US007027393B1

(12) **United States Patent**
Cheriton

(10) **Patent No.: US 7,027,393 B1**
(45) **Date of Patent: Apr. 11, 2006**

(54) **TCP OPTIMIZED SINGLE RATE POLICER**

(75) Inventor: **David R. Cheriton**, Palo Alto, CA (US)

(73) Assignee: **Cisco Technology, Inc.**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1124 days.

(21) Appl. No.: **09/798,648**

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

(51) **Int. Cl.**
H04J 1/16 (2006.01)
H04J 3/14 (2006.01)
H04L 12/26 (2006.01)

(52) **U.S. Cl.** **370/230.1**; 370/232; 370/235.1
(58) **Field of Classification Search** 370/235, 370/232, 229, 230, 230.1, 231, 233, 234, 370/235.1, 253, 252; 709/224, 225, 232, 709/233, 234, 235; 109/223
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,734,825 A * 3/1998 Lauck et al. 709/233
6,144,639 A * 11/2000 Zhao et al. 370/235
6,147,969 A * 11/2000 Benmohamed et al. 370/230
6,324,165 B1 * 11/2001 Fan et al. 370/232
6,400,684 B1 * 6/2002 Benmohamed et al. .. 370/230.1
2003/0223370 A1 * 12/2003 Jain et al. 370/235
2004/0057378 A1 * 3/2004 Gronberg 370/230
2004/0170127 A1 * 9/2004 Tanaka 370/235

* cited by examiner

Primary Examiner—Feben M. Haile
(74) *Attorney, Agent, or Firm*—Campbell Stephenson Ascolese LLP

(57) **ABSTRACT**

An extension to the conventional single rate microflow policer that provides dual rate policing with a minimum of extra resource utilization. Using the extended microflow policer, an aggressive TCP flow ramps up to exceed the policer rate, setting a burst drop flag. Once the flow rate exceeds the burst rate, a single packet is dropped and the burst drop flag is cleared. On seeing the single packet drop, the TCP sender is then expected to reduce its rate. Flows that do not back off will eventually exceed a higher, hard drop threshold and experience packet drop. An aggressive TCP rate thus oscillate around the burst rate, efficiently approaching the hard drop rate without exceeding it. The addition of only a single bit flag avoids the cost of a dual-rate policer and the tail drop behavior induced by a single rate policer.

