On the Performance Modelling and Optimisation of DOCSIS HFC Networks

Neelkamal P. Shah[1], Demetres D. Kouvatsos[1], Jim Martin[2], and Scott Moser[2]

[1] School of Computing, Informatics and Media
University of Bradford, UK
{N.P.Shah,D.Kouvatsos}@bradford.ac.uk
[2] School of Computing
Clemson University, USA
{Jim.Martin,S.Moser}@cs.clemson.edu

Abstract. The DOCSIS protocol defines the MAC and physical layer operations governing two-way transmission of voice, video and multimedia data over HFC cable networks, thus constituting a complex system with many interdependent parameters. This tutorial employs simulation and analytic methodologies for the performance modelling and optimisation of DOCSIS 1.1/2.0 HFC networks with particular focus on the contention resolution algorithm, upstream bandwidth allocation strategies, flow-priority scheduling disciplines, QoS provisioning and TCP applications. In this context two performance evaluation case studies are reviewed in detail - based respectively on two open queueing network models of DOCSIS 1.1/2.0 HFC networks. The first study evaluates, via '*ns*' simulation, the effect of carrying TCP/IP traffic on network performance whilst the second optimises analytically the upstream network bandwidth allocation. It is expected that many of the performance affecting operational behaviours exhibited by former releases of DOCSIS-based HFC networks will also exist under the latest DOCSIS 3.0 protocol and future extensions.

Keywords: Data-over-cable service interface specification (DOCSIS), hybrid fibre coax (HFC), Media Access Control (MAC), Quality of Service (QoS), broadband access, scheduling, contention resolution algorithms, bandwidth allocation, Transmission Control Protocol (TCP), performance modelling and evaluation, queueing network models (QNM's), network decomposition.

1 Introduction

Community antenna television (CATV) systems were introduced as a way to deliver television content to households located in hilly terrain that could not receive broadcast television. Over the years CATV companies began offering Internet access, data and telephony services to their customers in addition to television channels as a means of increasing revenue. Initially cable operators deployed proprietary systems. To stay competitive with other access technologies such as the digital subscriber line (DSL), it was decided to open the cable modem (CM) market by creating a single standard hoping to make CM's commodity items. The industry converged on the DOCSIS

D. Kouvatsos (Ed.): Next Generation Internet, LNCS 5233, pp. 682–715, 2011.

standard [1] which defines the MAC and physical layers that are used to provide high speed data communication over HFC cable networks. By pushing fibre further to the subscriber, fewer amplifiers are needed, noise is less of a problem and two-way data communication is possible.

In the early 1990's, the cable industry developed a large number of schemes for supporting two-way data over cable and several competing standards emerged:

1.1 IEEE 802.14

In 1994 the IEEE 802.14 working group was chartered to develop a MAC layer that would support both the asynchronous transfer mode (ATM) and IP over HFC networks. This working group has since disbanded. The upstream channel was time-slotted to enable multiple access with a slot size of 8 bytes. ATM's constant bit rate (CBR), variable bit rate (VBR), available bit rate (ABR) and unspecified bit rate (UBR) services were supported over the HFC network. Primarily due to time constraints, the standard did not obtain vendor support.

1.2 Multimedia Cable Network System's (MCNS's) DOCSIS

In response to competition from DSL, key multiple system operators (MSO's) in the early 1990s formed the MCNS to define a standard system for providing data and services over a CATV infrastructure. In 1997 MCNS released version 1 of DOCSIS (DOCSIS 1.0). The upstream channel was time-slotted with a configurable slot size (referred to as a minislot). This standard was quickly endorsed by the cable industry. The DOCSIS standard is now managed by CableLabs, a non-profit research and development group funded by cable industry vendors and providers.

