Cable Modem Bandwidth Provisioning and Quality of Service

Shlomo Rakib
Terayon Communication Systems

Abstract

Cable operators today are challenged to provide data services to a broad range of users with differing bandwidth and Quality of Service requirements. The nature of a shared cable network requires systems that allow operators to provision and allocate the bandwidth to multiple tiers of users. Only in this way can operators match bandwidth to the needs of both residential and commercial users, and to build a service model that maximizes revenue. To support such a model, the cable operator should deploy a cable modem system that provides adequate bandwidth capacity in both directions, and the ability to provision the bandwidth to support a broad range of user application needs.

This paper analyzes the performance of Terayon's TeraComm system, operating in its UBR (unspecified bit rate) mode, which along with CBR (constant bit rate), represents the system's two primary MAC layer modes. These test results clearly demonstrate the sophisticated bandwidth management capabilities of the TeraComm system, which support multi-tiered data services. The Terayon system achieves maximum bandwidth utilization and maintains fair allocation of bandwidth among users, maintaining minimum latency for access and data transfer. This capability enables a new generation of data services and broadband applications.

PHYSICAL LAYER

The TeraComm system uses PHY and MAC layers which are both unique in the cable modem industry. With Terayon's Synchronous Code Division Multiple Access PHY implementation, a set of 144 orthogonal codes are used to allow simultaneous transmission from up to 144 individual data streams on each channel. The system defines a separate channel each for upstream and downstream; thus there are 288 simultaneous transmissions streams allowed. The transport rate of each transmission stream is 72kbps. The system reserves 128 data streams for user data. Allowing for the overhead of data cell framing mechanisms, the capacity for data is 64kbps per stream. The remaining 16 data streams are used for system management and access control. This means user data traffic never contends for bandwidth with management or access request traffic (see figure 1).

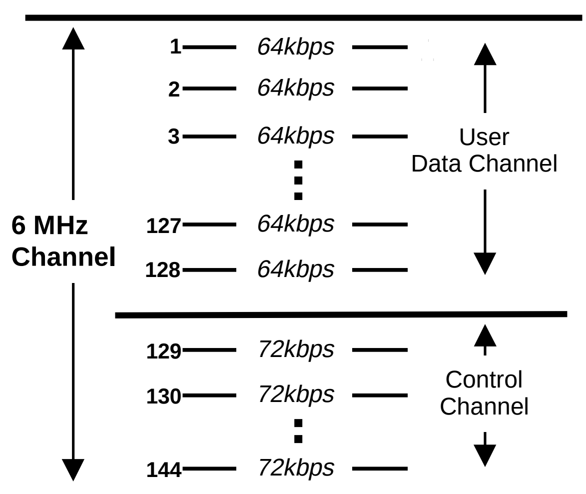

Figure 1

MEDIA ACCESS CONTROL LAYER

The 128 users' data stream codes are managed and assigned to modems by the Bandwidth Manager at the TeraLink headend controller. Figure 2 demonstrates how one or more codes can be assigned on a permanent basis (CBR mode) or on a dynamic basis (UBR mode). Modems requiring access to one of the 128 user data streams will contend for access using 4 slots (see figure 2) in the control channel. Providing the higher transmission rate of 4 slots, rather than a single slot, lowers the probability of corrupting an access request due to impulse noise. These time slots also lessen the probability of a collision during the access process. Each modem will randomly choose one of these time slots to issue its access request. After a successful access has been achieved, additional codes may be allocated based on measured usage within a small time window, rather than via additional access requests.

