The University of Kansas

Information and Telecommunication Technology Center

Technical Report

Preliminary Performance Evaluation of QoS in DOCSIS 1.1

Gayathri Chandrasekaran, Mohammed Hawa and David W. Petr

ITTC-FY2003-TR-22736-01

January 2003

Sponsor:
Sprint Corporation

Contents

1. Introduction

Hybrid Fiber Coax (HFC) cable networks have been used in the past primarily to deliver broadcast-quality TV signals to homes. The wide availability of such systems and their extremely wide bandwidth allows extending their functionality to deliver high-speed broadband data signals to end-users. Data over Cable System Interface Specifications (DOCSIS), a MAC protocol elaborated under the leadership of Cable Television Laboratories, Inc., has been established as the major industry standard for two-way communications over HFC cable plants. The first specifications, referred to as DOCSIS 1.0, have been largely deployed in cable networks. The second specifications (DOCSIS 1.1) are now in the early phase of deployment. The physical layer of DOCSIS 1.1 specifications is the same as that of DOCSIS 1.0. The difference between the two sets of specifications lies on the MAC layer, which includes Quality of Service (QoS) capabilities, in DOCSIS 1.1 specifications. This has been adopted as the *de facto* standard for delivering broadband services for HFC networks. Although DOCSIS is designed for interoperability among cable modems and related products, with a few modifications it can be used in wireless Multipoint, Multichannel Distribution Service (MMDS) and Local Multipoint Distribution Service (LMDS) systems. Such Broadband Wireless Access (BWA) networks attempt to carry high speed Internet traffic to subscriber systems, are easy to implement and can be installed without extensive infrastructure. Hence, The IEEE802.16 standard developed for BWA systems was based on two HFC MAC protocols, IEEE 802.14a and DOCSIS. This is due to the striking similarities between the cable medium and the wireless environment.

2. DOCSIS 1.1 Overview

2.1 DOCSIS 1.1 Basics

DOCSIS 1.1 is an IP centric point-to multipoint standard that was developed for broadband Internet access applications over cable TV networks. DOCIS specifies both the physical layer and the Media Access Control (MAC) layer. Its major components are the Cable Modems (CM) at the customer premises and the Cable Modem Termination System (CMTS) at the head-end of the cable plant. DOCSIS 1.1 evolved from the DOCSIS 1.0 specs, but with Quality of Service (QoS) mechanisms and algorithms implemented at the MAC layer. To summarize, DOCSIS 1.1 builds upon 1.0, but also includes the following features:

- Quality of Service
- Dynamic Services
- Concatenation
- Fragmentation
- Payload Header Suppression
- IP Multicast
- CM Authentication
- SNMPv3
- CM Account Management
- Fault Management
- Secure Software

This report focuses on the Quality of Service, Concatenation and Fragmentation features. The medium between the CMTS and the different CMs is a two-way shared medium, in which the downstream channels carry signals from the head-end to users and upstream channels carry signals from users to head-end. Upstream and downstream channels are separated using Frequency Division Duplex (FDD). A CM is normally tuned to one upstream channel and the associated downstream channel. The upstream channel is an inherently shared medium while the downstream is a broadcast dedicated link from the CMTS to the CM.

DOCSIS has a reservation-based, centralized approach for allocating bandwidth on the upstream channel. The upstream is modeled as a stream of mini-lots, with a dynamic mix of contention and reservation-based transmission opportunities. CMs may use the contention mini-slots for transmitting their requests, and the CMTS allocates transmission opportunities for the CMs in the next frame if capacity is available. Periodically, the CMTS sends a Bandwidth Allocation Map (MAP) message over the downstream channel to indicate to the CMs the specific mini-lots allocated to them. The Allocation MAP is a MAC Management message that describes some slots as grants for particular stations to transmit data in, other slots as available for contention transmission and other slots as an opportunity for new stations to join the link. As a result of bandwidth reservation, the CMs are guaranteed collision-free data transmission. But collisions may occur during the contention (request) period, and this is resolved using a Contention Resolution Algorithm (CRA). Specifically DOCSIS uses the Truncated Binary Exponential Backoff as its primary CRA.

Each CM has a unique MAC address. To support QoS, DOCSIS further introduces the notion of Service Flows ID (SID). Each SID is unique and defines a particular service class or flow mapping between a CM and the CMTS. The CMTS may assign one or more SIDs to a particular CM, which is essentially used for bandwidth request and allocation. ***An upstream service flow*** in DOCSIS 1.1 may be one of the following

- Unsolicited Grant Service (UGS)
- Real-time Polling Service (rtPS)
- Non-Real-time Polling Service (nrtPS)
- Best Effort (BE)
- Unsolicited Grant Service with Activity Detection (UGS-AD)

Other key enhancements made in DOCSIS 1.1 are

- *Support for multiple service flows per cable modem* allows a single modem to support a combination of video, voice and data packets.
- *Dynamic service establishment* allows MAC messages to dynamically create, modify and delete traffic flows.

- *Payload Header Suppression* (PHS) conserves link-layer bandwidth by suppressing unnecessary packet headers on both upstream and downstream traffic flows.
- *Layer 2 fragmentation* allows fragmenting larger data packets
- *Concatenation* allows CMs to send multiple MAC frames in the same transmission

1.2 Upstream Bandwidth Allocation:

There are a number of ways by which the CM can obtain bandwidth for data transmission from the CMTS:

a. By making explicit requests through contention, piggybacking or unicast opportunities
b. Unsolicited grants

Contention:

Portions of the upstream bandwidth are open to all modems (contention) for requesting upstream bandwidth and for initial ranging. The requests transmitted through contention are subject to collision, and these collisions are resolved by a **Contention Resolution Algorithm.**

Piggybacking:

Piggyback is a request for additional bandwidth sent in a data transmission. Piggybacking obviates contention, since the requests are transmitted with the data packets.

Unicast Request Polls:

Periodic unicast request opportunities are sent as a means of real-time polls regardless of network congestion. These opportunities in the Allocation MAP (see below) may be used by the stations to transmit their request packets, avoiding contention.

The allocation MAP is a varying length MAC Management message that is transmitted by the CMTS to define the transmission opportunities on the upstream. It includes a fixed-length header followed by variable number of Information Elements (IEs). Each IE defines the allowed usage for a range of minislots. Each IE consists of a 14-bit SID, a 4-bit type code and a 14-bit starting offset. The four defined types of SIDs are *broadcast* (intended for all stations), *multicast* (intended for a group of stations), *unicast* (intended for a particular station) and a *null address* (intended for no station). The various IEs that the CMTS may send in the MAP are:

- Request IE – provides intervals in which requests may be made for bandwidths for upstream data transmission
- Request/Data IE – provides intervals in which requests for bandwidth or short data packets may be transmitted
- Initial Maintenance IE – provides intervals in which new stations may join the network
- Station Maintenance IE - provides intervals in which stations are expected to perform some aspect of routine maintenance
- Short and Long Data Grant IEs - provide an opportunity for a CM to transmit one or more upstream PDUs (Protocol Data Unit which follows a MAC Header in a MAC frame)

Many different scheduling algorithms may be implemented in the CMTS by different vendors. The specification does not mandate a particular algorithm.

1.3 Contention Resolution Algorithm (CRA) Overview

The CMTS controls assignments on the upstream channel through the MAP and determines which minislots are subject to collision. The mandatory method of contention resolution that must be supported is based on a ***Truncated Binary Exponential Back-off,*** with the initial back-off window and the maximum back-off window controlled by the CMTS. The values are specified as a part of the MAP and represent a power-of two value. When a CM has information to send and wants to enter the contention resolution process, it sets its internal back-off window equal to the Data Back-off start defined in the MAP currently in effect. The CM randomly selects a value within its back-off window - $[0, 2^{\text{Backoff}}]$. This random value indicates the number of contention opportunities that the CM must defer before transmitting. A CM must only consider the contention transmit opportunities for which this transmission would be eligible. These are defined by either Request IEs or Request/Data IEs in the MAP. After a contention transmission, the CM waits for a Data Grant (Data Grant Pending) or Data Ack in a subsequent MAP. Once either is received, the contention resolution is complete. The CM determines that the contention transmission was lost when it finds a MAP without a Data Grant (Data Grant Pending) or Data Ack for it. The CM must now increase its back-off window by a factor of two, as long as it is less than the maximum back-off window (Back-off End). The CM now randomly selects a number within its new back-off window and repeats the deferring process described above. This re-try process continues until the maximum number of retries (16) has been reached, at which time the PDU must be discarded.

