

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
13 September 2007 (13.09.2007)

PCT

(10) International Publication Number  
**WO 2007/103841 A2**


- (51) International Patent Classification:  
*H04L 12/56* (2006.01)
- (21) International Application Number:  
PCT/US2007/063252
- (22) International Filing Date: 5 March 2007 (05.03.2007)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
11/370,141 6 March 2006 (06.03.2006) US
- (71) Applicant (for all designated States except US): CISCO TECHNOLOGY, INC. [US/US]; 170 West Tasman Drive, SJC/10/2/1, San Jose, California 95134-1706 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): CHAPMAN, John T. [US/US]; 13250 Pierce Road, Saratoga, California 95070 (US). LIU, Xiaomei [CN/US]; 1072 South Blaney Ave., San Jose, CA 95129 (US). PARANDEKAR, Harsh [IN/US]; 903 McBride Loop, San Jose, California 95125 (US). BOMMAREDDY, Satish [IN/US]; 5843 Comanche Dr., San Jose, California 95123 (US).

- (74) Agents: LAFFERTY, Wm. Brook et al.; Scientific-Atlanta, Inc., Intellectual Property Dept., 5030 Sugarloaf Parkway, Lawrenceville, Georgia 30044 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, 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, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**  
— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: RESOURCE RESERVATION AND ADMISSION CONTROL FOR IP NETWORK



(57) Abstract: A resource reservation and admission control scheme uses pseudowires to reserve bandwidth over a layer-2 and/or layer-3 network. The pseudowires are associated with ports on different network processing devices. During a resource reservation and admission control session, the physical links used by the pseudowire are selected and reserved to more effectively allocate network bandwidth. The negotiated pseudowire is then used to transport content for a communication session over the network. In one example application, the resource reservation and admission control scheme is used during a Downstream External PHY Interface (DEPI) session for pseudowires established between a Modular Cable Modem Termination System (M-CMTS) Core and an Edge Quadrature Amplitude Modulation Interface (EQAM) device. However, the reservation protocol can be used in any application that needs to reserve bandwidth over an Internet Protocol (IP) network.

WO 2007/103841 A2



---

*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.*

## RESOURCE RESERVATION AND ADMISSION CONTROL FOR IP NETWORK

5

## Background

A new Modular Cable Modem Termination System (M-CMTS) architecture has been developed for the Data Over Cable Service Interface Specification (DOCSIS) environment that is described in co-pending application No. 11/134,818, filed May 20, 2005, entitled “TIMING SYSTEM FOR MODULAR CABLE MODEM TERMINATION SYSTEM” which is herein incorporated by reference.

One characteristic of the Modular CMTS architecture is that the DOCSIS Media Access Control (MAC) and Physical Interface (PHY) are located in different chassis and connected together through any type of packet switched network, such as a Gigabit Ethernet (GE), 10GE, or Multi-Protocol Label Switching (MPLS). In one embodiment, the DOCSIS MAC is located in the M-CMTS Core and the PHY is located in an Edge Quadrature Amplitude Modulation (EQAM) device. The interface in the downstream direction between the M-CMTS Core and the EQAM device is referred to as the Downstream External PHY Interface (DEPI), and as described above, may be established over a Gigabit Ethernet network or some other type of packet switched Internet Protocol (IP) network.

The intervening GE network between the MAC and PHY is alternatively referred to as a Converged Interconnect Network (CIN) and typically comprises one or more Layer-2 (L2) bridges and/or Layer-3 (L3) routers. In the M-CMTS architecture, multiple different Ethernet connections may be connected over the CIN between one or more M-CMTS Cores and one or more EQAM devices.

Bridging and routing protocols try and choose the best links on intermediate hops between these different layer-2 and layer-3 switching devices. This can result in the over subscription of some links and the under subscription of other links. For example, the intermediate CIN network may have a small number of source and destination MAC  
5 addresses that may not evenly distribute traffic due to hashing on the limited number of MAC addresses. Similar uneven load distribution may be created by layer-3 load balancing algorithms. In either case, bandwidth capability in the CIN can be negatively impacted by the packet switching devices between the M-CMTS Core and the EQAM device.

The present invention addresses this and other problems associated with the prior art.

10

#### Summary of the Invention

A resource reservation and admission control scheme uses pseudowires to reserve bandwidth over a layer-2 and/or layer-3 network. The pseudowires are associated with ports on different network processing devices. During a resource reservation and admission  
15 control session, the physical links used by the pseudowire are selected and reserved to more effectively allocate network bandwidth. The negotiated pseudowire is then used to transport content for a communication session over the network. In one example application, the resource reservation and admission control scheme is used during a Downstream External PHY Interface (DEPI) session for pseudowires established between a Modular Cable Modem  
20 Termination System (M-CMTS) Core and an Edge Quadrature Amplitude Modulation Interface (EQAM) device. However, the reservation protocol can be used in any application that needs to reserve bandwidth over an Internet Protocol (IP) network.

The foregoing and other objects, features and advantages of the invention will become more readily apparent from the following detailed description of a preferred embodiment of the invention which proceeds with reference to the accompanying drawings.

5

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a cable modem system that uses a resource reservation and admission control scheme.

FIG. 2 is a more detailed diagram of the cable modem system described in FIG. 1.

FIG. 3 is a flow diagram describing the operation of a resource reservation controller  
10 located in a Modular CMTS (M-CMTS) Core.

FIG. 4 is a flow diagram describing the operation of a resource reservation controller located in a Quadrature Amplitude Modulation (QAM) device.

FIG. 5 is a diagram of a packet used for establishing a pseudowire across an Ethernet network between the M-CMTS Core and the EQAM device.

15 FIG. 6 is a block diagram showing how the resource reservation and admission control scheme can be used over an IP network for other non-cable applications.

FIG. 7 is a diagram showing how the resource reservation and admission control scheme can be used to provide redundant communication sessions.

