

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
23 December 2004 (23.12.2004)

PCT

(10) International Publication Number
WO 2004/112335 A1

(51) International Patent Classification⁷: H04L 12/66

(21) International Application Number:
PCT/US2004/018784

(22) International Filing Date: 14 June 2004 (14.06.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/477,970 12 June 2003 (12.06.2003) US
60/547,208 24 February 2004 (24.02.2004) US

(71) Applicant (for all designated States except US): CAMI-ANT, INC. [US/US]; 5 Mount Royal Avenue, Marlborough, MA 01752 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): RILEY, Yusun, Kim [US/US]; 630 Hemenway Street, Marlborough, MA 01752 (US). NAIR, Ajit [US/US]; 7 Crest Drive, Westford, MA 01886 (US). MERCURIO, Michael [US/US]; 6 Haverford Street #1, Jamaica Plain, MA 02130 (US).

(74) Agents: DEMSHER, Ronald, R. et al.; Hale and Dorr LLP, 60 State Street, Boston, MA 02109 (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, 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, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, 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, IT, LU, MC, 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:
— with international search report
— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

[Continued on next page]

(54) Title: PCMM APPLICATION MANAGER



(57) Abstract: A method of assigning network resources between network endpoints (100 and 102) includes providing a session initiation request to an application manager. The application manager is logically and physically decoupled from an application server associated with the network endpoints. This request initiates communication between the network endpoints through the set of network resources.

WO 2004/112335 A1



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.

PCMM APPLICATION MANAGER

TECHNICAL FIELD

[0001] This invention relates generally to policy-based network resource management.

BACKGROUND

[0002] In general, a policy server manages network resources according to a set of rules, i.e., a policy. While some policy servers simply carve out network resources for undifferentiated data between network routing or switching elements (e.g., from router A to router B), an application-focused policy server manages network resources to fulfill the needs of applications. For example, in a Voice over Internet Protocol (VoIP) application, two network endpoints may wish to establish a VoIP call with one another via the network. The network endpoints request network resources from the policy server for the VoIP call.

[0003] FIG. 1 shows one way to implement this VoIP example. A first network endpoint 10 wishes to establish a VoIP call to a second network endpoint 12 via an access network 14. The first endpoint 10 specifies the session attributes in its session setup request to the application server. The application server 18 includes the VoIP application that is responsible for coordinating and setting up the session or call between the endpoints. The application server 18 relays the request for resources associated with the VoIP session to a policy server 20 using PacketCable Multimedia (PCMM) protocol (PKT-TR-MM-ARCH-V01-030627, V01, June 27, 2003; and, PKT-SP-MM-I01-030627, June 27, 2003). The application server 18 is synonymous with an application manager 22 that provides the functionality necessary to reserve and manage QoS for the session via the policy server. The policy server in turns controls the network resources to ensure proper treatment is applied to all the media streams associated with the call or session.

[0004] Having the functionality of the application manager 22 located in the application server 18 complicates the design of the application server 18. Since some application servers do not need application managers, an application manager constitutes unnecessary cost and complexity for those installations.

SUMMARY OF THE INVENTION

[0005] In one aspect, a method of assigning a set of network resources between network endpoints using a session initiation request to an application manager for reserving resources between network endpoints. The application manager is logically and physically (or only logically) decoupled from an application server associated with the network endpoints. The method further includes providing a PCMM message from the application manager to a policy server as a result of the session initiation request, wherein the PCMM message includes at least some information embedded in the session initiation request. Suitable protocols other than PCMM may also be used to convey this information between the application manager and the policy server. The method further includes selecting, via the policy server, a set of network resources for creating a path connecting the network endpoints, wherein the policy server selects the set of network resources based on the PCMM message.

[0006] The session initiation request may be communicated via Session Initiation Protocol (SIP – see J. Rosenberg, et. al., "SIP: Session Initiation Protocol", RFC 3261, June 2002), using the Session Description Protocol (SDP – see Handley, M. and V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998) field in the SIP message to perform the necessary QoS reservation of resources in the network. The QoS information may be included in the session description protocol fields within the session initiation request. The application manager may be inserted between one of the network endpoints and an application server. The application manager transparently relays the session initiation request between the network endpoint and the application server, without requiring functional modifications of the network endpoints or the application server.

[0007] In one embodiment, the method combines high-level QoS information from the session initiation request, with understanding of underlying network technologies, to produce detailed QoS information for use in the PCMM message. In another embodiment, the method embeds QoS information either via SDP or via RSVP FlowSpecs (as defined in IETF documents Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997. Wroclawski, J., "Specification of the Controlled-Load Network Element Service", RFC 2211, September 1997. Shenker, S.,

Partridge, C., Guerin, R., "Specification of Guaranteed Quality of Service", RFC 2212, September 1997) in an eXtensible Markup Language message (XML), and uses the XML message to communicate the session initiation request to an application server, which subsequently transfers the QoS information to the application manager in an XML message. The method may alternatively use HTTP to communicate the XML message to an application server, which subsequently transfers the QoS information to the applications manager.

[0008] In one embodiment, the application server remains stateless through the use of a unique session identifier for communicating with the application manager, and the association between the session and the resources are maintained in the decoupled application manager. The session identifier is constructed from the SIP signaling messages between the end points and the application server.

[0009] In one embodiment, the method includes selecting RSVP FlowSpec parameters based on SDP media attributes, and providing the ability to configure FlowSpec parameter selection behavior for the purposes of performing resource reservation in a PCMM framework. The method further includes ensuring interoperability with DOCSIS Service Flow QoS parameters at the CMTS for providing QoS by normalizing the selected FlowSpec parameters across different CMTS vendor implementations. The method also includes ensuring interoperability between network endpoints that implement different optional features of SDP media attributes.

[0010] In one embodiment, the method includes providing configurable settings per direction (upstream and downstream), using these setting in FlowSpec parameter selection, along with the default values for each setting. The configurable settings and associated default values include:

Delay, 100 milliseconds (downstream and upstream)

BandwidthAdjust, 1.125 (downstream)

BandwidthAdjust, 1.0 (upstream)

RateFactor, 4.0 (downstream and upstream)

MinPacketSize, 640 (downstream and upstream)

MaxPacketSize, 1500 (downstream and upstream)

SlackTerm, 800 (downstream and upstream)

RateLimit, 96000 (upstream)

RateLimit, 0 (downstream)

PollIntervalLimit, 6000 microseconds (upstream)

PollIntervalSet, empty (upstream)

[0011] When the "ptime" media attribute (packetization time) is not given in the SDP, the method includes selecting RSVP FlowSpec parameters such that sufficient resources are reserved for all packetization times mandated by the PacketCable Audio/Video Codecs Specification.

