US005802106A

United States Patent [19]

Packer

[11] Patent Number: 5,802,106

[45] Date of Patent: Sep. 1, 1998

[54] **METHOD FOR RAPID DATA RATE DETECTION IN A PACKET COMMUNICATION ENVIRONMENT WITHOUT DATA RATE SUPERVISION**

[75] Inventor: **Robert L. Packer**, Los Gatos, Calif.

[73] Assignee: **Packeteer, Inc.**, Cupertino, Calif.

[21] Appl. No.: **762,828**

[22] Filed: **Dec. 6, 1996**

[51] **Int. Cl.**[6] **H04B 3/46**; H04B 17/00; H04Q 1/20

[52] **U.S. Cl.** **375/225**; 370/229; 370/230; 370/233; 370/253

[58] **Field of Search** 375/225; 370/229, 370/233, 234, 235, 236, 253

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,887,280	12/1989	Reisenfeld	375/121
4,912,765	3/1990	Virupaksha	381/31
5,072,407	12/1991	Gutz et al.	364/514
5,301,303	4/1994	Abraham et al.	395/500
5,440,298	8/1995	Kuramatsu	340/825.44
5,689,511	11/1997	Shimazaki et al.	370/545

Primary Examiner—Stephen Chin
Assistant Examiner—Mohammad Ghayour
Attorney, Agent, or Firm—Townsend and Townsend and Crew LLP; Kenneth R. Allen

[57] **ABSTRACT**

A method for data flow rate detection useful with data flow rate control is introduced into a TCP packet communication environment which does not have data rate supervision by computing presumed speed based on latency between packets during the initial interchange of synchronization (SYN) packet and the acknowledgement (ACK) packet, presuming to know the initial length of each. This information may be utilized to determine potential rate of data flow for further use in making bandwidth allocation and rate enforcement decisions.