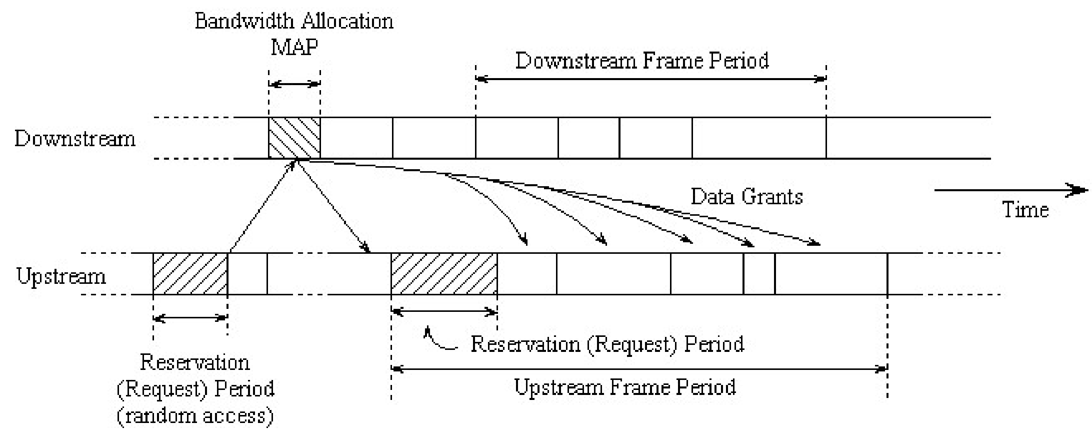

Fig. 1 Allocation MAP Structure

3. Quality of Service in DOCSIS 1.1

The DOCSIS 1.1 Specification includes several new Quality of Service (QoS) related concepts not present in DOCSIS 1.0. These include:

- Packet Classification & Flow Identification
- Service Flow QoS Scheduling
- Fragmentation and Concatenation

3.1 Theory of Operation

The various DOCSIS protocol mechanisms can be used to provide Quality of Service for both upstream and downstream traffic through the CM and the CMTS.

The requirements for Quality of Service include:

- A configuration and registration function for pre-configuring CM based Service Flows and traffic parameters
- A signaling function for dynamically establishing QoS- enabled Service Flows and traffic parameters
- Utilization of MAC scheduling and traffic parameters for upstream Service Flows
- Utilization of QoS traffic parameters for downstream Service Flows
- Classification of packets arriving from upper layer service interface to a specific active Service flow.

3.2 Service Flows

A Service Flow is a MAC-layer transport service that provides unidirectional transport of packets either to upstream packets transmitted by the CM or to downstream packets transmitted by the CMTS. A Service Flow is characterized by a set of QoS Parameters such as latency, jitter and throughput assurances. In order to standardize operation between the CM and CMTS, these attributes include details of how the CM requests upstream minislots and the expected behavior of the CMTS upstream scheduler.

The CM and CMTS provide this QoS by shaping, policing, and prioritizing traffic according to the QoS Parameter Set defined for the Service Flow. The flow in which a packet is transmitted is based on the content of the IP header fields, allowing every application to receive a different service flow. Multiple data flows (each flow corresponding to a service and identified by a service identification descriptor [SID]) concurrently exist in a cable modem. A transmission request in the upstream and the corresponding grant includes the SID as the flow identifier. The cable modem and the CMTS negotiate the QoS for each flow upon allocation and dynamically as the service requirement changes. QoS is then achieved by the implementation of sophisticated scheduling mechanisms in the CMTS. A classification function is applied to every packet.

3.3 QoS Service Flows in DOCSIS 1.1

Scheduling services are designed to improve the efficiency of the poll/grant access. By specifying a scheduling service and its associated QoS parameters, the CMTS can anticipate the throughput and latency needs of the upstream traffic and provide polls and/or grants at the appropriate times. Each service is tailored to a specific type of data flow described. The basic services comprise:

- *Unsolicited Grant Service (UGS)* is designed to support real-time service flows that generate fixed size data packets on a periodic basis, such as Voice Over IP (VoIP). The service offers fixed size grants on a real-time basis, which eliminate the overhead and latency of CM requests and assure that grants will be available to meet the flow's real-time needs. The CMTS must provide fixed size data grants at periodic intervals to the Service Flow. The CM is prohibited from using any contention request and the CMTS should not provide unicast data opportunities. Piggyback requests are also prohibited. The key service parameters are:
    - Unsolicited Grant Size (bytes): Used by the CMTS to compute the size of the unsolicited grant in minislot units.
    - Grants Per Interval: The actual number of data grants per Nominal Grant Interval
    - Nominal Grant Interval: Specifies the nominal interval between successive data grant opportunities for this Service Flow.
    - Tolerated Grant Jitter: Specifies the maximum amount of time that the transmission opportunities may be delayed from the nominal periodic schedule for a particular Service Flow
- *Real-Time Polling Service (rtPS)* is designed to support real-time service flows that generate variable size data packets on a periodic basis, such as MPEG video.

The service offers real-time, periodic, unicast request opportunities, which meet the flow's real-time needs and allow the CM to specify the size of the desired grant. This service requires more request overhead than UGS, but supports variable grant sizes for optimum data transport efficiency. The CMTS provides periodic unicast request opportunities. When the source becomes inactive, the transmission reservations are released to other flows. The CM is prohibited from using any contention request or request/data opportunities, and also from sending piggybacked requests. The CM uses only the unicast request opportunities in order to obtain upstream transmission opportunities. Key service parameters:

  - Nominal Polling Interval: Specifies the nominal interval between successive unicast request opportunities for this Service Flow on the upstream channel
  - Minimum Reserved Traffic Rate (bits/s): The CMTS should be able to satisfy bandwidth requests for a service flow up to its Minimum Reserved Traffic Rate.

- *Non-Real-Time Polling Service (nrtPS)* is designed to support non-real-time flows that require variable size data grants on a regular basis, such as high bandwidth FTP. The service offers unicast polls on a regular basis, which assures that the flow receives request opportunities even during network congestion. The CMTS typically polls nrtPS SIDs on an (periodic or non-periodic) interval on the order of one second or less. The CMTS must provide timely unicast request opportunities. The CM is allowed to use the contention request opportunities as well as the unicast request opportunities. Key service parameters:
  - Nominal Polling Interval: Specifies the nominal interval between successive unicast request opportunities for this service flow on the upstream channel.
  - Minimum Reserved Traffic Rate (bits/s): The CMTS should be able to satisfy bandwidth requests for a service flow up to its Minimum Reserved Traffic Rate.
  - Traffic Priority: The value of this parameter specifies the priority assigned to a service flow (Valid range: 0-7)

- *Best Effort (BE)* service provides efficient service to best-effort traffic. The CM is allowed to use contention request opportunities as well as piggybacking of requests. This service flow has limited QoS support. The key service parameters are:
    - Minimum Reserved Traffic Rate (bits/s): The CMTS should be able to satisfy bandwidth requests for a service flow up to its Minimum Reserved Traffic Rate.
    - Traffic Priority: The value of this parameter specifies the priority assigned to a service flow (Valid range: 0-7)
- *Unsolicited Grant Service with Activity Detection (UGS/AD)* is designed to support UGS flows that become inactive for substantial portions of time, such as VoIP with silence suppression. The service provides unsolicited grants when the flow is active and unicast polls when the flow is inactive. This combines the low overhead and low latency of UGS with the efficiency of rtPS. Though UGS/AD combines UGS and rtPS, only one scheduling service is active at a time. The CMTS must provide periodic unsolicited grants when the flow is active, but reverts to providing unicast request opportunities when the flow is inactive. It detects inactivity by unused grants. The CM is prohibited from using contention requests or request/data opportunities. Piggybacking requests are also prohibited. The key service parameters are Nominal Polling Interval, Nominal Grant Interval, and Unsolicited Grant Size.