[0012] When multiple codecs are negotiated in the SDP fields, the method includes selecting RSVP FlowSpec parameters such that sufficient resources are reserved for all packetization times mandated by the PacketCable Audio/Video Codecs Specification.

[0013] In one embodiment, the method includes using the DOCSIS QoS MIB to accurately determine ptime (packetization time) when a ptime media attribute is unknown, and modifying reserved resources if the ptime media attribute is subsequently determined. In another embodiment, the method includes using the DOCSIS QoS MIB to accurately determine ptime (packetization time) when negotiating multiple codecs to determine a particular codec being used, and modifying reserved resources if the particular codec is subsequently determined.

[0014] In one embodiment, the application manager learns the contact IP address and port number of one or both endpoints from the REGISTER messages so that the application manager can forward INVITE messages to the endpoint. In another embodiment, the application manager queries the application server to obtain the contact IP address and port number of one or both endpoints so that the application manager can forward INVITE messages to the endpoint.

[0015] In another aspect, a system for assigning a set of network resources between network endpoints includes an application manager for (i) receiving a session initiation request for initiating communication between the network endpoints through the set of network resources, and (ii) generating a PCMM message including at least some information embedded in the session initiation request. The communication between the

network endpoints employs an application residing on an application server. The system further includes a policy server for receiving the PCMM message, and for selecting the set of network resources for creating a path connecting the network endpoints, wherein the policy server selects the network resources based on the PCMM message.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] FIG. 1 shows a prior art architecture for establishing a VoIP link between two network endpoints; and,

[0017] FIG. 2 shows the network architecture of the described embodiment;

[0018] FIG. 3 shows a sample call flow for two network endpoints in the same network domain; and

[0019] FIG. 4 shows a sample call flow for two network endpoints in different network domains. .

DETAILED DESCRIPTION

[0020] The described embodiment is a network architecture that uses an application server, an application manager, and a policy server to implement application-specific policies for network resource allocation. This embodiment abstracts the application manager away from the application server, so that the application server does not have to implement complex network specific signaling interfaces such as PCMM in order to perform admission control or resource requests over the network. This embodiment describes two endpoints implementing a high quality video teleconference session, although these concepts can apply also to other applications, such as video streaming, Voice over Internet Protocol (VoIP) communications, and networked gaming. The abstraction of the application manager from the application server mandates a new signaling interface between the application manager and the application server. This interface is now a much simpler interface than PCMM, and can use either SIP signaling, or another protocol such as eXtensible Markup Language (XML), Real Time Streaming

Protocol (RTSP) or HyperText Transfer Protocol (HTTP) to communicate between the application server and the application manager. The now abstracted application manager can interpret the much simpler 'language' and translate it into the more complex and network specific language in order to communicate with the Policy Server for reserving and managing network resources.

[0021] The CableLabs PacketCable Multimedia (PCMM) protocol specification (PKT-TR-MM-ARCH-V01-030627, V01, June 27, 2003), defines an application manager function embedded in an application server. The application manager function is responsible for communicating with the policy server to implement policy decisions, e.g., admission control and Quality of Service (QoS), through a PCMM signaling interface. By abstracting the application manager from the applications themselves (i.e., away from the application server), this embodiment relieves the application server from needing to know about the PCMM architecture and its associated protocols. In this embodiment, the application server can therefore interface to the PCMM environment in a much more generic way (i.e., via the abstracted applications manager), while delegating the task of maintaining state (regarding gate tokens) to the application manager. Now, the application manager can either be standalone, or be embedded in the Policy Server; this is an optimization to minimize the number of network elements involved in the delivery of the service.

[0022] FIG. 2 shows one particular embodiment of the network architecture, including a first network endpoint 100, a second network endpoint 102, an access network 104, an application server 106, a policy server 108 and an application manager 110. This embodiment has a broken-line box around the policy server 108 and the application manager 110 to show that the functionality of both of these components resides on the same physical platform, although as we will describe later, the application manager 110 can reside on a physical platform separate from the policy server 108 (i.e., in a stand-alone configuration). In this embodiment, the first network endpoint 102 uses a SIP exchange to request a video teleconference session through the access network 106 to the second network endpoint 104. This video teleconference session uses a video teleconference application residing on the application server 106.

[0023] The application manager 110 sits in the SIP signaling path 112. Although the SIP connections 112 between the endpoints 100, 102 and the application server 106 are shown in FIG. 2 as direct connections for simplicity, it should be understood that all connections from the endpoints pass through the access network 104.

[0024] Every SIP session begins with a three-way handshake between two SIP endpoints. The endpoint that is initiating the session (the “caller”) first sends an INVITE message to the receiving endpoint (the “callee”). The callee responds with an OK message back to the caller. The caller then responds with an ACK message. In FIG. 2, the two SIP endpoints are the first network endpoint 100 and the second network endpoint 102, where the first network endpoint 100 is the caller, and the second network endpoint 102 is the callee. The application server 106 functions in this embodiment as a SIP proxy server, receiving and forwarding SIP messages from and two the endpoints 100 and 102.

[0025] While the application manager 110 sits in the SIP signaling path, the application manager 110 transparently passes the SIP handshake exchanges, so that the SIP endpoints or the proxy server (application manager 106) do not perceive the application manager 110 in the SIP signaling path. As the application manager 110 passes the first SIP INVITE message from the first network endpoint 100 to the application server 106, the application manager 110 extracts the Session Description Protocol (SDP) fields from the INVITE message for information about the first network endpoint 100. When the application manager 110 receives the OK message from the application server 106, the application manager extracts the SDP fields from the OK message for information about the application server. Before forwarding the OK message to the first network endpoint 100, the application manager 110 determines whether to dynamically create Quality of Service (QoS) for only the first network endpoint 100 or for both the first network endpoint 100 and the second network endpoint 102. Based on this determination, the application manager sends Gate Set messages via the PCMM link 114 to the policy server 108. Details regarding how the application manager translates the SDP field in the SIP exchanges into parameters for the Set Gate messages are described later for both voice and video communications.

[0026] For every media type listed in the SDP, the application manager 110 creates two gates on the policy server: one in the “upstream” direction, and one in the

“downstream” direction. Thus, for the video teleconference session in this embodiment, the application manager 110 creates four gates per endpoint. One gate is for upstream audio, one gate for upstream video, one gate for downstream audio, and one gate for downstream video.

[0027] The application manager 110 determines whether to create QoS by examining at the Via header in the INVITE and OK messages. Based on the number of entries in the Via header, the application manager 110 can determine whether the first network endpoint 100 and the second network endpoint 102 are both using the same applications manager as the first hop. In a multiple-domain environment, it is possible that the caller is using one application manager as its first hop, and the callee is using another application manager as its first hop. In such a multiple-domain case, each application manager is only responsible for creating QoS for the network endpoint that uses that application manager for its first hop. To relate INVITE messages and OK messages to corresponding PCMM messages and related gates, the application manager uses the CallID header in the SIP messages. Using this technique, a decoupled application manager can transparently exist in a network requiring no modifications to the endpoints or to the application server functionality (i.e., no changes to the actual design or functional implementation of these components).

