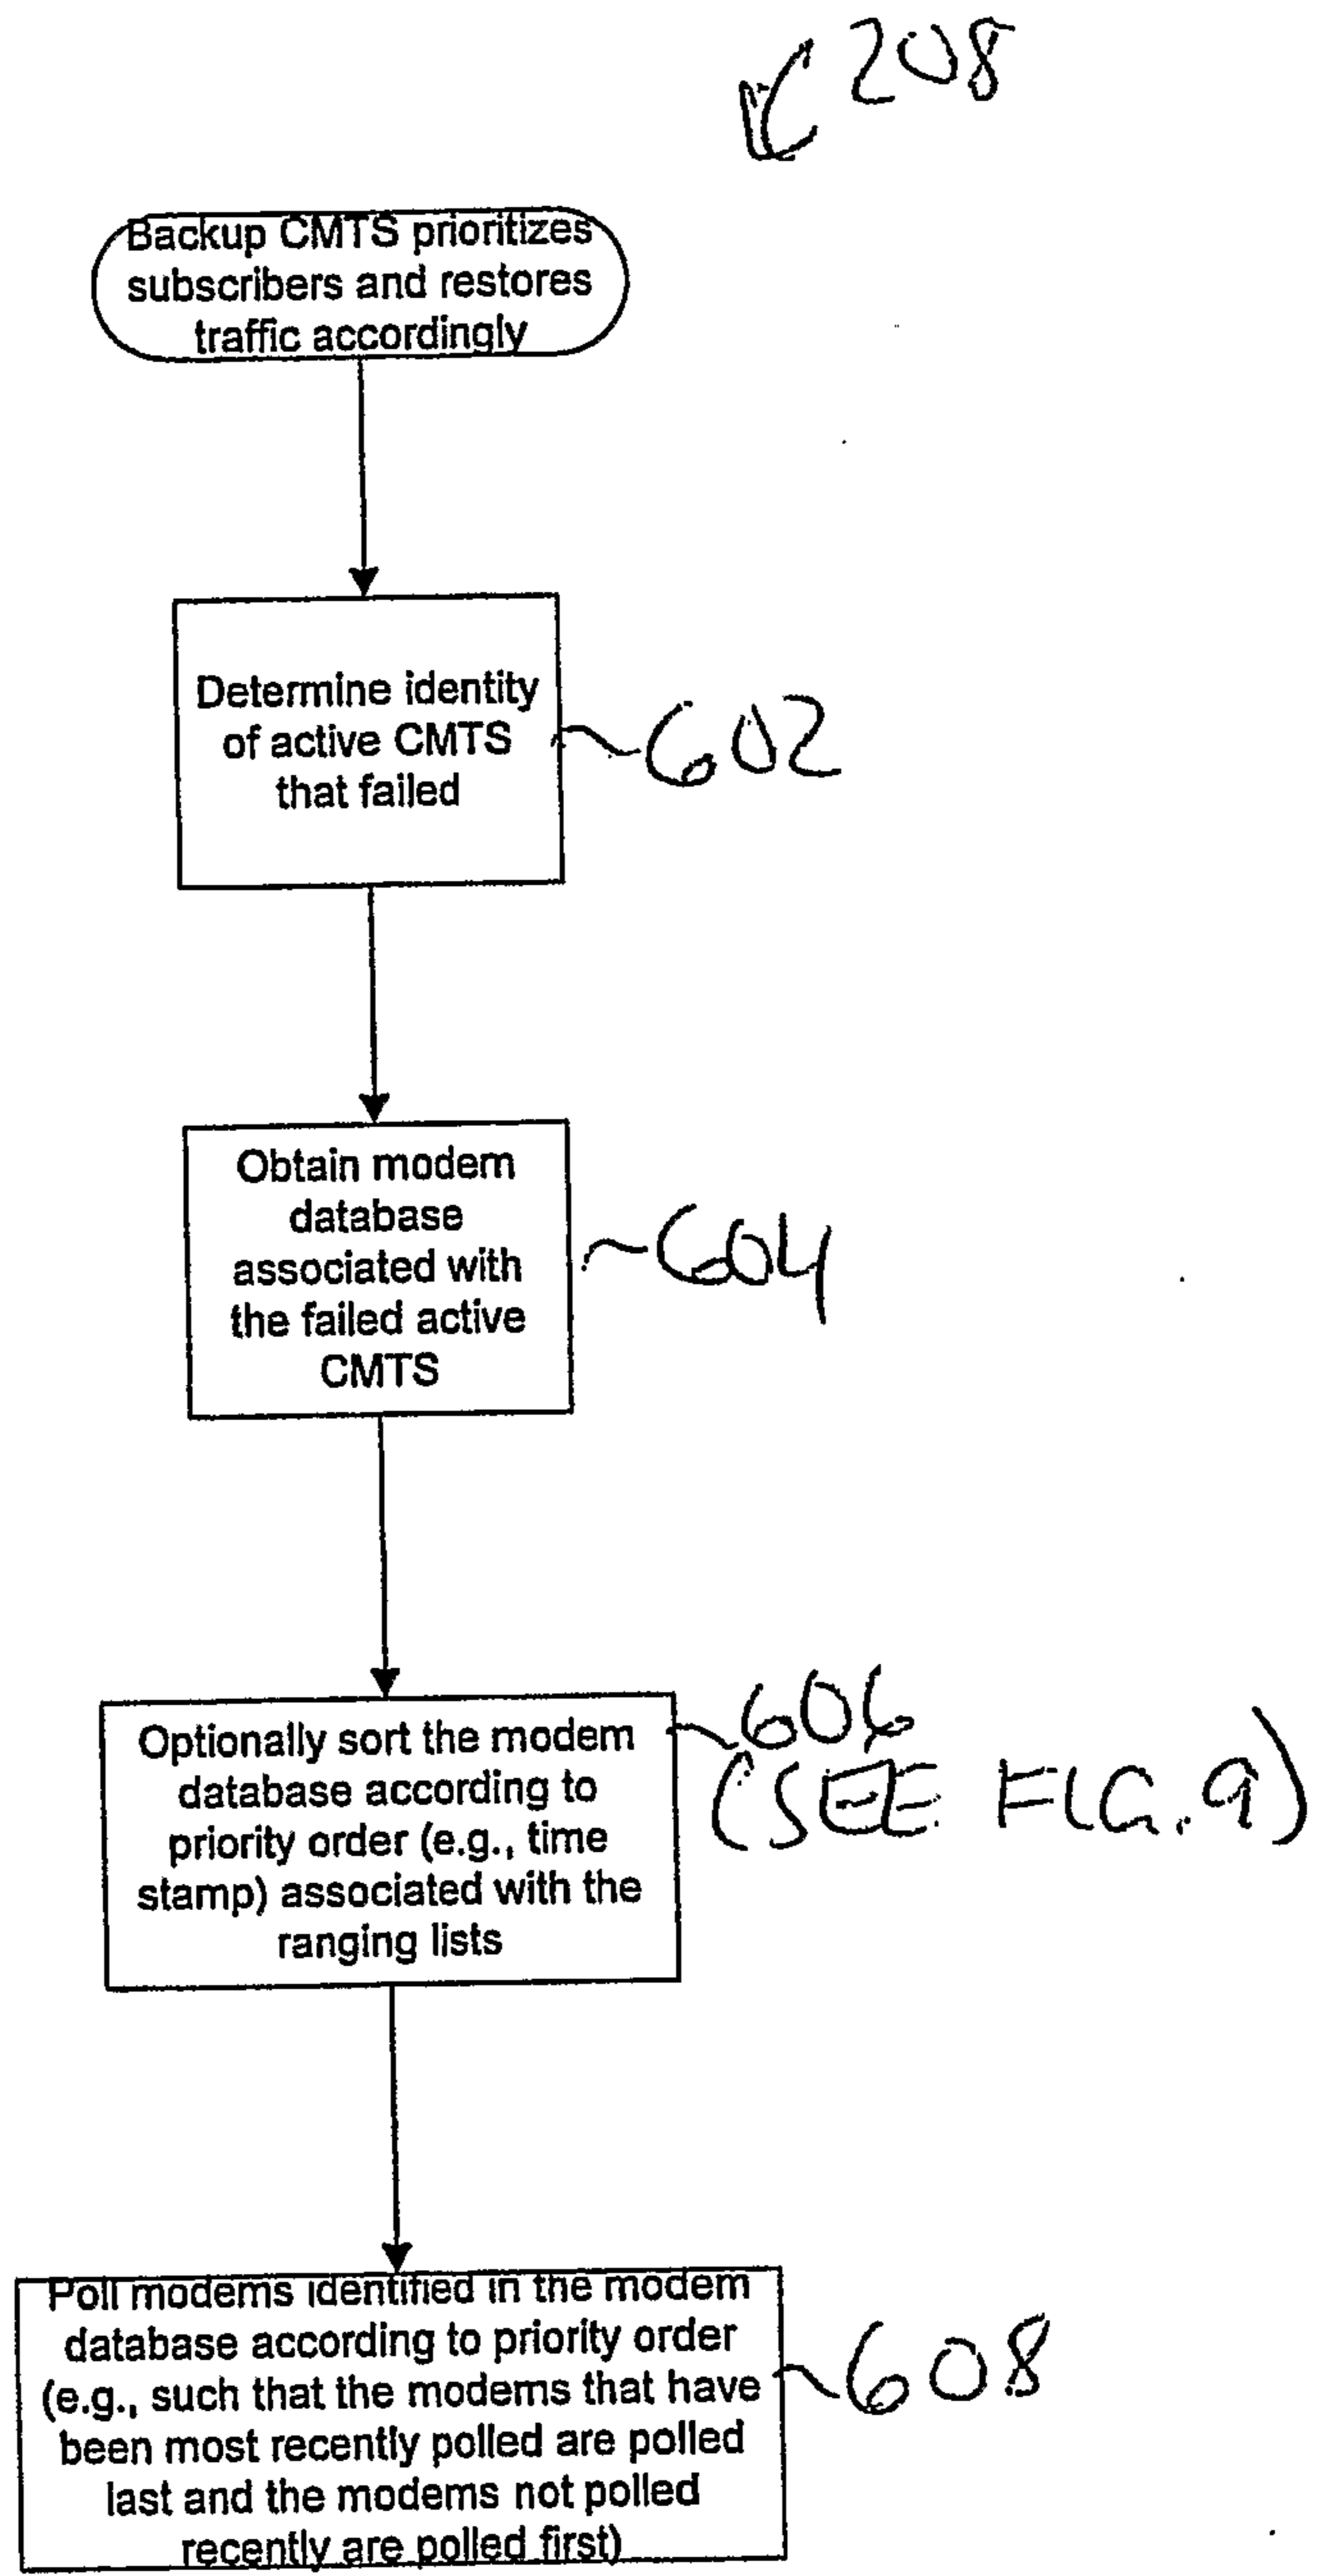


FIG. 6




FIG. 7




FIG. 8






FIG. 9

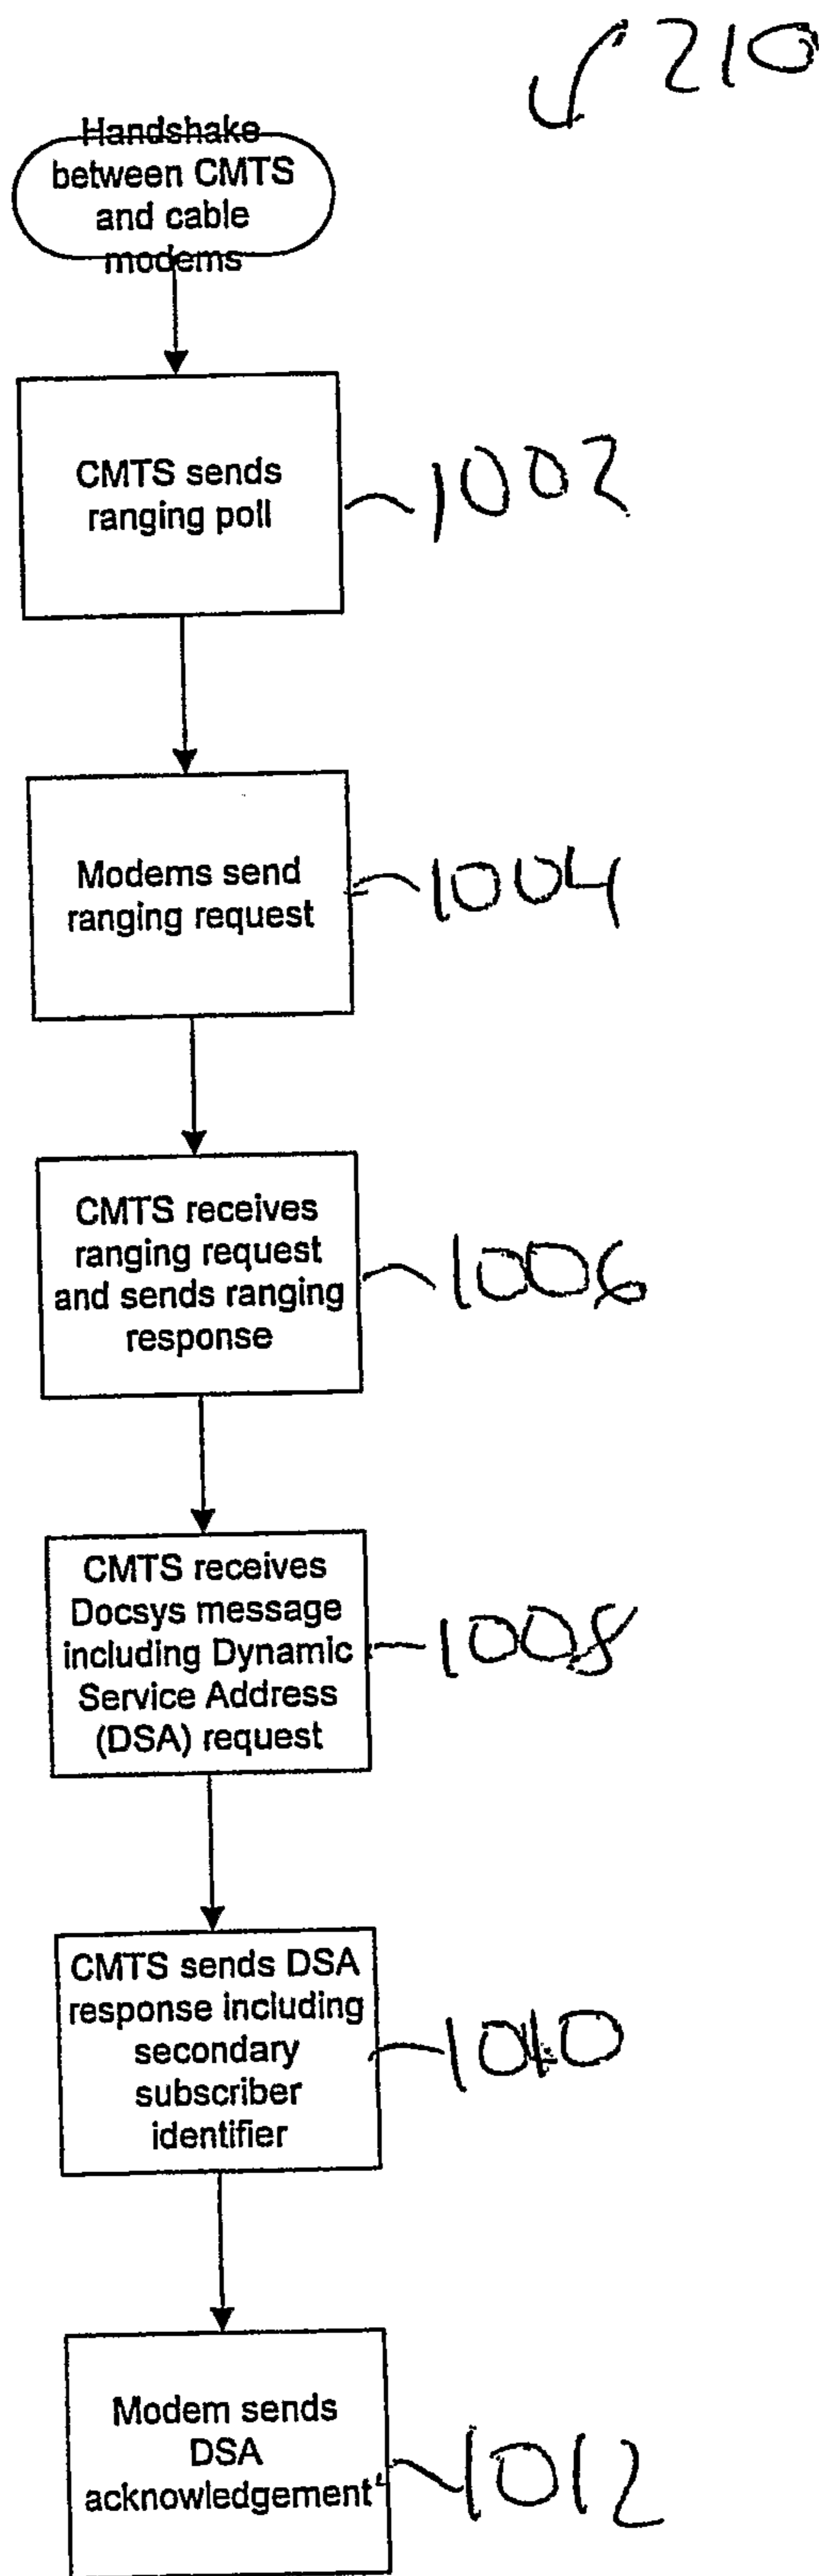


