(12) **United States Patent**
Ashwood Smith

(10) Patent No.: **US 7,296,087 B1**
(45) **Date of Patent:** **Nov. 13, 2007**

(54) **DYNAMIC ALLOCATION OF SHARED NETWORK RESOURCES BETWEEN CONNECTION-ORIENTED AND CONNECTIONLESS TRAFFIC**

(75) Inventor: **Peter J. Ashwood Smith**, Hull (CA)

(73) Assignee: **Nortel Networks Limited**, St. Laurent (CA)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/527,584**

(22) Filed: **Mar. 17, 2000**

(51) **Int. Cl.**
G06F 15/173 (2006.01)

(52) **U.S. Cl.** **709/238**; 709/226; 709/229; 709/235; 709/242

(58) **Field of Classification Search** 370/229–240; 709/241
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,363,053 B1 * 3/2002 Schuster et al. 370/230

6,363,319 B1 * 3/2002 Hsu 701/202
6,473,404 B1 * 10/2002 Kaplan et al. 370/238

OTHER PUBLICATIONS

RFC2702: Requirements for Traffic Engineering Over MPLS (D. Awduche; J. Malcolm; J. Agogbua; M. O'Dell; and J. McManus—Sep. 1999.

* cited by examiner

Primary Examiner—Syed A. Zia
(74) *Attorney, Agent, or Firm*—Matthew M. Roy; Ogilvy Renault LLP

(57) **ABSTRACT**

Resources of a shared physical network element of a communications network are dynamically allocated between connection-oriented traffic and connectionless traffic. For each shared physical network element of the network, a resource requirement of the connection-oriented traffic is determined; and a respective traffic metric to be used for routing connectionless traffic is dynamically adjusted based on the determined resource requirement of the connection-oriented traffic. As a result, resources of the shared physical network element can be efficiently utilized, and congestion of connectionless traffic being routed through the shared physical network element is avoided.

30 Claims, 2 Drawing Sheets

Figure 1

Fig 2

START

RECEIVE MPLS RESERVATION REQUEST — 100

A

PROCESS REQUEST — 102

TOTAL RESOURCES ALLOCATED TO MPLS — 104

QUERY IGP TABLE — 106

108 — OBTAIN IGP METRIC

No — WAIT

RESPONSE RECEIVED? — 110

No — TIME OUT?

Yes — ALARM

Yes — UPDATE LOCAL PATH DATABASE — 112

PROPAGATE LINK STATE PACKET — 114

A

1

DYNAMIC ALLOCATION OF SHARED NETWORK RESOURCES BETWEEN CONNECTION-ORIENTED AND CONNECTIONLESS TRAFFIC

CROSS-REFERENCE TO RELATED APPLICATIONS

This is the first application filed for the present invention.

MICROFICHE APPENDIX

Not Applicable.

TECHNICAL FIELD

The present invention relates to resource management in communications networks, and in particular to dynamic allocation of shared network resources between connection-oriented and connectionless traffic in a communications network.

BACKGROUND OF THE INVENTION

In the modern network space, packetized data traffic of various different protocols (e.g. internet protocol, frame relay, asynchronous transfer mode, etc.) is transported over a common network infrastructure. Each protocol provides its own packet (or frame) size and format standards. Additionally, some protocols (e.g. IP) are specifically designed to allow packets having widely varying lengths. New routing protocols, for example the multi-protocol label switching (MPLS) protocol have been proposed to facilitate multi-protocol traffic across a common network infrastructure.

Under the MPLS protocol, label switched paths (LSPs) are propagated across the network hop-by-hop along a path that is set up at the beginning of a communications session. In a general, the label assigned to the LSP can be different for each hop, with the label conversion being performed by the node serving the respective hop. Resources of each hop (i.e. the node serving the hop) of the path are reserved during set-up of the path, and normally will not be available for carrying other traffic until the path is released.