8 Claims, 3 Drawing Sheets

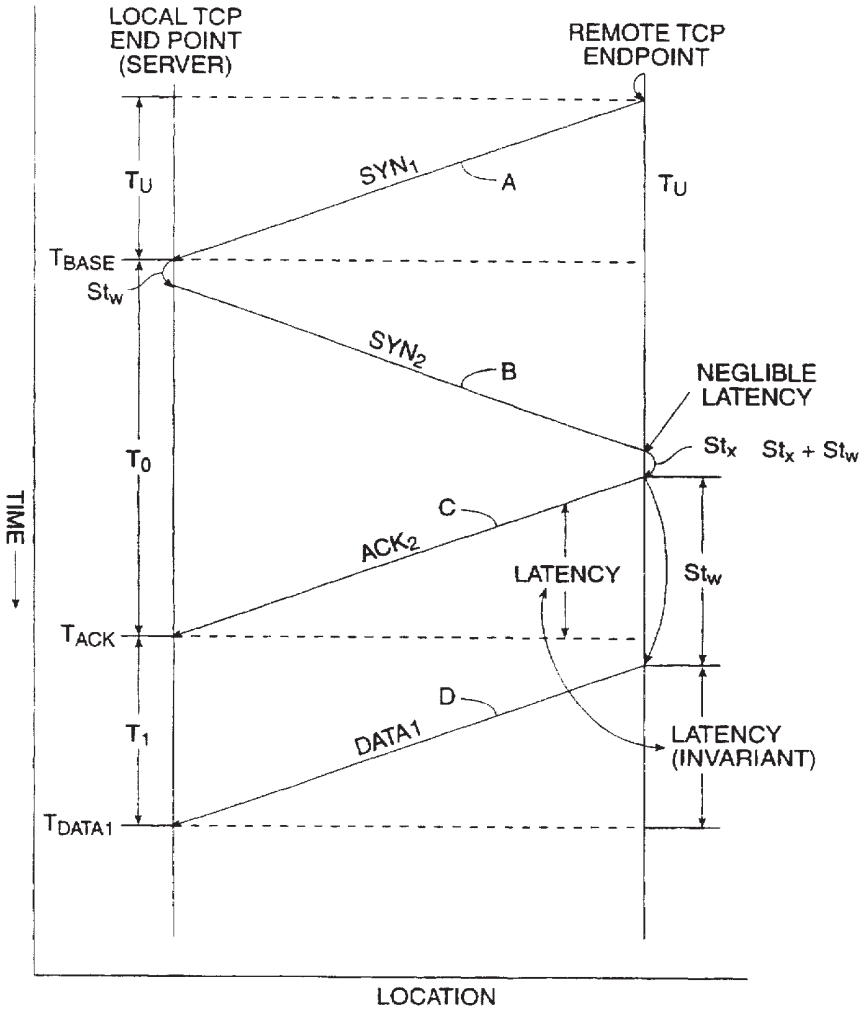

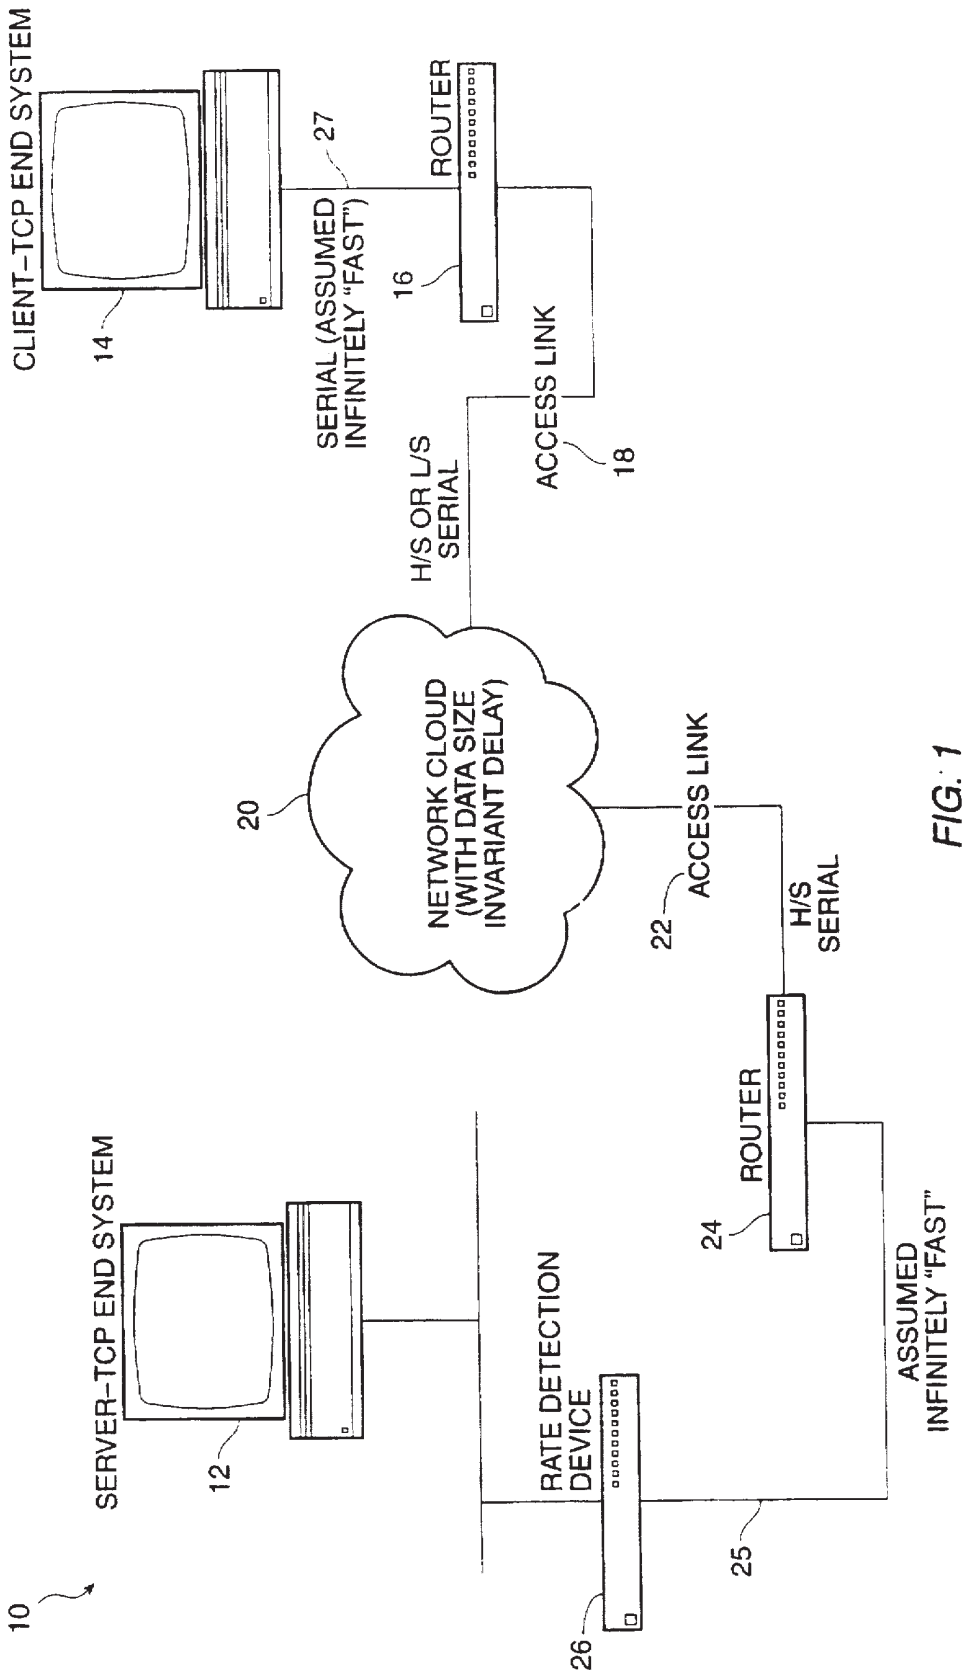

FIG. 1

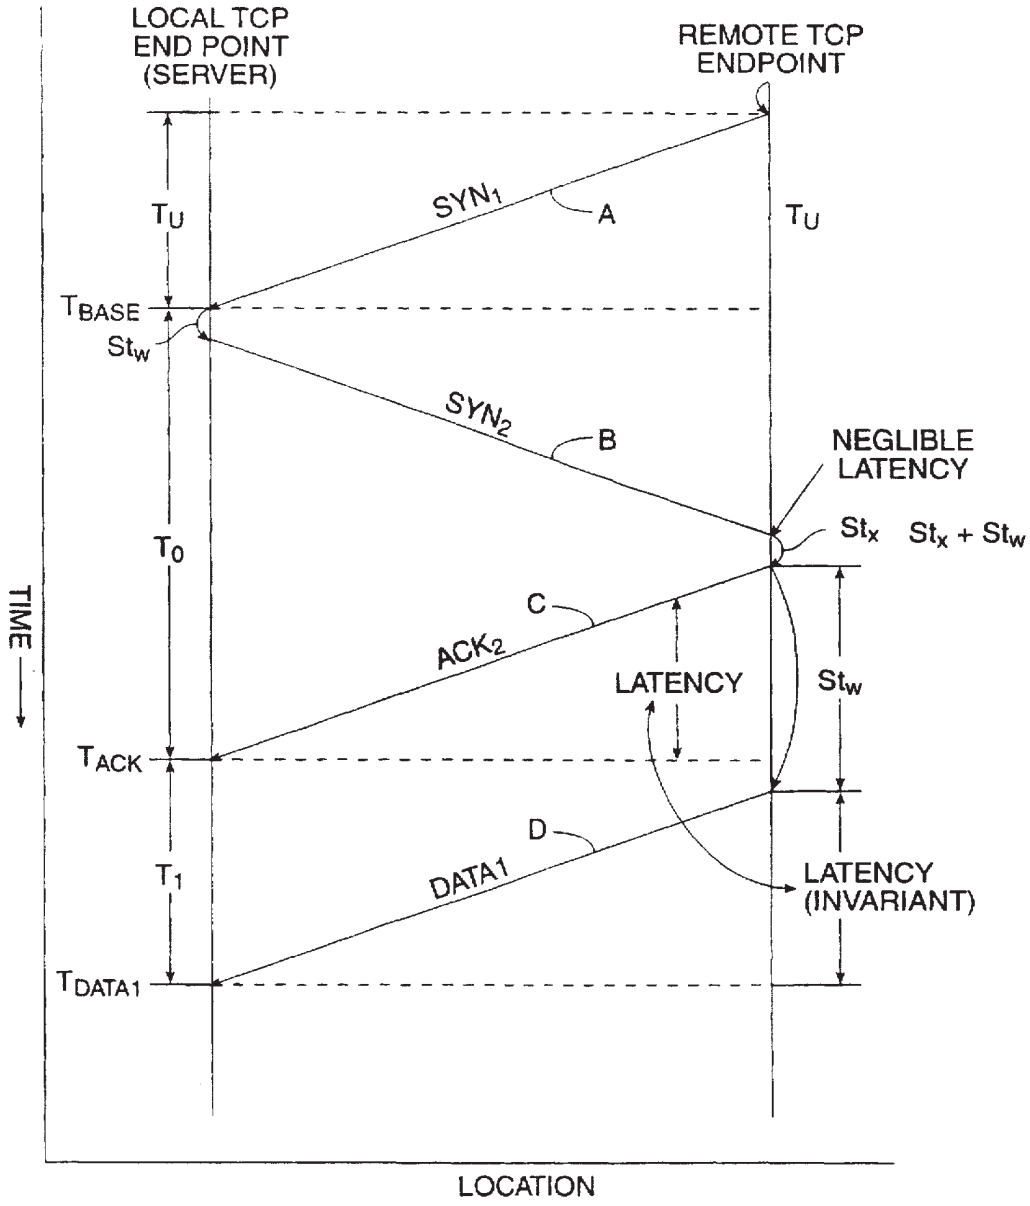

FIG. 2

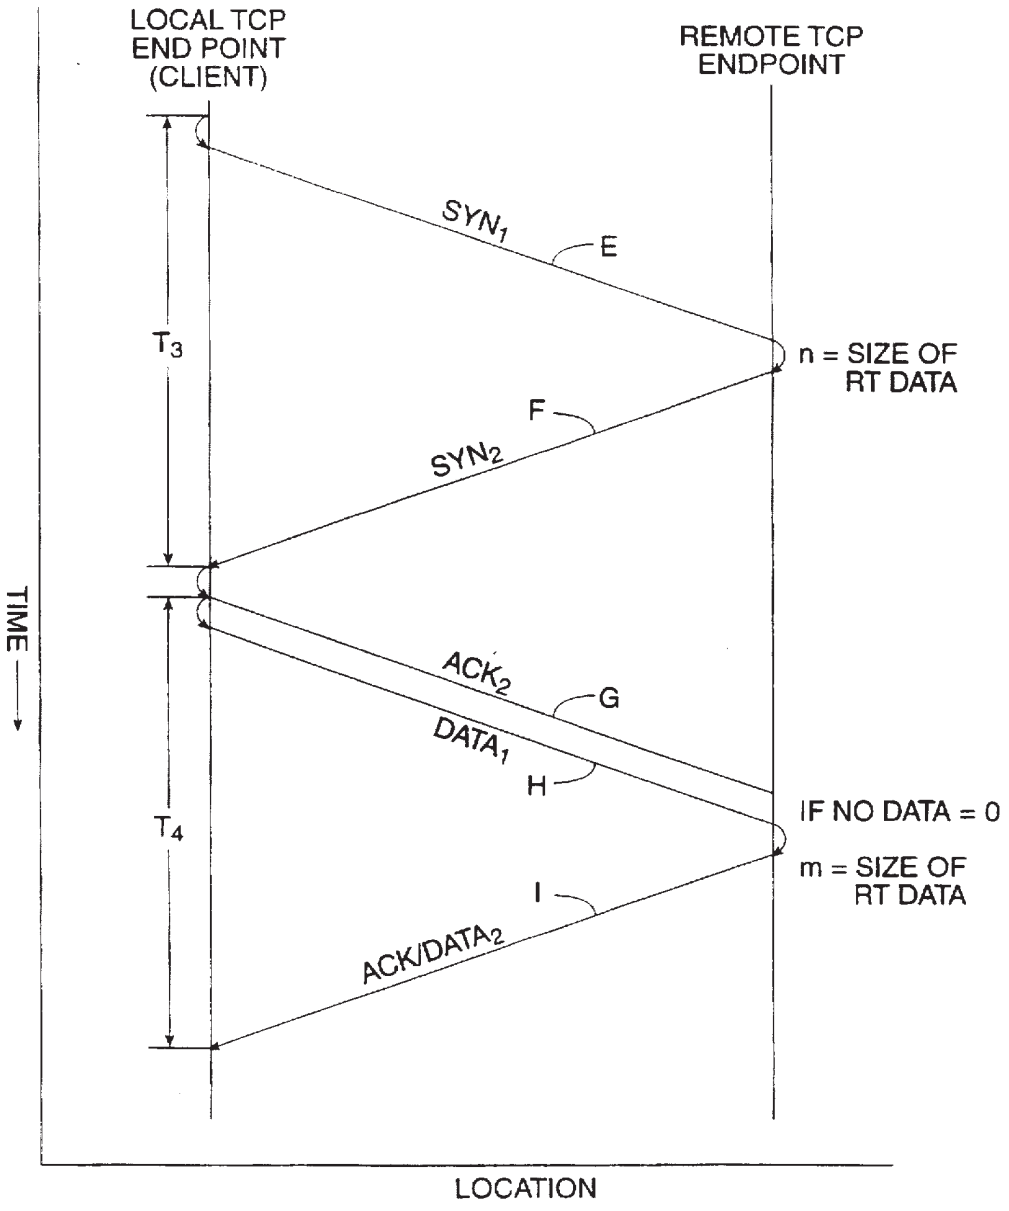

FIG. 3

METHOD FOR RAPID DATA RATE DETECTION IN A PACKET COMMUNICATION ENVIRONMENT WITHOUT DATA RATE SUPERVISION

COPYRIGHT NOTICE

BACKGROUND OF THE INVENTION

This invention relates to digital packet telecommunications, and particularly to data flow rate detection. It is particularly useful in conjunction with data flow rate control at a particular layer of a digitally-switched packet telecommunications environment normally not subject to data flow rate control wherein data packets are communicated at a variety of rates without supervision as to rate of data transfer, such as under the TCP/IP protocol suite.

The widely-used TCP/IP protocol suite, which implements the world-wide data communication network environment called the Internet and is employed in local networks also (Intranets), intentionally omits any explicit supervisory function over the rate of data transport over the various media which comprise the network. While there are certain perceived advantages, this characteristic has the consequence of juxtaposing very high-speed packet flows and very low-speed packet flows in potential conflict, which results in inefficiencies. Certain loading conditions can even cause instabilities which could lead to overloads that could stop data transfer temporarily. It is therefore considered desirable to provide some mechanism to optimize efficiency of data transfer while minimizing the risk of data loss. It is extremely useful to obtain an early indication of the rate of data flow which can or must be supported. In fact, a knowledge of data flow rate capacity is a key indicator for use in resource allocation decisions. For example, if a particular path is of inadequate capacity to handle a high rate of data flow, an alternative route may be sought out. To this end, some form of data flow rate detection is needed, preferably one which operates extremely fast, which operates on a packet flow as early as possible during a link up between nodes and which does not require a measurement on the substantive data to be communicated in order to determine the data flow rate.