1.3 DAVIC/DVB

The non-profit Swiss organisation Digital Audio Visual Council (DAVIC) was formed in 1994 to promote the success of digital audio-visual applications and services. The organisation produced the DAVIC 1.2 and the very similar Digital Video Broadcast Return Channel for Cable (DVB-RCC) radio frequency (RF) CM standards that defined the physical and MAC layers for bidirectional communications over CATV HFC networks. The DVB-RCC standard was popular in Europe for several years. However, to benefit from the economies of scale, the European cable industry moved towards the EuroDOCSIS standard.

Fig. 1 illustrates a simplified DOCSIS cable network. A Cable Modem Termination System (CMTS) interfaces with hundreds or possibly thousands of CM's. A Cable Operator allocates a portion of the RF spectrum for data usage and assigns a channel to a set of CM's[1]. A downstream RF channel of 6MHz (8MHz in Europe) is shared by all CM's in a one-to-many bus configuration (i.e. the CMTS is the only sender).

---

[1] A group of CM's that share an RF channel connect to an Optical/Electrical (O/E) node with a coaxial cable using a branch-and-tree topology.

The cable industry is undergoing a period of rapid change. Fuelled primarily by demand for voice over IP (VoIP) (requiring a more symmetric service) and IPTV services (requiring greater downstream bandwidth), the operations support systems of MSO's are being upgraded. In DOCSIS 1.0, only one QoS class was supported, that is, 'best effort' (BE), for data transmission in the upstream direction. Upstream data rates were limited to 5.12Mbps. DOCSIS 1.1 provides a set of ATM-like QoS guarantees. In addition, the physical layer supports an upstream data rate of up to 10.24 Mbps. DOCSIS 2.0 further increases upstream capacity to 30.72 Mbps through more advanced modulation techniques and by increasing the RF channel allocation to 6.4MHz. DOCSIS 3.0 supports hundreds of Mbps in both the upstream and downstream channels through channel bonding techniques.

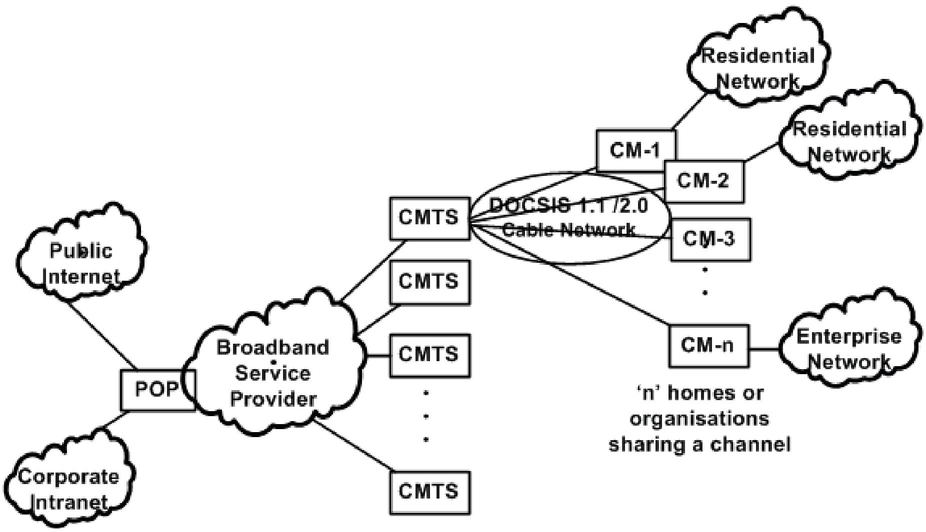

Fig. 1. Simplified DOCSIS cable network

In DOCSIS HFC networks, there are many configuration parameters and it is difficult to know a priori how particular combinations of parameters and different traffic mixes impact the network performance. This tutorial is a development of [2] and is motivated by the need to unveil the aforementioned mysteries by solving the numerous intriguing operational design, bandwidth allocation and performance evaluation problems and opportunities presented by HFC networks particularly under the complex DOCSIS protocol. The performance modelling reviewed herein relates to HFC cable networks operating under the DOCSIS 1.1 or 2.0 protocols (shortened to 'DOCSIS 1.1/2.0 HFC networks' throughout the manuscript) and it is expected that in many situations the models and their solutions will be applicable to the cases of DOCSIS 3.0 and future releases. In addition to performance modelling characterised by the underlying aims of response, throughput and/or utilisation improvement, DOCSIS has received attention in the context of monetary cost-reduction efforts. This has occurred via proposals to couple cable broadband networks to fixed wireless networks in order to eliminate last mile cabling or couple mobile communication systems such as 3G UMTS networks with these wired networks to minimise the need to build backhaul networks for the mobile systems [3, 4].