3.4 Fragmentation & Concatenation

Fragmentation is sending a portion of a packet frame during a reserved slot time.
Fragmentation is an upstream CM "modem capability". The CMTS must enable or disable this capability on a per-modem basis. The per-modem basis provides compatibility with DOCSIS 1.0 CMs. Once fragmentation is enabled for a DOCSIS 1.1 modem, fragmentation is enabled on a per Service Flow basis. When enabled for a Service Flow, fragmentation is initiated by the CMTS when it grants bandwidth to a particular CM with a grant size that is smaller than the corresponding bandwidth request from the CM. This is known as a **Partial Grant.**

Concatenation refers to sending more than a frame during a transmission opportunity i.e. it allows the cable modem to make a single time slice request for multiple packets and send all packets in a single large burst on the upstream. Fragmentation and concatenation also make better use of the scarce upstream resource and improve throughput. There was no fragmentation in DOCSIS 1.0 and concatenation was optional. Figure 2 is a state diagram summarizing the request/grant process.

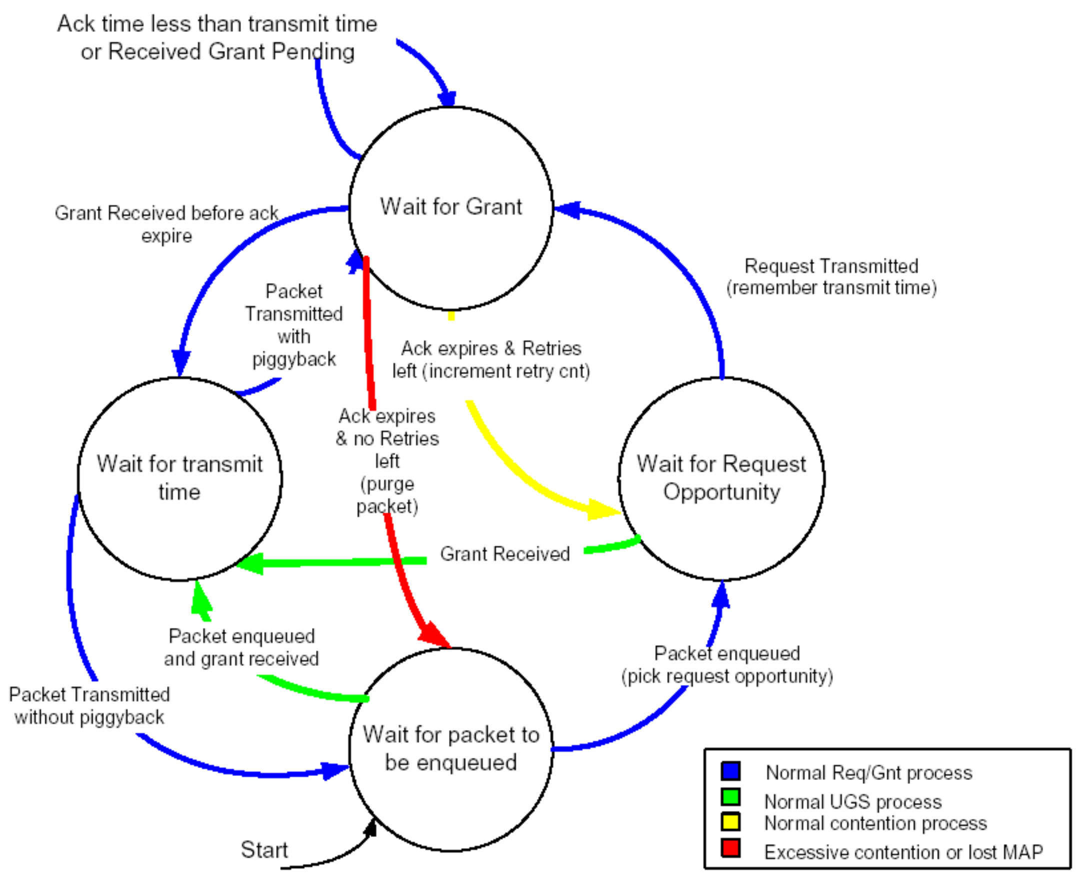

Fig. 2 State Diagram of Request/Grant Process

4.OPNET’s DOCSIS Models

OPNET's Model Library is recognized industry-wide as the most advanced suite of models of network protocols, technologies, and applications available. OPNET's DOCSIS model suite has been developed jointly by OPNET and the CableLabs' Bandwidth Modeling and Management Vendor Forum. The model is based on the DOCSIS 1.1 specification (as established by CableLabs), and includes features relevant to both DOCSIS 1.1 and 1.0 system. The model enables testing configurations and architectures before building expensive prototypes. It also allows creation of virtual representations of proposed cable modem networks so as to evaluate the capacity and quality of service (QoS) characteristics of alternative designs.

Model suite

4.1 CM Nodes

Cable modem hosts can send and receive application traffic from any of the standard network applications (voice, video conferencing, HTTP, etc).

The model has three node models that simulate CM functionality.

docsis_ethernet_cable_modem	-	a cable modem connected to an ethernet port
docsis_cm_wkstn	-	a combined Ethernet workstation and cable modem
docsis_cm_server	-	a combined Ethernet server and cable modem

Fig.3 CM models

4.2 CMTS nodes

The CMTS nodes model cable network head-end nodes. Two types of CMTS nodes are available: a CMTS gateway and a CMTS server.

- **Router and gateway nodes**
  Router and gateway nodes typically connect the bus to the cable network. Each RF interface is managed by a separate CMTS process and is independent of other RF interfaces.
- **CMTS server model**
  This node incorporates all of the CMTS functionality, but can also send and receive application data. Other differences between the CMTS router and the CMTS server are listed below:
  - the CMTS server has only one RF interface (routers have at least two interfaces)
  - the CMTS server contains a server module that allows the node to send and receive application data (routers do not usually have a server)

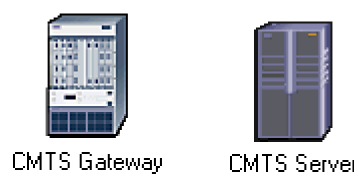

Fig.4 CMTS models

4.3 Link models

A DOCSIS bus link model connects the CMTS and cable modem hosts.
Taps connect individual host and CMTS nodes to the bus. The link is modeled as a bus with a single channel that uses different frequencies for CMTS (downstream) and CM (upstream) data.

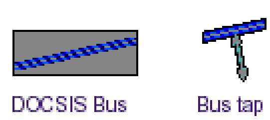

Fig.5 CMTS models

4.4 DOCSIS attribute configuration object

This global configuration object allows configuration of general DOCSIS parameters that can be applied to DOCSIS nodes in the network. This object is used to configure profiles for MAP, modulation, physical media overhead, upstream physical properties, and payload header suppression parameters. These profiles can then be deployed on CMTS nodes in the network.

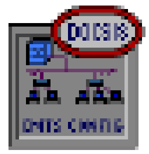

Fig.6 Attribute Configuration Object

4.5 Application Traffic (Application and Profile objects)

The custom application is a generic model that is extremely versatile in representing a broad class of applications. It can be used when the application of interest does not correspond to any of the standard applications. The custom application provides attributes that allow one to configure various aspects of the application in detail. A custom application can be used to represent any number of tiers, including two-tier applications. The custom application has been configured so that only the clients communicate with the server.

4.6 Bugs

- DOCSIS RTP process does not add unicast requests to the MAP in the required intervals

  The docsis_cmts process calls the RTP process to find out whether or not a unicast request should be added to the MAP every time it creates a new MAP. The RTP process should be basically checking whether the time since the last poll is greater than the real-time polling interval or not. If yes, a unicast request must be added to the MAP. There seems to be a logical error due to which this is not happening. As a result performance studies were not done with the rtPS QoS class.