26 Claims, 3 Drawing Sheets

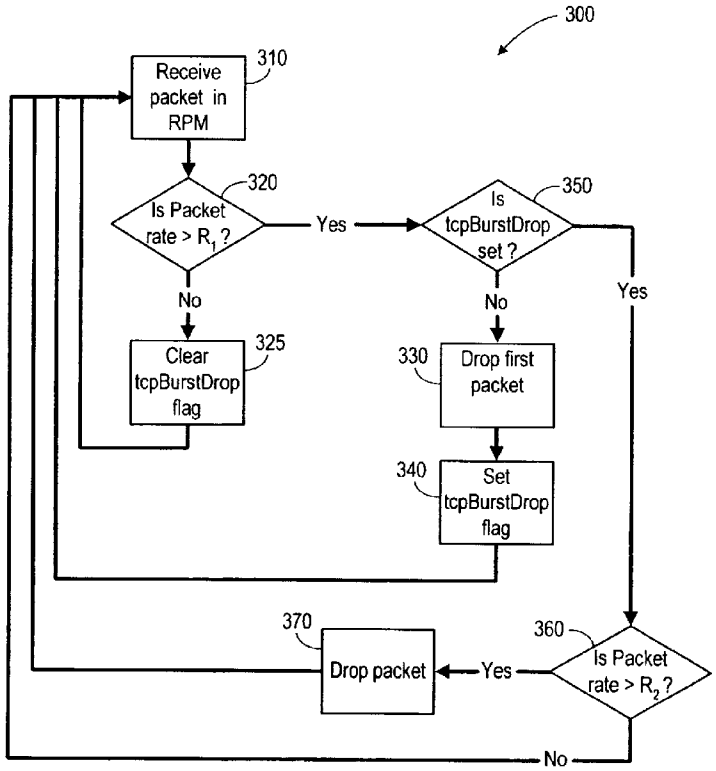

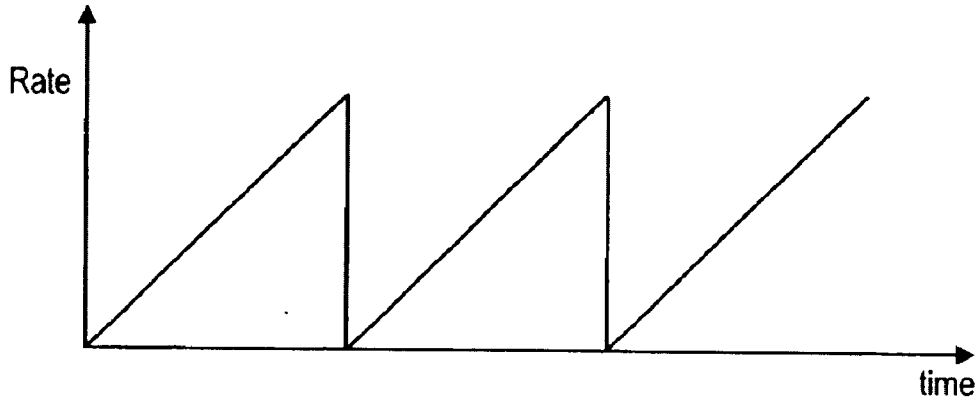

Fig. 1A
(PRIOR ART)