FIG. 10

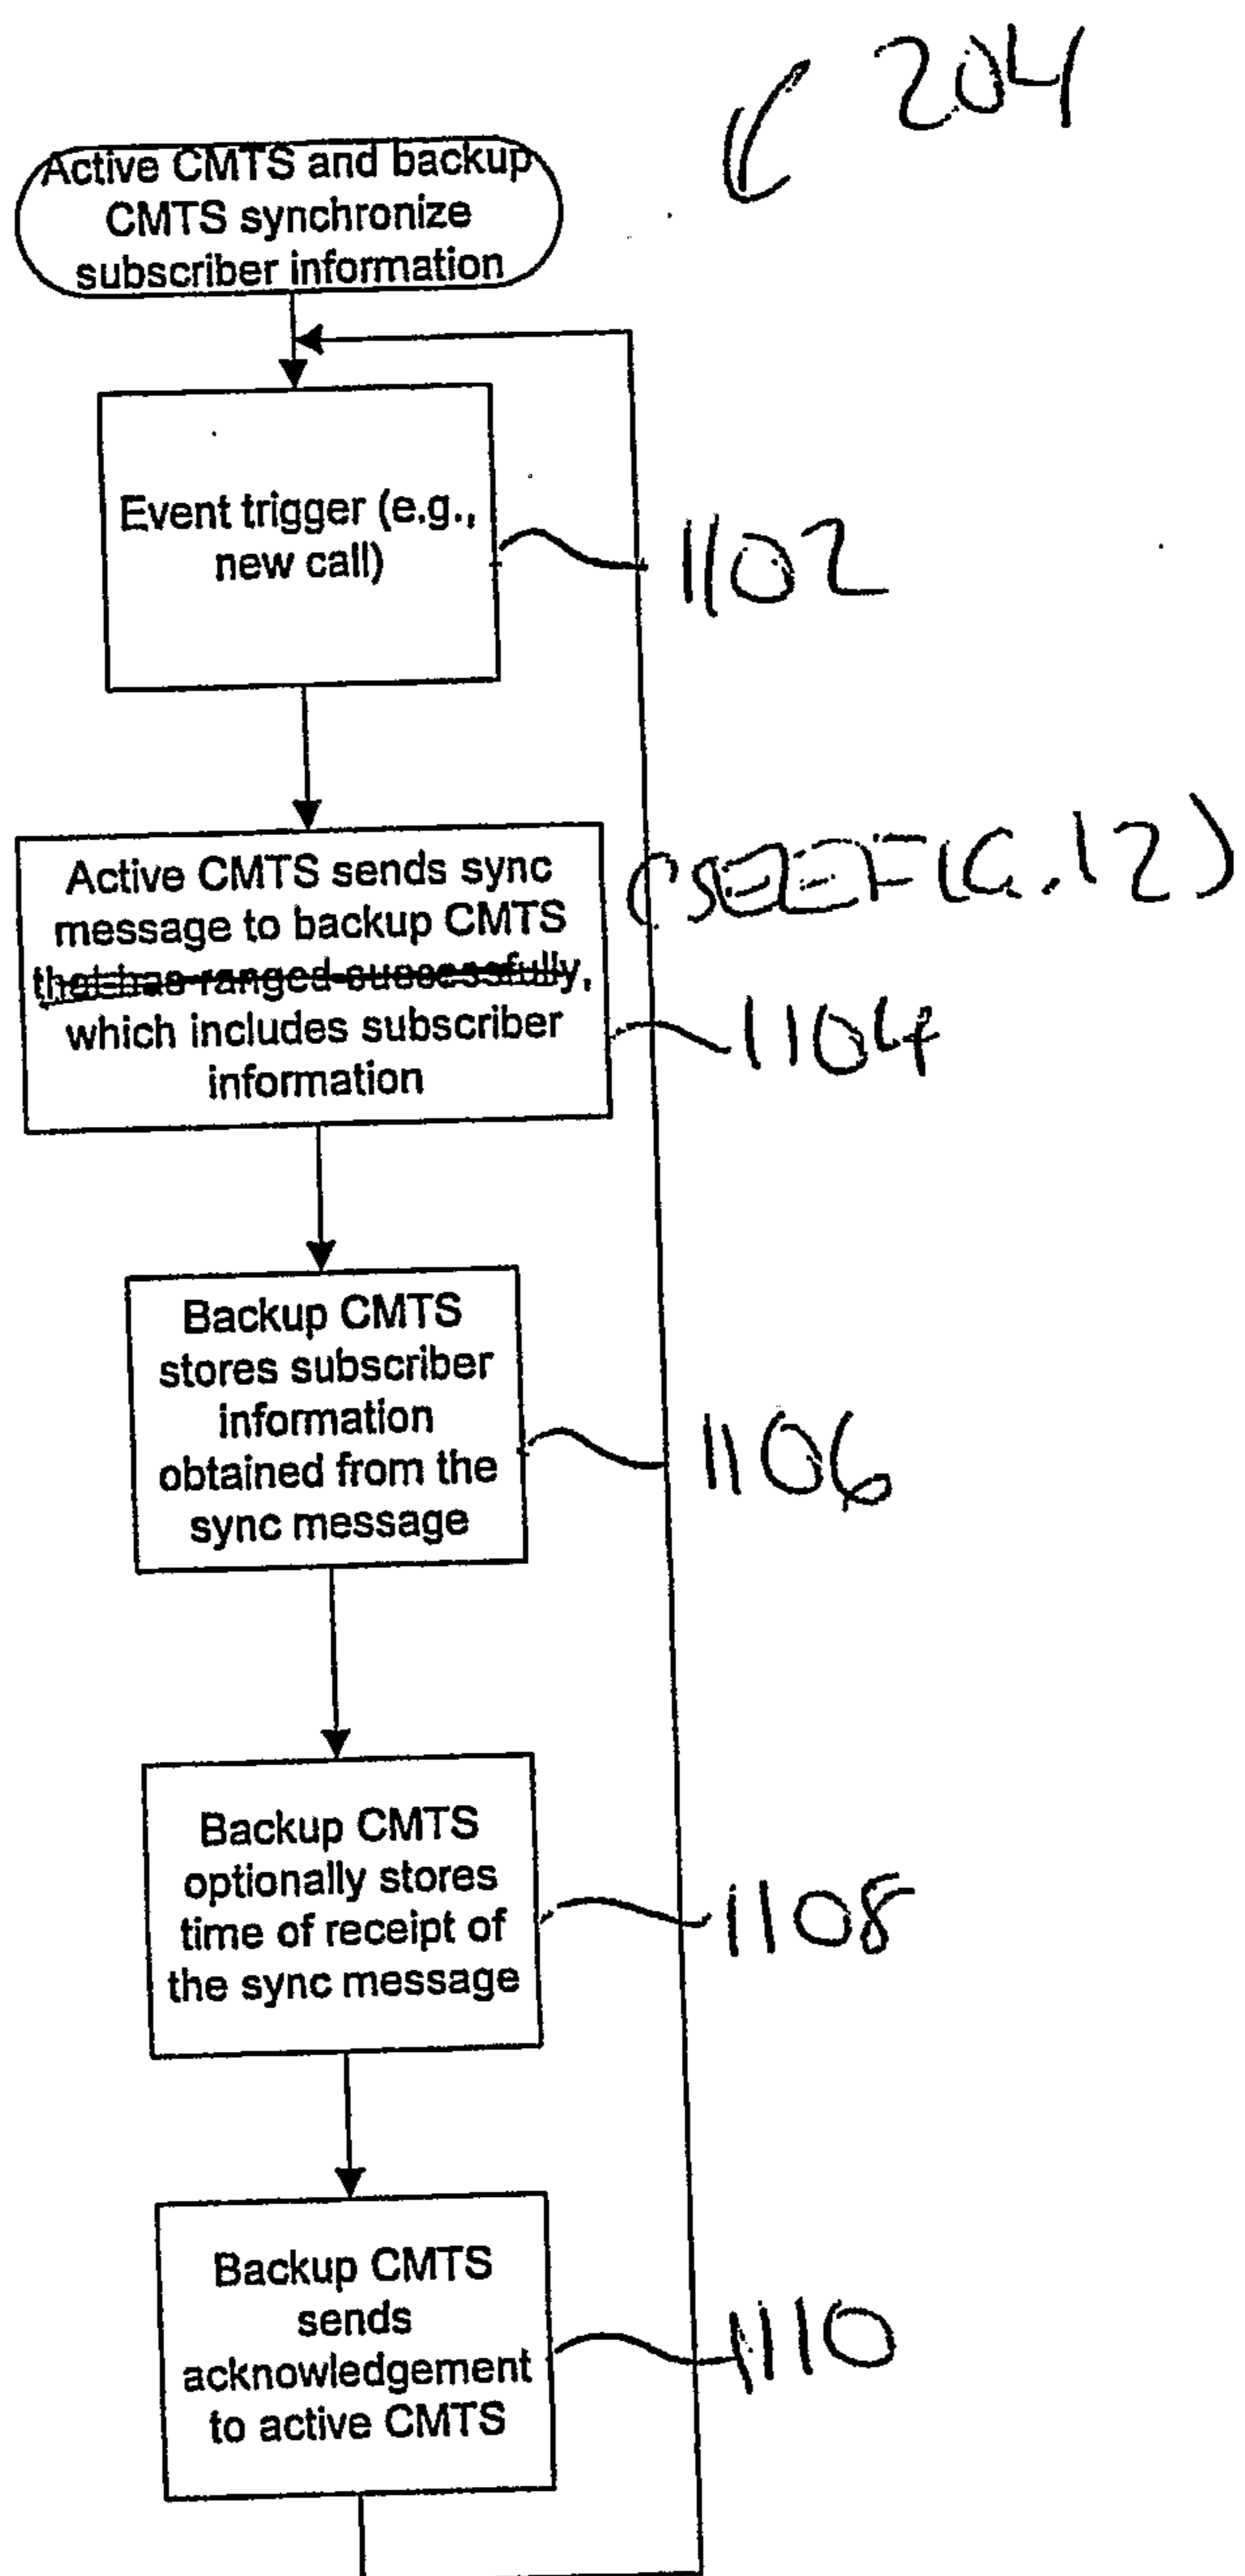


FIG. 11

11/17

1200

Sync message

Primary subscriber ID	1202	Secondary subscriber ID	1204
Service template including QOS requirements (e.g., scheduling type)	1206		
Packet classifier	1208		

