Performance Analysis of Real-time Streaming Media over Shared Medium Access Networks

James Martin
Michele Weigle
James Westall

Department of Computer Science
Clemson University
Clemson, South Carolina 29634-0974

1 Introduction

The objective of the proposed study is to evaluate the impact of shared medium access networks on the performance of streaming media applications having real-time constraints. Applications such as IP based telephony (VoIP) and video conferencing have such constraints, and the performance of VoIP over shared medium access networks will be the focus of this research. Specifically, we will characterize the effect of increasing levels of competing traffic on the quality of a VoIP connection in which one or both endpoints reside on shared medium access networks in which QoS support at the MAC layer is limited or non-existent.

1.1 Evolution of the access network

Twenty years ago shared medium access networks in the form of shared bus Ethernet and token ring networks were pervasive. These wired LANs employed CSMA-CD and token passing protocols to arbitrate access to a shared bus having aggregate bandwidth on the order of 10 Mbps. By the mid 1990's, shared medium access networks were rapidly disappearing as existing installations were replaced by fully switched Ethernet and ATM LANs. These switched LANs provided point-to-point bandwidth on the order of 100 Mbps and aggregate LAN bandwidth of 10Gbps or higher. In the past 10 years wired LANs have experienced another order of magnitude growth in both point-to-point and aggregate capacity, and the shared medium access network has reappeared in the form two important new and evolving technologies: cable and wireless.

1.1.1 Cable access networks

Cable networks employ a new and complex MAC protocol called DOCSIS (Data over Cable Service Interface Specification). DOCSIS itself is evolving rapidly, has not been widely studied by academic researchers, and will be a central focus of the proposed study. The physical layer of a DOCSIS system is a shared access cable. Medium access is controlled by a head-end device called the Cable Modem Terminating System (CMTS). Packet flow is always between the CMTS and the Cable Modems (CMs), never CM to CM. Downstream bandwidth is shared in an asynchronous TDM fashion under the control of the CMTS using a possibly prioritized form of round robin scheduling.

Multiple CMs must contend for access to the upstream channel, and the channel allocation procedure is quite complex. The upstream channel is subdivided into transmission slots referred to as mini-slots. The capacity in bytes of a mini-slot on a given DOCSIS network is fixed and is typically near that of an ATM cell. Permission to transmit data in a block of one or more mini-slots must be granted to a CM by the CMTS. The CMTS grants mini-slot ownership by periodically sending a frame called the MAP on the downstream channel. In addition to ownership grants, the MAP also typically identifies some mini-slots as contention slots in which CM's may bid for for quantities of future mini-slots. To minimize collisions in the contention slots, a non-greedy backoff procedure is employed. Each CM is required to randomly select the contention slot in which it transmits a bid for mini-slots. When collisions do occur in contention slots, all parties that collide are required to employ an exponential backoff, doubling the size of the window of slots in which the next bid is randomly placed. The MAP may also indicate that a mini-slot is reserved for a particular CM to bid for future mini-slots if it wishes to do so.

DOCSIS also provides mechanisms for assigning desired QoS attributes to specific *service flows*. The unsolicited grant service (UGS) is an example. In this premium service, the CMTS grants mini-slots to a service flow of a CM at a prearranged rate, effectively providing that flow a CBR channel. This facility can clearly support large scale deployments of toll quality VoIP, but at a premium cost to the cable subscriber. Therefore, for the foreseeable future, it seems likely that non-cable operators such as Vonage and Sprint will continue offer lower cost VoIP access to the PSTN using best effort DOCSIS. Our focus will be on characterizing the performance of VoIP using the DOCSIS best effort service, and to investigate the use of service flow prioritization in DOCSIS 1.1 to provide better than best effort QoS that is short of the CBR service available with UGS.

1.1.2 Wireless access networks

IEEE 802.11 wireless networks are widely deployed, and their use is rapidly growing. Although 802.11 networks have been the subject of considerable study by academic research community, both the 802.11 standards themselves and the applications using the wireless networks are rapidly evolving, thus opening new avenues for investigation. Like CMs, 802.11 transmitters cannot detect collisions in which they participate, and the 802.11 MAC protocol is commonly called Carrier Sense Multiple Access with Collision Avoidance (CSMA-CA). The collision avoidance algorithm is also very similar in spirit to that used in DOCSIS. Following a successful transmission, all stations are required to use a slotted time model. A station with traffic to send must randomly select a slot, wait until that time slot is reached, sense the channel, and transmit only if the channel remains available.