- CMTS does not add Pending Grant IEs to the MAP

  When the CMTS identifies that it can't satisfy a grant, it has to add a Pending IE to the MAP to inform the CM that it’s request has been received and that it would be satisfied in the subsequent MAP. This does not happen, due to a bug in the implementation. This results in the CM contending again which it wouldn't have done had the pending grant IE been added to the MAP. This would have a small but definite impact on the results obtained. The fact that the CMs contend again would have led to an increase in the number of contentions and hence the collision probability. But this is true only for the experiments for which concatenation and fragmentation had been disabled, because if fragmentation had been enabled, the CMTS does not deny a grant, but instead gives a partial grant.
- DOCSIS nrtPS does not add unicast requests to the MAP

  DOCSIS 1.1 offers unicast request opportunities to nrtPS to avoid contention during times of network contention. But due to a logical bug in the implementation, the CMTS was not adding the unicast request polls to the MAP. This is the only feature that differentiates nrtPS from BE. Again because of this we were not able conduct performance studies with nrtPS.

5. Simulation Scenarios

The figure shows an OPNET simulation scenario, with 2 CMs and the CMTS. Also the various configuration objects for configuring DOCSIS parameters and traffic parameters can be seen.

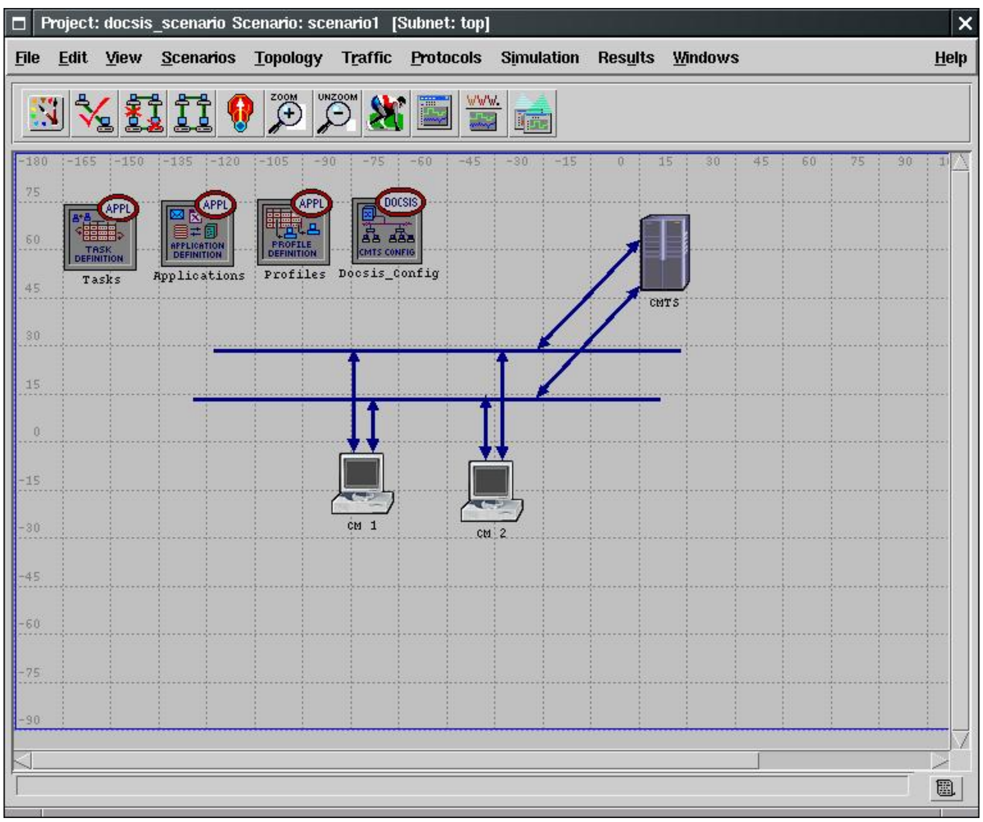

Fig. 7 Simulation Scenario in OPNET

6. Performance Evaluation of DOCSIS 1.1 QoS

6.1 DOCSIS CMTS Parameters:

- Physical Media Profile:

  This defines profiles that allow the user to configure the RF specifications (modulation, channel width, data rate, and center frequency) and interleave latency of downstream channels and upstream channels. Parameters used for these simulations were:

  Upstream Bandwidth: 640 Kb/s
  Downstream bandwidth: 27Mb/s
  Upstream modulation scheme: QPSK
  Downstream modulation scheme: 64 QAM

- MAP Profile:

  This attribute defines the bandwidth allocation MAP profiles that are applied to the upstream channel parameters on CMTS nodes.

  MAP inter-arrival time: 0.2s
  Short data grant limit: 128 bytes
  Long data grant limit: 10000 bytes
  Data back off start: 7
  Data back off end: 16
  Number of contention slots: 64
  Bytes per minislot: 4 bytes

6.2 Experiment 1: a. Comparison of Best Effort & UGS Delays

b. Effect of Piggybacking on Best Effort Delay

Configuring the UGS station:

The total upstream bandwidth is 640Kbps. 64 minislots have been allotted for contention. The remaining bandwidth can be calculated in the following way:

Duration of one minislot = 50 us

Number of minislots in MAP: MAP time / Duration of one minislot

= 0.2/0.00005

= 4000 minislots

Number of minislots available

for data transmission	= 4000- (32+2) [2 slots for station maintenance]
Size of one minislot	= 4 bytes
Remaining bandwidth	= (3966 * 4 * 8) / 0.2
	= 634.56 Kbps

Of the remaining bandwidth ~ 10% has been allotted to UGS. Since UGS is generally used for Constant Bit Rate Traffic, i.e. for flows that generate fixed packets in regular intervals, the station has been configured in the following manner.

<u>UGS Station Traffic Statistics</u>:

- Request Packet size = 1000 bytes
- Packet Inter-arrival time = 0.21s
- Packet size distribution: Constant
- Inter-arrival times distribution: Constant

<u>UGS QoS Parameters:</u>

The Nominal Grant interval is set so that there is a grant every MAP and it is also approximately equal to the packet inter-arrival time.

- Grant Size = 1300 bytes
- Nominal Grant Interval = 0.2s

Physical overhead calculation:

The following parameters are used for the physical overhead calculation. These are configured as a part of the Physical Media Overhead within the MAP Profiles. Messages can be short or long data frames. Short data frames have a 100-byte limit while long frames can go up to 10,000 bytes. Since all the packets lengths used in the simulation exceeds 100 bytes, the long data frame physical parameters are listed below. Before transmitting a packet received from the higher layer, the CMTS divides the packet into code words. It then adds FEC error correction bytes to each code word and preamble length bits to each frame.

Preamble Length: 56 bits

FEC Error Correction (FEC Parity): 16 bytes

FEC Code Word Length: 226 bytes

Guard Time (bits): 40

Last Code Word Mode: Fixed

A sample overhead calculation for a packet size of 1300 bytes is shown below.

The MAC headers are added to the IP frame (includes 20 bytes of IP header and 20 bytes of TCP header).

Total message size before upstream physical overhead is added = Size of IP frame + DOCSIS_MSG_PDU_HDR_SIZE (20 bytes)

Message size = Payload size + TCP/IP headers + MAC Overhead

= 1300 + 40 + 20

= 1360 bytes

No of code words = (message size)/(FEC Code word length- FEC Parity)

= 1360 /(226-16)

= 6.47

= 7 (rounded off to the highest integer)

Now the final message size = number of code words x FEC code word

= 7 x 226

= 1582

Total Frame Size = Message size + (Preamble + Guard Time)/8

= 1582 + (56 + 40)/8

= 1582 + 12

= 1594 bytes.

The total bandwidth allocated to the UGS station is (1594 * 8)/0.2 =63.76 Kbps.

Configuring the BE stations:

The rest of the bandwidth was distributed among the BE stations. The stations had identical parameters and also traffic statistics. The load was increased by adding stations. Fragmentation and concatenation were disabled.

