Optimizing Network Bandwidth Utilization With Flow Rate Control

Internet Traffic and the Bandwidth Management Challenge

The exponential growth in global Internet traffic volume is being fueled by continued broadband service deployments, enterprises turning towards cloud computing services, the rapid growth of smartphones and other intelligent mobile devices with always-on Internet connections, and the change in traffic mix from email and other bulk data applications to rich media and time-sensitive services. TeleGeography's Global Internet Geography study reveals that international Internet traffic grew 62% in 2010 after realizing 74% growth recorded in 2009. Furthermore, it has been forecasted that global Internet traffic will increase more than four-fold by 2014. To accommodate such increased traffic demand, significant network capacity has been added by broadband service providers. In 2010 carriers added 13.2Tbps of new international capacity, up from 9.4Tbps in 2009, and 6Tbps in 2008.

A phenomenon and problem that just about every carrier or institution worldwide has seen is that although no matter how much capacity is added to their network, it is used almost instantly as it is the nature of Internet Protocol (TCP/IP) to consume all the bandwidth that is available. Capacity added one day could easily be fully consumed the next day. Furthermore, although the traffic volume growth is high, typically a very small percentage of users consume the vast majority of capacity on any given network. Most Broadband Service Providers (BSP's) have found that as little as 4-5% of their users are consuming up to 80% of all their network capacity, leaving a limited amount of capacity for the vast majority of subscribers. In an environment where growth of quality sensitive applications such as video, voice, and gaming are accelerating at exponential rates with online video alone (Internet, TV, VOD and P2P) forecasted to exceed 90% of all consumer traffic by 2014, the negative implications to service quality and profitability are huge.

Thus the challenge for BSP's, be they wireline, wireless or cable network operators, is how to optimally improve network utilization and improve the users' quality of experience (QoE) in the most efficient, cost effective manner.

The Underlying Technology Problem: TCP Behavior and the Impact of Queues On Network Congestion

TCP/IP packet networks, including the Internet, operate as best effort networks with no assured QoS, typically displaying considerable delay, delay jitter and major bandwidth rate variances especially at network congestions points. These attributes are caused by the interaction of TCP with queue based network elements. The problems are aggravated for the rapidly growing Real Time Traffic (video and voice) being carried by wireless and wireline packet networks.

As packet traffic traverses the Internet, TCP grows the data rate, sending larger and larger blocks of packets in bursts until a packet is reported lost by the receiver, at which point the TCP sender drops to half rate and then starts to grow again. This causes (1) the rate to increase until the path overloads and packets are randomly discarded during periods of network congestion or (2) all the flows in the path grow until they have reached their "maximum rate" when the network is not congested.

The current generation of routers and other network equipment use queues to avoid the overload of the next network link or trunk. The queues are simple buffers which drop packets if the queue is full. In addition, Weighted Random Early Discard (WRED) is currently used where a small percentage of packets are discarded if the queue is partially full, and all packets are discarded if the queue is full. All such methods create multiple impacts on TCP with several undesirable effects of handling network congestion.

Delay: Because TCP rate is dependent on the Round Trip Time (RTT) of the path, one immediate impact is the delay caused by the queue which slows down TCP flows. Typically, bottlenecks in any path occur at one or both edges of the network. Thus, if a TCP flow crosses two edges where overload tends to occur, then the delay in the two network edge queues add to the round trip transit time, potentially in both directions. Often the queuing delay is greater than the transit time, adding 40 ms at each edge queue for a total of up to 200 ms.

Delay Jitter: Perhaps the most harmful aspect of the added queuing delay is that it varies as the load varies. Thus in the prior example, the delay could vary anywhere from 40 ms to 240 ms, creating a 200 ms delay jitter. From a technology perspective, such large delay jitter is unacceptable for real time applications such as VoIP and video. Given the desire to improve the users' quality of experience, coupled with the increased revenue opportunity associated with the delivery of high quality VoIP and video services, there is a substantial incentive to eliminate the jitter.

Stalled TCP: When a queue gets full and must drop all the new packets, some TCP flows suffer from losing more than one packet. Since TCP bursts a group of packets each RTT, the flows' whose bursts just arrived at the queue generally lose two, three or more packets from one burst. TCP rules then tell the sender to drop to slow start (one packet per RTT). Even more harmful are the cases where the TCP process sees the loss of more packets in one RTT and goes into a stall. The stall is a period with no transmission and this period is extended longer and longer if additional losses are detected.