Initial deployments of 802.11 supported no QoS capabilities at all, and that remains true of low-end access points (APs) and end systems today. However, QoS support is gradually appearing in 802.11 networks in the form of both emerging standards and vendor specific extensions to the standard. High-end APs now commonly support 802.1q VLANs. This permits priority based packet scheduling *at the AP itself*, but the facility is not sufficient to ensure QoS to the endpoints because of the possibility for unbounded delay in the CSMA-CA network. To address this issue, the emerging E-DCF standard allows both the AP and end systems to use backoff windows in the CSMA-CA protocol that are dependent on the QoS attribute of the VLAN to which the packet belongs. When non-overlapping windows are used, absolute priority of VoIP over non-voice traffic is guaranteed. Nevertheless, as with DOCSIS, our focus will be on characterizing the performance of VoIP using IEEE 802.11 best effort service, and to investigate the use of flow prioritization capabilities of 802.1Q E-DCF to provide better that best effort QoS that is still short of absolute priority of VoIP over competing traffic.

2 Proposed research

In this project, we will evaluate, under a variety of access network configurations and competing traffic loads, the effect of competing traffic on a VoIP connection. In each configuration we will provide a quantitative characterization of the loss of call quality as function of the increasing intensity of competing traffic. Our objectives are to characterize access network configurations and load levels for which the use of VoIP over best effort is problematic and to quantify the way in which the VoIP connection is degraded.

We will also explore the degree to which the negative impact of competing traffic may be mitigated by incremental application of differentiated services at the MAC layer. This can be done in the DOCSIS environment by assigning different service characteristics to the service flows associated with VoIP endpoints. In the 802.11 environment it may not be possible to obtain configurable E-DCF phones or NICs, but the effect of incremental differentiated services on a simplex flow can be explored via the configurable backoff window on a Cisco AP. We understand that market forces make it likely that only best effort service and premium services with *guaranteed* QoS will be offered by equipment vendors and operators. Nevertheless, from a purely scientific perspective, we are interested in better understanding the effect of providing incremental differentiated services at the MAC layer.

The research will be conducted on a private network testbed. Switching elements in the testbed will consist of a Cisco uBR7114 CMTS, a Cisco Catalyst 3750 switch, and two Cisco 802.11a/b/g access points. The CMTS and twenty cable modems are on presently on loan from Cisco. The other devices are being requested in this proposal. Two Cisco wireless phones and two Cisco Ethernet phones are also requested in the proposal.

2.1 Access network configurations

We will study VoIP over DOCSIS, VoIP over 802.11b, and VoIP over 802.11b over DOCSIS. The latter configuration is becoming an increasingly popular way to provide low-cost Internet access in the home and small business. For cost reasons, best effort provisioning over all of these configurations will remain common for the foreseeable future.

In each configuration, exactly one VoIP connection will be active. One endpoint of the connection will always reside on the shared medium access network under study. Depending upon the objective of the particular study, the other endpoint may reside on the same shared medium access network, a different shared medium access network, or a private gigabit LAN.

2.2 Competing traffic sources

In previous published work done by us and others, it has been found that realistic competing traffic loads may be achieved with no more than twenty competing hosts. We will use twenty traffic sources, of which one will be an endpoint of the measured VoIP connection.

The other nodes will operate as stochastically identical sources with parameterized offered load. Details of the load model are not yet fully defined but are likely to include (in increasing level of complexity): constant bit rate (CBR) sources; *on/off* sources with heavy-tailed state holding times; synthetic Web traffic sources such as the *Surge* system developed at Boston University.

2.3 Performance measurement

The Mean Opinion Score (MOS) is a numeric quantification of the average call quality as perceived by a large collection of listeners. Several published studies have proposed algorithms for mapping measurable attributes of the traffic arrival process of a VoIP connection to MOS. These algorithms are driven by measurements of packet loss, delay, and variation in delay or jitter.

Delay in audio delivery arises at each stage in the pipeline: encoding, packetization, network transmission; playout buffer; and decoding. All except network transmission occur within the VoIP endpoint and will be assumed to be small and deterministic. Thus we will capture and use measures of loss and network transmission delay.

Loss and delay variation can be readily captured via packet sniffers. Absolute packet delay may be computed by equipping the sniffer system with multiple NICs and simultaneously connecting it to source and destination networks.