This tutorial is organised as follows: An overview of the DOCSIS 1.1/2.0 protocol is presented in Section 2. A review of selected DOCSIS HFC network performance models and their solutions is carried out in Section 3. Two quantitative DOCSIS HFC network performance evaluation studies (simulation and analytic) are detailed in Section 4. Conclusions and areas of future work follow in Section 5. Finally a list of the acronyms used throughout the manuscript is provided in Appendix I after the references.

2 The DOCSIS 1.1/2.0 Protocol

DOCSIS 1.1/2.0 is presented successively through descriptions of its general operation, QoS provision and security issues and finally an illustration of its MAC transmission layer through a QNM.

2.1 General Operation

Once powered on, the CM establishes a connection to the network and maintains this connection until the power to it is turned off. Registration of the CM onto the network involves acquiring upstream and downstream channels and encryption keys from the CMTS and an IP address from the ISP. The CM also determines propagation time from the CMTS in order to synchronise itself with the CMTS (and in effect the network) and finally logs in and provides its unique identifier over the secure channel. Due to the shared nature of these cable networks, transmissions are encrypted in both the upstream and downstream directions [5].

DOCSIS 1.1/2.0 specifies an asymmetric data path with downstream and upstream data flows on two separate frequencies. The upstream and downstream carriers provide two shared channels for all CM's. On the downstream link the CMTS is a single data source and all CM's receive every transmission. On the upstream link all CM's may transmit and the CMTS is the single sink.

Packets sent over the downstream channel are broken into 188 byte MPEG frames each with 4 bytes of header and a 184 byte payload. Although capable of receiving all frames, a CM is typically configured to receive only frames addressed to its MAC address or frames addressed to the broadcast address. In addition to downstream user data, the CMTS will periodically send management frames. These frames control operations such as ranging, channel assignment, operational parameter download, CM registration and so on. Additionally, the CMTS periodically sends MAP messages over the downstream channel that identify future upstream time division multiple access (TDMA) slot assignments over the next MAP time. The CMTS makes these upstream CM bandwidth allocations (bandwidth grants) based on CM requests and QoS policy requirements.

The upstream channel is divided into a stream of time division multiplexed 'minislots' which, depending on system configuration, normally contain from 8 to 32 bytes of data. The CMTS must generate the time reference to identify these minislots. Due to variations in propagation delays from the CMTS to the individual CM's, each CM must learn its distance from the CMTS and compensate accordingly such that all CM's will have a system wide time reference to allow them to accurately identify the proper location of the minislots. This is called ranging and is part of the CM initialisation process.

Ranging involves a process of multiple handshakes between the CMTS and each CM. The CMTS periodically sends sync messages containing a timestamp. The CMTS also sends periodic bandwidth allocation MAPs. From the bandwidth allocation MAP the CM learns the ranging area from the starting minislot number and the ranging area length given in the message. The CM will then send a ranging request to the CMTS. The CMTS, after evaluating timing offsets and other parameters in the ranging request, returns to the CM a ranging response containing adjustment parameters. This process allows each CM to identify accurately the timing locations of each individual minislot.

In addition to generating a timing reference so that the CM's can accurately identify the minislot locations, the CMTS must also control access to the minislots by the CM's to avoid collisions during data packet transmissions. Fig. 2 illustrates a hypothetical MAP allocation that includes allocated slots for contention requests, user data and management data. For BE traffic, CM's must request bandwidth for upstream transmissions. There are several mechanisms available: contention bandwidth requests, piggybacked bandwidth requests and bandwidth requests for concatenated packets.