The mapping of an end-to-end path at the beginning of a communications session characterizes the MPLS protocol as "connection oriented". Other protocols, (such as IP) which do not transport data over predefined end-to-end paths are referred to as "connectionless". Typically, connectionless traffic is routed across a network using a shortest-path or least-cost-path routing protocol, such as, for example, the Interior Gateway Protocol (IGP). In general, a metric (e.g. a link distance vector, or a link cost factor) is assigned to each link and used within each router for mapping packet destination addresses to downstream links. The metric is normally provisioned for traffic engineering, and reflects not only geographic distances, but also provisioned bandwidth of each link. A higher metric on a particular link makes that link less attractive for carrying connectionless traffic, so that the IGP will normally operate to route connectionless traffic away from that link. Both connection-oriented and connectionless traffic may be carried over shared network infrastructure. This situation is normally accommodated by adjusting the provisioned IGP metric to reflect an average anticipated amount of bandwidth allocated to the connection-oriented traffic. However, this raises a difficulty in that

2

the resources reserved for connection-oriented traffic. Accordingly, during periods of heavy demand for connection-oriented traffic, the provisioned IGP metric for a link may provide an inflated indication of the amount of band-width actually available for connectionless traffic. This can easily result in undesirable congestion on the link. Conversely, during periods of low demand for connection-oriented traffic, the provisioned IGP metric for a link may provide a deflated indication of the amount of band-width actually available for connectionless traffic. This can result in undesirable under-utilization of the link.

A technique which allows connection-oriented and connectionless traffic to efficiently share network resources is therefore highly desirable.

SUMMARY OF THE INVENTION

An object of the present invention is to provide a technique for efficiently allocating resources of a shared network element between connection oriented and connectionless traffic.

A further object of the present invention is to provide a method of allocating resources between connection oriented and connectionless traffic, by adjusting an IGP metric in accordance with MPLS resource reservations.

Accordingly, an aspect of the present invention provides a shared network element operative within a communications network capable of end-to-end transport of connection-oriented traffic and connectionless traffic through the shared network element. The shared network element comprises: means for determining a resource requirement of the connection-oriented traffic; and means for dynamically adjusting a respective traffic metric to be used for routing connectionless traffic based on the determined resource requirement of the connection-oriented traffic.

Another aspect of the present invention provides a method of managing an allocation of resources between connection-oriented traffic and connectionless traffic being routed through a shared physical network element of a communications network. The method comprises the steps of: determining a resource requirement of the connection-oriented traffic; and dynamically adjusting a respective traffic metric to be used for routing connectionless traffic based on the determined resource requirement of the connection-oriented traffic.

In embodiments of the invention, the connection-oriented traffic is multi-protocol label switched (MPLS) traffic. In such cases, the step of determining the resource requirement of the connection-oriented traffic comprises the steps of: receiving MPLS reservation requests in respect of the shared physical network element; and dynamically adjusting a total amount of resources required to satisfy the received MPLS reservation requests.

In embodiments of the invention, the connectionless traffic includes internet protocol (IP) packet traffic. In such cases, routing of the connectionless traffic may be controlled using an interior gateway protocol (IGP) routing system adapted to calculate a shortest path route of the connectionless traffic through the communications network, the shortest path routing being based on a respective metric concerning each physical network element forming the network.

The step of dynamically adjusting the respective metric preferably comprises the steps of: increasing the respective metric as the determined resource requirement of the connection-oriented traffic increases; and decreasing the respec-

In some embodiments of the invention, the respective metric may be a link distance vector associated with a respective link connected to a node of the communications network. In such cases, the step of dynamically adjusting the respective metric comprises the steps of: determining an updated value of the link distance vector; and updating a mapping table maintained by the node with the updated value of the link distance vector.

An updated value of the link distance vector may be determined by querying a resource allocation table comprising a plurality of characteristic resource allocation values and a respective link distance vector value corresponding to each characteristic resource allocation value. Querying the resource allocation table may include the steps of: identifying the characteristic resource allocation value which most closely matches the determined resource requirement of the connection-oriented traffic; and selecting the corresponding link distance vector as the updated link cost factor.

In other embodiments of the invention, the respective metric may be a link cost factor associated with a respective link connected to a node of the communications network. In such cases, the step of dynamically adjusting the respective metric comprises the steps of: determining an updated value of the link cost factor; updating a PATH database maintained by the node with the updated link cost factor value; and propagating a link state packet containing the updated link cost factor value to neighboring nodes within the network.