20

### DETAILED DESCRIPTION

Referring to FIG. 1, a Modular CMTS Core 14 is connected to a Wide Area Network (WAN) 12 and to one or more Edge Quadrature Amplitude Modulation Interface (EQAM) devices 22 through an Internet Protocol (IP) network 18. In one embodiment, the IP network 18 may include one or more Layer-2 and/or Layer-3 packet switching devices 20A and/or

20B. The EQAM device 22 is connected to a Hybrid Fiber Coaxial (HFC) cable plant 26 that is used for communicating to multiple Cable Modems (CMs) 28. The CMs 28 are connected to any type of equipment that may need to transmit or receive DOCSIS IP data. For example, a Personal Computer (PC) 30, Set-Top Box (STB), television, IP phone, or any other type of  
5 IP endpoint or customer premise equipment.

### Resource Reservation

A resource reservation and admission control scheme is performed by a controller 16 in the M-CMTS Core 14 and a controller 24 in the EQAM device 22. The controllers 16 and  
10 24 set-up pseudowires 32 over the IP network 18. The pseudowires 32 are alternatively referred to as tunnels or virtual connections and are used by the controllers 16 and 24 to more effectively reserve and distribute bandwidth over the IP network 18. In one embodiment, the pseudowires (PW) 32 set up a logical connection between two end points. When the PW is set up, it is then associated with a Layer-2 (L2) tag, such as a Virtual Local Area Network  
15 (VLAN) tag. The VLAN tags are reserved and allocated and then used in packets for a particular communication session between the M-CMTS Core 14 and EQAM device 22.

Some VLAN tables are configured for some or all of the egress ports 34 and/or ingress ports 36 in the M-CMTS Core 14 and EQAM device 22, respectively. The VLAN tables associate one or more VLAN tags with each port. The VLAN table can also identify  
20 the bandwidth configuration of the port. For example, some ports may be 1 GE ports and some may be 10GE ports. In another embodiment, an egress port 34 may be associated with one or more VLAN tags while the ingress port 36 may accept any VLAN tag.

The controllers 16 and 24 in the M-CMTS Core 14 and EQAM device 22, respectively, negotiate which VLAN tags are used for a particular communication session.

The selected VLAN tag is associated with ports having sufficient available bandwidth for transporting content 10 from WAN 12 over IP network 18.

The EQAM device 22 and the M-CMTS Core 14 can select which physical link to place a PW 32 onto and have it stay there. Since the max bandwidth of each PW 32 is known  
5 and the bandwidth of the physical links are known, the PW 32 can be admitted based upon available bandwidth (admission control) and then directly associated with a physical link (resource reservation).

The pseudowires 32 established over the IP network 18 use the negotiated VLAN tag to direct packets out particular egress ports 34 in M-CMTS Core 14 and into particular  
10 ingress ports 36 in EQAM device 22. Because bandwidth for these ports is reserved, bandwidth utilization is more effectively controlled over particular physical links in the IP network 18.

In one embodiment, the IP network 18 is a Gigabit Ethernet (GE) network and the egress ports 34 on the M-CMTS Core 14 and the ingress ports 36 in the EQAM device 22 are  
15 GE ports that may each contain one or more VLAN IDs. For example, GE egress port 34A and GE ingress port 36A may each be associated with one or more of the same VLAN ID values. Similarly, GE egress port 34B and GE ingress port 36B may both be associated with one or more of the same VLAN ID values. Of course, different egress ports 34 or ingress ports 36 can be assigned to different combinations of VLAN ID values.

20 The VLANs may connect between any combinations of physical ports (34A to 36A or 34A to 36B) since the intervening switches 20A and 20B can provide the desired connectivity. For example, the ports in FIG. 1 may have the following VLAN ID associations:

Port 34A: VLAN ID 1, VLAN ID 3;

Port 34B: VLAN ID 2, VLAN ID 4;

Port 36A: VLAN ID 1, VLAN ID 4; and

Port 36B: VLAN ID 2, VLAN ID 3.

A virtual connection between port 34A and 36A may be associated with VLAN ID 1 or

5 between port 34A and port 36B with VLAN ID 3 depending on bandwidth availability.

Similarly, a virtual connection between port 34B and 36A may be associated with VLAN ID 4 or between port 34B and port 36B with VLAN ID 2. Further, VLAN ID values associated with other M-CMTS Cores can also be associated with either ports 36A and 36B and VLAN IDs associated with other EQAM devices may be associated with port 34A and 34B.

10 In an alternative embodiment, Ethernet addresses associated with the different ports 34 and 36 may be used to assign different ports to different communication sessions. In yet another embodiment, Multi-Protocol Label Switching (MPLS) tags can be used to reserve ports on the M-CMTS Cores 14 and the EQAM devices 22.

The pseudowires 32 can be viewed by the controllers 16 and 24 as variable bit rate  
15 connections with a maximum burst rate placed on a physical link in the CIN 18 between the MCMTS 14 and the EQAM 22. The controllers 16 and 24 then use the resource reservation protocol described below to decide which physical links (physical ports) to place the pseudowires 32.

As an example, the pseudowires 32 may have a maximum burst rate of around 40  
20 Million bits per second (Mbps) and 24 of these pseudowires could be built into a same physical GE connection. There may be multiple MAC framers, QAM channels (PHYs), and physical GE connections in the communication system (see FIG. 2). The controllers 16 and 24 use the resource reservation protocol to efficiently pack each physical connection with the maximum number of logical pseudowire connections 32. For example, as mentioned above,

the EQAM device 22 may have one ingress port 36 for every 24 QAM output channels. Accordingly, the controllers 16 and 24 may establish 24 logical pseudowires 32 over each ingress port 36.

The resource reservation can be used for any system that needs to allocate ports on a common subnet as further described below in FIG. 7. For example, the resource reservation and admission control scheme can be used as an extension to DEPI, RSVP, or any other resource reservation protocol that needs to establish pseudowires, tunnels, or virtual connections across a layer-2 and/or layer-3 IP. The resource reservation system can also be used with multiple different subnets by configuring the L-2 and L-3 packet switches in the different subnets to conduct the same resource reservation operation conducted by the M-CMTS Core 14 and EQAM device 22.

