US007095715B2

(12) **United States Patent**
Buckman et al.

(10) **Patent No.: US 7,095,715 B2**
(45) **Date of Patent: Aug. 22, 2006**

(54) **SYSTEM AND METHOD FOR PROCESSING NETWORK PACKET FLOWS**

(75) Inventors: **Charles R. Buckman**, Austin, TX (US); **Dennis J. Cox**, Austin, TX (US); **Donovan M. Kolbly**, Austin, TX (US); **Craig S. Cantrell**, Austin, TX (US); **Brian C. Smith**, Fort Worth, TX (US); **Jon H. Werner**, Austin, TX (US); **Marc Willebeek-LeMair**, Austin, TX (US); **Joe Wayne Blackard**, Austin, TX (US); **Francis S. Webster, III**, Austin, TX (US)

(73) Assignee: **3Com Corporation**, Marlborough, MA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 956 days.

(21) Appl. No.: **09/897,189**

(22) Filed: **Jul. 2, 2001**

(65) **Prior Publication Data**

US 2003/0012147 A1 Jan. 16, 2003

(51) **Int. Cl.**
G01R 31/08 (2006.01)
H04L 12/28 (2006.01)

(52) **U.S. Cl.** **370/230**; 370/392; 370/401

(58) **Field of Classification Search** 370/389, 370/392, 400, 401, 428, 474, 466, 412, 230, 370/235, 263; 713/150, 153, 154, 160; 709/224, 230, 236
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

651,636	A	6/1900	McCarthy	52/54
5,101,402	A	3/1992	Chiu et al.	709/224
5,648,965	A	7/1997	Thadani et al.	370/241
5,742,772	A	4/1998	Sreenan	709/226
5,787,253	A	7/1998	McCreery et al.	709/231
5,835,726	A	11/1998	Shwed et al.	709/229
5,845,267	A	12/1998	Ronen	705/40
6,028,842	A *	2/2000	Chapman et al.	370/235
6,091,709	A *	7/2000	Harrison et al.	370/235
6,104,700	A *	8/2000	Haddock et al.	370/235
6,108,700	A	8/2000	Maccabee et al.	709/224
6,172,990	B1 *	1/2001	Deb et al.	370/474
6,212,559	B1	4/2001	Bixler et al.	709/221
6,262,983	B1 *	7/2001	Yoshizawa et al.	370/389
6,286,030	B1	9/2001	Wenig et al.	709/203
6,292,489	B1 *	9/2001	Fukushima et al.	370/401

(Continued)

OTHER PUBLICATIONS

PacketShaper Enterprise Solutions; Packeteer, Inc. (www.packeteer.com).
Network-Based Application Recognition, Release 12.0(5)XE2, 12.1(1)E and 12.1(2)E, Cisco Systems (www.cisco.com).

Primary Examiner—Doris H. To
Assistant Examiner—Ian N. Moore
(74) *Attorney, Agent, or Firm*—Hamilton & Terrile, LLP

(57) **ABSTRACT**

A system and method provides a broadband network node for a best effort network such as the Internet or intranets which supports the inexpensive and rapid deployment of services to the best efforts network. Separate data path and control path mechanisms allow high-speed data transfers with parallel processing flows for the data path that are controlled across data flows by the control path. Packets are classified, modified and shaped to enable the service on the network with an accountant to track packet traffic for control and billing purposes. A series of processing blades perform a modification function for each blade that processes packets according to classifications. The processing blades are modular and scalable for insertion in the broad band switch to rapidly adapt the broadband network node for new services.