[0028] Once the application manager 110 has set up a dialog between the endpoints, and the QoS is established for the endpoints, one of the endpoints may send another INVITE message altering its QoS requirements. The application manager 110 handles such re-INVITE messages by altering the gates created earlier for that dialog.

[0029] When the dialog has to be torn down, the endpoint initiating the tear down sends a BYE message. After passing along the BYE message, the application manager 110 clears the QoS that it created for that dialog by sending PCMM Gate Delete messages to the Policy Server 108. The application manager 110 maintains a database for storing information regarding which gates were created for a dialog, so the application manager can identify which gates need to be deleted when the dialog is being torn down.

[0030] FIG. 3 shows a sample call flow for two network endpoints in the same network domain, and FIG. 4 shows a sample call flow for two network endpoints in

different network domains. The call flow for each figure includes the flow of SIP signals and PCMM messages between endpoints in two domains. For FIG. 4, each domain has its own application manager 110 and policy server 108, but a central proxy server connects the two domains.

[0031] In another embodiment, the application manager requests QoS via a two-phased approach. When the application manager receives an INVITE message from the caller, the application manager will not only learn the caller's requirements, but also send GateSet messages to the policy server. These GateSet messages will only "reserve" the QoS for the caller and/or the callee, i.e., it does not yet "commit" the necessary gates. As the call setup progresses and the application manager receives the corresponding OK message from the callee, the application manager sends a Gate Set messages for the callee (if needed) and also send Gate Set messages for the caller. This time the application manager "commits" the caller and callee's gates, rather than only reserving them.

[0032] This embodiment provides better control of the gates on the policy server. In other words, by first "reserving" gates, the application manager does not consume valuable resources on the PEP (policy enforcement point). The application manager only commits these resources when the application manager receives the OK, at which point the application manager knows that the callee is available for completing the session establishment.

[0033] In another embodiment, the application manager learns contact information from REGISTER messages coming from endpoints. This feature allows the application manager to learn the IP address and port number of the endpoints for the purposes of forwarding an INVITE message to an endpoint. In cases where the Registration server cannot handle multiple hops between the endpoint and itself, the Registration server will only forward INVITE messages to the application manager. The application manager then uses this information to deliver the INVITE to the correct endpoint.

[0034] In another embodiment, the application manager obtains contact information from the application server by querying it with a REGISTER message containing an empty Contact header. This feature allows the application manager to learn the IP address and port number of the endpoints for the purposes of forwarding an INVITE message to

an endpoint. In cases where the Registration server cannot handle multiple hops between the endpoint and itself, the Registration server will only forward INVITE messages to the application manager. The application manager then queries the application server to obtain the contact information and deliver the INVITE to the correct endpoint.

[0035] Although the described embodiment performs session initiation between the endpoint and the application server is SIP, other even simpler and more flexible protocols can also be used for communication between the endpoint and the application server. For example, one embodiment uses eXtensible Markup Language (XML) instead of SIP for an endpoint to communicate session setup embedded with QoS information to the application server. Other embodiments use Real Time Streaming Protocol (RTSP) or HyperText Transfer Protocol (HTTP) for session initiation and resource requests. The QoS information can be in embedded or in the format of SDP or IETF RSVP FlowSpec

[0036] Although the described embodiment performs session initiation with SIP, and the interface between the application server and the application manager is also SIP, other even simpler protocols can also be used for communication between the application server and the application manager. For example, one embodiment uses eXtensible Markup Language (XML) for a application server to communicate resource requests to the application manager. Other embodiments use Real Time Streaming Protocol (RTSP) or HyperText Transfer Protocol (HTTP) for session initiation and resource requests. The QoS information can be in embedded or in the format of SDP or IETF RSVP FlowSpec (as defined in IETF documents Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997. Wroclawski, J., "Specification of the Controlled-Load Network Element Service", RFC 2211, September 1997. Shenker, S., Partridge, C., Guerin, R., "Specification of Guaranteed Quality of Service", RFC 2212, September 1997)

[0037] When using XML, the application server inserts the SDP field into the XML message sent to the application manager. This application server extracts the SDP field out of the SIP messages used to communicate with the end points, and puts it in a XML message and sends it to the application manager, which then performs the necessary

translations to appropriate PCMM messages. In addition, the application server communicates to the application manager the context for the session, through a unique session identifier. This session identifier can be constructed from the SIP signaling messages between the end points and the application server. Because the application server only needs to communicate with the application manager as a result of messaging activity between the end points and the application server, the application server can construct the session ID during this process to issue messages to the application manager. This enables the application server to remain "stateless." The application server does not have to keep state of the application session and the resource mapping, thus simplifying the implementation and functionality of an application server. When resources need to be relinquished, the Application server uses the session identifier constructed from the SIP signaling messages to communicate to the application manager. The application manager maintains state of the session and the resource association, and thus is "stateful."

[0038] As with the described embodiment, the application manager may be embedded in the policy server, or it may be resident on a stand-alone platform. In either case, the application manager transforms the QoS information in the SDP into more detailed QoS by combining the SDP with knowledge of the access network, and forwards the more detailed QoS information to the policy server via a PCMM message.

VIDEO SDP TRANSLATION

[0039] The paragraphs that follow describe the details of translation from SDP fields to PCMM data for digital video traffic. The PCMM architecture uses the IETF's RSVP FlowSpec parameters to communicate Quality of Service (QoS) requirements between the Policy Server and the Cable Modem Termination System (CMTS). The following paragraphs describe some ways to select FlowSpec parameters appropriately for use in PCMM for the purposes of enabling high quality video teleconference sessions for SIP User Agents (i.e., network endpoints) based on SDP parameters.

[0040] Three challenges in selecting FlowSpec parameters include:

1. The SDP specifications continue to evolve and many parameters (e.g., media attributes) are not well established; most are optional. There exist SIP/SDP implementations that do not include all the optional media attributes that would result in optimal FlowSpec parameter selection.
2. Compressed digital video traffic is Variable Bit Rate (VBR), meaning packets vary in size and arrive at random intervals.
3. QoS as specified by DOCSIS and provided by the CMTS has many optional and CMTS implementation-dependent features. (DOCSIS refers to the set of Data-Over-Cable Service Interface Specifications, which defines how to transmit data over cable networks in a standard fashion. For example, DOCSIS RFI MAC protocol defines how to connect to cable modems over the access network. RFI is the DOCSIS Radio Frequency Interface specification defining MAC and Physical Layer interfaces between CMTS and CM network elements.) This means that while one CMTS vendor may decide to implement a particular feature, another may not, or may decide to implement it differently. Interoperability issues between CMTS, Cable Modem, and other equipment further complicate the process of selecting optimal parameters to ensure QoS.