<u>BE Stations Traffic Statistics</u>:

- Mean packet Size = 800 bytes
- Mean packet inter-arrival time = 0.22s
- Packet size distribution: Exponential
- Inter-arrival time distribution: Exponential

Length of Simulation (min): 37

Approximate number of packets generated by each station: 10,000

Since Best Effort flows are similar to Internet traffic, exponential arrivals and exponential packet sizes were set. The inter-arrival time must be greater than the MAP inter-arrival time for the queue size to remain bounded. This is because fragmentation and concatenation were disabled (to make the scenario simple). Hence only one packet at a time can be sent in an upstream transmission interval. Hence the mean inter-arrival time should be set so that there is at most one arrival during the MAP inter-arrival time.

Results:

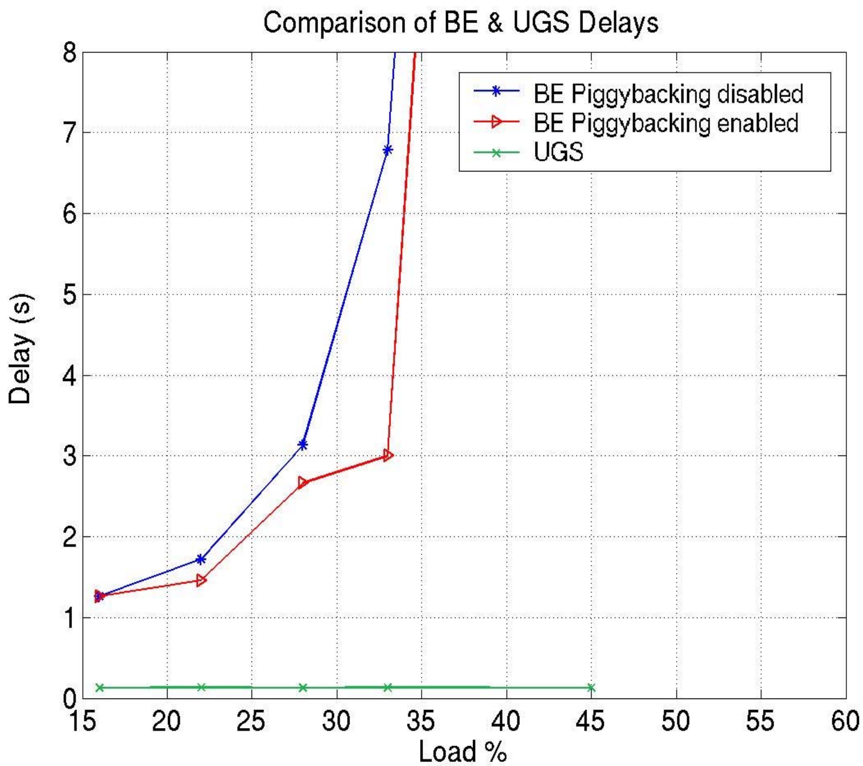

Fig. 8 Comparison of UGS and BE Delays

Analysis:

The performance of BE traffic was compared with UGS. UGS flows are allowed to reserve a certain portion of the bandwidth. No requests for transmission are needed; hence UGS has a low bounded constant delay, since it receives a grant every MAP and thus queue size is low and bounded. BE on the other hand has "request grant, request

grant" pattern. Stations have to contend for sending requests and have to wait for grants in the subsequent MAPs. Contentions may result in collision and thus there may be increased delay due to retransmissions. The effect of piggybacking on delay is also studied. With piggybacking enabled, requests are piggybacked to outgoing data and thus the frequency of contention is reduced and delay is reduced. We see that the system saturates quickly. This is because fragmentation and concatenation were disabled; hence only one frame can be sent per MAP. This leads to relatively high delay values, since the queue size builds up quickly. This is because as load increases, the CMTS would not be able to satisfy all the requests from all the stations. Also in this experiment and in all the experiments that follow, the load value is the actual load value and not the calculated value. The load value was calculated from the simulations as
Reserved slot time in MAP/ Total MAP duration. We report actual loads because with the exponential packet lengths and exponential inter-arrival times, it is very difficult to predict the load values from the mean packet size and inter-arrival time. The packet lengths are exponentially distributed and hence they would be distributed in a wide range of values. Also, as explained before, there is also some amount of complexity involved in the computation of the upstream physical overhead.

6.3 Experiment 2: Effect of Fragmentation & Concatenation

Configuring the UGS station: As before ~ 10% of the total bandwidth was allotted to the UGS Station. The MAP configuration parameters were retained for all the experiments.

<u>UGS Station Traffic Statistics</u>:

- Request Packet size = 1000 bytes
- Packet Inter-arrival time = 0.21s
- Packet size distribution: Constant
- Inter-arrival times distribution: Constant
- Number of packets generated by each station: 10,000

<u>UGS QoS Parameters:</u>

- Grant Size = 1300 bytes
- Nominal Grant Interval = 0.2s

The total bandwidth allocated to the UGS station is (1594 * 8)/0.2 =63.76 Kbps.

Configuring the BE stations:

The rest of the bandwidth was distributed among the 11 BE stations. The stations had identical parameters and also traffic statistics. The load was increased by adding stations to the scenario. The experiment was performed to study the effect of fragmentation and concatenation on performance i.e., on mean delay. Hence one set of simulations were carried out with fragmentation and concatenation disabled, and another with enabled. Piggybacking was enabled and data back-off start was maintained at 7.

<u>BE Stations Traffic Statistics</u>:

- Mean packet Size = 800 bytes
- Packet size distribution: Exponential
- Inter-arrival time distribution: Exponential
- Mean packet inter-arrival time: variable
- Number of packets generated by each station: 10,000

Results:

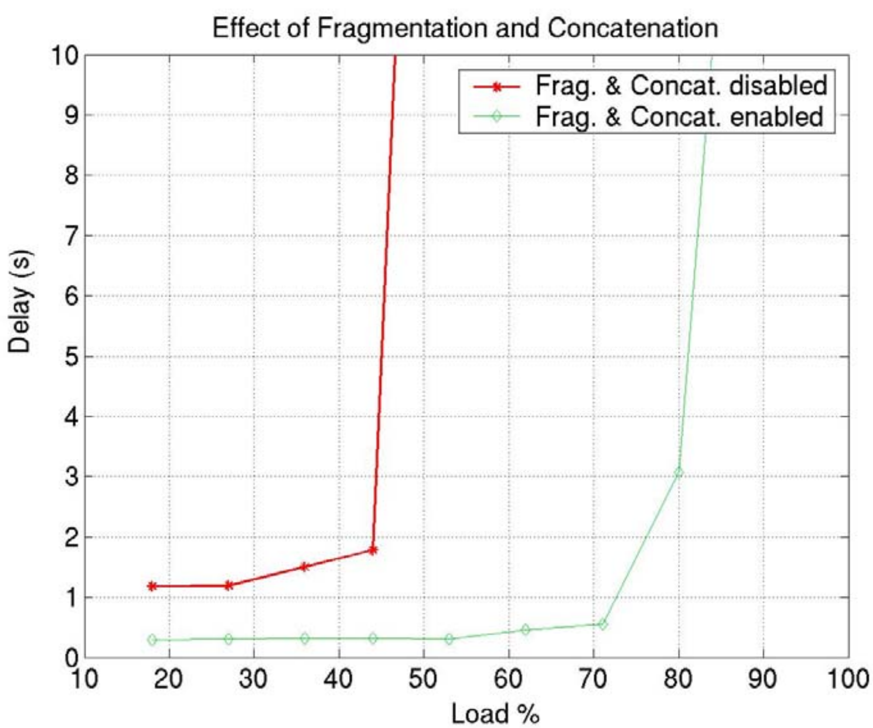

Fig. 9 Effect of Fragmentation & Concatenation

Analysis:

DOCSIS concatenation combines multiple upstream packets into one packet to reduce packet overhead (as explained already, the physical upstream overhead is considerable)` and overall latency, and to increase transmission efficiency. Using concatenation, a DOCSIS cable modem makes only one bandwidth request for multiple packets, as opposed to making a different bandwidth request for each individual packet; this technique is especially effective for bursty real-time traffic, such as voice calls. Since a single bandwidth request is made for a multiple frames, the queuing time is reduced, and this brings down the overall access delay. This also saves upstream bandwidth to an extent since sending of many small request packets is avoided. Also the need for frequent contention is also reduced, resulting in fewer collisions. This also contributes to an improvement in performance. When enabling concatenation, fragmentation should also be enabled; otherwise packets that are too large for transmission in a single MAP are not transmitted. That is, when concatenation is enabled the request sizes are generally big, and the CMTS does not issue a grant at all if fragmentation is disabled and there are not enough minislots in the current MAP. Thus fragmentation and concatenation together help in better utilization of the upstream bandwidth compared to the case when both were disabled. In that case we are only able to reach loads of 45%, since after that the queue builds up and delay increases indefinitely. But when fragmentation and concatenation were enabled, we are able to reach much higher load values with lesser access delays.