5 Claims, 1 Drawing Sheet

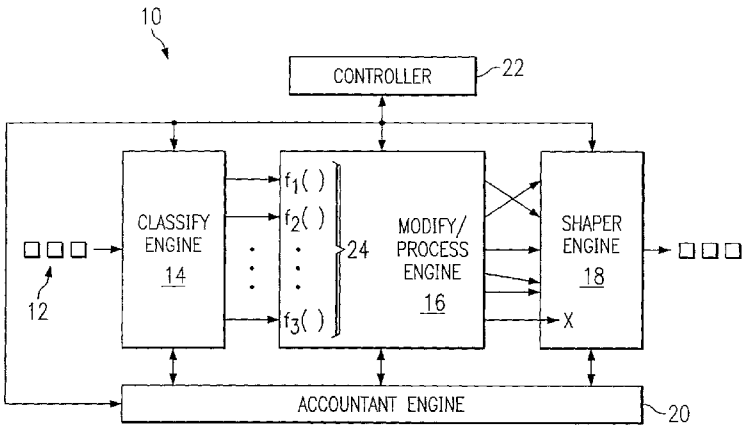

U.S. PATENT DOCUMENTS

6,320,848	B1 *	11/2001	Edwards et al.	370/255
6,418,125	B1 *	7/2002	Oran	370/266
6,452,915	B1	9/2002	Jorgensen	370/38
6,542,466	B1 *	4/2003	Pashtan et al.	370/235
6,560,233	B1 *	5/2003	Hatanaka et al.	370/401
6,590,885	B1	7/2003	Jorgensen	370/338
6,594,246	B1	7/2003	Jorgensen	370/338
6,628,617	B1 *	9/2003	Karol et al.	370/237
6,628,629	B1	9/2003	Jorgensen	370/322
6,636,481	B1 *	10/2003	Yamaguchi et al.	370/230
6,640,248	B1	10/2003	Jorgensen	709/226
6,680,922	B1	1/2004	Jorgensen	370/328
6,697,368	B1 *	2/2004	Chang et al.	370/395.1
6,711,165	B1 *	3/2004	Tzeng	370/392
6,714,517	B1 *	3/2004	Fawaz et al.	370/236
6,772,223	B1	8/2004	Corl et al.	709/238
6,801,530	B1 *	10/2004	Brandt et al.	370/392
6,804,240	B1 *	10/2004	Shirakawa et al.	370/392
6,826,147	B1 *	11/2004	Nandy et al.	370/229
6,831,893	B1 *	12/2004	Ben Nun et al.	370/235
6,865,602	B1	3/2005	Nijemcevic et al.	709/223
6,892,233	B1	5/2005	Christian et al.	709/223
2002/0085560	A1 *	7/2002	Cathey et al.	370/392
2002/0099854	A1	7/2002	Jorgensen	709/249
2002/0107908	A1 *	8/2002	Dharanikota	709/203
2002/0152303	A1	10/2002	Dispensa	709/224
2002/0191622	A1 *	12/2002	Zdan	370/401
2002/0194369	A1 *	12/2002	Rawlins et al.	709/238
2003/0005144	A1	1/2003	Engel et al.	709/235
2003/0014627	A1 *	1/2003	Krishna et al.	713/153
2003/0028624	A1	2/2003	Hasan et al.	709/220
2003/0067903	A1	4/2003	Jorgensen	370/338
2003/0076855	A1	4/2003	Chamberlain	370/447
2004/0022237	A1	2/2004	Elliott et al.	370/356
2004/0088646	A1	5/2004	Yeager et al.	715/500
2004/0098447	A1	5/2004	Verbeke et al.	709/201