The following paragraphs set forth how the described embodiment overcomes these challenges.

[0041] VBR Video Traffic Characteristics – Providing QoS for the transmission of real-time video over digital networks is particularly difficult because of the VBR nature of compressed digital video traffic. Video compression codecs (e.g., MPEG, H.263, H.261) use temporal compression techniques, in which the compression ratio changes from frame to frame and is dependent on the change or movement across subsequent frames from a periodic snapshot of the scene. This leads to high peak burst rates and generally high but variable peak-to-mean ratios. These characteristics make it more difficult to efficiently estimate resource utilization and make accurate reservations than for constant bit rate traffic such as voice.

[0042] Without an accurate estimate of required resource reservations, either too many resources are reserved, leading to low utilization and a high cost of service, or not enough resources are reserved, leading to higher loss and higher delay. Either case results in poor quality of service. Providing high quality multimedia services at the lowest cost

possible requires an efficient and accurate method to predict resources required for video transmission on a per-session basis.

[0043] DOCSIS Service Flow QoS Parameters – The following few paragraphs describe DOCSIS QoS parameters and their relationship to the IETF's FlowSpec parameters for transmission of digital video. The FlowSpec to DOCSIS parameter mapping is specified in the PacketCable Multimedia Specification (PKT-SP-MM-I01-030627, June 27, 2003). The DOCSIS Service Flow QoS Parameters applicable for video transmission utilize the DOCSIS Real-Time Polling Service (rtPS) in the Upstream combined with Minimum Reserved Rate and Maximum Sustained Rate in both the Upstream and Downstream. The Real-Time Polling Service works well for bursty VBR traffic such as digital video because it provides a cable modem with contention-free, timely request opportunities without the need to allocate upstream bandwidth until the modem (also referred to herein as the network endpoint) has data to send.

[0044] The first two QoS parameters, Nominal Polling Interval and Tolerate Polling Jitter, are only used in the upstream. The last parameter, Maximum Downstream Latency, is only used in the downstream. The remaining parameters are used in both the upstream and downstream.

[0045] Nominal Polling Interval – The DOCSIS Nominal Polling Interval specifies the time interval, in microseconds, in which the cable modem will receive periodic unicast (contention free) request opportunities. For VBR video traffic, it is important that the Nominal Polling Interval be short enough that the modem receives request opportunities and subsequent grants often, enabling it to transmit at its peak rate during burst periods. The only way to ensure the modem enough request opportunities during peak bursts is to set the polling interval as if a flow were sending traffic at its peak rate continuously. If the polling interval only allows for the average rate, packets will be dropped or delayed during the burst period.

[0046] For example, consider a video source with average rate of 256 Kbps (32000 bytes/second) and average packet size of 640 bytes (50 packets per second on average). Also consider that during burst periods, it will peak at a rate of 768 Kbps (150 packets per second). If the modem were only allowed 50 request opportunities per second, during

burst periods it wouldn't receive request opportunities fast enough, and packet loss will occur. In order to minimize packet loss during peak bursts, the modem must receive 150 request opportunities per second.

[0047] It is true that during non-burst periods, where the Service Flow is sending data at a lower rate, the unused request opportunities will be wasted. However, the amount of resources used to schedule and transmit a request opportunity downstream by the CMTS is insignificant to the amount of data being transmitted upstream, as well as the amount of time required for the modem to use a contention request (because of exponential back-off and the possibility for collision). The DOCSIS Real-Time Polling Service was designed with this in mind, which is why it is ideal for bursty VBR real-time traffic such as video.

[0048] Because the Polling Interval is derived from the RSpec's Reserved Rate R for rtPS, the Reserved Rate R will typically be set equal to Peak Rate p, and both R and p will be greater than the Bucket Rate r. This is consistent with the IETF usage of these parameters for Guaranteed Service.

[0049] Tolerated Polling Jitter – The DOCSIS Tolerated Polling Jitter parameter specifies the amount of time, in microseconds, a unicast request opportunity for Real-Time Polling may be delayed. For real-time video sessions, this value should be relatively small. PCMM specifies a default of 800 microseconds and is conveyed to the CMTS via the IETF RSpec Slack Term S parameter.

[0050] Maximum Sustained Traffic Rate – The DOCSIS Maximum Sustained Traffic Rate parameter specifies the maximum sustainable rate over time (i.e., the average rate) in bits per second of DOCSIS MAC layer datagrams, and includes all MAC layer header and CRC overhead. In DOCSIS, the upstream overhead is slightly larger due to extended headers in the upstream. For VBR video traffic, this should be equal to the video's maximum sustainable rate after accounting for MAC layer overhead.

[0051] Maximum Traffic Burst – The DOCSIS Maximum Traffic Burst parameter specifies the largest burst size in MAC layer bytes, the flow may transmit at its peak rate. For VBR video traffic, it is important that the Maximum Traffic Burst is set high enough to allow the flow to burst at its peak rate. The Maximum Traffic Burst is set equal to the

IETF Token Bucket Depth b after accounting for MAC layer overhead. There is the additional requirement that the Traffic Burst must not be less than the DOCSIS specified minimum value of 1522 bytes.

[0052] Minimum Reserved Traffic Rate – The DOCSIS Minimum Reserved Traffic Rate specifies a guaranteed minimum rate the flow will receive and closely resembles the IETF RSpec Reserved Rate R after accounting for MAC layer overhead. However, for the purposes of PCMM, the conversion from RSVP FlowSpec to DOCSIS specifies the use of the TSpec Bucket Rate r parameter. This conversion works well for VBR video traffic, because the RSpec Reserved Rate R is typically the same as (or greater than) the TSpec Bucket Rate r . That is to say, at the DOCSIS MAC layer, the reservation is made using the average rate of the video session. Test results show that setting the Minimum Reserved Traffic Rate using the bandwidth as advertised in the SDP media attributes works well in the upstream, however the downstream experiences enough packet loss to significantly affect the video quality at some bit rates. The described embodiment provides for a small, selectable percentage increase in the bandwidth to overcome this packet loss.

[0053] Assumed Minimum Reserved Rate Packet Size – The DOCSIS Assumed Minimum Reserved Rate Packet Size, closely resembles the IETF TSpec Minimum Policed Unit m . Like the IETF parameter it is used to estimate per packet overhead. If the Service Flows sends packets smaller than this size, the packets will count as being of this size for the purposes of maintaining the reserved rate. For VBR video traffic where there is a large variance of packet size, our implementation will use a configured fixed packet size parameter. Note that although DOCSIS specifies this parameter for both upstream and downstream Service Flows, PCMM only specifies it for downstream. The DOCSIS default for this parameter is not defined, and is CMTS-implementation dependent.