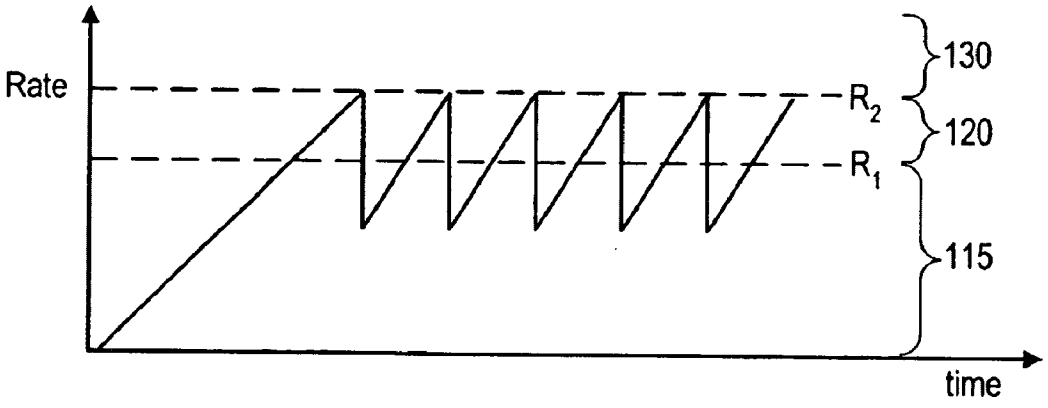

Fig. 1B

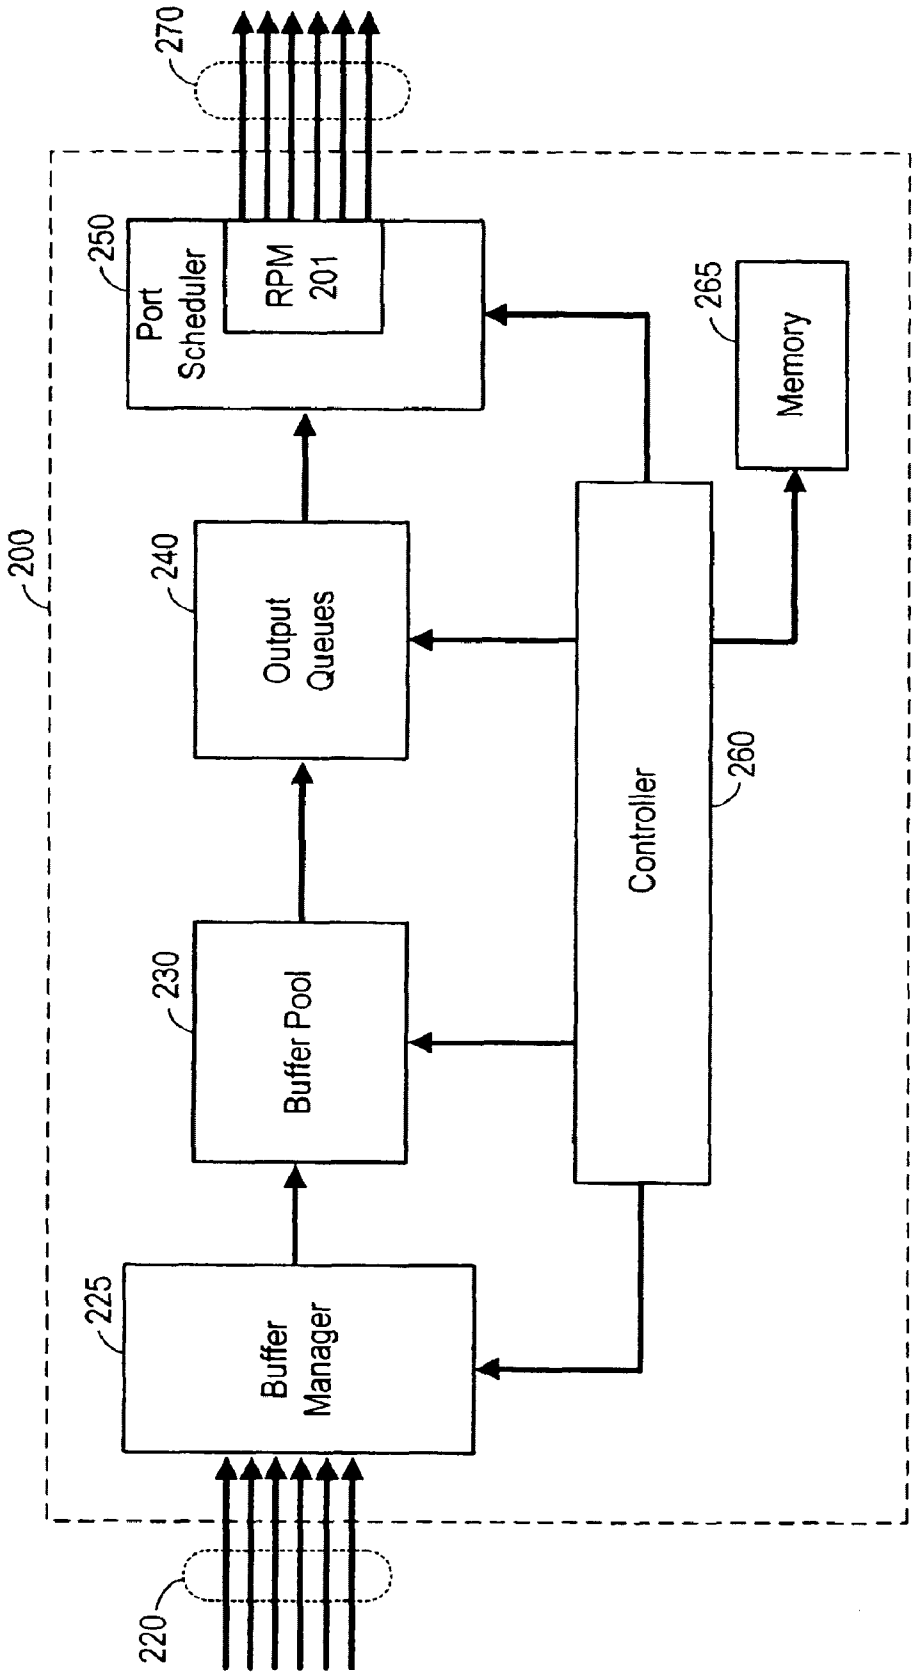

Fig. 2

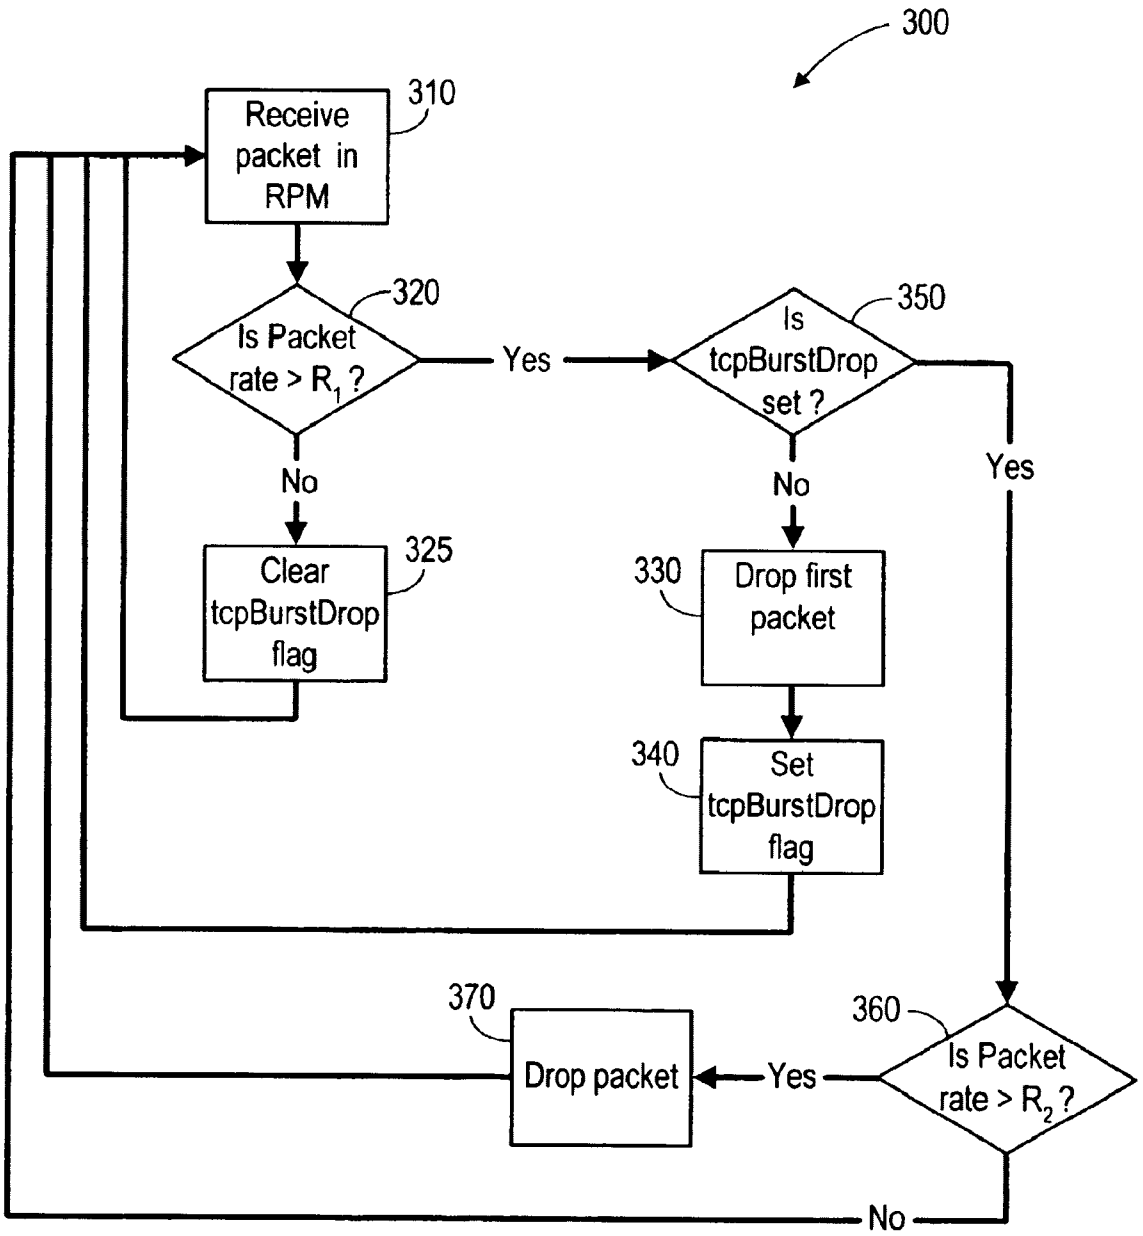

Fig. 3

TCP OPTIMIZED SINGLE RATE POLICER

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to digital communications systems, in particular computer networking, and specifically data flow rate control.

2. Description of the Related Art

In the field of computer networking, one area of concern is maintaining and supplying a pre-negotiated quality of service (QoS) and/or a guaranteed packet rate. Further discussion of the general quality of service problem can be found in James F. Kurose and Keith W. Ross, *Computer Networking: A Top Down Approach Featuring the Internet* (Addison Wesley 2000), Chapter 6.6, incorporated herein by reference in its entirety.

Many systems attempt to provide a guaranteed bit rate or packet rate for designated flows through a switching or routing system. A "flow" is here defined as a unique data connection between a certain designated source address and a designated destination address. Generally speaking, a "flow" is a defined subset of the packet cell traffic between designated endpoints, not merely a transport connection.