* cited by examiner

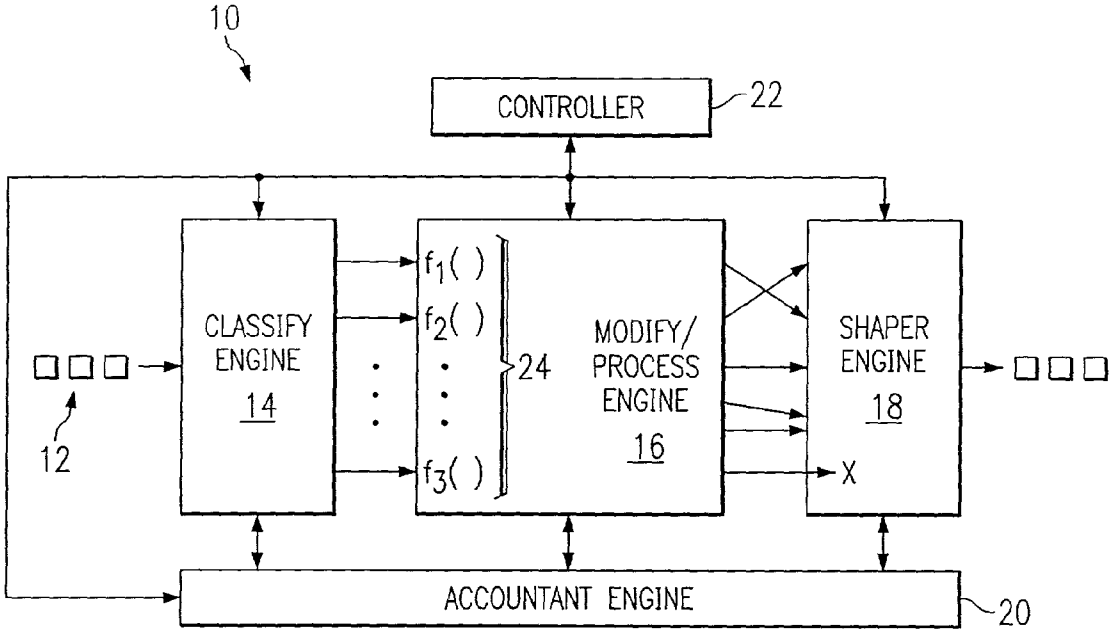

FIG. 1

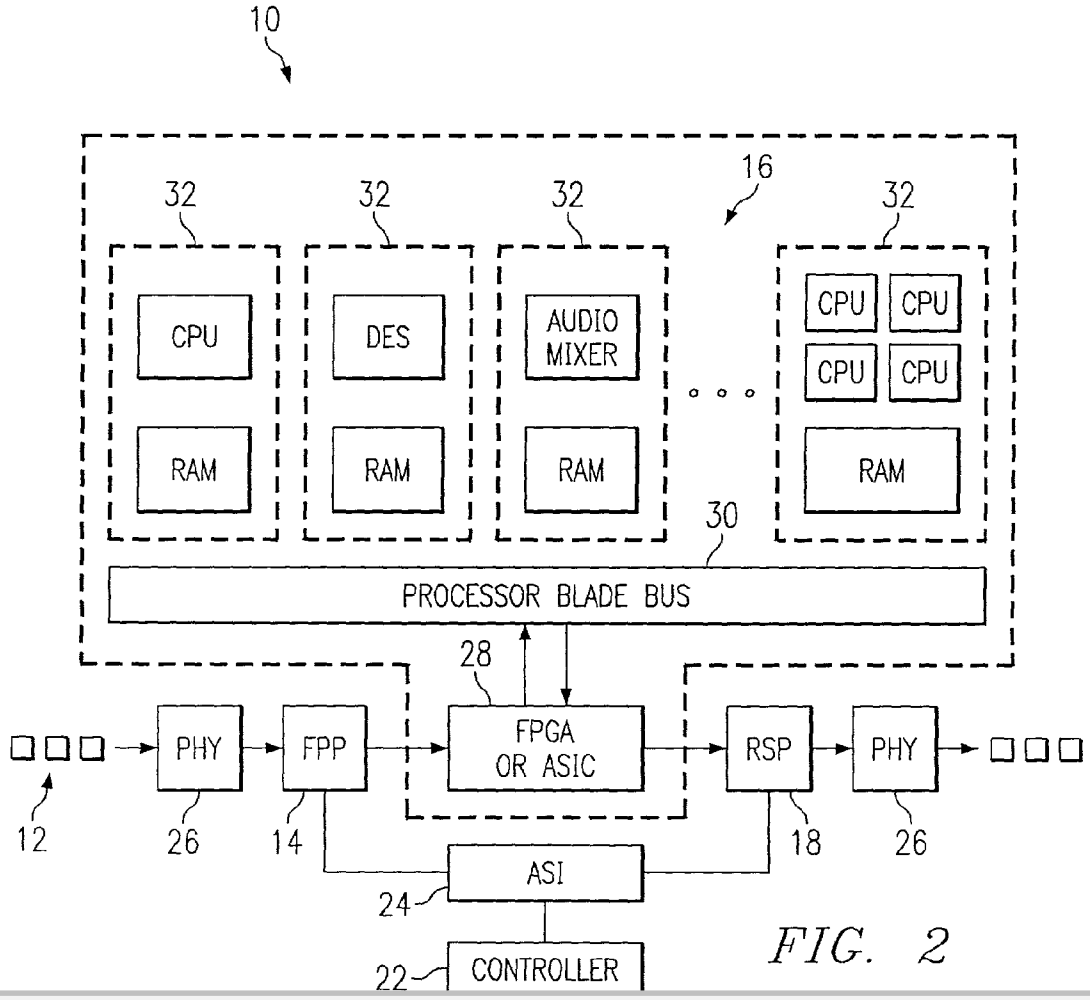

FIG. 2

SYSTEM AND METHOD FOR PROCESSING NETWORK PACKET FLOWS

TECHNICAL FIELD OF THE INVENTION

This invention relates generally to computer networks, and more specifically relates to a system and method for processing packet flows through a network.

BACKGROUND OF THE INVENTION

The Internet has grown in popularity largely because it provides a simple and uniform underlying packet protocol for exchanging data that in turn enables more complex applications to occur. The relative simplicity of the Internet has lead to an explosion of growth in data traffic for business and personal usage. An industry of Internet service providers (ISP) has developed to provide access to the Internet for businesses and homes. These ISPs have invested heavily in infrastructure with the general goal of improving data transfer rates for customer end points. Thus, ISPs have evolved from low speed analog modem Internet connections provided through dial-up service across the plain old telephone system (POTS) into broadband access provided by cable modems, DSL, and ADSL that are capable of transferring data at substantially higher rates.