[0054] Maximum Downstream Latency – The DOCSIS Maximum Downstream Latency parameter specifies the maximum latency (in microseconds) a packet may experience between being received at the CMTS Network Interface and being transmitted downstream on the HFC. This parameter only applies to flows that have a downstream

Minimum Reserved Rate configured and the flow is not exceeding this rate. For PCMM, this parameter is set from the RSpec Slack Term S.

[0055] Mapping SDP Video Bandwidth to RSVP FlowSpec – If the SDP does not include a bandwidth media attribute parameter, all FlowSpec parameters are selected using preconfigured defaults. Otherwise, the application manager 110 uses the SDP bandwidth media attribute as input to map SDP to FlowSpec parameters. Other settings that affect the selection of RSVP FlowSpec parameters are configured and described in this section. In general, most of the following parameters the application manager uses are configured for each direction (i.e., upstream and downstream).

[0056] Delay (milliseconds) – This is the maximum worst-case queuing delay for video, and corresponds to the maximum burst size (the amount of time a flow may burst at its peak rate). An initial default value of 100 milliseconds is suggested.

[0057] BandwidthAdjust (floating point multiplier) – This is a “fudge factor” used to multiply original SDP requested bandwidth to reserve more bandwidth for clients that underestimate video bandwidth or to solve other interoperability issues. An initial default value of 1.0 is suggested for upstream and 1.125 for downstream. During initial testing it was observed that at all rates below 512 Kbps, additional bandwidth was required in the downstream for video quality to be acceptable, due mostly to excessive packet loss and delayed packets

[0058] RateFactor (floating point multiplier) – This is used to compute peak and reserved rate given a maximum sustained throughput (bandwidth). An initial default value for RateFactor is 4.0.

[0059] MinPacketSize (bytes) – This is the estimated IP packet size used for minimum policed unit and other calculations requiring packet size. An initial (conservative) default value for MinPacketSize is 640.

[0060] MaxPacketSize (bytes) – This is the maximum IP packet size expected to conform to the IETF TSpec. An initial default value for MaxPacketSize is 1500.

[0061] SlackTerm (microseconds) – This is the value used for IETF RSpec Slack Term S. For PCMM it corresponds to the DOCSIS Upstream Tolerated Polling Jitter in

the upstream and the DOCSIS Downstream Latency in the downstream. An initial default value for SlackTerm is 800 microseconds, which is the default in PCMM.

[0062] RateLimit (bytes per second) – This is the upper limit on the Reserved Rate R and Peak Rate p, parameters specified in bytes per second. An initial default value for RateLimit is 96000 for the upstream, and 0 for the downstream (a value of 0 means there is no limit).

[0063] The following parameters are specified for upstream only.

[0064] PollIntervalLimit (microseconds) – This is the minimum polling interval to request from the CMTS. An initial default value for PollIntervalLimit is 6000 microseconds (a value of 0 means there is no limit).

[0065] PollIntervalSet is a list of acceptable polling intervals at the CMTS. See description below. The suggested initial default set is empty.

[0066] The following are relevant TSpec Parameters from the PCMM specification.

[0067] Bucket Depth (b) – The TSpec Token Bucket Depth combined with the TSpec Token Bucket Rate bounds the maximum burst size (at the peak rate) of a flow, as well as the maximum worst case queuing delay the flow will experience. As such, both the maximum worst case queuing delay and the peak rate of the flow must be considered.

[0068] The TSpec Bucket Depth b in bytes is given by:

$$b = (\textit{Delay} * p) / 1000$$

where *Delay* is the configured maximum queuing delay in milliseconds, and *p* is the TSpec Peak Rate in bytes per second.

Example:

$$\textit{Delay} = 100 \text{ msec}$$

$$p = 32000$$

$$b = (100 * 32000) / 1000 = 3200$$

[0069] Bucket Rate (r) – The TSpec Token Bucket Rate specifies the average rate over time to which the IP flow will conform. The TSpec Bucket Rate r in bytes per second is:

$$r = (\textit{Bandwidth} * \textit{BandwidthAdjust}) / 8$$

where *Bandwidth* is the advertised bandwidth from the SDP in bits per second, and *BandwidthAdjust* is the configured adjustment multiplier.

Example:

$$\begin{aligned} \textit{Bandwidth} &= 512000 \\ \textit{BandwidthAdjust} &= 1.0 \\ r &= 512000 * 1.0 / 8 = 64000 \end{aligned}$$

[0070] Maximum Datagram Size (M) – The TSpec Maximum Datagram Size M in bytes is set to the configured *MaxPacketSize*.

[0071] Minimum Policed Unit (m) – The TSpec Minimum Policed Unit m in bytes is set to the configured *MinPacketSize*.

[0072] Peak Rate (p) – The TSpec Peak Rate is the maximum rate the flow is expected to burst data onto the network. The TSpec Peak Rate p in bytes per second is always equal to RSpec Reserved Rate. This statement is an indirect result of using Reserved Rate R for the DOCSIS rtPS polling interval. The Polling Interval must be set using the Peak Rate, which effectively means the Peak Rate must be equal to the Reserved Rate.

[0073] The following are relevant RSpec Parameters from in the PCMM specification.

[0074] Reserved Rate (R) – The RSpec Reserved Rate is the guaranteed rate the flow is entitled to, provided that its traffic conforms to the TSpec. In PCMM, the Reserved Rate is used to determine the DOCSIS Nominal Polling Interval. In the described embodiment, the derivation of the Reserved Rate is controlled by configured limits. This is necessary to solve interoperability issues between the CMTS and the Policy Server/application manager. Specifically, hard limits on the maximum value for the Reserved Rate R parameter (*RateLimit*), as well a limit on the derived DOCSIS Upstream Nominal Polling Interval (*PollIntervalLimit*) are configured.

RSpec Reserved Rate R in bytes per second is:

$$R = r * \textit{RateFactor}$$

$$\textit{RateLimit} = 0$$

$$R = \min(r * RateFactor, RateLimit) \quad RateLimit > 0$$

where r is the TSpec Token Bucket Rate, $RateFactor$ is the configured rate multiplier, and $RateLimit$ is the configured upper limit on the Reserved Rate parameter. If $RateLimit$ is zero, then there is no upper limit on the Reserved Rate.

[0075] In the upstream there are two ways to configure the Nominal Polling Interval, which is derived from the Reserved Rate: either setting a limit on the Polling Interval using $PollIntervalLimit$ setting, or only allowing Polling Intervals in a predetermined set of intervals by configuring $PollIntervalSet$. The $PollIntervalSet$ is a list of acceptable Nominal Polling Intervals supported at the CMTS. If this list is non-empty, then the second method will be used, otherwise the first method is used.