6.4 Experiment 3: Effect of Backoff Start and Piggybacking

Configuring the UGS station: As before ~ 10% of the total bandwidth was allotted to the UGS Station. The MAP configuration parameters were retained for all the experiments.

<u>UGS Station Traffic Statistics</u>:

- Request Packet size = 1000 bytes
- Packet Inter-arrival time = 0.21s
- Packet size distribution: Constant
- Inter-arrival times distribution: Constant
- Number of packets generated by each station: 10,000

<u>UGS QoS Parameters:</u>

- Grant Size = 1300 bytes
- Nominal Grant Interval = 0.2s

The total bandwidth allocated to the UGS station is (1594 * 8)/0.2 =63.76 Kbps.

Configuring the BE stations:

The rest of the bandwidth was distributed among the BE stations. The stations had identical parameters and also traffic statistics. In this experiment the number of BE stations were maintained constant at 11 and the load was changed by changing the inter-arrival times. In order to make load values predictable and to simplify protocol operation, fragmentation and concatenation were disabled and packet lengths were made constant. This was done since with constant packet lengths, we would know for sure the exact load values since there would be no randomness involved in the packet lengths. Also disabling fragmentation and concatenation helps study the effect back-off values better. The experiment was done first disabling piggybacking, and then enabling it. Back-off end value was maintained at 16.

BE Stations Traffic Statistics:

- Mean packet Size = 800 bytes
- Packet size distribution: Constant
- Inter-arrival time: variable
- Inter-arrival time distribution: Exponential
- Number of packets generated by each station: 10,000

Case a: Piggybacking disabled

Results:

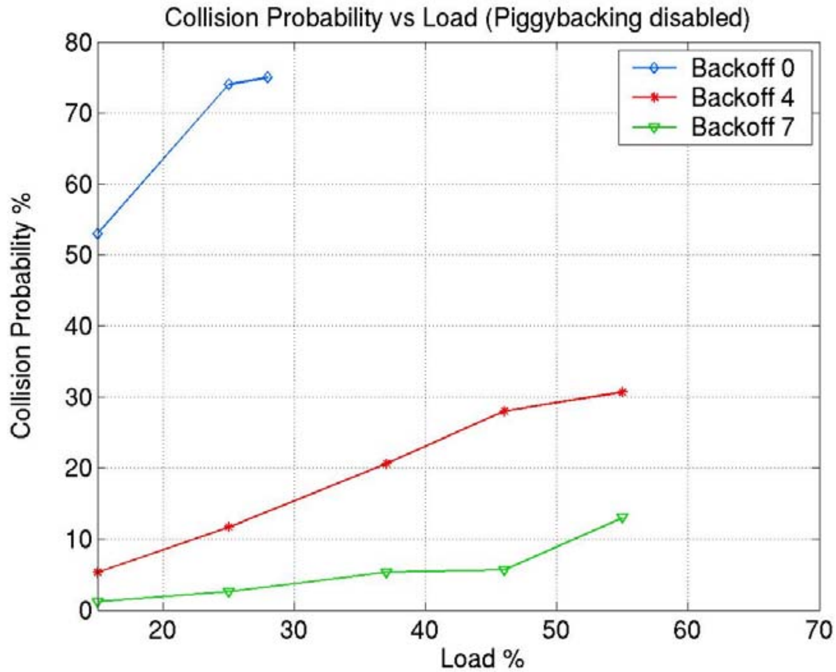

Fig. 10 Effect of Back-off on Collision Probability (Piggybacking disabled)

Analysis:

We see the effect of the data back-off start value on collision probability. For a data back-off zero we see that the collision probability is very high and it increases with increase in load. This is because since the initial value is zero, the back-off window is zero (no random wait). Hence the probability that two stations would pick the same minislot for contention is very high. This would result in a collision and the contention requests of both the stations would be lost. Also since piggybacking has been disabled, there are more frequent contentions as load increases. This leads to increased probability of collisions. We see that the collision probability decreases as the back-off value increases. This is because the station defers its transmission by a greater number of minislots (since its window size is big). Also we see that we are able to go up to much higher loads for back-off 4 & 7. This is because with a lower back-off value, there would be increased number of retransmissions since the collision probability is more. Hence the queue builds up quickly. But with greater back-off values, the collision probability is comparatively less and hence we are able to go up to comparatively high load values.

Case b: Piggybacking enabled

Result:

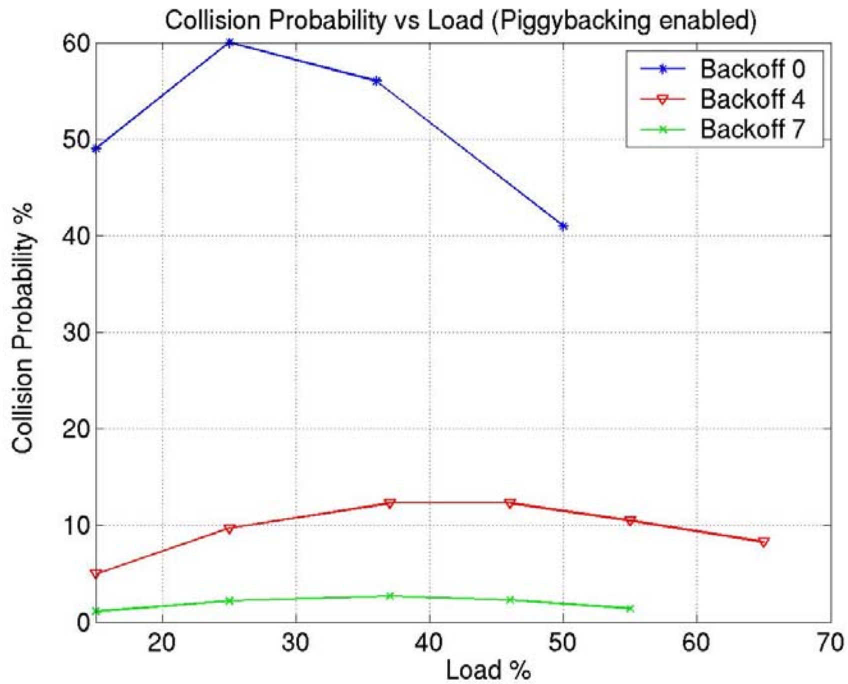

Fig. 11 Effect of Back-off on Collision Probability (Piggybacking enabled)

Analysis:

The same experiment was repeated with piggybacking enabled. Here the collision probability increases with load for lower load values, and then decreases. This is because with higher load values, the contention load decreases, since most of the requests can by piggybacked with the data transmissions. But at lower loads, the probability that there may be a packet in the queue while data is being transmitted is less. Hence the station might be forced to use contention more often, even if piggybacking is enabled. Hence the collision probability increases. But at higher loads, the piggybacking feature may be used more effectively and hence the number of contentions and hence the collisions may be reduced. Also it can be observed that we are able to reach higher loads compared to the case when piggybacking was disabled.