Internet access offered by ISPs has become a commodity service with different ISPs typically offering similar access speeds and competing on price points. However, as ISPs gain subscribers to their high-speed broadband access services, the reliability of ISP service, i.e. the ability to transfer data at predictable rates, available at end points has suffered and indeed is often less than the capacity of broadband customer modems. For instance, surges in subscriber use tends to create bottlenecks that slow data transfer rates and use ISP infrastructure inefficiently. This difficulty relates generally to the Internet's architecture which transfers data on a "best effort" basis in which TCP/IP packets are generally transferred between routing points without prioritization. This "best effort" architecture is attractive for its simplicity, but creates difficulties for ISPs who desire to distinguish themselves based on the services offered compared to other ISPs.

One manner in which ISPs are working to distinguish themselves is to offer subscribers different types of services. In general terms, a service is the processing of certain data on the network in a predetermined manner with associated billing. For instance, one service typically offered by telephone network providers is call waiting, in which customers receive notice of an incoming telephone call when the customer is on the phone for a monthly fee. A typical service offered by ISPs is e-mail service, although e-mail service is not typically billed. Typically e-mail service is provided by routing e-mails to an ISP server that stores the e-mails until the subscriber retrieves them.

More complex services are generally difficult to implement on a best effort network architecture, such as the Internet, since best effort networks are generally designed to route packets to destinations on a first come first served basis. An ISP that seeks to provide a new service to its subscribers has to design and install hardware and software that generally require a considerable capital investment and time to develop. Even after investing considerable time and capital in the development of a new service, subscriber needs and technology often advance more rapidly than a

revenue source or market differentiator for the ISP. Thus, ISPs who attempt to lead the market by developing and offering new services face considerable risk if a large investment in a new service fails to result in greater revenue or customer satisfaction.

SUMMARY OF THE INVENTION

Therefore a need has arisen for a system and method which processes network packet flows in a best effort network with scalable and flexible hardware and software architecture.

A further need has arisen for a system and method which supports rapid deployment of services in a best effort network with reduced expense and complexity.

In accordance with the present invention, a system and method is provided that substantially eliminates or reduces disadvantages and problems associated with previously developed systems and methods for processing network packet flows in a best effort network. A network node provides a data path that performs repetitive operations at high speeds to classify, process and shape packets according to a slower speed control path that performs complex logic operations associated with predetermined services. Processing blades interfaced with the data flow receives predetermined packets based on classification and modifies the packets for transmission across the network.

More specifically, packets received by the network node are classified according to a set of classification rules. The classification rules sort the packets into packet groups and each packet group is sent to a predetermined input port of a modify\process engine. The modify\process engine performs multithreaded processing with a function associated with each input port. This highly parallel implementation in which internal state is not shared among functions allows the network node switch to perform repetitive operations at high speeds for processing packet flows. The modify\process engine sends packets to a shaper engine which controls the transmission of the packets with queuing disciplines for transmission to the network. An accountant engine tracks statistics in the classify, modify\process and shaper engines for control and billing purposes. A controller programs the classify, modify\process, shaper and accountant engines to achieve a desired packet prioritization and flow. For instance, the controller installs or modifies rules in the classify engine, installs or modifies functions in the modify\process engine, establishes queues in the shaper engine and determines the variables or tracking in the accountant engine.

In one embodiment, packet flows are controlled by packet processing behaviors that are specified with a dataflow program and reflective programming. The dataflow program specifies interconnected objects for the data to flow through. For instance, one packet processing behavior of a service in a dataflow representation creates a queue for a new TCP connection. The dataflow program detects SYN packets of new TCP connections and, when a new connection is detected, a new dataflow program is constructed using reflective programming that contains elements to detect the host/port quadruple of the new flow. Thus, bandwidth may be dynamically assigned for new connections.

The modify\process engine includes an interface with a blade bus adapted to accept processor blades programmed to modify packets. Each processor blade performs a function on packets received and then forwards the modified packet

complex operations with minimal delay. The processing blade architecture provides a scalable broadband network node that is easily adapted to support new services with reduced complexity and increased speed. For instance, a broadband network node programmed to classify packets for supporting a voice over internet protocol (VOIP) is adapted to support a conference call service by interfacing a processing blade with the processor blade bus so that packets classified as conference call VOIP packets are mixed as audio streams with the processing blade.