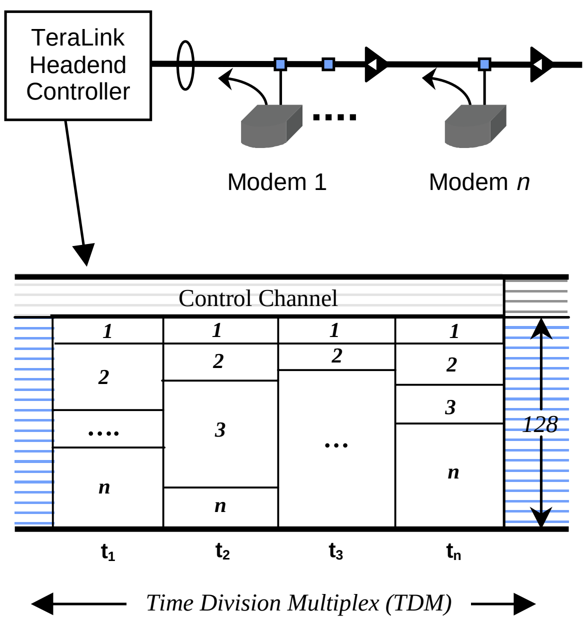

Figure 2

This means that some calculable latency will be incurred only during initial access. In the event that the channel subscription is higher than the available channel rate, data streams will be distributed among multiple modems without requiring further access requests. For this reason, there is zero latency incurred at the MAC layer after initial access in an "undersubscribed" system. MAC layer latency in an "oversubscribed" system is discussed in the next section. While this MAC technique is much like that used for standard Ethernet or other collision-based contention techniques, there are some important differences. By using orthogonal codes for data streams that are separate from those used for access contention ensures that even a high rate of access collisions has no effect on channel efficiency. A standard rule of thumb for Ethernet network scaling is to allow no more than 30% utilization on average, due to the fast break-down of the channel under high contention rates. The actual traffic being sent is used for contention in an Ethernet network, therefore the amount of data (channel capacity) lost and to a collision event can be large. The effect of this data loss is that much of the channel capacity is spent resending data. The TeraComm system can always run at 100% utilization, independent of the amount of access requests being processed.

Finally, within an Ethernet network, the probability of a collision is proportional to the size of the datagram sent. Also, in an Ethernet network all traffic is open to collisions. With the TeraComm MAC, the smallest possible station identifiers are used for access requests, so there is no data lost in a collision event. Due to its small size, the probability of collision is small. Finally, no access requests are required for bandwidth after the initial access is granted in an "undersubscribed system." The Bandwidth Manager task may automatically grant more bandwidth based on usage, system load and operator-defined Quality of Service provisioning.

RESPONSE LATENCY OF AN OVERSUBSCRIBED SYSTEM

A system in which the requested bit rate is higher than the available bit rate is said to be "Oversubscribed.” In a system that supports guaranteed Quality Of Service, over-subscription occurs when the sum of bandwidth reserved for Constant Bit Rate plus the maximum bandwidth allocated to Unspecified Bit Rate applications, (minimum bit rate = 0) is greater than the total bit rate of the system. In a purely contention-based system, over-subscription is merely implied by an access latency which on average is not tolerable by the applications which are running over the system. As collisions reduce the capacity of the data channel in a purely contention-based system, which in turn increases the required access time, access latency will grow non-deterministically.

In an undersubscribed system, the TeraComm MAC uses time division multiplexing of codes to reduce the number of access requests, thus avoiding non-deterministic access latency. Under normal operation, modems that are no longer utilizing granted bandwidth will lose all but their last data code. After the initial code is granted, access latency is zero from then on. In the event that the system is "oversubscribed,” that last code will also be revoked in the event of an access request from another modem. While the system is in an "oversubscribed" state, multiple modems may share a set of unreserved codes which are periodically multiplexed from modem to modem. As the number of UBR modems increases in comparison to the number of unreserved codes, the duty cycle of allowable transmission time decreases.

For example, while 129 modems sharing 128 codes may each transmit 99% of the time, 256 modems sharing 128 unreserved codes may only transmit 50% of the time. The time spent awaiting re-allocation of a code is considered response latency. The average response latency of a time division multiplexed system is then deterministic and is derived from the product of the multiplexing period and the ratio of transmitting UBR modems to unreserved codes.