Policers are a critical component in providing quality of service in data networks. Policers are used to hold a packet flow to a target rate in the presence of burst traffic. Token bucket and leaky bucket mechanisms are well known approaches to policing packet streams. See, for example, Kurose and Ross, cited above. In addition, there are "virtual time" based approaches to policing such as that described in the ATM Forum Traffic Management Specification, (version 4.0, af-tm-0056.000, June 1996) as the theoretical arrival time (TAT) algorithm. The ATM Forum Traffic Management Specification is incorporated herein by reference in its entirety. However all of these approaches have the same drawbacks seen in packet buffering, namely tail dropping. Tail dropping, as that term is understood in the art, refers to the complete drop of all packets in a transmission burst after the bursting flow exceeds its designated maximum flow rate.

The problem of tail dropping in packet buffers is described in S. Floyd, and V. Jacobson, *Random Early Detection Gateways for Congestion Avoidance, IEEE/ACM Transaction on Networking*, vol. 1, No. 4, August 1993, p. 397–413 and in V. Jacobson, K. Nichols, and K. Podhuri, *RED in a Different Light*, Technical Report, April 1999. Both of these papers are incorporated herein by reference in their entireties.

Generally speaking, bandwidth management on the links between routers and switches is the key element in maintaining quality of service. As noted in Kurose and Ross, there are three aspects of a flow's packet rate among which one could choose to implement a policing scheme. These three important policing criteria, which differ from each other according to the time scale over which the packet flow is policed, are as follows:

Average Rate. The network may wish to limit the long term average rate (i.e., packets per time interval) at which a flow's packets can be sent into the network. A crucial issue here is the interval of time over which the average rate will be policed. For example, a flow whose average rate is limited to 100 packets per second is more constrained than a flow that is limited to 6,000

latter constraint would allow a flow to send 1000 packets in a given second-long interval of time (subject to the constraint that the rate be less than 6,000 packets in a minute), while the former constraint would disallow this sending behavior entirely.

Peak Rate. While the average rate constraint limits the amount of traffic that can be sent into the network over a relatively long period of time, a peak rate constraint limits the maximum number of packets that can be sent over a shorter period of time. Using the example above, the network may police a flow at an average rate of 6,000 packets per minute, while limiting the flow's peak rate to 1,500 packets per second.

Burst Size. The network may also wish to limit the maximum number of packets (i.e., the burst packets) that can be sent into the network in an extremely short interval of time. As this interval length approaches zero, the burst size limits the number of packets that can be instantaneously sent into the network. While it is physically impossible to instantaneously send multiple packets (after all, every link has a physical transmission rate that cannot be exceeded), the abstraction of a maximum burst size is a useful one.

One model that can be used to characterize different policing schemes is known as the "leaky bucket" mechanism (sometimes called the leaky bucket algorithm). A leaky bucket consists of a bucket (a logical container) that can hold up to b tokens.

In the leaky bucket mechanism, tokens are added to the bucket as follows: new tokens (which may potentially be added) are always generated at a rate of r tokens per second. If the bucket is filled with less than b tokens when a token is generated, the newly generated token is added to the bucket. Otherwise, the newly generated token is ignored and the token bucket remains full to its capacity of b tokens. The "leak" arises from the fact that tokens are removed from the bucket according to a defined rule representing the act by which the parameter policed (here, packet transmission).

The leaky bucket mechanism can be used to police a packet flow in the following manner: suppose that before a packet is transmitted into the network it must first remove a token from the token bucket. If the token bucket is empty, the packet must wait for a token. In this way, packets cannot enter the network until a token is available for them. This is analogous to requiring a ticket to enter a freeway.

Alternatively, rather than waiting for a token, a packet that arrives at an output queue looking for a token could be dropped if there are insufficient tokens to allow it to be enqueued. This is an example of a leaky bucket mechanism employed as an output queue control device.

The virtual time policing scheme, also well-known in the art, can also be used, as virtual time policers are generally considered an alternate to leaky bucket algorithms. In the virtual time scheme, the process first determines the "next time" that a flow is allowed to send a packet. When the next packet in that flow arrives, its time of arrival is compared to the "next time." If the packet has arrived earlier than the "next time," it needs to be policed or perhaps dropped. If the packet arrived later than the "next time," it is allowed. A burst parameter is usually associated with each policer to indicate how much earlier than the "next time" a packet can arrive before it is policed.

The question now becomes, "How does the network behave in response to packet that is either dropped or held (i.e., buffered)?" Adaptive flows, such as TCP, typically

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.