[0076] First Method: Using $PollIntervalLimit$ – In the upstream, there is the additional requirement that the Minimum Policed Unit m (bytes) divided by the Reserved Rate R (bytes/second) MUST NOT be less than the configured $PollIntervalLimit$ (microseconds). That is to say, the following inequality must always hold true when the configured $PollIntervalLimit > 0$:

$$m/R * 10^6 \geq PollIntervalLimit$$

For the Upstream, Reserved Rate R in bytes per second is:

$$R = \text{same as equations above for } R \quad PollIntervalLimit = 0$$

$$R = \min(r * RateFactor, m * 10^6 / PollIntervalLimit) \quad PollIntervalLimit > 0$$

Example (limits on Rate and Polling Interval):

$$r = 64000$$

$$m = 640$$

$$RateFactor = 4.0$$

$$RateLimit = 96000$$

$$PollIntervalLimit = 10000$$

$$R = \min(64000 * 4.0, 96000) \quad PollIntervalLimit = 0 \\ = \min(256000, 96000)$$

$$R = \min(96000, 640 * 10^6 / 10000) \quad PollIntervalLimit > 0 = 10000 \\ = \min(96000, 80000) \\ = 80000$$

[0077] Second method: Using PollIntervalSet – When the *PollIntervalSet* is non-empty, the derived Nominal Polling Interval must be selected from the set of configured values in microseconds. When using this method, the setting, *RateLimit* and *PollIntervalLimit* have no meaning.

[0078] Algorithm to choose R:

1. Compute the possible values of R ($R_1 \dots R_n$) based on m and the set of configured polling intervals, $I_1 - I_n$.

In general, $R = (m * 10^6) / \text{NominalPollingInterval}$

$$R_i = (m * 10^6) / I_i \quad \text{for all } i [1..n]$$

2. Select one of the values for R above ($R_1 \dots R_n$) that is closest to, but not greater than $r * \text{RateFactor}$.

Example (PollIntervalSet):

$$r = 12000$$

$$m = 640$$

$$\text{RateFactor} = 4.0$$

$$\text{PollIntervalSet} = \{10000, 20000, 30000\}$$

$$R_1 = (640 * 10^6) / 10000 = 64000$$

$$R_2 = (640 * 10^6) / 20000 = 32000$$

$$R_3 = (640 * 10^6) / 30000 = 21333.3333\dots$$

$$r * \text{RateFactor} = 48000$$

$$R = 32000$$

[0079] Slack Term (s) – The Slack Term S in microseconds is set from the configured *SlackTerm* parameter.

AUDIO SDP TRANSLATION

[0080] The paragraphs that follow describe the details of translation from SDP fields to PCMM data for digital audio traffic. This description supplements the description for video traffic.

The audio codecs at the network endpoints, as well as the inter-packet arrival time, determine the FlowSpec parameters for audio data. The PacketCable Codec Specification specifies three inter-packet arrival times (10, 20, 30 milliseconds). The PCMM Audio/Video Codecs Specification (PKT-SP-CODEC-I05-0401.13) also specifies each of the FlowSpec parameters for all audio codecs in the specification for each of the three packet times. The supported audio codecs include:

- G.729E
- G.729A
- G.728
- G.726.40
- G.726.32
- G.726.24
- G.726.16
- G.711 (PCM encoding) a-Law
- G.711 (PCM encoding) μ -Law

[0081] When a network endpoint in a SIP exchange specifies an audio codec using the SDP (Session Description Protocol), it should also use the optional "p-time" packetization time media attribute.

[0082] Two issues exist with respect to determining the correct FlowSpec parameters to use for audio codecs by SIP/SDP sessions. The first issue is that the packet time (i.e., "p-time") SDP media attribute is optional, and is therefore not used by all network endpoints. When the packet time is not known, the system can select FlowSpec parameters that will work at all three packet times: 10, 20 or 30 milliseconds. This method will always work, but will not always provide the efficient resource reservation. The smallest packet interval, combined with the largest packet size is used to accommodate all combinations of packet size and packet rate. When the packet time is not known, the system can also use a configured default packet time. The PacketCable default packet time is specified at 20 milliseconds. This or any configured default may be used when the behavior of the network endpoints in use are known.

[0083] The second issue with respect to determining the correct FlowSpec parameters is that there is no way to know, in general, which codec a SIP network endpoint will use when multiple codecs are available.

[0084] When a SIP network endpoint has multiple audio codecs available for use, and includes multiple codecs for the same media type (in this case, audio), it may use any

of the available codecs negotiated between two network endpoints. Either endpoint must be prepared to receive any of the negotiated codecs, and there is no signaling to let the other end point know which codec is being used until the RTP (Real Time Protocol) media stream starts.

[0085] The described embodiment uses a combination of the following three techniques for determining which of multiple specified codecs on which to base resource reservation decisions.

[0086] First Technique – When the SDP field specifies multiple codecs for the same media type, the application manager initially bases resource reservation decisions on the codec utilizing the greatest amount of network resources. This ensures that enough resources have been reserved so that any of the negotiated codecs will work properly.

[0087] Second Technique – This technique uses the codec ordering in the SDP for each media type to select a particular codec. The SDP specification states that the order of the codecs in the SDP should determine the codec preference of the network endpoint. In most cases the first codec in the list will be the one used by the network endpoint. When network endpoints are known to adhere to this standard, this configuration will result in a more stringent QoS reservation, resulting in a lower cost of service.

[0088] Third Technique – This technique uses the DOCSIS QoS MIB (Patrick, M., Murwin, W., “Data Over Cable System Interface Specification Quality of Service Management Information Base (DOCSIS-QOS MIB)”, Internet-Draft (expires April 2004), <http://www.ietf.org/internet-drafts/draft-ietf-ipcdn-qos-mib-09.txt>, October 27, 2003) at the CMTS to determine the audio packet size that the network endpoints are sending. Since audio traffic is CBR (Constant Bit Rate), the same packet size will be used for all packets. Examining the DOCS-QOS MIB will reveal the number of packets and bytes forwarded by the DOCSIS Service Flows used to carry the audio data, as well as the amount of time the flow is active. Dividing the total number of bytes by number of packets used to send those bytes yields the packet size. The rate is determined by dividing the total number of bytes forwarded by the amount of time that the flow is active.

[0089] The described embodiment makes an initial resource reservation using either the first technique or the second technique. After a few seconds of data transfer, the application manager determines the exact packet size and rate, and adjusts the resource reservation using the third technique, to more accurately represent what the audio codec requires. If dynamic codec switching is a known possibility, the described embodiment uses only the first technique, which will always work for any of the available codecs.

[0090] Other aspects, modifications, and embodiments are within the scope of the claims.