In order to understand the exact context of the invention, an explanation of technical aspects of the Internet/Intranet telecommunications environment may prove helpful.

Internet/Intranet technology is based largely on the TCP/IP protocol suite, where IP is the network level Internet Protocol and TCP is the transport level Transmission Control Protocol. At the network level, IP provides a "datagram" delivery service. By contrast, TCP builds a transport level service on top of the datagram service to provide guaranteed, sequential delivery of a byte stream between two IP hosts.

TCP has flow control mechanisms operative at the end stations only to limit the rate at which a TCP endpoint will emit data, but it does not employ explicit data rate control.

The basic flow control mechanism is a 'sliding window', a range of bytes beyond the last explicitly-acknowledged byte, which by its sliding operation essentially limits the amount of unacknowledged transmit data that a transmitter can emit.

Another flow control mechanism is a congestion window, which is a refinement of the sliding window scheme involving a conservative expansion to make use of the full, allowable window. A component of this mechanism is sometimes referred to as "slow start".

The sliding window flow control mechanism works in conjunction with the Retransmit Timeout Mechanism (RTO), which is a timeout to prompt a retransmission of unacknowledged data. The timeout length is based on a running average of the Round Trip Time (RTT) for acknowledgment receipt, i.e. if an acknowledgment is not received within (typically) the smoothed RTT+4*mean deviation, then packet loss is inferred and the data pending acknowledgment is retransmitted.

Data rate flow control mechanisms which are operative end-to-end without explicit data rate control draw a strong inference of congestion from packet loss (inferred, typically, by RTO). TCP end systems, for example, will 'back-off', i.e., inhibit transmission in increasing multiples of the base RTT average as a reaction to consecutive packet loss.

Bandwidth Management in TCP/IP Networks

Bandwidth management in TCP/IP networks is accomplished by a combination of TCP end systems and routers which queue packets and discard packets when some congestion threshold is exceeded. The discarded and therefore unacknowledged packet serves as a feedback mechanism to the TCP transmitter. (TCP end systems are clients or servers running the TCP transport protocol, typically as part of their operating system.)

The term "bandwidth management" is often used to refer to link level bandwidth management, e.g. multiple line support for Point to Point Protocol (PPP). Link level bandwidth management is essentially the process of keeping track of all traffic and deciding whether an additional dial line or ISDN channel should be opened or an extraneous one closed. The field of this invention is concerned with network level bandwidth management, i.e. policies to assign available bandwidth from a single logical link to network flows.

Routers support various queuing options. These options are generally intended to promote fairness and to provide a rough ability to partition and prioritize separate classes of traffic. Configuring these queuing options with any precision or without side effects is in fact very difficult, and in some cases, not possible. Seemingly simple things, such as the length of the queue, have a profound effect on traffic characteristics. Discarding packets as a feedback mechanism to TCP end systems may cause large, uneven delays perceptible to interactive users.

A mechanism is needed to measure packet flow rate in order to control packet traffic efficiently.

It is particularly challenging to deal with TCP protocols because TCP protocols ignore the fact that delay is proportional to packet size or data length. The load issue means that there is a need to distinguish between high speed serial links over long distances and low speed serial links over short distances. TCP protocols cannot distinguish these two situations.

SUMMARY OF THE INVENTION

According to the invention, in a packet communication environment in which there is no data rate supervision to control assignment of available bandwidth from a single logical link to network flows, such as an environment using TCP protocols, a method is provided for making an early

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.