The ramifications of such stalls can be significant. For instance, a typical web access may have 40 flows, each delivering 20-60 kb of data. In a network with even moderate overload, these stalls will occur on many of the 40 flows. The slowest flow determines the display time for the web page. Queues spread out the distribution of flow rates so much that web pages are typically three times slower to complete than would be true if the flows kept about the same rate. In fact, experiments have shown that if the stalls are eliminated one can improve web page response time by a factor of 3X.

Synchronization: When the TCP traffic on a trunk approaches 50% utilization it tends to synchronize. Initially a few flows peak at once and thus the queue drops some of their packets. If the flows' Round Trip Delays are similar they will tend to peak again at the same time. When this occurs additional flows get caught in the peak and the aggregate of all flows peak together. At some point all the TCP flows with similar RTT's will peak together and the channel will suffer considerable delay jitter and packet loss. This is why most trunks need to be upgraded when they exceed 45% utilization in the peak hour. However, if the packet drops are organized to avoid synchronization utilization can be increased to about 95%, thus doubling the effective capacity of the trunk.

The Solution: Flow Rate Control (FRC) Internet Traffic Management Elements

To briefly summarize the balance of this paper, flow rate control (FRC) systems accurately manage the rate of every flow with the composite rate of all flows in any subset of the traffic to be maintained within 90% - 95% of the capacity of a multiplicity of network congestion points. Three critical benefits are realized: 1) significant improvement in the users' quality of their Internet experience, 2) material improvement in network capacity utilization, lowering both Service Provider's CAPEX and OPEX, and 3) increased revenue generation opportunities through the delivery of tiered services.

Flow Rate Control: From a macro perspective, a flow rate control-enabled network element looks into a flow table and finds a match for the five IP packet fields that define a flow (source address, destination address, protocol, source port and destination port). These IP packet fields determine where and how the new packet should be sent on the same route computed for the first packet of the flow.

The flow rates, as well as other behavioral parameters, are collected that allow the TCP flow rates to be precisely controlled at the input to a system. If the flow rate is excessive, one and only one packet is dropped such that the TCP flow operates at precisely the rate that the flow should be allowed. Clearly if all flows could be managed at an aggregation point, the total rate of any group of flows could be controlled such that when that group arrived at a congestion point there would be no need for delaying packets or for dropping packets randomly at the congestion point. Figure 1 shows the structure of a flow rate control system.

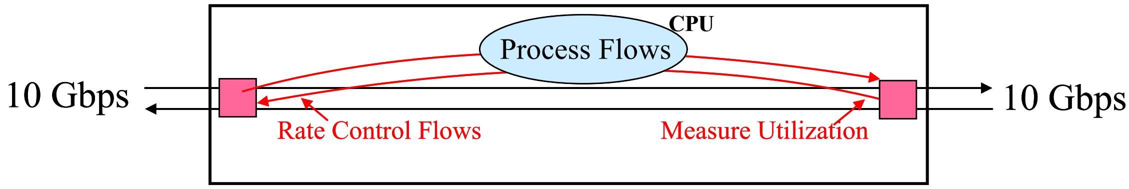

Figure 1. Flow Rate Control System

Flow Rate Priority: The historical concept of capacity allocation for a group of flows was to divide capacity equally between all flows. This is the natural outcome of TCP and queues and was acceptable when the Internet got started since typically each user used one flow and this made all users equal. If one flow was sending more packets than another, the random discard at the queue would on the average discard more packets from the larger flow, slowing it down such that all flows tend to the same rate. Today, however, a more flexible allocation of rates per flow, flow rate prioritization, is required to allow each flow or subset of flows to receive different rates. Flow rate priorities can be used so that the rates received are proportional to the priorities.