Effect of Back-off on Delay:

Results:

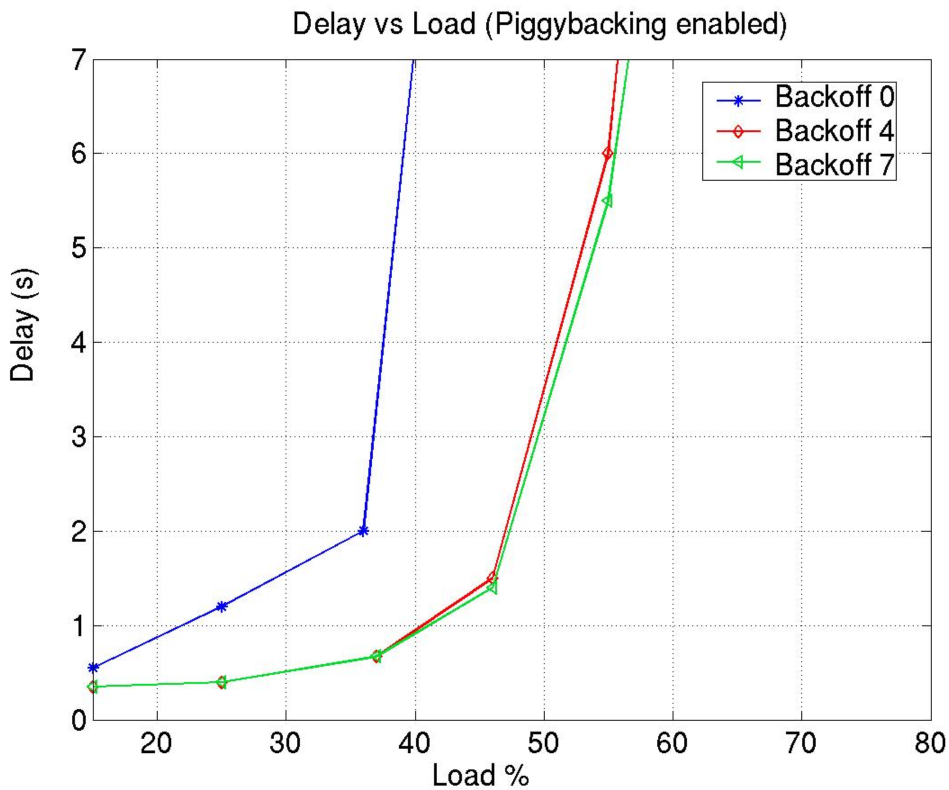

Fig. 12 Effect of Back-off on Delay (Piggybacking enabled)

Analysis:

We see that increasing the data back-off value above 0 produces a marked improvement in access delays. This is because if the back-off value is high, the collision probability reduces and hence delay decreases. We also can observe that there is not much difference in delays for back-off 4 and 7. Though the collision probability is much lesser than what it is for 4, the corresponding decrease in delay is offset by the increase due to increased delay in contention. This is because the stations have to back off by a larger number of minislots, and this increases the average waiting time to send the request packets.

6.5 Experiment 4: Effect of Back-off Start and Piggybacking – A more realistic scenario

Configuring the UGS station: As before ~ 10% of the total bandwidth was allotted to the UGS Station. The MAP configuration parameters were retained for all the experiments. The rest of the attributes were same as that of the previous experiments.

Configuring the BE stations:

The rest of the bandwidth was distributed among the BE stations. The stations had identical parameters and also traffic statistics. In this experiment the number of BE stations was maintained constant at 11 and the load was changed by changing the inter-arrival times. Here fragmentation and concatenation were enabled and packet lengths were made exponentially distributed. Hence it was not possible to predict the load values ahead and they were measured from the simulation. The experiment was done first disabling piggybacking, and then enabling it.

<u>BE Stations Traffic Statistics</u>:

- Mean packet Size = 800 bytes
- Packet size distribution: Exponential
- Inter-arrival time: variable
- Inter-arrival time distribution: Exponential
- Number of packets generated by each station: 10,000

Results:

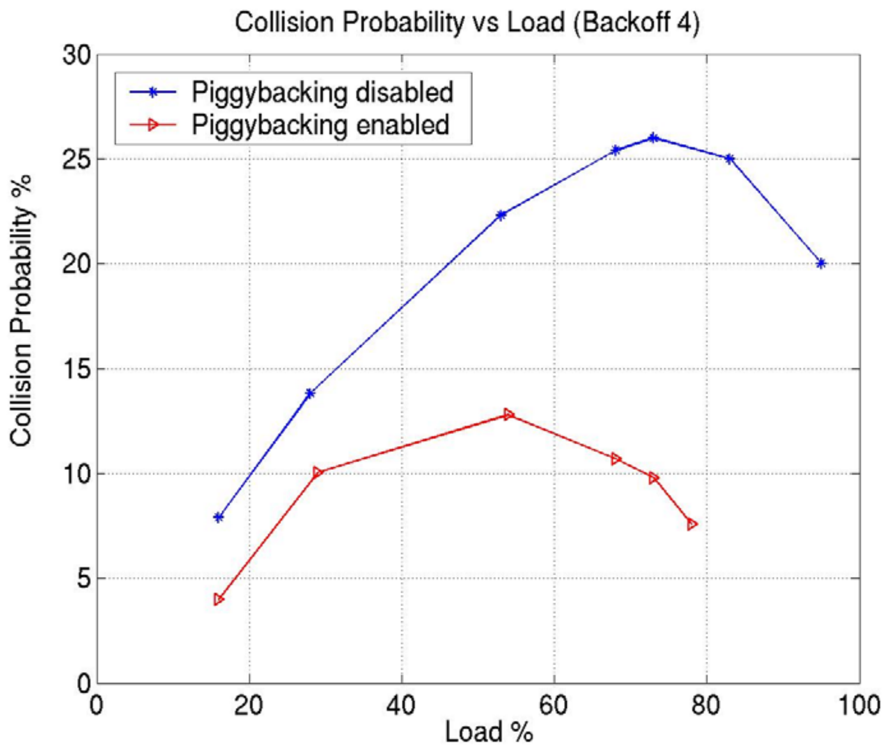

Fig. 13 Effect of Piggybacking on Collision Probability

Analysis:

The simulations were carried out to study the effect of back-off on collision probability, for back-offs 0, 4 & 7. The results are plotted in Figure 12 showing the effect of piggybacking for back-off 4. We see that enabling piggybacking has reduced the collision probability since the need for contention is reduced with piggybacking enabled. But we see that the trend is same whether piggybacking has been enabled or not. This is because concatenation has been enabled. So the effect of concatenation is more prominent. Since with concatenation enabled, requests for multiple frames can be made, there is less need for contention. Hence the collision probability first increases with load, for lower load values. This is because there are a large number of intervals during which there are no packet arrivals, and hence the stations are forced to use contention, with piggybacking the need for contention is a little less. But as load increases, the buffer is almost never empty and hence the contention load keeps decreasing and hence the collision probability decreases.

Effect on delay:

For the same scenario, the effect of back-off on delay is discussed for piggybacking enabled case. The behavior is the same for the piggybacking disabled case, except that the delay is slightly more.

Results:

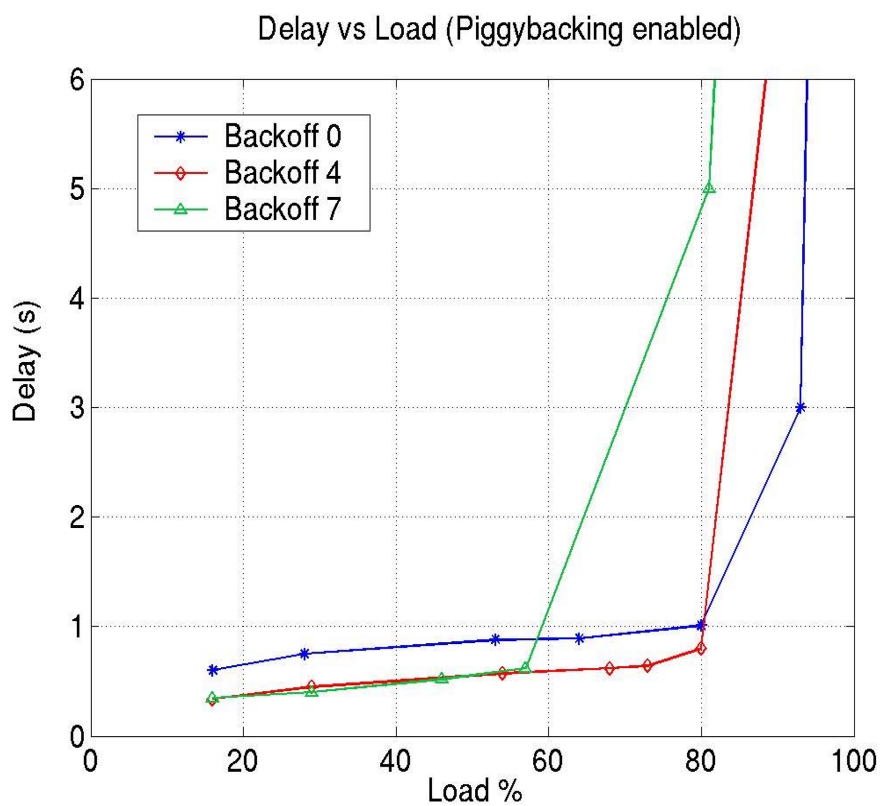

Fig. 14 Effect of Back-off on Delay (Piggybacking enabled)