The present invention provides a number of important technical advantages. One important technical advantage is the inexpensive and rapid deployment of new and modified services in a best effort network. The classification, modification and shaping of packets is performed at line speed by software residing on function specific hardware. Flexibility is provided to add and modify services through software programming and the addition of appropriate hardware processing blades. Thus, the broadband network node provides a flexible and scalable solution for routing packets in a best effort network according to the packet's associated applications.

Another important technical advantage of the present invention is that bandwidth is dynamically allocated for packet flows using reflective programming techniques in the context of dataflow programs. This allows the network to adapt as the type of data across the network changes over time, resulting in improved efficiency and flexibility in bandwidth usage.

Another important technical advantage of the present invention is that it provides separate mechanisms for a data path and a control path in routing packets. The data path is highly parallel, processing each packet flow independently to perform repetitive operations at high speeds for routing packets including complex modifications of the packets. The control path looks across data flows to implement the data path with the parallel processing so that complex routing operations are performed in an efficient manner at line speeds.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the present invention and advantages thereof may be acquired by referring to the following description taken in conjunction with the accompanying drawings, in which like reference numbers indicate like features, and wherein:

FIG. **1** depicts a block diagram of the logical architecture of a broadband network node; and

FIG. **2** depicts a block diagram of a hardware architecture for modifying packets with a processing blade architecture.

DETAILED DESCRIPTION OF THE INVENTION

Preferred embodiments of the present invention are illustrated in the figures, like numerals being used to refer to like and corresponding parts of the various drawings.

Building a service for a network presents a substantial task which is often time consuming and expensive. For instance, years often pass from the development and approval of a business case through the design and provisioning of a service on a network. The conventional development and provisioning of a service on a best effort network, such as the Internet or intranets that use TCP\IP, are deployed on a best effort network, modification of the service over time presents a continuing problem.

One solution that eases the development and deployment of services on a best effort network is to deploy broadband network nodes on the edges of the network, such as the advanced traffic processor disclosed in U.S. patent application Ser. No. 09/875,639, filed Jun. 6, 2001, entitled "A System and Method for Allocating Bandwidth Across a Network," incorporated herein by reference. A broadband network node is a hardware device that processes network packet flows at line speeds at the ingress and egress points of an intranet. For instance, services are provisioned to the intranet by programming the broad band network nodes to classify, modify and shape packets to accomplish the desired service.

Referring now to FIG. **1**, a block diagram depicts the logical architecture of a broadband network node **10** for provisioning and operating a service on a best effort network. Generally, a service is a packet processing behavior or behaviors and associated billing rules which provides value to subscribers to the service. Broadband network node **10** provides packet processing behaviors through software that classifies, processes and shapes packet traffic flowing through broadband network node **10**. Broadband network node **10** allows flexible and scalable deployment of services through software additions and modifications and, in some cases, the addition of hardware modules, that operate the service with separate data path and control path mechanisms. A highly parallel data path processes packet flows independently to perform repetitive operations at high speeds which reduce the impact of the service on network data transfer rates, while the control path performs more complex logic at lower speeds to allow control of one or more services by looking across data flows.

The data path mechanism is the flow of packets **12** through a classify engine **14**, a modify\process engine **16** and a shaper engine **18**. The data path mechanism uses a dataflow program which specifies a set of interconnected objects through which data flows and is processed. An accountant engine **20** tracks statistics in classify engine **14**, modify\process engine **16** and shaper engine **18** to evaluate the data path usage and for billing purposes. The control path mechanism is established by controller **22** which interfaces with classify engine **14**, modify\process engine **16**, shaper engine **18** and accountant engine **20**. Controller **22** programs the dataflow programs of each of the engines. For instance, installing or modifying rules in classify engine **14**, installing or modifying functions in modify\process engine **16**, establishing queues in shaper engine **18** and determining the variables to track in accountant engine **20**. Controller **22** also controls the routing of the packets between the engines.

Classify engine **14** accepts inbound packets **12** and applies a set of classification rules in a well-defined order determined by controller **22** so that the first matching rule defines the packet group. Classification into packet groups is based on any part of the packet, including the datalink, network, and transport header fields as well as packet content or payload. For example, classify engine **14** groups packets based on VCI\VPI\CLP (ATM fields), IP source or destination addresses, port numbers, protocol type such as TCP or UDP, application protocol such as DNS, DHCP, HTTP, or the MIME content type such as text\html.

In one embodiment, classify engine **14** classifies groups of packets by flow statistics maintained by classify engine **14** or accountant engine **20**, such as flow lifetime, average

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.