Flow rate control systems also provide benefits to User Datagram Protocol (UDP) and layer 2 protocols. By precisely controlling TCP rates, non-TCP traffic can be protected and prioritized, minimizing the effects of delay jitter and reducing the pressure on higher latency TCP queues. This improves the experience and responsiveness of streaming traffic, VPN sessions, or other quality-sensitive traffic. Alternatively, non-TCP traffic can be controlled just as TCP can, giving an equitable amount of bandwidth to each flow. This protects other flows from aggressive applications that use UDP or tunnels attempting to circumvent TCP back off mechanisms and use more bandwidth than they would normally receive. FRC's are also capable of transparently performing UDP-based Call Admission Control (CAC), only accepting new sessions if sufficient bandwidth exists to support it and existing sessions without reducing their quality.

Impact of Flow Rate Control on Network Congestion at the Broadband Edge

In normal usage of the Internet the primary point where sufficient overload occurs is at the broadband edge between the BSP and the customer. For enterprise customers, traffic management will be most effective at the corporate edge given the worst congestion point is generally the connection to the Internet. For residential broadband customers, the congestion point is usually at the BSP's service concentration points for subscribers, i.e. the DSLAM, CMTS or GGSN backhaul links.

Net Neutral Management of Abusive Bandwidth Users

One of the major problems facing BSP's today is that a small fraction of the users consume the majority of the pool of network bandwidth capacity intended for a entire community. To date, P2P applications have been the primary offenders. By its very nature a P2P application opens many flows to download music or videos. Each flow tends to receive equal bandwidth due to TCP's interaction with queue-based congestion management. P2P applications having 100 flows receive 100 times the capacity of other single flow applications. As application designers have determined that the ability to open several flows increases overall application throughput, other non-P2P applications are starting to utilize multiple flows.

The FCC in the U.S. has ruled that Service Providers should not discriminate against a particular application but can exercise fair and equitable traffic management to maintain good service levels. Anagran's patent-pending FRC technology is uniquely suitable for providing fair bandwidth allocation in an application agnostic method. Anagran FRC Internet Traffic Managers (ITMs) measure all the traffic and flows per subscriber and in cases of network congestion, rate control all flows such that multi-flow users receive the same capacity as the other normal subscribers. This is referred to as “subscriber equalization”. FRC is not fooled by encryption but is application agnostic thus meeting the FCC requirements for a net neutral method of traffic and bandwidth management.

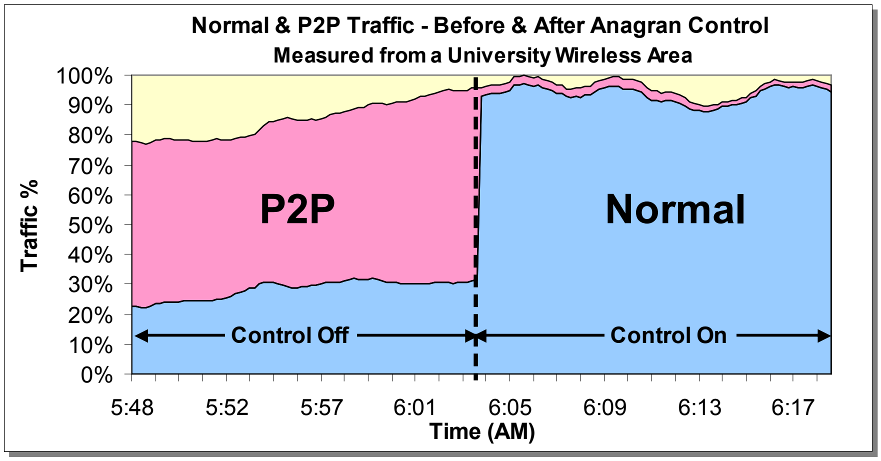

Figure 2: Example of Subscriber Equalization

A specific example of FRC control of P2P traffic overloading a network is shown in Figure 2. P2P bandwidth hogging is the most aggravating problem caused by the "equal capacity per flow" rule. In measurements made at a University which installed Anagran FRC Internet Traffic Managers, P2P dominates the capacity until subscriber equalization was turned on at 6:03AM. At that point the P2P traffic received its fair share of bandwidth and the average user's capacity increased fourfold.

The solution for Service Providers is to deploy FRC technology, implementing "equal capacity for equal pay". Where Service Providers have deployed Anagran FRC Internet Traffic Managers for subscriber equalization, Quality of Service (QoS) and Quality of Experience (QoE) radically improves for all users, with multi-flow users still receiving their fair share of capacity.