FIG. 2 shows the system in FIG. 1 in more detail. The M-CMTS Core 14 includes a packet switching engine 40 that receives different content A, B and C over the same or different WAN connections 41. The content A, B, and C are contained in IP packets that are directed to different associated cable modems 62A, 62B, and 62C, respectively.

The packet switching engine 40 forwards the packets for content A, B, and C to different MAC framers 42 associated with the QAM channels 60 for the associated destination cable modems 62. For example, the content A is directed to a MAC framer 42A that communicates with CM 62A via a QAM channel 60A. The content B and C is directed to a MAC framer 42B that communicates with CMs 62B and 62C via a QAM channel 60B. The MAC framers 42 in this example operate as resource reservation and admission controllers 16 and the QAM channels 60 operate as the controllers 24 previously shown in FIG. 1. In one embodiment, the MAC framers 42 and the QAM channels 60 include

processors that are programmed to perform the resource reservation and admission control scheme.

The VLAN IDs for separate different M-CMTS Cores 14 and separate MAC framer 42 are permitted to converge on a single ingress port in the EQAM device 22 (the PHY box).

5 The converse is also true. One M-CMTS Core 14 can have a single GE port 34 with 2 VLAN IDs that cross the network 18 with each VLAN ID terminating on a different EQAM device 22. The association of VLAN IDs and GE ports can be configured by a network administrator.

FIGS. 3 and 4 in combination with FIG. 2 will be referenced to describe the resource  
10 reservation and admission control scheme in more detail. Referring first to FIG. 3, resource reservation and admission control may take place when M-CMTS Core 14 and EQAM device 22 first initiate a communication session. In one example, this may be the initiation of a communication session in operation 70 by the MAC framer 42A. Any type of communication session may kick off the resource reservation operation, such as initiation of a  
15 DEPI session in the M-CMTS environment, the initiation of a communication session between two gateways, etc.

In operation 71, the MAC framer 42A announces in message 50 how much bandwidth the M-CMTS Core 14 requires from both the QAM channel 60A and the IP network 18 for transmitting content A. In operation 72, the controller 16 lists in the same message 50, or in a  
20 different message 52, all of the VLAN IDs that can connect from the MAC framer to QAM channel 60 that also have sufficient bandwidth capacity for transmitting content A. As described above, this may comprise any VLAN IDs associated with physical egress ports 34 that have the required connectivity and available bandwidth.

A VLAN ID may simultaneously participate in multiple session negotiations. In one embodiment, the M-CMTS Core 14 may provisionally make a reservation of bandwidth for any or all of the VLAN IDs identified in message 52. This reservation can prevent bandwidth from the listed VLAN IDs from being allocated to other communication sessions while a VLAN ID is being negotiated with the QAM channel 60A. In this embodiment, after the QAM channel 60A selects one of the identified available VLAN IDs, and the selected VLAN ID is accepted, the M-CMTS Core 14 may release the bandwidth reservation for any of the non-selected VLAN IDs that were provisionally reserved.

Alternatively, the M-CMTS Core 14 may not provisionally reserve any bandwidth on any of the currently available VLAN IDs. In this embodiment, the QAM channel 60A in response message 54 selects one of the VLAN IDs from the list in message 52. If the selected VLAN ID is no longer available when message 54 is received, the M-CMTS Core 14 may send a failure message back to the QAM channel 60A indicating that bandwidth for the selected VLAN ID is no longer available. The failure message could then list a new set of currently available VLAN IDs for reselection by QAM channel 60A.

Alternatively, the M-CMTS Core 14 could make a temporary reservation on all the VLAN IDs it publishes to the EQAM device 22. When the EQAM response is received, the M-CMTS Core 14 could then release all the temporary reservations except the one selected by the EQAM device 22, which it would then make permanent.

Either way, in operation 74, the M-CMTS Core 14 receives a response 54 back from the QAM channel 60A identifying one of the VLAN IDs previously presented in message 52. Alternatively, the QAM channel 60A could send back an error message indicating none of the listed VLAN IDs are available.

When a VLAN ID is identified, the framer 42A in operation 76 reserves the bandwidth for the egress port 34A associated with the identified VLAN ID. A connect message is then sent back to the QAM 60A in operation 78. The framer 42A then uses the selected VLAN ID in packets containing content A to establish a pseudowire 48 with the  
5 QAM channel 60A.

FIG. 4 shows the operations performed by the QAM channel 60A. In one embodiment, the EQAM device 22 and QAM channels 60 can be connected to multiple different M-CMTS Cores 14. Thus, in this configuration, the EQAM device 22 is given the responsibility of selecting and arbitrating a particular VLAN ID. The QAM channel 60A in  
10 operation 82 receives session initiation messages from the M-CMTS Core 14. In operation 84, the QAM 60A receives the message 50 that identifies the bandwidth requirements for the communication session and in operation 86 receives the message 52 that contains the list of available VLAN IDs for the communication session.

In operation 88, the QAM 60A chooses one of the identified VLAN IDs associated  
15 with an ingress port 36 that has sufficient available bandwidth to handle the connect speed identified in message 50. The QAM 60A chooses the VLAN ID based upon any variety of internal criteria, such as available bandwidth that has not been reserved yet and path connectivity. For example, if the listed VLAN IDs are associated with more than one physical port, the VLAN ID associated with the port having the most available bandwidth  
20 may be selected. The QAM 60A reserves the bandwidth on the port 36A associated with the selected VLAN ID in operation 90 and then returns the selected VLAN ID to the framer 42A in operation 92.

After a connect message is received back from the framer 42A in operation 94, the negotiated VLAN ID is then associated with the pseudowire 48 that is used in operation 96

for conducting the communication session. The framer 42B may conduct a similar resource reservation and admission control operation for the pseudowires 56 used for transmitting content B and C.

FIG. 5 shows an example of a packet 120 that may be sent from the M-CMTS Core 14 to the EQAM device 22. The packet 120 includes a VLAN ID field 122 that contains the VLAN ID value that is negotiated between the M-CMTS Core 14 and EQAM device 22. In one embodiment, the destination address 124 may be used instead of, or in conjunction with, the VLAN ID value to direct the packet 120 over a particular pseudowire between M-CMTS 14 Core and EQAM device 22.