Analysis:

For low loads, increasing the back-off value reduces the collision probability and hence delay decreases significantly due to reduced retransmissions. There is not much of a difference between back-off 7 and 4 for lower loads, since the decrease in delay as a consequence of decreased collision is offset by the increased delay incurred during contention. This is because as the back-off value is more, the back-off window is larger and hence the station defers by a greater number of minislots before transmitting its request. This increases delay since for a specified value of priority, the CMTS issues grants in the order in which they were received. But it is observed that the load value at which the system saturates, i.e. the buffer builds up, decreases as the back-off value increases. This is because each station defers its transmission by a greater number of minislots, and hence it is possible that the request may not reach the CMTS before the next MAP is transmitted. Hence it may so happen that the station is forced to wait for two MAP periods. This leads to the queue building up and delay rising indefinitely. But for back-off zero, the window starts with a small value and slowly increases each time a collision occurs and hence it is more adaptive to the load, and helps utilizing the entire bandwidth.

Comparing the results illustrated in Figure 14 with those in Figure 12, we see that the system saturates quickly for higher back-off values in Figure 14, as opposed to Figure 12. This is because for the experiment of Figure 14, fragmentation and concatenation were enabled, whereas for experiment 3 (figure 12) both were disabled. So with fragmentation and concatenation enabled, it is possible to achieve a much higher throughput and utilize the entire bandwidth available in one MAP period. Also if the back off is high, with the piggybacking enabled, it is possible that certain requests reach the CMTS only after the subsequent MAP is sent. This leads to an increased waiting time, and also building up of the queue. On the other hand, with the previous experiment (Figure 12), since fragmentation and concatenation were disabled, the collisions had a prominent effect and it wouldn't be possible to saturate the entire available bandwidth. So the system saturates quickly for the lower back-off value.

6.6 Experiment 5: Effect of Priorities on Delay

Simulation Scenario:

The set up consisted of 2 UGS stations and 8 Best Effort stations. The packet attributes and other UGS parameters were the same as that of previous experiments.

The total bandwidth allocated to the UGS station is (1594 * 8 *2)/0.2 = 127.72 Kbps.

There were 8 Best Effort stations with priorities [0-3], i.e. 2 stations per priority. The number of stations was maintained constant throughout the experiment. Load was changed by changing the packet inter-arrival times. Since fragmentation and concatenation were enabled, there was no constraint on the inter-arrival times (i.e. they are no longer required to be greater than the MAP duration.). Again load values were measured from the simulation.

BE Stations Traffic Statistics:

- Mean packet Size = 800 bytes
- Mean packet inter-arrival time = variable
- Packet size distribution: Exponential
- Inter-arrival time distribution: Exponential
- Number of packets generated by each station: 10,000

Results:

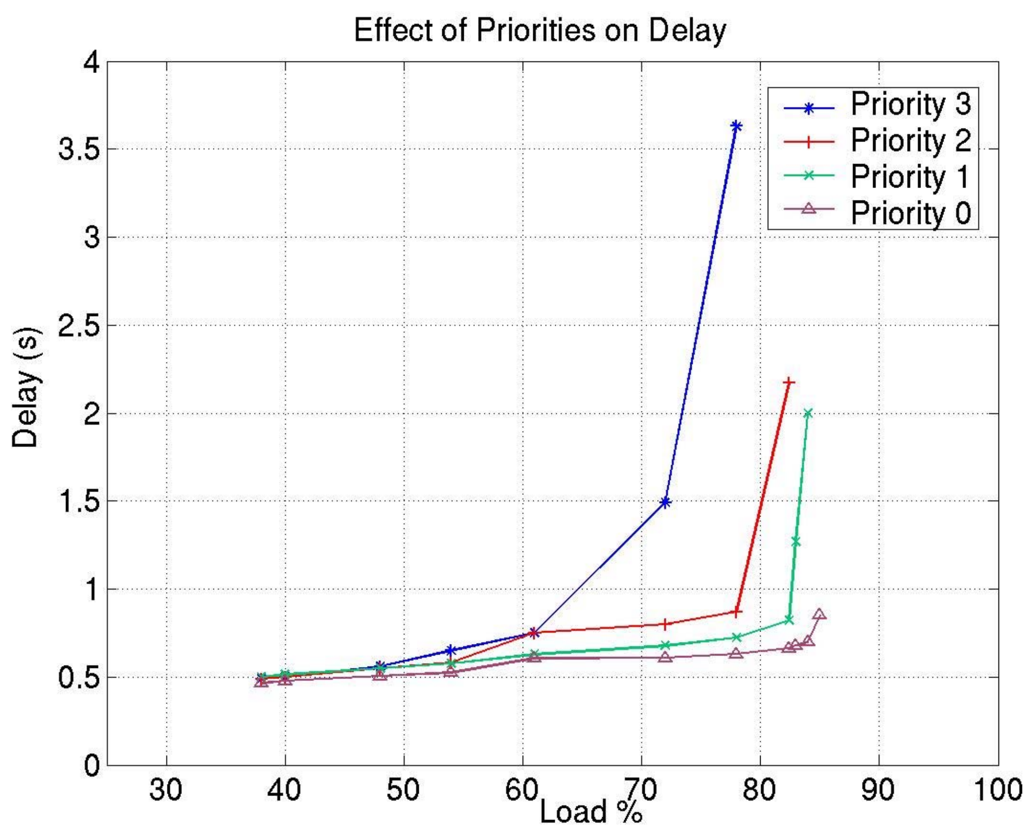

Fig. 15 Effect of Priorities on Delay

Analysis:

The effect of Traffic Priorities on delay is illustrated in Figure 15. The standard recommends a Traffic Priority parameter for the BE and the nrtPS flows. There is no priority for the UGS and the rtPS effort stations, since UGS flows are guaranteed bandwidth, and the rtPS flows get unicast request opportunities. Thus these flows are implicitly assigned a high priority. Only BE flows were set up in the simulation experiment for studying the effect of priorities. The CMTS uses the Traffic Priority attribute for determining precedence in grant generation. Priorities do not affect contention process, though. Traffic Priorities range from 0-7 applicable to BE and nrtPS, 0 being the highest. It is clearly seen that the higher priority stations have lower access delay, since it is given preference in transmitting data. The grants in the MAP are scheduled in order of priority, and thus a high priority station gets to transmit its packets first and thus suffers lower access delays.

7. Conclusions & Future Work

7.1 Conclusions

A detailed evaluation of the performance of some of the DOCSIS 1.1 QoS features has been done. The performance of Best Effort and Unsolicited Grant Service were analyzed and also comparative studies were made. Also, fragmentation, piggybacking, concatenation, and contention back-off parameters were studied. Detailed analysis of how these affected performance of the broadband access system, especially with respect to access delay, was done. It was found that nominal back off values like 4 or 5 provide near-optimal access delays. It was also found that fragmentation, concatenation and piggybacking improve throughput and reduce delay greatly.

7.2 Future Work

The performance evaluation of the other two QoS classes rtPS and nrtPS remains to be done. The framework for carrying out these simulations has been designed. Also the implementation of the new scheduling architecture * in OPNET remains to be done.

* **Quality of Service Scheduling in Cable and Broadband Wireless Access Systems**
Mohammed Hawa and David W. Petr, IWQOS 2001.