BANDWIDTH MANAGER RESPONSE IN AN UNDERSUBSCRIBED SYSTEM

In an undersubscribed system, where the number of active modems is less than the number of unassigned time slots, the bandwidth manager will fairly distribute bandwidth among the active modems. Any modem that utilizes more than 60% of its assigned bandwidth is eligable for an increase in allocation. To efficiently use the channel in an environment where data is TCP based, the rate at which capacity is moved between active modems must compliment TCP Slow Start implementations.

MAC TEST RESULTS

The TeraComm MAC implementation performance is superior to other collision-based contention techniques, providing deterministic latency for access and data transfer and for data channel efficiency. This remains true for undersubscribed as well as oversubscribed channel conditions.

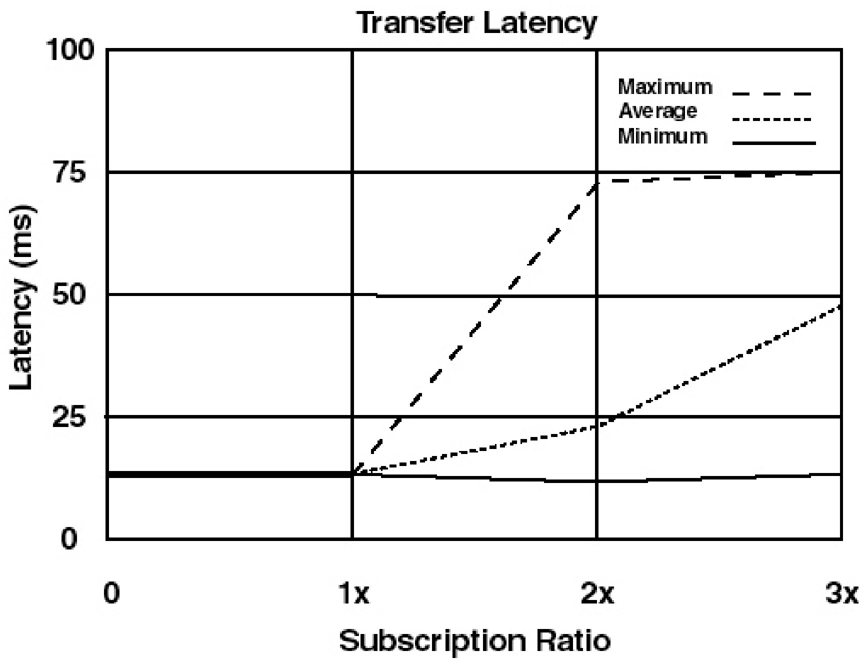

Figure 3: The transfer latency with access in the Terayon system is far superior to alternative solutions due to the architecture of the Terayon MAC and capacity of the channel supported by S-CDMA.