In one embodiment, the M-CMTS Core 14 and the EQAM device 22 may not care what particular path is taken within network 18 for routing or switching the packet 120 from the negotiated egress port 34 to ingress port 36. Accordingly, the packet switches 20A and 20B (FIG. 1) in the CIN may not be configured with VLAN ID/port associations and the only network processing elements that are configured with VLAN ID/port associations are the M-CMTS Core 14 and the EQAM device 22. This still ensures that packets 120 for the negotiated pseudowire will be output on a previously reserved egress port 34 and input to a previously reserved ingress port 36.

Alternatively, the pseudowire may be directed through particular ports on one or more packet switches 20 in IP network 18. In this embodiment, the one or more intermediate packet switches are also configured with different VLAN ID values associated with physical ports. In this alternative embodiment, the intermediate packet switches may send reply messages back to the M-CMTS Core identifying available VLAN tags. Any common VLAN tags could then be sent to the EQAM device 22. Alternatively, the intermediate nodes may not be aware or be configured for operating this feature within the DEPI protocol. In this

situation, the negotiation of VLAN tags by the intermediate packet switches could be integrated with some other existing VLAN discovery or assignment protocol such as MPLS.

#### Attribute Value Pairs

5 Any type of messaging protocol and message structure can be used for the resource reservation and admission control negotiation. As mentioned above, one application for the resource reservation scheme is for use during a Modular CMTS DEPI session. In this embodiment, the resource reservation negotiations between the M-CMTS Core 14 and the EQAM device 22 (FIGS. 1 and 2) may be conducted using Attribute Value Pairs (AVPs)  
10 during initiation of the DEPI session. One example of these AVPs are described below. Of course this is only one example and other message formats could also be used depending on the particular reservation application.

The concept of using a L3 protocol to manage physical links with L2 tags could be applied to other Layer 3 protocols, such as RSVP, and could be applied to other Layer 2 tags  
15 such as MPLS.

#### **Transmit Connect Speed AVP**

In this embodiment, the message 50 shown in FIG. 2 may be a transmit connect speed AVP. This AVP can be used in the DEPI session to indicate the maximum transmission rate  
20 of the QAM channel that will be used by the M-CMTS Core 14. If a QAM channel 60 (FIG. 2) is shared between two services, such as non-DOCSIS video and DOCSIS, this AVP could represent the bandwidth of the QAM channel being used for DOCSIS. The AVP could also reflect the DOCSIS channel being derated below 100% to prevent queue buildup from jitter.

The Tx Connect Speed may be calculated according to the maximum bit rate for MPEG-TS packets. The Tx Connect Speed can alternatively be calculated according to the maximum bit rate for DOCSIS frames. The EQAM device 22 can also account for the PHY level overhead including Forward Error Correction (FEC) and Trellis encoding if present.

5 Any reconfiguration change in the bandwidth rating can be communicated to each associated M-CMTS Core 14 and/or EQAM device 22. The controllers 16 and 24 in the M-CMTS Core 14 and EQAM device 22 would then automatically update the parameters used for conducting the resource reservation scheme. For example, the bandwidth rating could be changed from 98% to 96%. A connection may already be established prior to receiving an  
10 updated bandwidth rating. The current port may be overbooked as a result of the new bandwidth rating. The current session may then be torn down and then be re-established using the new rating. Alternatively, an error code may be generated.

#### **VLAN Available IDs**

15 A L2TPv3 Control Connection (CC) is specific to a pair of IP addresses. An IP subnet may span several physical connections in a bridged environment, such as the CIN 18 in FIG. 2. If that environment is managed with VLANs, the available VLAN ID AVP may be used for sending the available VLAN ID message 52 (FIG. 2). This allows the EQAM 22 to choose an ingress port 36 for a DEPI session by specifying one of the identified VLAN  
20 IDs. If this AVP is used, then during an Incoming Call Request message (ICRQ), the M-CMTS Core 14 supplies a list of VLAN IDs that can reach the EQAM device 22 that also have sufficient bandwidth to support the new DEPI session.

If a reconfiguration changes the VLAN assignments in the M-CMTS Core 14, the M-CMTS Core can issue an L2TP Setup Link Information (SLI) command with the updated VLAN ID information. The VLANs may be assigned such that each VLAN ID has only one end point at the M-CMTS Core 14 and only one end point at the EQAM device 22. A  
5 physical egress port on a M-CMTS Core 14 may have more than one VLAN ID assigned to it. A physical DEPI ingress port 36 on the EQAM device 22 may have more than one assigned VLAN ID.

#### **DEPI VLAN Assigned ID**

10 An assigned AVP may be used for response message 54 (FIG. 2) sent by the EQAM 22 during a L2TP Incoming Call Reply message (ICRP) in response to a valid DEPI VLAN available ID AVP sent during the L2TP ICRQ from the M-CMTS Core 14. The list of VLAN IDs from the M-CMTS Core 14 in the DEPI VLAN available ID AVP can be invalid. Alternatively, the VLAN available ID AVP may be valid, but none of the VLANs specified  
15 have sufficient bandwidth at the ingress port 36 of the EQAM 22 to admit an additional session. In either case, the EQAM 22 may return the appropriate error code message. Otherwise, the EQAM device 22 selects one of the VLAN IDs from the DEPI VLAN Available IDs AVP and returns it in the DEPI VLAN Assigned ID AVP in an ICRP message. If a configuration change occurs pertaining to the VLAN assignment at the EQAM 22, the  
20 EQAM 22 issues an SLI command with the updated VLAN ID information.

### Generic Layer-2 and Layer-3 Network Reservation

As described above, a mixture of L3 signaling and L2 resource tags can be used to reserve network bandwidth on any multi-hop, multi-path, L2 and/or L3 network. This mixture of L3 and L2 concepts is somewhat unconventional, but achieves the goal of controlled network resource reservation and allocation within the scope of L3 protocols on a L2 network. By putting the choice of paths under the control of the two endpoints, the resource reservation allows the intervening network 18 to be fully booked with pseudowire traffic without the fear of lost packets due to errors in load balancing or forwarding algorithms.