What is claimed is:

1. A method of assigning a set of network resources between network endpoints, comprising:
 - using a session initiation request to an application manager for reserving resources between network endpoints, wherein the application manager is logically and physically decoupled from an application server associated with the network endpoints;
 - providing a PCMM message from the application manager to a policy server as a result of the session initiation request, wherein the PCMM message includes at least some information embedded in the session initiation request;
 - selecting, via the policy server, a set of network resources for creating a path connecting the network endpoints, wherein the policy server selects the set of network resources based on the PCMM message.
2. The method of claim 1, further including communicating the session initiation request via Session Initiation Protocol.
3. The method of claim 2, further including inserting the application manager between one of the network endpoints and an application server, wherein the application manager transparently relays the session initiation request between the network endpoint and the application server, without requiring functional modifications of the network endpoints or the application server.
4. The method of claim 1, using the SDP in a SIP message to perform the necessary QoS reservation of resources in the network.
5. The method of claim 4, further including providing the QoS information in session description protocol fields within the session initiation request.

6. The method of claim 1, further including combining high-level QoS information from the session initiation request with understanding of underlying network technologies to produce detailed QoS information for use in the PCMM message.

7. The method of claim 1, further including embedding QoS information either via SDP or via RSVP flowspecs in an eXtensible Markup Language (XML) message, and using the XML message to communicate the session initiation request to an application server, which subsequently transfers the QoS information to the applications manager in an XML message.

8. The method of claim 7, further using HTTP to communicate the XML message to an application server, which subsequently transfers the QoS information to the applications manager.

9. The method of claim 1, wherein the application server remains stateless through the use of a unique session identifier for communicating with the application manager, and the association between the session and the resources are maintained in the decoupled application manager, and the session identifier is constructed from the SIP signaling messages between the end points and the application server.

10. The method of claim 4, further including selecting RSVP FlowSpec parameters based on the SDP media attributes.

11. The method of claim 10, further including selecting of RSVP FlowSpec parameters and providing the ability to configure FlowSpec parameter selection behavior for the purposes of performing resource reservation in a PCMM framework;

ensuring interoperability with DOCSIS Service Flow QoS parameters at the CMTS for providing QoS by normalizing the selected FlowSpec parameters across different CMTS vendor implementations;

ensuring interoperability between network endpoints that implement different optional features of SDP media attributes.

12. The method of claim 11, further including providing configurable settings for an upstream direction and a downstream direction of the session between the endpoints, and using these setting in FlowSpec parameter selection, along with the default values for each setting, wherein the configurable settings and associated default values include:

Delay, 100 milliseconds (downstream and upstream)

BandwidthAdjust, 1.125 (downstream)

BandwidthAdjust, 1.0 (upstream)

RateFactor, 4.0 (downstream and upstream)

MinPacketSize, 640 (downstream and upstream)

MaxPacketSize, 1500 (downstream and upstream)

SlackTerm, 800 (downstream and upstream)

RateLimit, 96000 (upstream)

RateLimit, 0 (downstream)

PollIntervalLimit, 6000 microseconds (upstream)

PollIntervalSet, empty (upstream)

13. The method of claim 11, further including, when the "ptime" media attribute (packetization time) is not given in the SDP, selecting RSVP FlowSpec parameters such that sufficient resources are reserved for all packetization times mandated by the PacketCable Audio/Video Codecs Specification.

14. The method of claim 11, further including, when multiple codecs are negotiated in the SDP fields, selecting RSVP FlowSpec parameters such that sufficient resources are reserved for all packetization times mandated by the PacketCable Audio/Video Codecs Specification.

15. The method of claim 11, further including using the DOCSIS QoS MIB to accurately determine ptime (packetization time) when a ptime media attribute is unknown, and modifying reserved resources if the ptime media attribute is subsequently determined.

16. The method of claim 11, further including using the DOCSIS QoS MIB to accurately determine ptime (packetization time) when negotiating multiple codecs to determine a particular codec being used, and modifying reserved resources if the particular codec is subsequently determined.

17. The method of claim 3, further including the application manager learning the contact IP address and port number of one or both endpoints from the REGISTER messages so that the application manager can forward INVITE messages to the endpoint.

18. The method of claim 3, further including the application manager to querying the application server to obtain the contact IP address and port number of one or both endpoints so that the application manager can forward INVITE messages to the endpoint.

19. A method of assigning a set of network resources between network endpoints, comprising:

using a session initiation request to an application manager for reserving resources between network endpoints, wherein the application manager is logically decoupled from an application server associated with the network endpoints;

providing a PCMM message from the application manager to a policy server as a result of the session initiation request, wherein the PCMM message includes at least some information embedded in the session initiation request;

selecting, via the policy server, a set of network resources for creating a path connecting the network endpoints, wherein the policy server selects the set of network resources based on the PCMM message.

20. A system for assigning a set of network resources between network endpoints, comprising:

an application manager for (i) receiving a session initiation request for initiating communication between the network endpoints through the set of network resources, wherein the communication employs an application residing on an application server, and (ii) generating a PCMM message including at least some information embedded in the session initiation request;

a policy server for receiving the PCMM message, and for selecting the set of network resources for creating a path connecting the network endpoints, wherein the policy server selects the network resources based on the PCMM message.

21. The system of claim 20, wherein the session initiation request is transmitted in Session Initiation Protocol (SIP).

22. The system of claim 21, wherein the application manager is disposed between one of the network endpoints and an application server, and the application manager transparently relays the session initiation request between the network endpoint and the application server.

23. The system of claim 22, wherein the application manager is a stand-alone component of the system.

24. The system of claim 22, wherein the application manager is incorporated within the policy server.

25. The system of claim 20, wherein at least one of the network endpoints provides QoS information to the application manager via the session initiation request.

26. The system of claim 25, wherein the session initiation request includes a session description protocol field containing the QoS information.

27. The system of claim 20, wherein the application manager combines high-level QoS information from the session initiation request with information about the network assets to produce detailed QoS information for use in the PCMM message.

28. The system of claim 20, wherein the application server or application Manager embeds QoS information within an eXtensible Markup Language (XML) message, and uses the XML message to communicate the session's required resources to the application manager, or Policy Server, respectively.

29. The system of claim 20, wherein the application server embeds QoS information within an eXtensible Markup Language (XML) message, and uses the XML message to communicate the set of network resources to the application manager, which subsequently translates the QoS information into a detailed QoS message to communicate with the Policy Server.