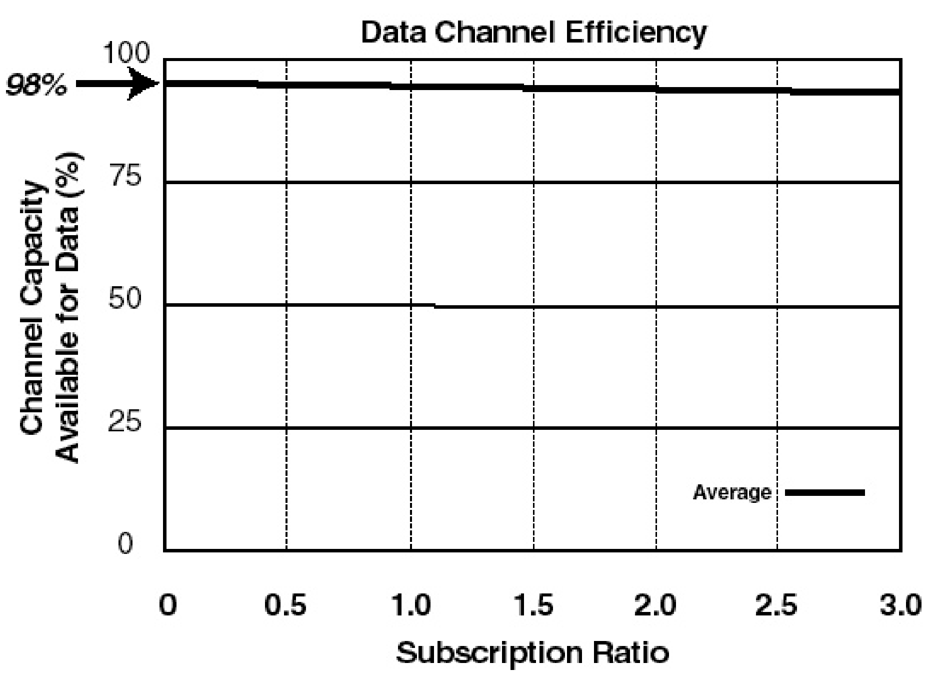

Figure 4: The low transfer latency in the Terayon system makes it well suited for time-sensitive applications, such as IP telephony, video conferencing, and online games. This is even true of best effort class of service systems with high over-subscription ratios.

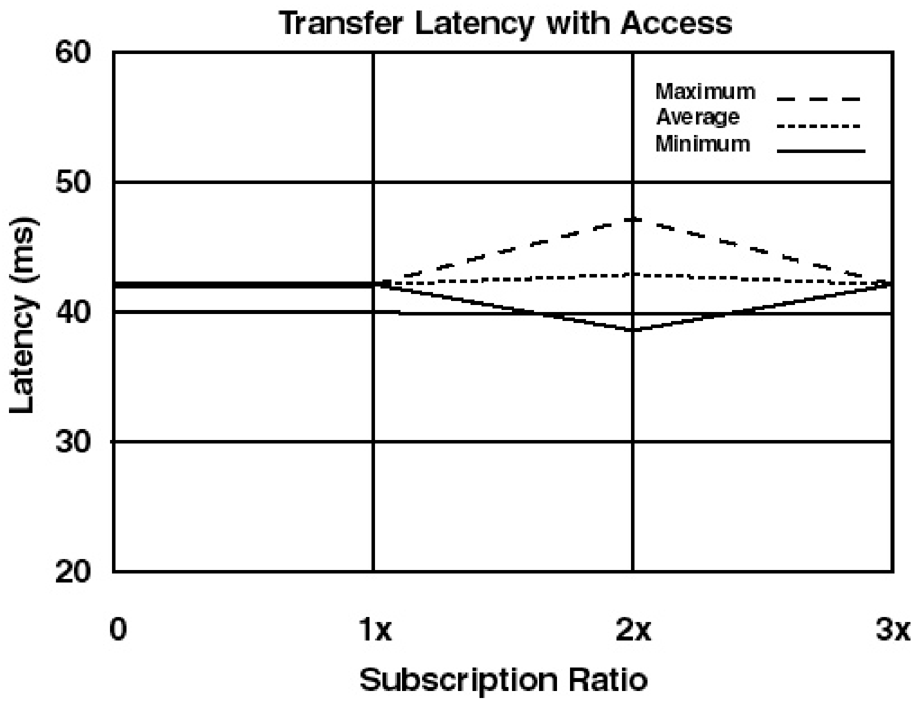

Figure 5Another important performance parameter is channel efficiency, or the impact on data channel capacity as more users are added and access requests to the channel increase. Because of separation of the data and control channels, channel efficiency is only slightly impacted by oversubscription.

AUTHOR'S BIOGRAPHY

Shlomo Rakib
President and Chief Technical Officer
Terayon Communication Systems
2952 Bunker Hill Lane
Santa Clara, California 95054
(408) 727-4400

Shlomo Rakib is President and Chief Technology Officer of Terayon Communication Systems, the cable modem access company he co-founded in 1993. Mr. Rakib developed the technical roadmap for Terayon, based on his vision of the growing need for broadband access to the home and office. Mr. Rakib invented and filed patents on S-CDMA (Synchronous Code Division Multiple Access) technology, which is optimized for robust performance over today's cable systems. Mr. Rakib has over 15 years experience designing products for the cable industry, including 9 years as Chief Engineer at the communications company, PhaseCom. He holds a bachelors degree in engineering from Technion University in Israel.