FIG. 12

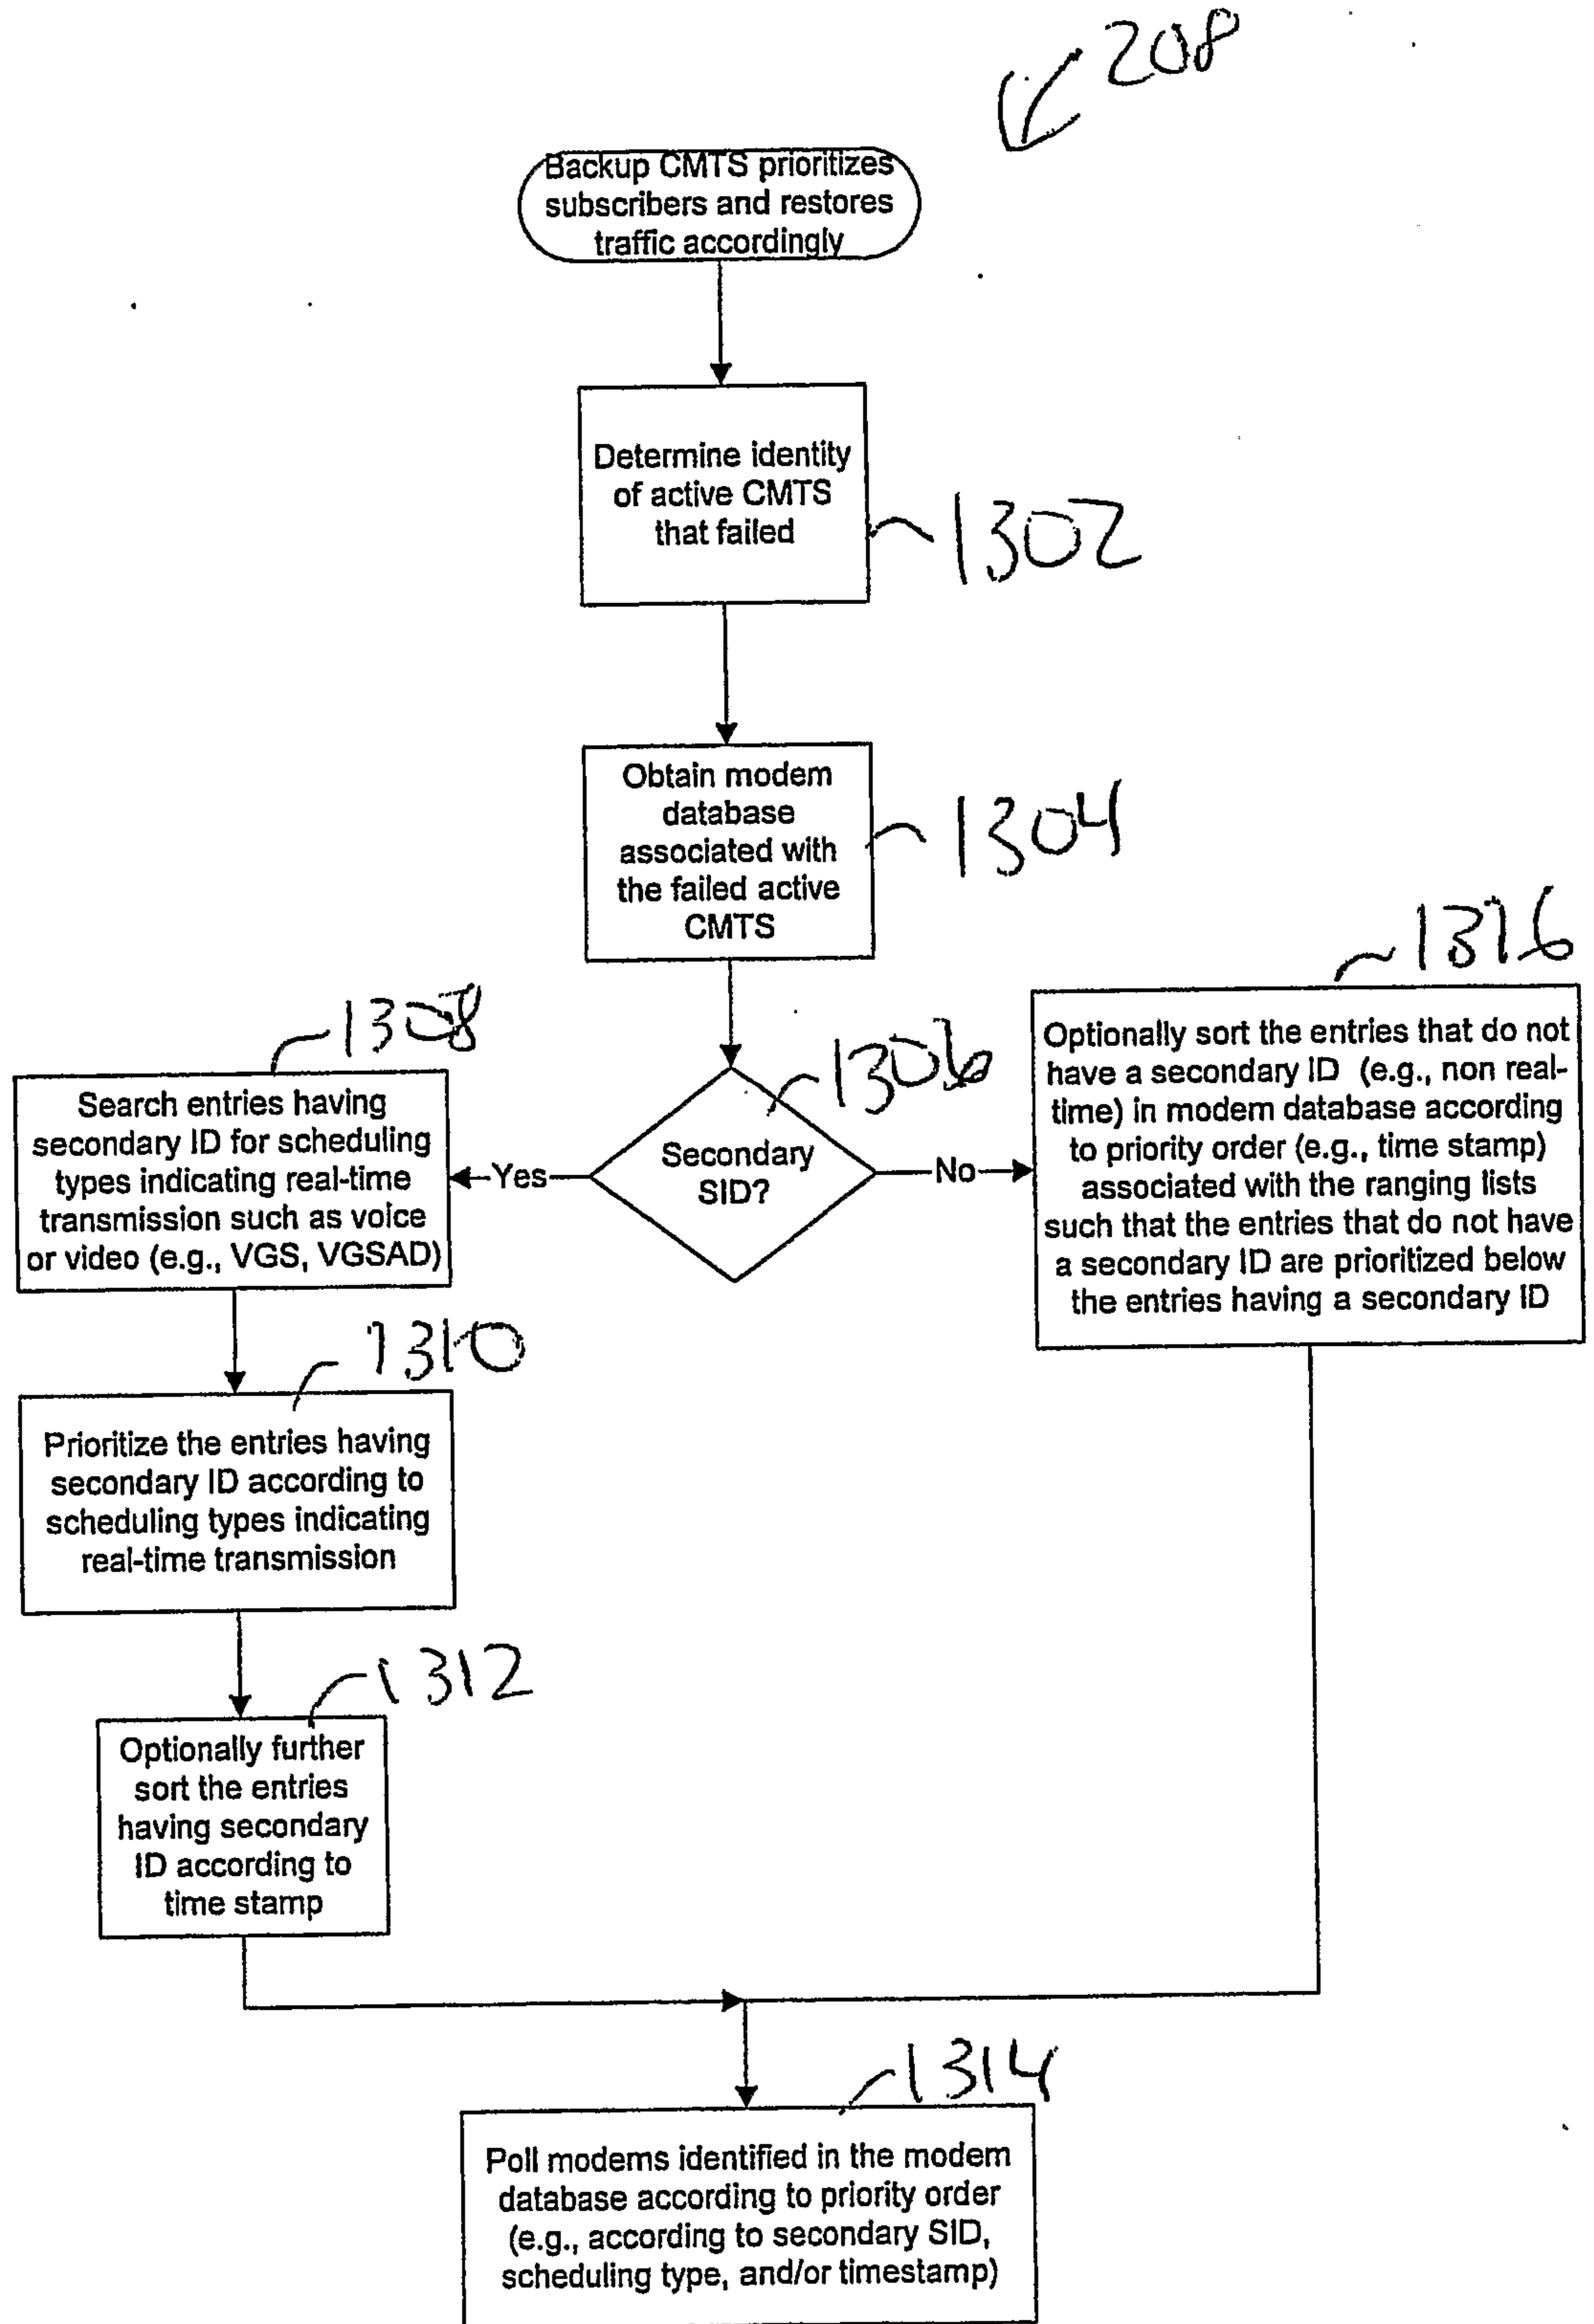


FIG. 13

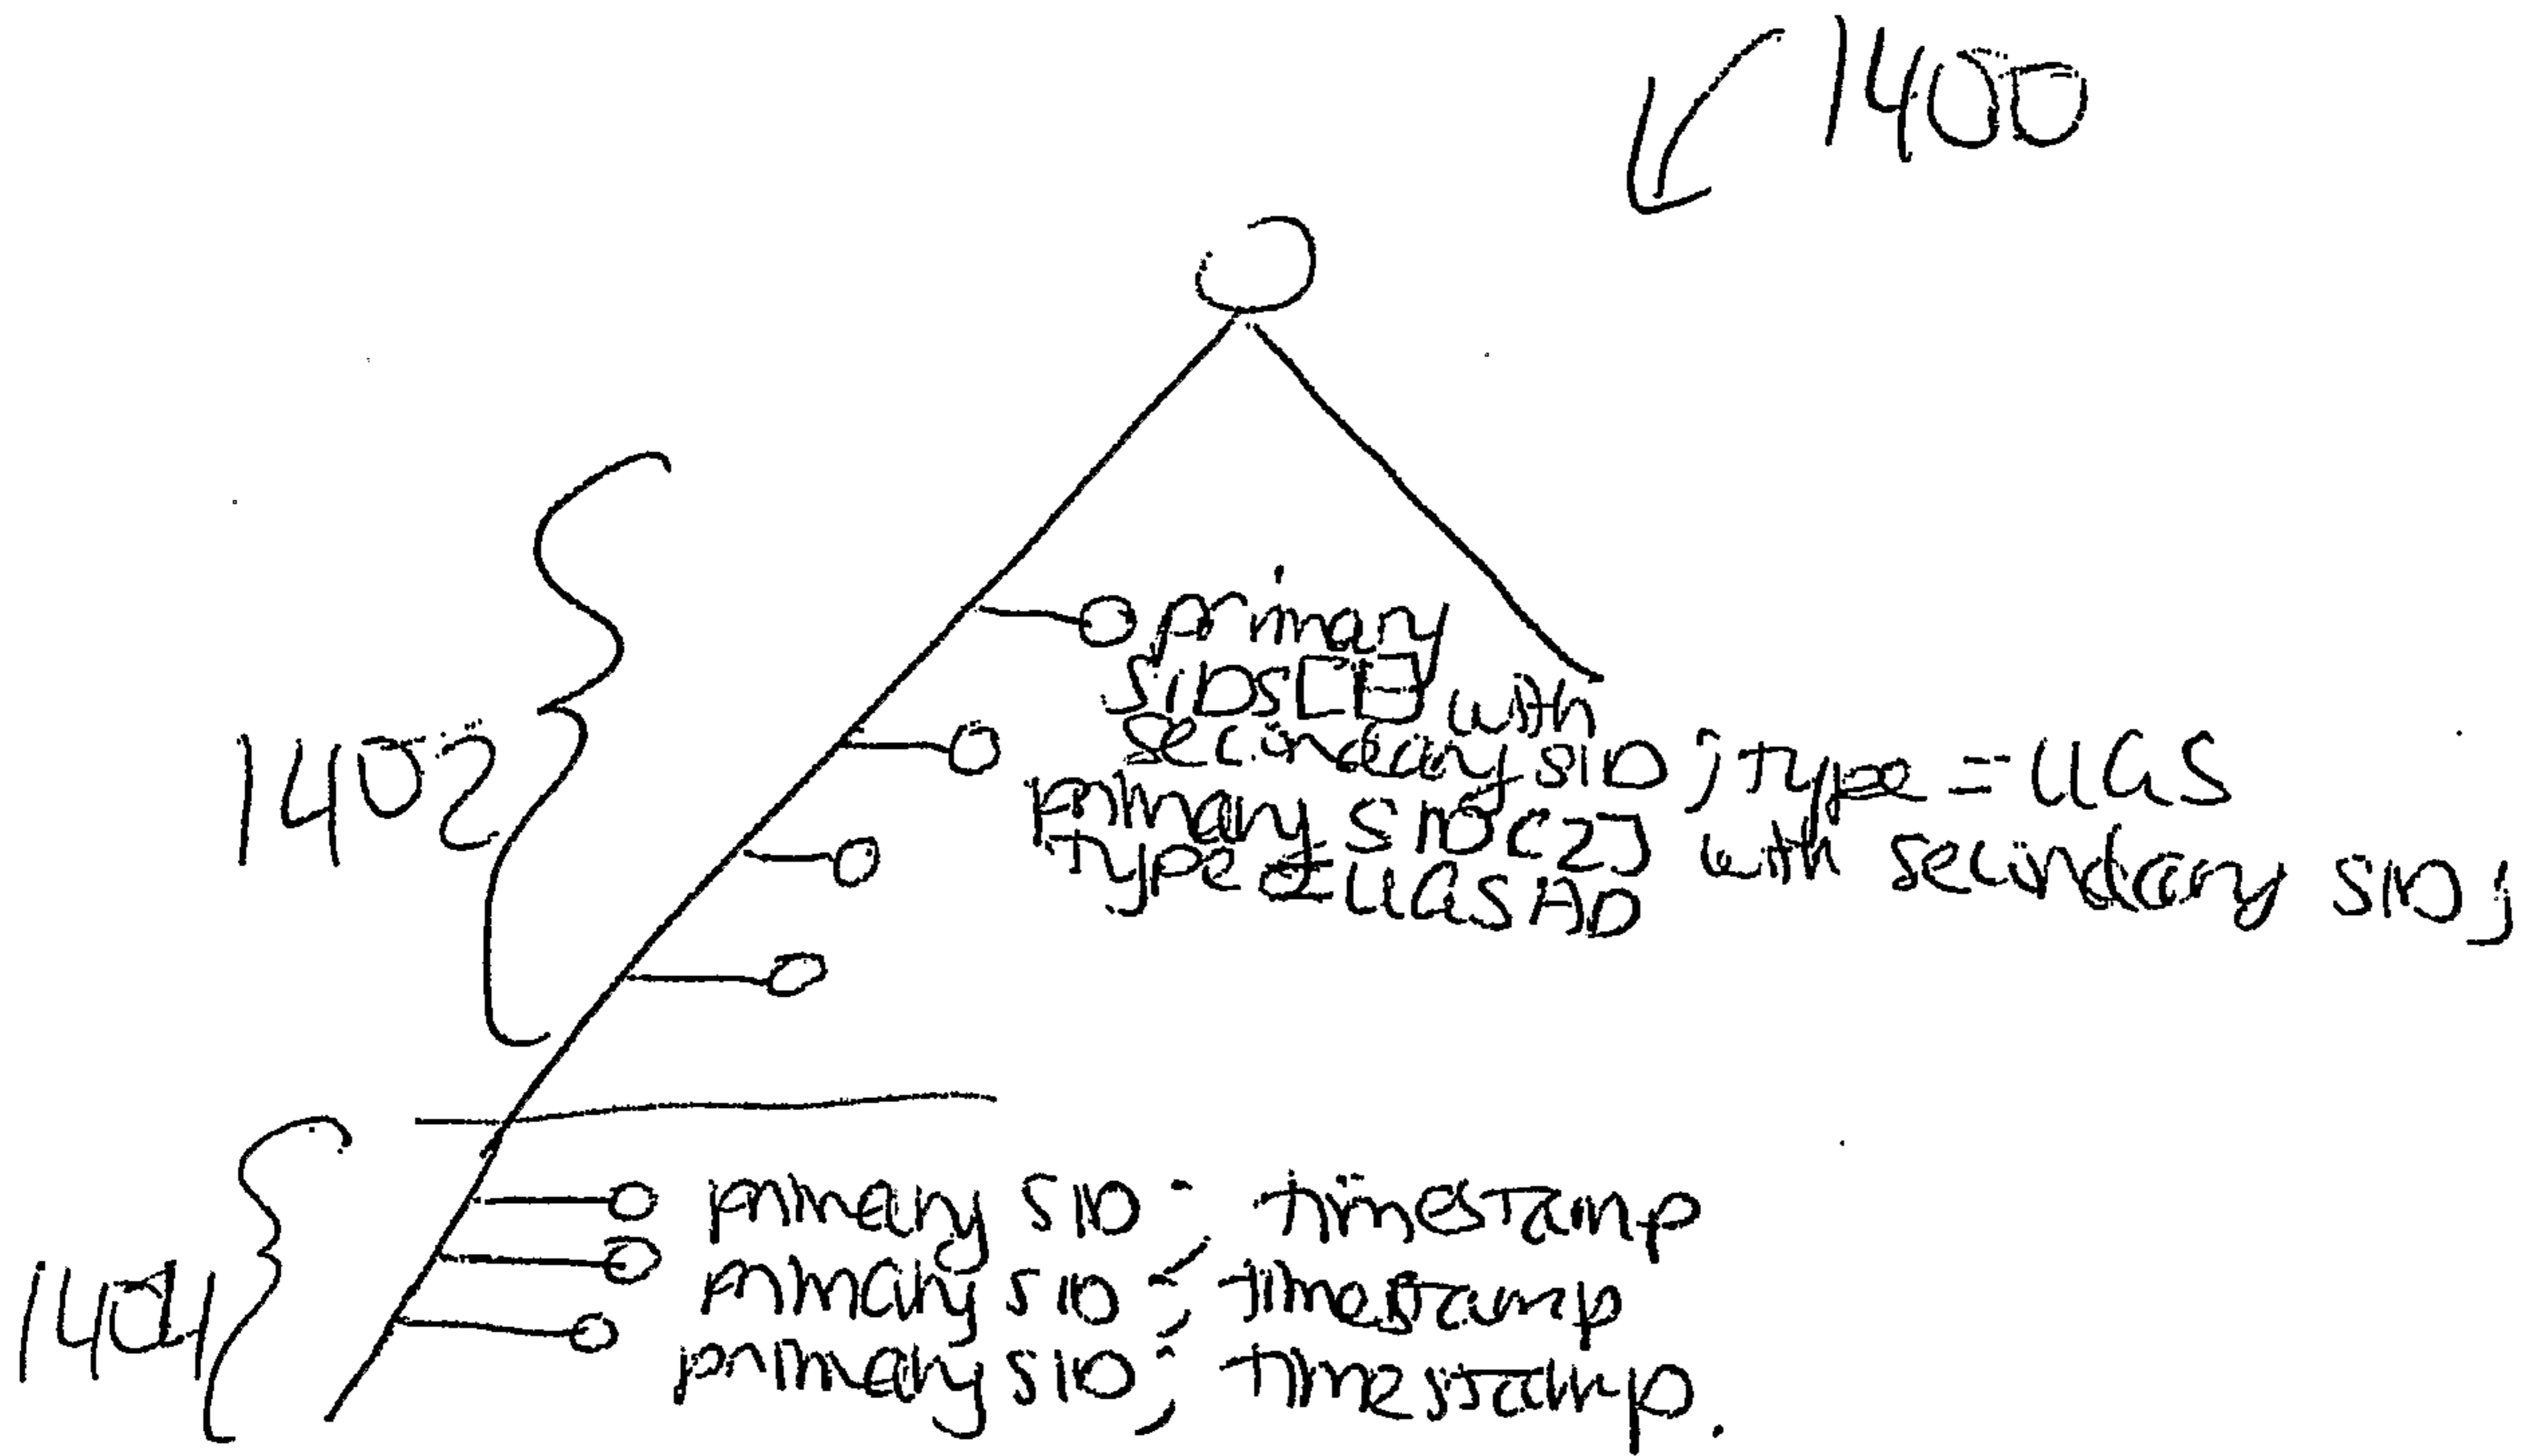


FIG. 14

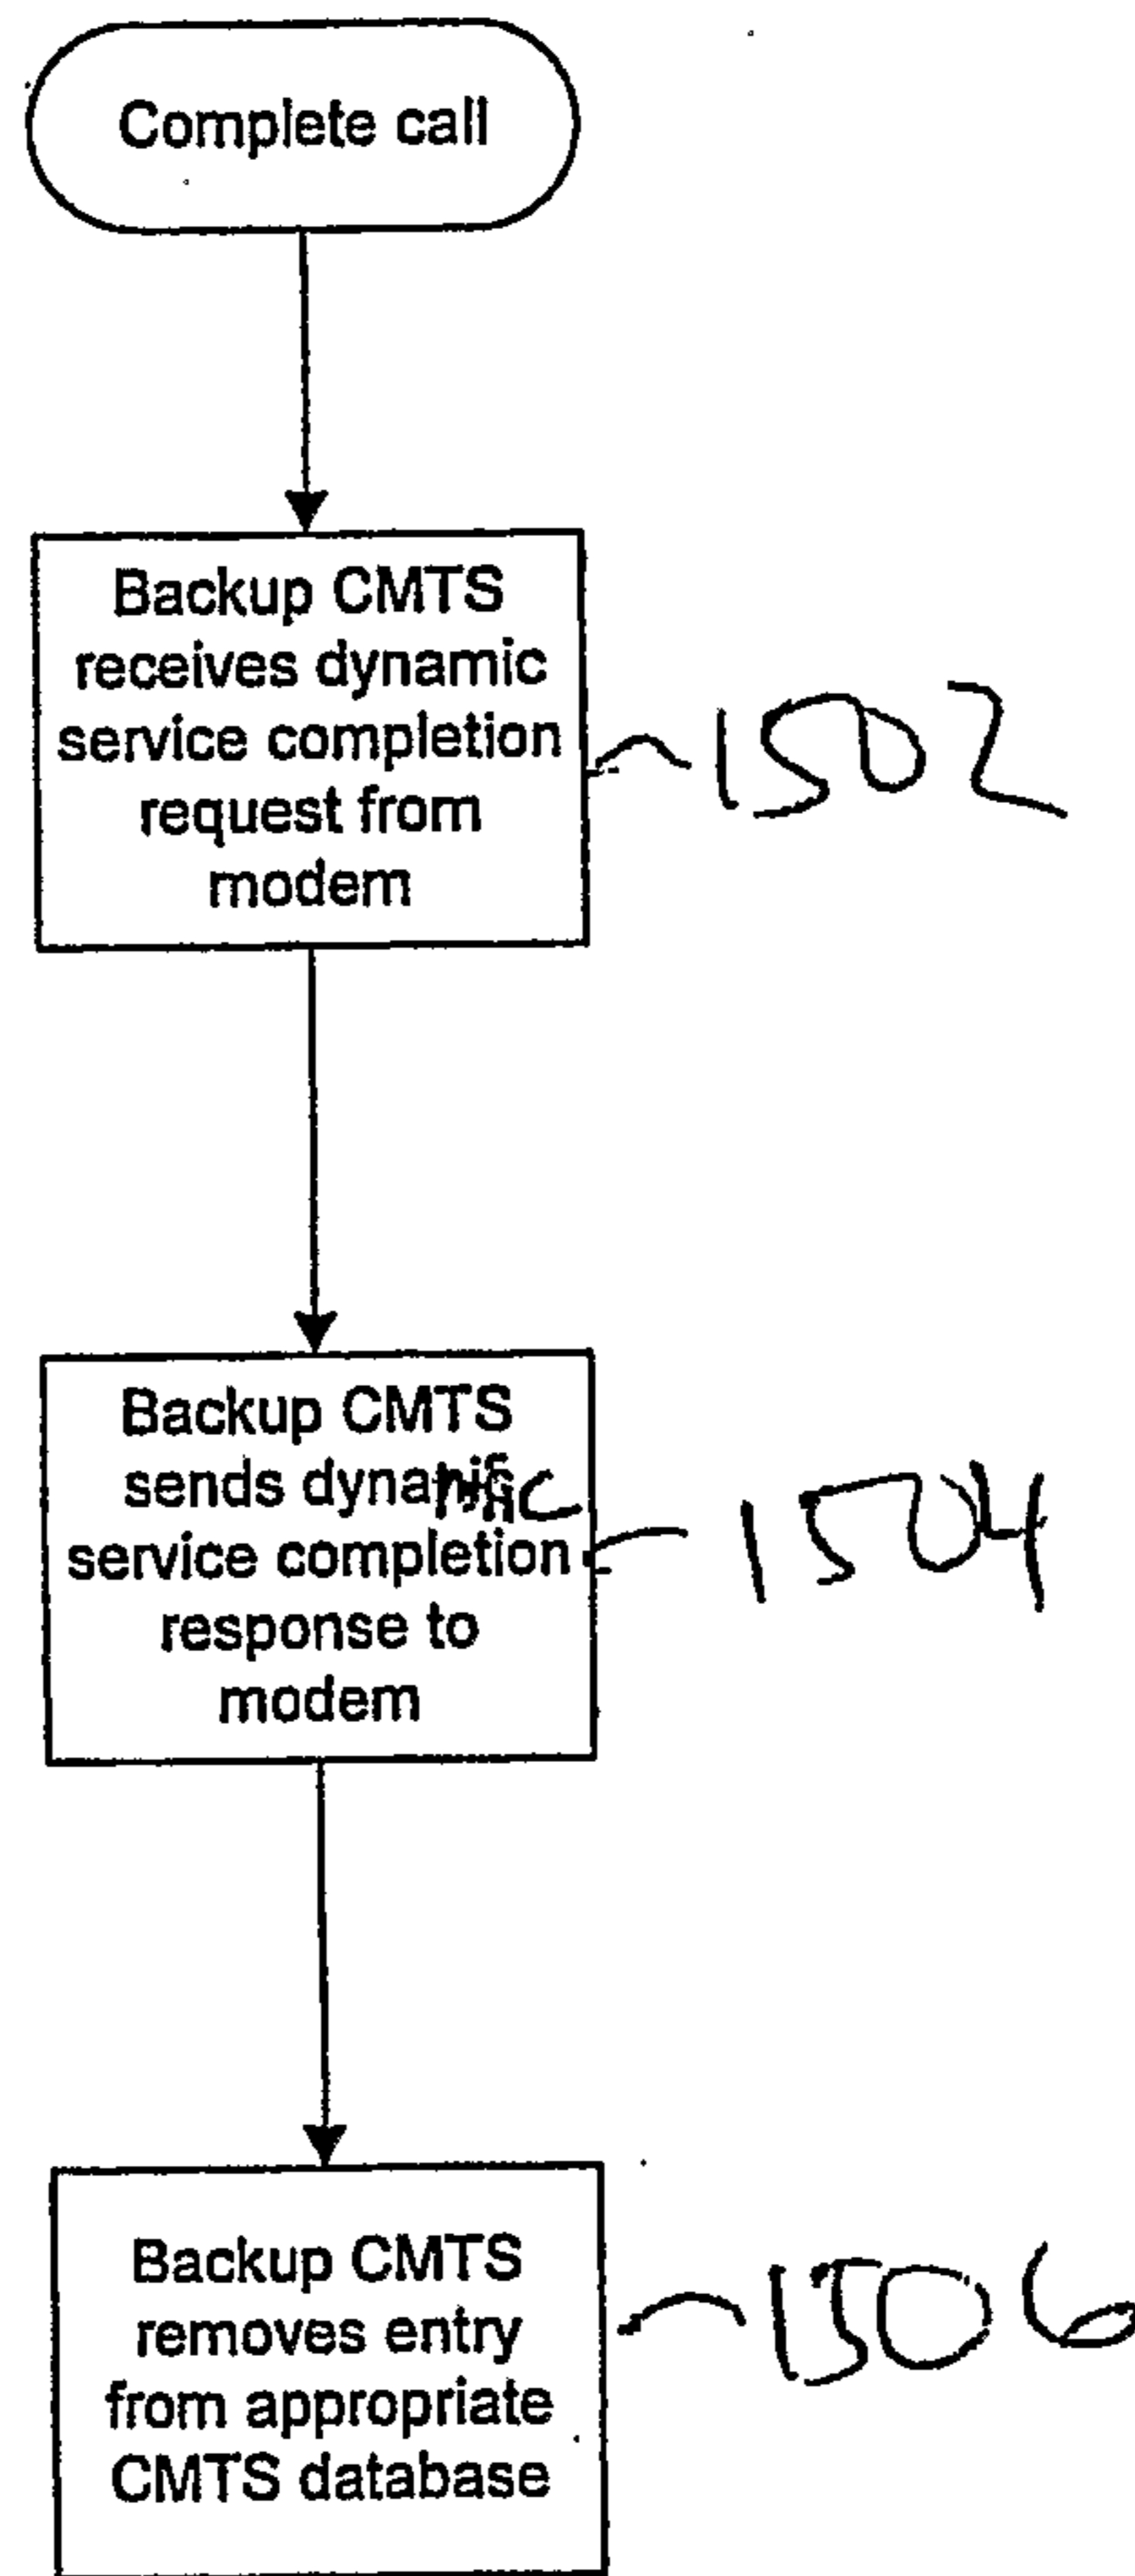


FIG. 15

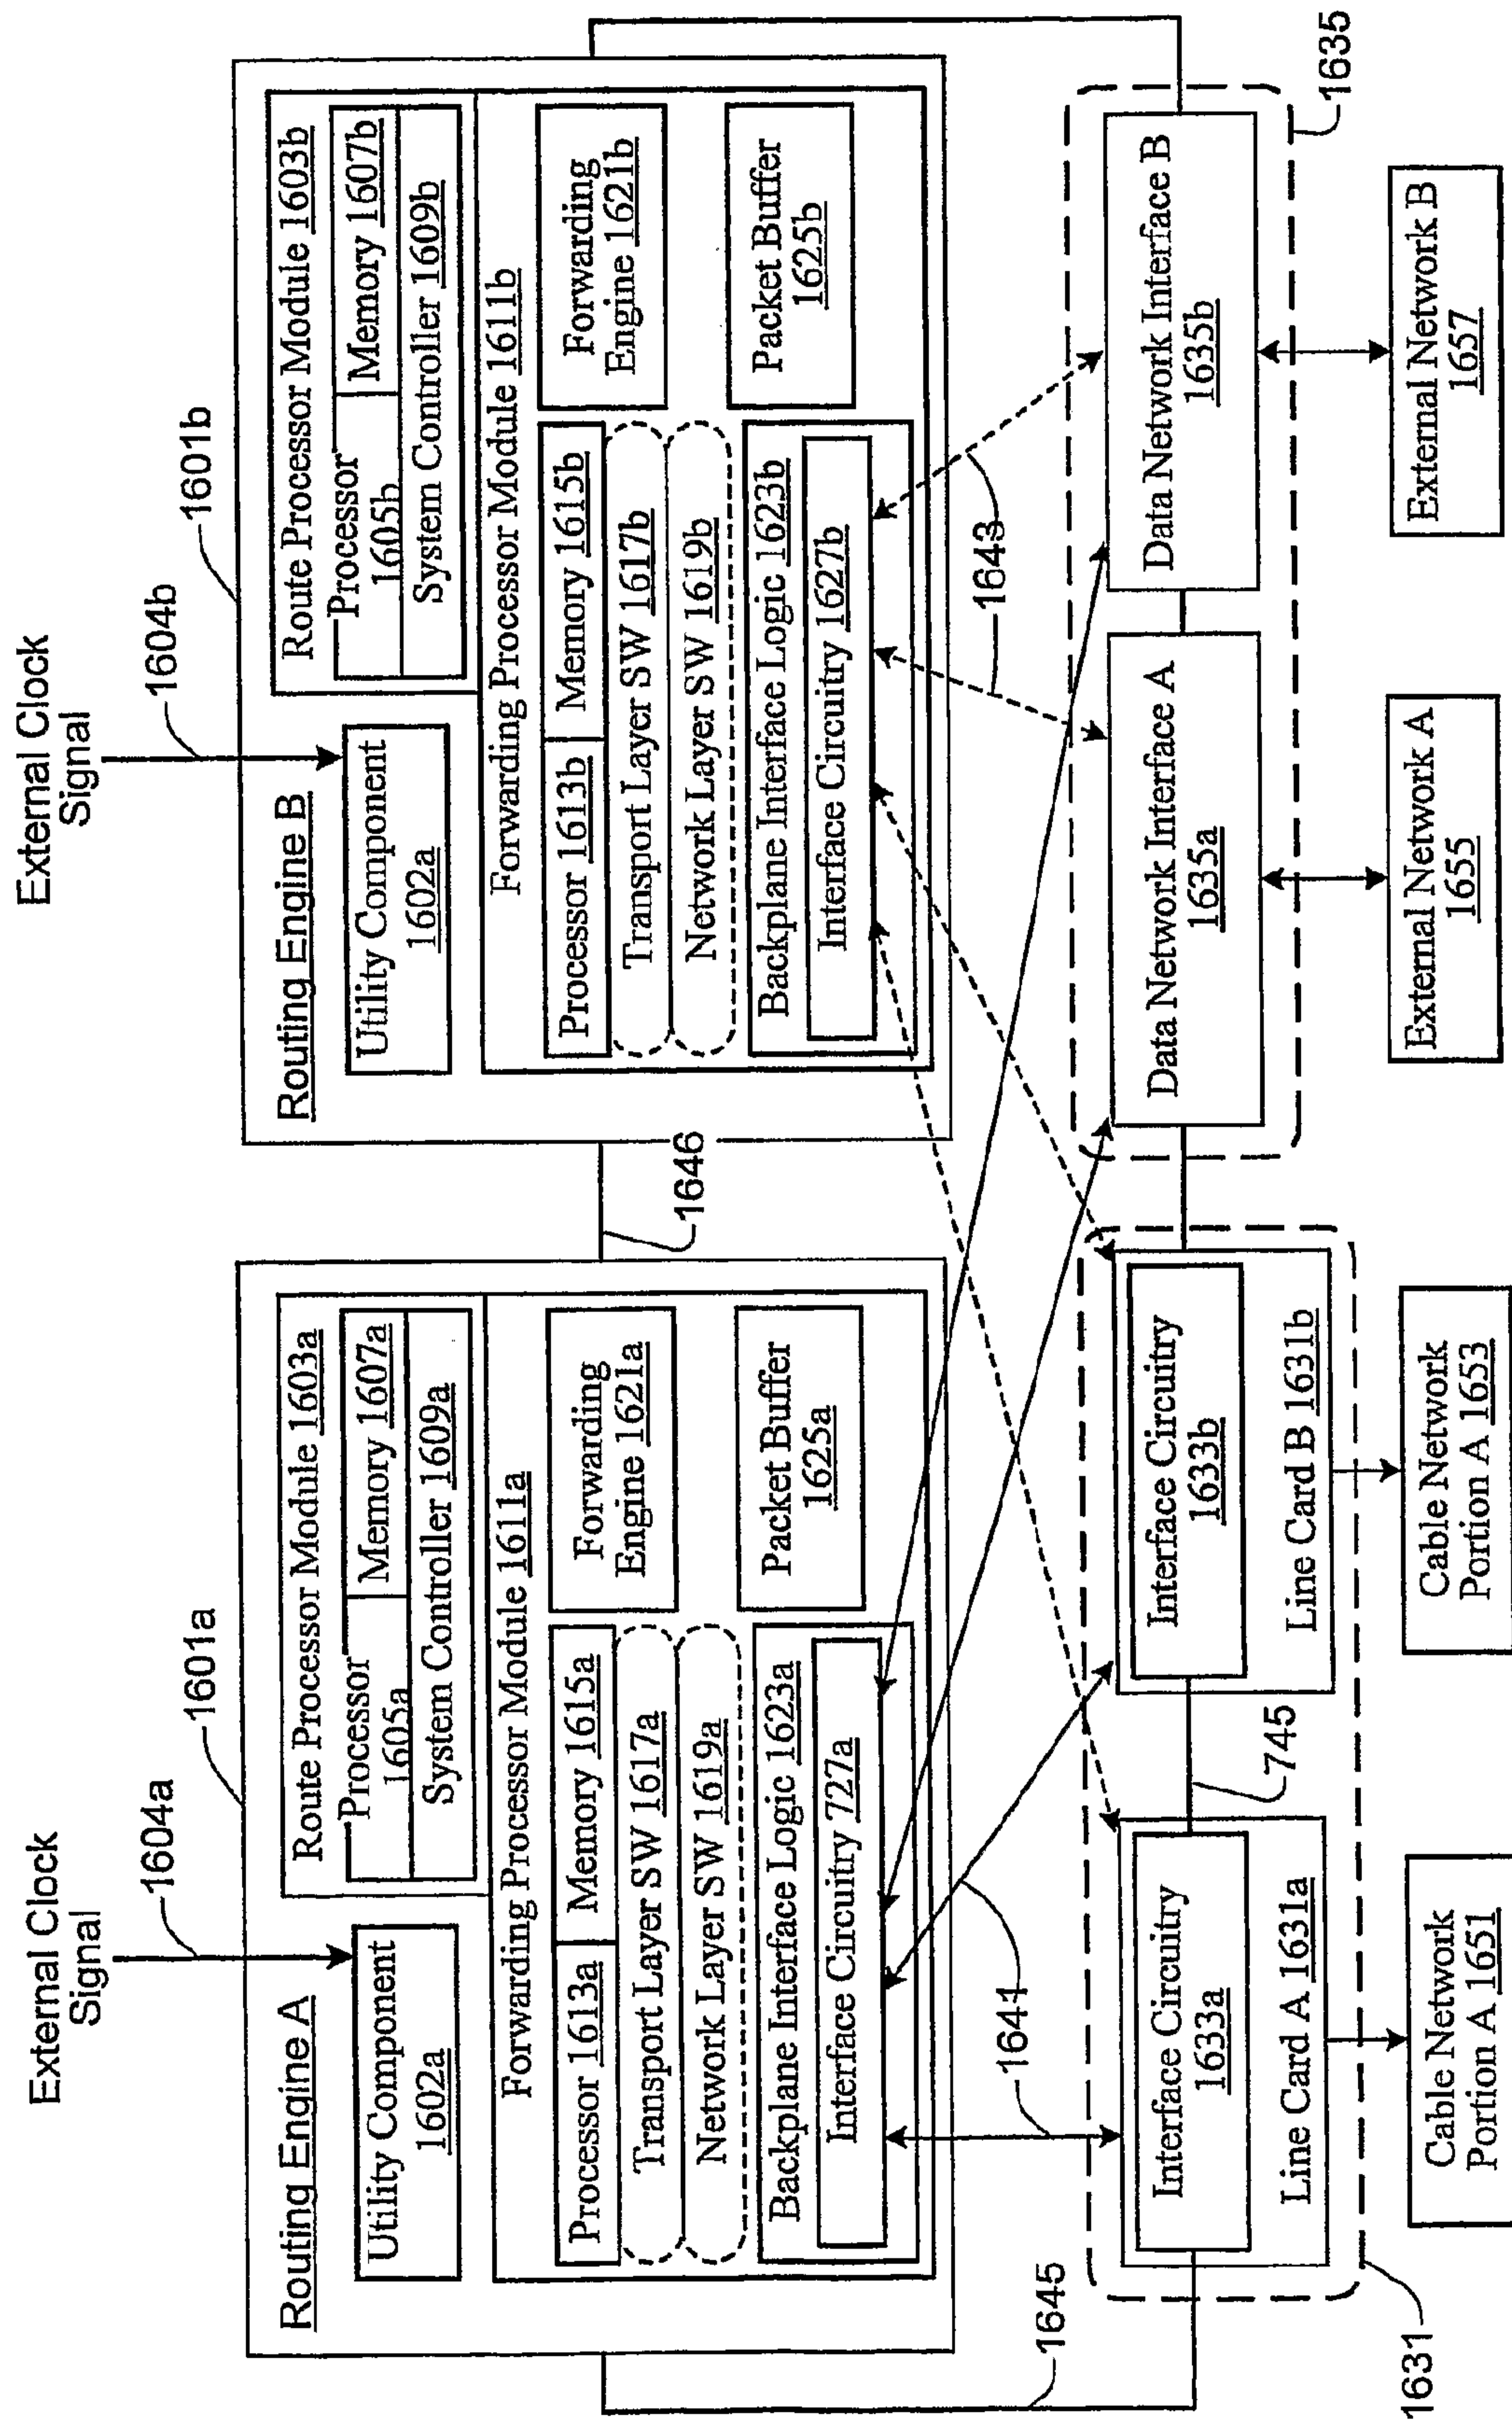


Figure 16 1600






Figure 17



208

Backup CMTS prioritizes subscribers and restores traffic accordingly

Determine identity of active CMTS that failed

1302

Obtain modem database associated with the failed active CMTS

1304

1316

Search entries having secondary ID for scheduling types indicating real-time transmission such as voice or video (e.g., VGS, VGSAD)

1308

Secondary SID?

1306

Optionally sort the entries that do not have a secondary ID (e.g., non real-time) in modem database according to priority order (e.g., time stamp) associated with the ranging lists such that the entries that do not have a secondary ID are prioritized below the entries having a secondary ID

Prioritize the entries having secondary ID according to scheduling types indicating real-time transmission

1310

Optionally further sort the entries having secondary ID according to time stamp

1312

Poll modems identified in the modem database according to priority order (e.g., according to secondary SID, scheduling type, and/or timestamp)

1314