An updated value of the link cost factor can be determined by querying a resource allocation table comprising a plurality of characteristic resource allocation values and a respective link cost factor value corresponding to each characteristic resource allocation value. Querying the resource allocation table may comprise the steps of: identifying the characteristic resource allocation value which most closely matches the determined resource requirement of the connection-oriented traffic; and selecting the corresponding link cost factor as the updated link cost factor.

An advantage of the present invention is that by adjusting the IGP metric for a link in accordance with changing resource requirements of connection oriented traffic, routing of connectionless traffic is automatically altered to make effective use of remaining resources while avoiding congestion. No modification of conventional (e.g. IGP) routing methodologies are required to accomplish this result.

BRIEF DESCRIPTION OF THE DRAWINGS

Further features and advantages of the present invention will become apparent from the following detailed description, taken in combination with the appended drawings, in which:

FIG. **1** is a block diagram illustrating a communications network usable in conjunction with an embodiment of the present invention; and

FIG. **2** is a flow chart illustrating exemplary steps in a process for managing resource allocation in accordance with an embodiment of the present invention.

It will be noted that throughout the appended drawings, like features are identified by like reference numerals.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention provides a system for allocating resources between connectionless and connection-oriented

the present invention generally comprises a plurality of routers **4** (four are shown in FIG. **1**) interconnected by links **6**. The links **6** may be fiber optic links. The routers **4** may be agile or non-agile optical routers, and may be configured for wave division multiplex (WDM) and/or dense wave division multiplex (DWDM) transport of packet data traffic. Communications devices **8**, for example end user personal computers (PCs) or local area network (LAN) servers may be connected to the communications network **2** via one or more access points **10**. The communications network **2** may also be connected to one or more federated networks **12**, for example an ATM or an IP network, through a respective gateway **14**.

In the example of FIG. **1**, connection-oriented traffic is conveyed through an end-to-end MPLS path **16** mapped across the communications network **2** between a source node and a destination node via one or more intervening routers **4**. The path **16** is divided into hops **18**, each of which is served a respective node (e.g. the source node or a router **4**) connected at the up-stream end of the respective hop **18**. In the illustrated example, the source and destination nodes are located at respective access points **10***a* and **10***b*, and two intervening routers **4***a* and **4***b* are incorporated into the path **16**. A first hop **18***a* of the path **16** extends between the source node at access point **10***a* and a first router **4***a*. A second hop **18***b* extends between the routers **4***a* and **4***b*. Finally, a third hop **18***c* extends between the router **4***b* and the destination node at access point **10***b*.

In addition, connectionless traffic is transported through the communications network between the gateway **14** and the destination node **10***b*. Routing of the connectionless traffic is handled, for example, in accordance with conventional least cost path routing using the Interior Gateway Protocol (IGP) to map destination addresses to downstream links. In the example illustrated in FIG. **1**, a least cost path **20** calculated using provisioned IGP metrics for each link follows the route indicated by dashed arrows. As may be seen in FIG. **1**, this least cost path **20** shares two hops (**18***b* and **18***c*) with the MPLS path **16**, which may lead to congestion on those hops. Thus the present invention provides a technique of managing the allocation of resources between the two traffic flows, in order to avoid congestion within the shared hops.

In accordance with the present invention, congestion within shared physical elements of the network (e.g. within the routers **4***a*, **4***b* and links **6** of the shared hops **18***b*, **18***c*) is avoided by adjusting the IGP metrics concerning each of the shared links **6** to reflect the resources (e.g. bandwidth) which have been allocated to the MPLS path **16**. Adjustment of the IGP metric in this manner makes each of the shared links **6** less attractive for either shortest-distance-path routing or least-cost-path routing. In cases where either the bandwidth allocated to the MPLS path **16**, or the volume of connectionless traffic is large, the adjusted IGP metrics on the shared links **6** may cause the routing protocol to favour an alternative path **22** (indicated by dotted arrows in FIG. **1**) which avoids sharing links **6** with the MPLS path **16**.

Adjustment of IGP metrics can be based on a resource allocation table **24** which may be co-resident within each router **4** or centrally located and accessible by each router **4** through the network **2**. The resource allocation table **24** generally operates to receive a query from a node (any of access points **10**, routers **4**, or gateway **14**) containing a value indicative of resources allocated to connection oriented traffic (in the present example bandwidth allocated to

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | sales@docketalarm.com | 1-866-77-FASTCASE

fastcase®
Smarter legal research.