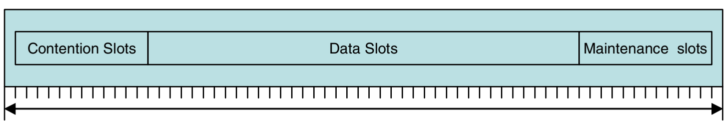

Fig. 2. Example upstream MAP allocation

Contention Bandwidth Requests. The CMTS must periodically provide transmission opportunities for CM's to send a request for bandwidth to the CMTS. As in slotted Aloha networks [6], random access bandwidth request mechanisms are inefficient as collisions will occur if two (or more) CM's attempt to transmit a request during the same contention minislot. Most implementations will have a minimum number of contention minislots to be allocated per MAP time, and in addition, any unallocated minislot will be designated as a contention minislot.

DOCSIS 1.1/2.0 specifies a truncated binary exponential backoff (tBEB) collision resolution algorithm (CRA) and it can be described as follows. When a packet arrives at the CM that requires upstream transmission, the CM prepares a contention-based bandwidth request by computing the number of minislots that are required to send the packet including all framing overhead. The contention algorithm requires the CM to randomly select a number of contention minislots to skip before sending this request (an initial backoff). This number is drawn from a range between 0 and a value that is provided by the CMTS in each MAP. The values sent are assumed to be a power of 2, so that a 5 would indicate a range of 0 – 31. After transmission, if the CM does not receive an indication that the request was received, the CM must randomly select another number of contention minislots to skip before re-sending the request. The CM is required to exponentially backoff the range following each collision with the maximum backoff specified by a maximum backoff range parameter contained in each MAP. The CM drops the packet after it has attempted to send the request 16 times.

As an example of tBEB, assume that the CMTS has sent an initial backoff value of 4, indicating a range of 0 – 15, and a maximum backoff value of 10, indicating a range of 0 – 1023. The CM, having data to send and looking for a contention minislot to use to request bandwidth, will generate a random number within the initial backoff range. Assume that an 11 is randomly selected. The CM will wait until eleven available contention minislots have passed. If the next MAP contains 6 contention minislots, the CM will wait. If the following MAP contains 2 contention minislots, a total of 8, the CM will still continue to wait. If the next MAP contains 8 contention minislots the CM will wait until 3 contention minislots have passed, 11 in total, and transmit its request in the fourth contention minislot in that MAP.

The CM then looks for either a Data Grant from the CMTS or a Data Acknowledge. If neither is received, the CM assumes a collision has occurred. The current backoff range is then doubled, i.e. the current value is increased from 4 to 5 making the new backoff range 0 – 31, and the process is repeated. The CM selects a random value within this new range, waits the required number of contention minislots, and resends its request. The backoff value continues to be incremented, doubling the range, until it reaches the maximum backoff value, in this example 10, or a range of 0 – 1023. The current backoff range will then remain at this value for any subsequent iterations of the loop. The process is repeated until either the CM receives a Data Grant or Data Acknowledge from the CMTS, or the maximum number of 16 attempts is reached.

Piggybacked Bandwidth Requests. To minimise the frequency of contention-based bandwidth requests, a CM can piggyback a request for bandwidth on an upstream data frame. For certain traffic dynamics, this can completely eliminate the need for contention-based bandwidth requests. This takes place via the Extended Header field in the MAC frames which can be used to request bandwidth for additional upstream transmissions during the current data transmission. Thus requests for bandwidth can be made outside of the contention process thereby reducing the frequency of collisions and consequently the access delay.