FIG. 1




FIG. 2

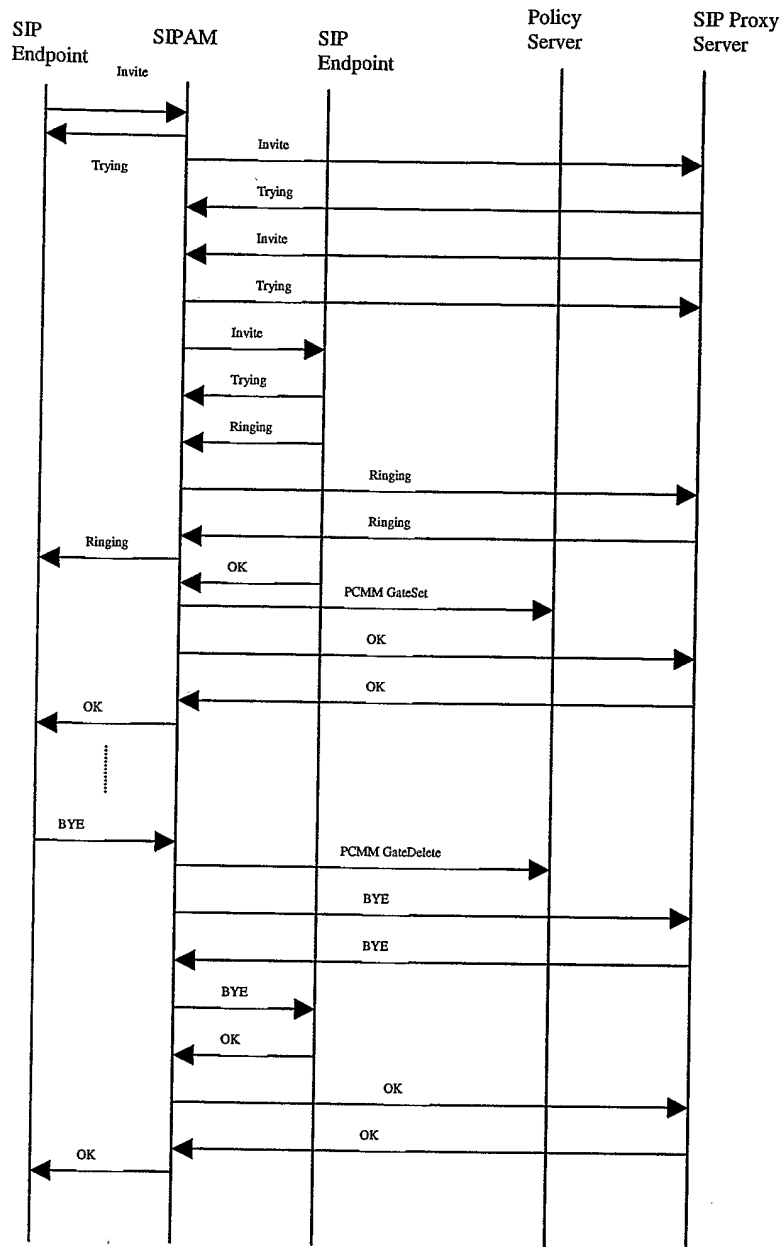


FIG. 3

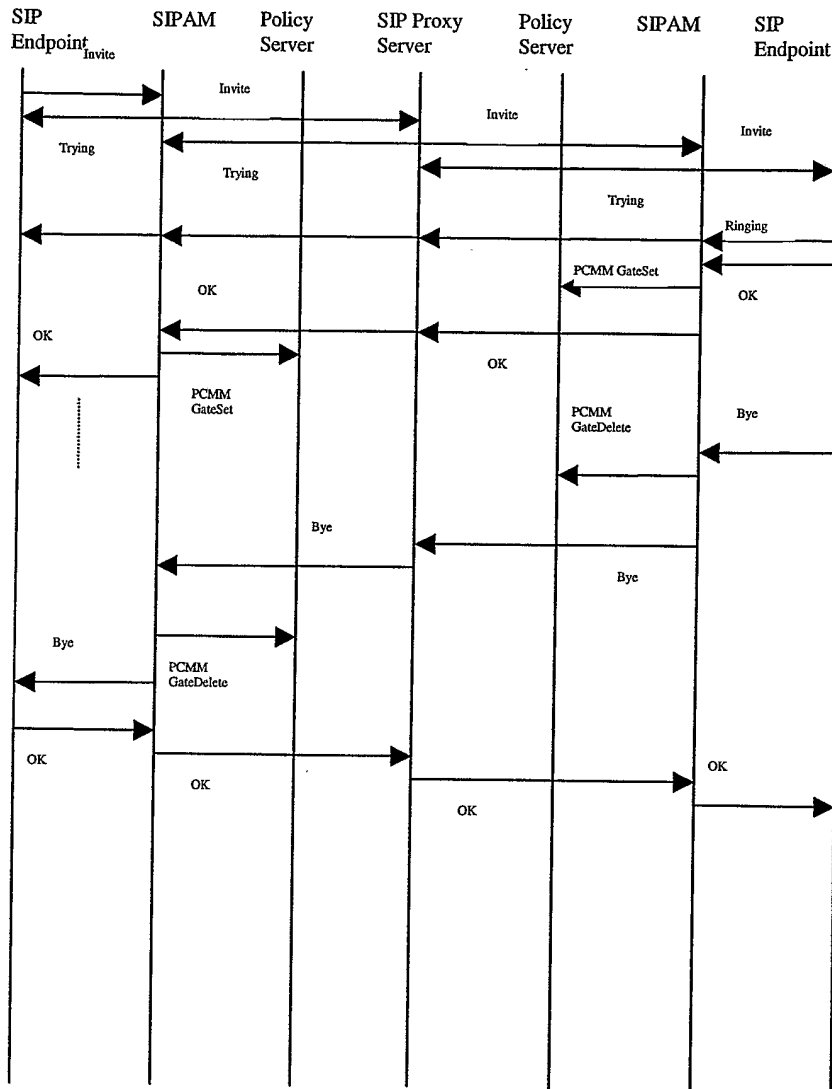


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US04/18784

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : H04L 12/66
 US CL : 370/351-356

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 U.S. : 370/351-356, 260, 264,265

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
 EAST

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y,E	US 6,771,639 B1 (HOLDEN) 03 August 2004, col 3 lines 17-59, col 4 lines 17-33, col 6 lines 54-67, col 7 lines 1-10 and 35-67, col 8 lines 1-7	1-2, 4-6, 10, 19, 20-22, 25, 27
Y,P	US 6,738,390 B1 (XU et al) 18 May 2004, col 5 lines 53-67, col 6 lines 1-10.	1-2, 4-6, 10, 19, 20-22, 25,27

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:	
"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 17 August 2004 (17.08.2004)	Date of mailing of the international search report 13 OCT 2004
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450 Alexandria, Virginia 22313-1450 Facsimile No. (703)305-3230	Authorized officer Douglas Olms Telephone No. (703) 305-4700 <i>Kiyoma Zogan</i>