Prioritizing Subscriber Traffic

One extremely valuable feature of the Anagran FRC Internet Traffic Manager is the ability to assign a rate priority to flows, subscribers and/or traffic classes.

Subscriber Priority: Depending on the service a given subscriber pays for (e.g., gold, silver, or bronze), Anagran's FRC system can set the designated flow rate priority for each subscriber. For example, assume gold customer gets a rate priority of "4", silver "2", and bronze "1". The gold user will have a flow rate four times the rate of the bronze subscriber's flows and twice the silver subscriber. Each class of subscribers can also be assigned a traffic rate cap which controls their peak rate during non-congestion periods. With both these controls a BSP can provide an extremely effective boost in their network performance and productivity.

Traffic Class Priority: One way to get much improved utilization from a fixed size Internet access link is to establish traffic classes, say bulk traffic and interactive traffic, and then prioritize such traffic accordingly. Anagran's FRC Internet Traffic Manager can identify bulk traffic with a "behavioral" command. For instance, the command would say: "If the number of bytes received on a flow exceeds 100,000, it is reclassified as bulk". Then if "bulk" is assigned a priority of ½ of interactive, the bulk traffic will flow rapidly if there is little interactive traffic but be suppressed when the interactive traffic peaks. This lets the bulk traffic fill the valleys between the interactive peeks rather than pushing the peak up even higher.

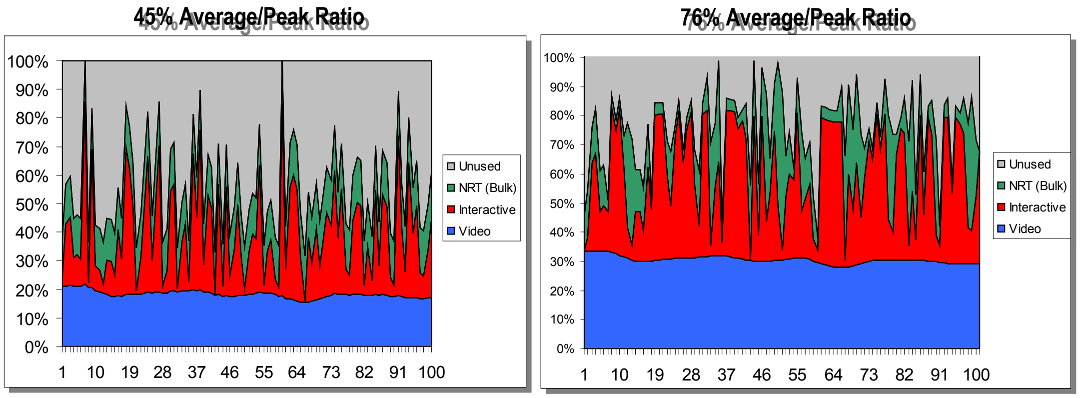

Figure 3: Bulk Traffic without (left) and with (right) Traffic Class Prioritization

Figure 3 is an example which in the left table indicates equal traffic class priorities as compared to the table on the right where interactive traffic has twice the priority of bulk traffic. The response time for interactive traffic is the same in both cases but the throughput of all three types of traffic has been increased 68%.

Summary

A fundamental challenge for network operators is the tremendous increase in wireline and wireless Internet traffic volume and the consequential effect on overall network performance, CAPEX and OPEX investments, and opportunities to increase ARPU via tiered services. As traffic is increasingly generated by applications sensitive to delay and jitter (voice and video), any effect even of temporary congestion on a network link results in a very noticeable decrease in the quality of the service by the end user.

Anagran's Flow Rate Control Internet Traffic Manager can be deployed transparently in the existing network infrastructure consisting of legacy packet-based networking devices to provide a significant improvement in network efficiency and quality of experience. Rather than managing individual packets with their associated increased delay and jitter, flow rate control technology manages flows per subscriber or traffic class. With a consistently predictable service quality, network operators have a greater opportunity to build profitable services for their customer base without the necessity of expensive capacity increases. With Anagran's unique flow rate control technology this approach immediately yields Service Providers improvements in the efficiency of link utilization, lowering CAPEX and OPEX, improving ARPU and their customers' overall user experience.