FIG. 6 shows how this resource reservation scheme is extendable past the DOCSIS environment. The network processing devices 130 and 132 can be any endpoints or network processing elements that reserve layer-2 network resources by establishing pseudowires 138 over an IP network 136. The IP network 136 may all be part of the same subnet that contains one or more packet switching devices 130, 132, and 134 that may include multiple physical ports 139 on a common IP subnet. The scheme described above allows bandwidth to be reserved in this network on a per physical port basis.

In one embodiment, the pseudowires 138 are used for transporting Moving Picture Experts Group Transport Stream (MPEG-TS) video or any other type of layer-2 or layer-3 network traffic. The MPEG market currently uses dedicated GE links between a video server and an EQAM device. If that network grows and/or when that network converges with the DOCSIS CIN network, the resource reservation scheme can be re-applied from the DEPI protocol to the protocols used in the video environment. This may also be generally useful for providing an IntServ like environment for pseudowires across L2 networks.

Other types of network applications can also use the resource reservation scheme. For example, the network processing devices 130 and 132 may be gateways that need to reserve bandwidth for different audio or media streams either separately or in conjunction with a Resource Reservation Setup Protocol (RSVP).

5

### Redundancy

FIG. 7 shows yet another embodiment of the resource reservation scheme that may be used to provide redundancy in the Modular CMTS environment or other non-cable applications. FIG. 7 shows the M-CMTS Core 14 operating both a working MAC framer 144 and a protect MAC framer 146. The working framer 144 is alternatively referred to as the primary framer and the protect framer 146 is alternatively referred to as the secondary framer. The M-CMTS Core 14 includes a packet switching engine 143 that receives content 142 over a connection 140 connected to the WAN network 12 (FIG. 1). The packet switching engine 143 can forward the content 142 either to the working framer 144 or the protect framer 146.

15 The DEPI protocol currently allows only one session between one of the MAC framers and the QAM channel 60 in EQAM 22. For redundancy, a second protection session is needed. However, it is undesirable to delay initiating a new session until after a primary session fails.

Accordingly, a working (primary) session ID 148 and a protect (secondary) session ID 150 are established for a same DEPI session. The working framer 144 establishes the working DEPI session 148 in a conventional manner with the QAM channel 60. The working session 148 has an associated MPEG2 Transport Stream ID (TSID), primary Control Connection (CC) ID and primary session ID 154. If the M-CMTS Core 14 is configured for operating in the protection mode, then information 145 associated with the working DEPI

session 148 is communicated to the secondary protect framer 146. For example, the TSID, primary CC ID, and primary session ID information 145 are communicated from working framer 144 to protect framer 146.

The protect framer 146 sends a message 152 to the QAM channel 60 requesting  
5 establishment of a secondary DEPI session 150. The request 152 identifies itself as a protect session for primary working session 148. The secondary request 152 may also include the TSID, primary CC ID, and primary session ID 154 for working session 148. The QAM channel 60 then verifies the working session information 145, and if authenticated, grants the secondary connection 150. Both the working MAC framer 144 and the protect MAC framer  
10 146 can then establish pseudowires with the QAM channel 60.

The working DEPI session 148 continues to operate as normal until there is a disruption in working framer 144 or in the associated pseudowire connection. At that time, the packet switching engine 143 starts sending the content 142 to the protect framer 146. The packets carrying the content 142 are assigned the VLAN tag associated with the pseudowire  
15 negotiated with protect session 150.

Because the protect session 150 is already established, the protect framer 146 can send the packets containing content 142 directly to QAM channel 60 without having to initiate another DEPI session. The packets received over protect session 150 are processed in the same manner as packets previously received over primary session 148 and forwarded to  
20 the same cable modem via the HFC plant 26 (FIG. 1).

The system described above can use dedicated processor systems, micro controllers, programmable logic devices, or microprocessors that perform some or all of the operations. Some of the operations described above may be implemented in software and other operations may be implemented in hardware.

For the sake of convenience, the operations are described as various interconnected functional blocks or distinct software modules. This is not necessary, however, and there may be cases where these functional blocks or modules are equivalently aggregated into a single logic device, program or operation with unclear boundaries. In any event, the

5 functional blocks and software modules or features of the flexible interface can be implemented by themselves, or in combination with other operations in either hardware or software.

Having described and illustrated the principles of the invention in a preferred embodiment thereof, it should be apparent that the invention may be modified in arrangement

10 and detail without departing from such principles. We claim all modifications and variation coming within the spirit and scope of the following claims.

## Claims

1. A processing device, comprising:

a processor to receive content for transporting over an Internet Protocol (IP) network and identifying both a local port and a remote port on a remote device having available  
5 bandwidth for transmitting the content, the processor then using the identified ports to establish a pseudowire for transporting the content over the IP network.

2. The processing device according to claim 1 wherein the processor sends a

required connect speed and a list of tags associated with local ports that can connect to the  
10 remote device and that have available bandwidth for transporting the content at the required connect speed, the processor receiving back a selected one of the tags that is also associated with a remote port on the remote device that also has available bandwidth for receiving the content at the required connect speed and then using the tag to establish the pseudowire over the IP network.

15

3. The processing device according to claim 1 wherein the local and remote ports are associated with Virtual Local Area Network (VLAN) tags and the processor sends a list of the VLAN tags associated with the local ports to the remote device, the processor then receiving back an identified one of the VLAN tags.

20

4. The processing device according to claim 4 wherein the processor operates in a Modular Cable Modem Termination System (M-CMTS) Core that negotiates the VLAN tag with an Edge Quadrature Amplitude Modulation (EQAM) device for use during a Downstream External PHY Interface (DEPI) session.

5. The processing device according to claim 4 including:

a first framer to negotiate a first VLAN tag for establishing a first pseudowire for carrying the content between the M-CMTS Core and the EQAM device over the IP network;

5 and

a second framer to negotiate a second VLAN tag with the EQAM for establishing a second backup pseudowire for carrying the content between the M-CMTS Core and the EQAM device over the IP network if the first pseudowire fails.

10 6. A method for reserving resources over an Internet Protocol (IP) network, comprising:

receiving content;

generating a list of tags associated with local links available for transporting the content to a remote device identified for receiving the content;

15 sending the list of tags to the remote device;

receiving a selected one of the tags back from the remote device associated with a remote link on the remote device available for receiving the content; and

using the selected tag to establish a virtual connection for carrying the content over the IP network.

20

7. The method according to claim 6 including:

identifying a required bandwidth for transporting the content over the IP network;

generating a list of tags associated with local links that have the required bandwidth available for transporting the content;

sending both the list of tags and the identified amount of required bandwidth to the remote device; and

receiving one of the tags back from the remote device also associated with a remote link having the required bandwidth available for transporting the content.

5

8. The method according to claim 7 including:

reserving the list of tags sent to the remote device;

using the tag received back from the remote device for transporting packets carrying the content; and

10 releasing the other tags from the list for use in establishing other virtual connections for carrying other content.

9. The method according to claim 6 including:

15 using the selected tag to establish a first virtual connection for carrying the content over the IP network;

generating a second list of tags associated with local links available for also transporting the content;

sending the second list of tags to the same remote device identified for receiving the content;

20 receiving a selected second one of the tags back from the remote device; and

using the selected second tag to establish a second backup virtual connection for carrying the content over the IP network when the first virtual connection fails.

10. A network processing device, comprising:

a processor to receive a list of tags associated with links on a remote device available for conducting a communication session over an Internet Protocol (IP) network and selecting one of the tags from the list that has an associated local link available for conducting the  
5 communication session, the selected tag then used for establishing a pseudowire over the IP network with the remote device.

11. The network processing device according to claim 10 wherein the processor receives a connect speed requirement for the communication session from the remote device  
10 and selects one of the tags from the list having an associated local link with sufficient bandwidth for the required connect speed.

12. The network processing device according to claim 11 wherein the processor is located in a Quadrature Amplitude Modulation (QAM) interface and receives the list of tags  
15 and the connect speed during initiation of a Downstream External PHY Interface (DEPI) session.

13. The network processing device according to claim 10 wherein the processor selects from a first list of tags for establishing a first pseudowire with the remote device for  
20 conducting the communication session and selects from a second list of tags for establishing a second pseudowire for conducting the communication session with the remote device when the first pseudowire fails.

14. A method for reserving resources on a network, comprising:  
negotiating with a remote device for reserving bandwidth for a communication  
session;  
identifying a layer-2 identifier during the negotiation associated with a physical link  
5 that has bandwidth available for supporting the communication session;  
establishing a virtual connection over the network using the negotiated layer-2  
identifier; and  
conducting the communication session over the virtual connection.

10 15. The method according to claim 14 including:  
identifying a Virtual Local Area Network (VLAN) tag that is associated with Gigabit  
Ethernet (GE) ports having sufficient bandwidth for conducting the communication session;  
and  
using the VLAN tag to establish the virtual connection over a GE network.

15  
16. The method according to claim 15 including:  
receiving a list of VLAN tags that are associated with GE ports on a remote device  
that have available bandwidth for conducting the communication session;  
selecting one of the VLAN tags from the list that also corresponds with a local GE  
20 port that has sufficient bandwidth for conducting the communication session;  
sending the selected VLAN tag back to the remote device; and  
establishing the virtual connection with the remote device using the selected VLAN  
tag.

17. The method according to claim 16 including:

receiving the VLAN tags from a Modular Cable Modem Termination System (M-CMTS) Core during initiation of a Downstream External PHY Interface (DEPI) session;

5 selecting one of the VLAN tags and sending the selected VLAN tag back to the M-CMTS Core;

using the selected VLAN tag to establish the virtual connection over the GE network between the M-CMTS Core and a Quadrature Amplitude Modulation (QAM) device; and conducting the DEPI session over the virtual connection.

10 18. A system for reserving resources over an Internet Protocol (IP) network, comprising:

means for receiving content;

means for generating a list of tags associated with local links available for transporting the content to a remote device identified for receiving the content;

15 means for sending the list of tags to the remote device;

means for receiving a selected one of the tags back from the remote device associated with a remote link on the remote device available for receiving the content; and

means for using the selected tag to establish a virtual connection for carrying the content over the IP network.

20

19. The system according to claim 18 including:

means for identifying a required bandwidth for transporting the content over the IP network;

means for generating a list of tags associated with local links that have the required

bandwidth available for transporting the content;

means for sending both the list of tags and the identified amount of required

bandwidth to the remote device; and

means for receiving one of the tags back from the remote device also associated with

5 a remote link having the required bandwidth available for transporting the content.

20. The system according to claim 19 including:

means for reserving the list of tags sent to the remote device;

means for using the tag received back from the remote device for transporting packets

10 carrying the content; and

means for releasing the other tags from the list for use in establishing other virtual

connections for carrying other content.

21. A system for reserving resources on a network, comprising:

15 means for negotiating with a remote device for reserving bandwidth for a  
communication session;

means for identifying a layer-2 identifier during the negotiation associated with a  
physical link that has bandwidth available for supporting the communication session;

means for establishing a virtual connection over the network using the negotiated

20 layer-2 identifier; and

means for conducting the communication session over the virtual connection.

22. The system according to claim 21 including:

means for identifying a Virtual Local Area Network (VLAN) tag that is associated

with Gigabit Ethernet (GE) ports having sufficient bandwidth for conducting the communication session; and

means for using the VLAN tag to establish the virtual connection over a GE network.

5           23.     The system according to claim 22 including:

means for receiving a list of VLAN tags that are associated with GE ports on a remote device that have available bandwidth for conducting the communication session;

means for selecting one of the VLAN tags from the list that also corresponds with a local GE port that has sufficient bandwidth for conducting the communication session;

10           means for sending the selected VLAN tag back to the remote device; and

means for establishing the virtual connection with the remote device using the selected VLAN tag.

15           24.     A computer readable medium containing instructions that when executed by a computer, comprise:

receiving content;

generating a list of tags associated with local links available for transporting the content to a remote device identified for receiving the content;

sending the list of tags to the remote device;

20           receiving a selected one of the tags back from the remote device associated with a remote link on the remote device available for receiving the content; and

using the selected tag to establish a virtual connection for carrying the content over the IP network.




FIG.1



3/7

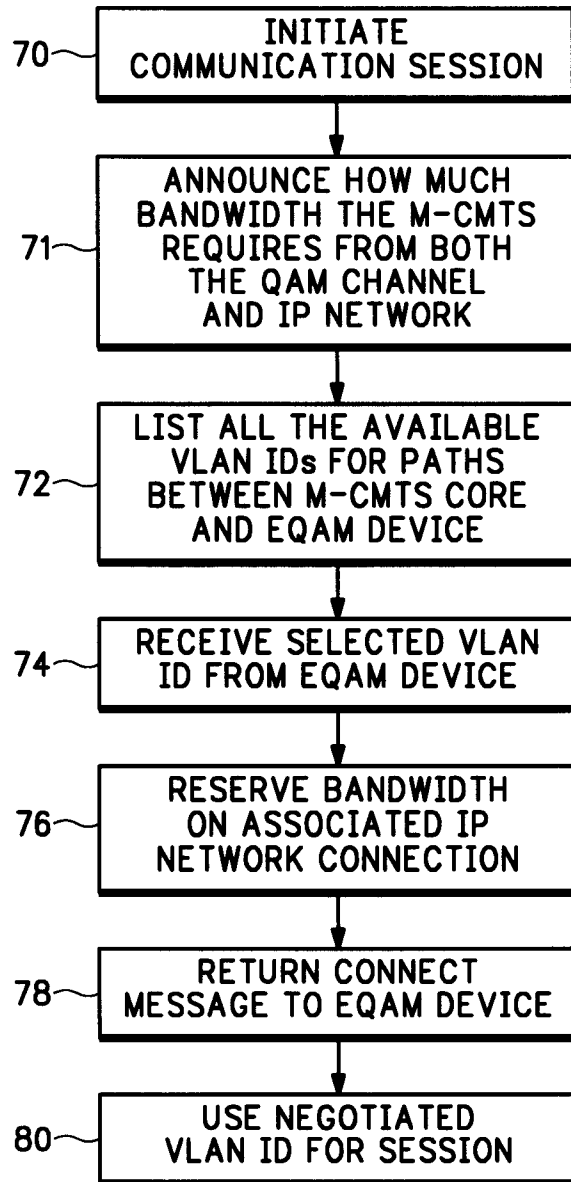


FIG.3

4/7

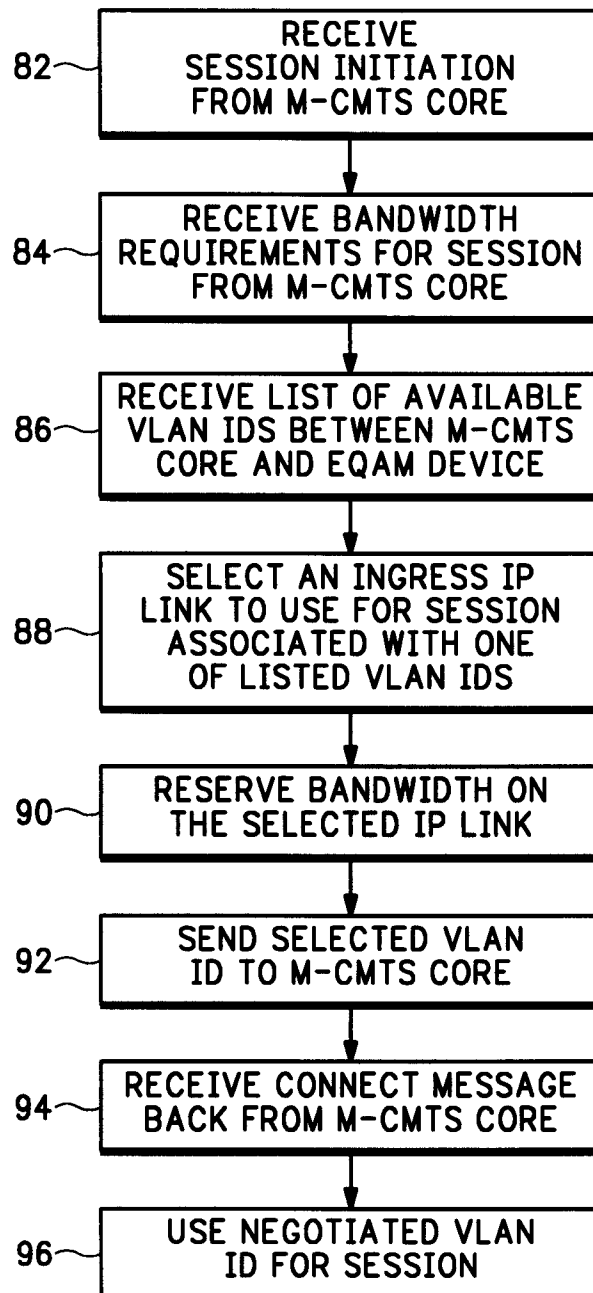


FIG.4

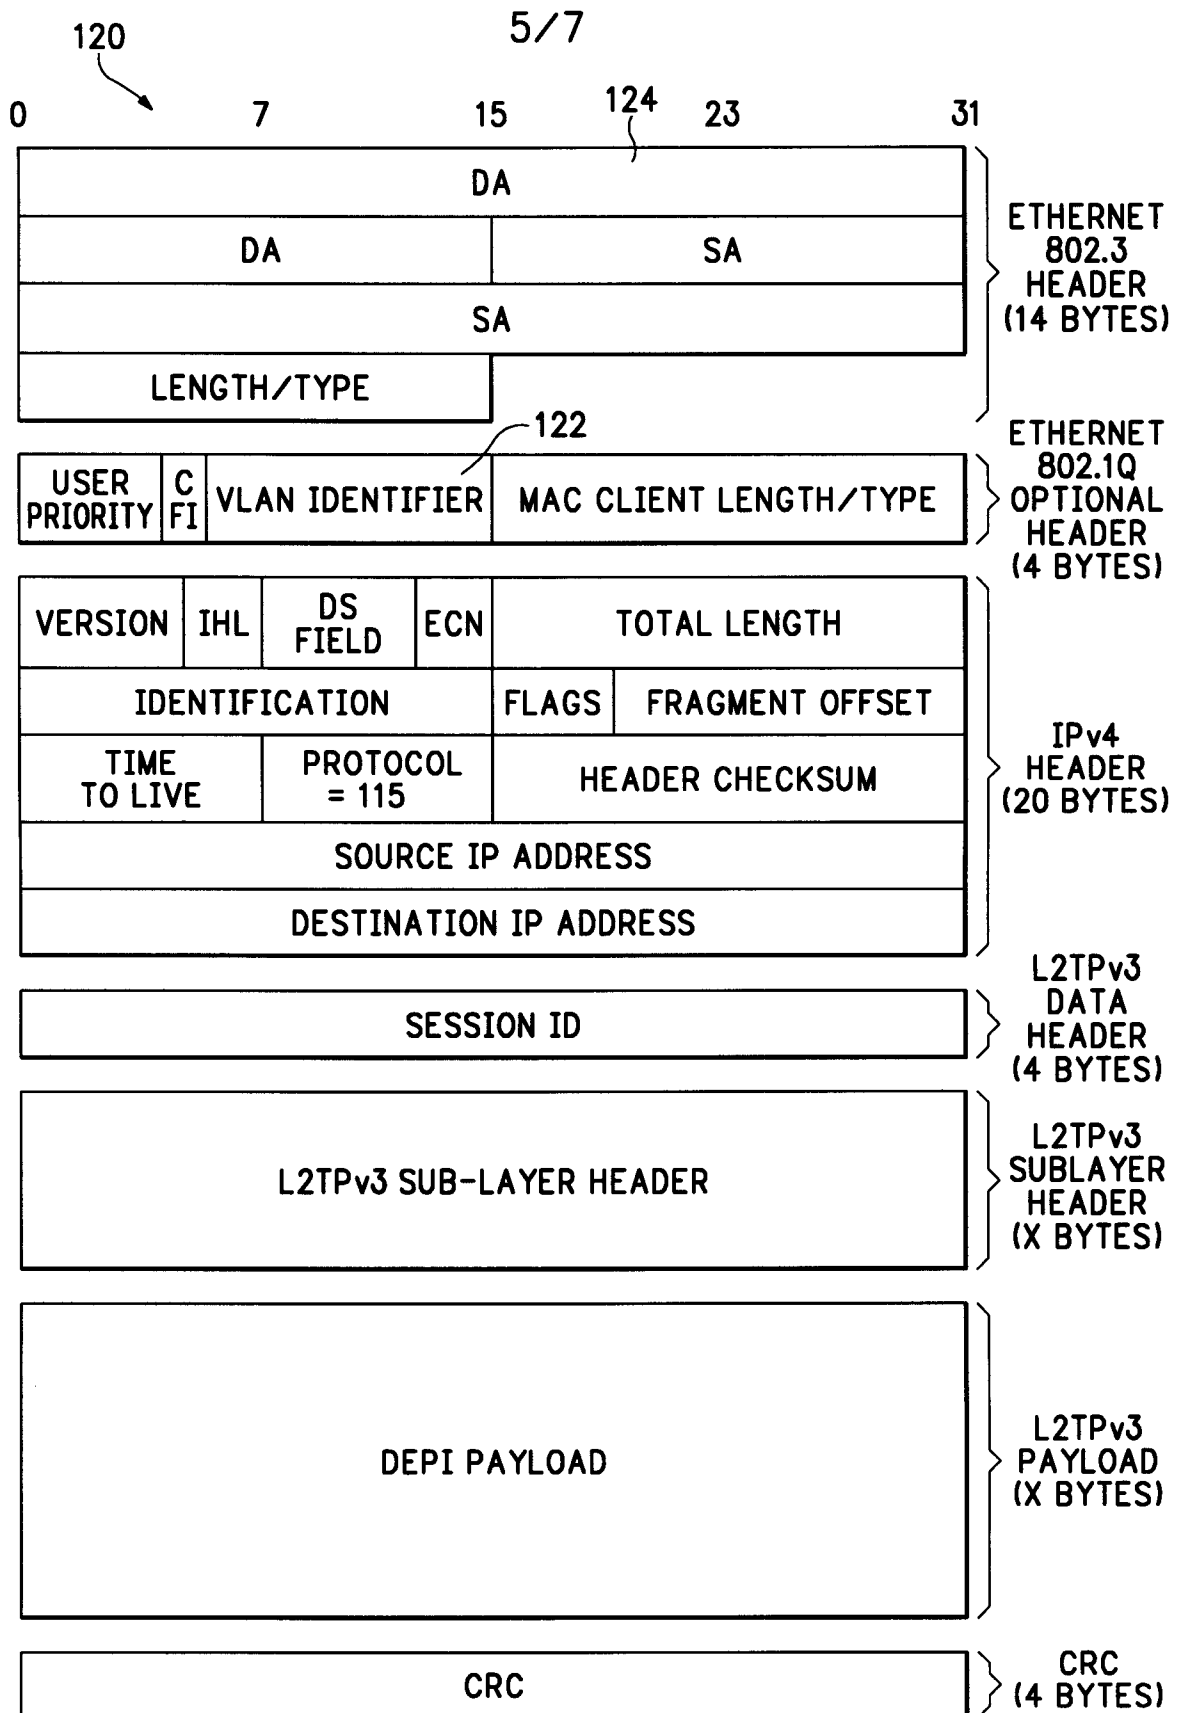


FIG.5




FIG.6




FIG.7