Bandwidth Requests for Concatenated Packets. DOCSIS 1.1/2.0 provides both Fragmentation MAC Headers, for splitting large packets into several smaller packets, and Concatenation MAC Headers, to allow multiple smaller packets to be combined and sent in a single MAC burst. One bandwidth request is presented to the CMTS for the group of packets undergoing concatenation. Concatenation can also be used to reduce the occurrence of collisions by reducing the number of individual transmission opportunities needed. Concatenation is the only method for transmitting more than one data packet in a single transmission opportunity. The CMTS, receiving the Concatenation MAC Header, must then 'unpack' the user data correctly. The Concatenation MAC Header precludes the use of the Extended Header field and therefore piggybacking of future requests cannot be done in a concatenated frame.

2.2 Quality of Service (QoS)

DOCSIS 1.1/2.0 manages bandwidth in terms of service flows that are identified by service flow IDs (SID's). Traffic arriving at either the CMTS or the CM for transmission over the DOCSIS 1.1/2.0 network is mapped to an existing SID and treated based on the profile. A CM will have at least 2 SID's allocated, one for downstream BE traffic and a second for upstream BE traffic. The upstream SID's at the CM are implemented as FIFO queues. Traffic-types extra to BE, such as VoIP, might be assigned to a different SID that supports a different scheduling service e.g. Unsolicited Grant Service (UGS) for toll quality telephony. The DOCSIS 1.1/2.0 protocol purposely does not specify the upstream bandwidth allocation algorithms so that vendors are able to develop their own solutions. DOSCIS requires CM's to support the following set of scheduling services: UGS, real-time polling service (rtPS), unsolicited grant service with activity detection (UGS-AD), non-real-time polling service (nrtPS) and BE service.

Unsolicited Grant Service (UGS). Designed to support real-time data flows generating fixed size packets on a periodic basis. For this service the CMTS provides fixed-size grants of bandwidth on a periodic basis. The CM is prohibited from using any contention requests. Piggybacking is prohibited. All CM upstream transmissions must use only the unsolicited data grants.

Real-Time Polling Service (rtPS). Designed to support real-time data flows generating variable size packets on a periodic basis. For this service the CMTS provides periodic unicast request opportunities regardless of network congestion. The CM is prohibited from using any contention requests. Piggybacking is prohibited. The CM is allowed to specify the size of the desired grant. These service flows effectively release their transmission opportunities to other service flows when inactive [1], demonstrating more efficient bandwidth utilisation than UGS flows at the expense of delay, which is worse.

Unsolicited Grant Service with Activity Detection (UGS-AD). Designed to support UGS flows that may become inactive for periods of time. This service combines UGS and rtPS with only one being active at a time. UGS-AD provides unsolicited grants when the flow is active and reverts to rtPS when the flow is inactive.

Non-Real-Time Polling Service (nrtPS). Designed to support non real-time data flows generating variable size packets on a regular basis. For this service the CMTS provides timely unicast request opportunities regardless of network congestion. The CM is allowed to use contention request opportunities.

BE Service. Designed to provide efficient service for the remaining flows. The CM is allowed to use contention or piggybacking to transmit bandwidth requests.

In the downstream direction, arriving packets are classified into a known SID and treated based on the configured service definition. For BE traffic, the service definition is limited to a configured service rate. For downstream traffic, the CMTS provides prioritisation based on SID profiles, where each SID has its own queue. Management frames, in particular MAP frames, are given highest priority. Telephony and other real-time traffic would be given next priority. BE traffic would share the remaining available bandwidth. There is also a single downstream transmission queue. Queuing occurs at the SID queues only if downstream rate control is enabled. All downstream queues are FIFO with the exception that MAP messages are inserted at the head of the transmission queue.

2.3 Security

Historically, cable systems have had an image as being insecure. The ‘always-on’ capability attracts attacks on subscriber networks. Subscriber networks with machines running Microsoft Windows with improper security settings have caused significant problems[2]. The security of cable networks has also been questioned since, as in a bus-based Ethernet LAN, data is received by all CM’s. By default, a CM is placed in non-promiscuous mode; however it is possible for a subscriber to change the configuration and to have the CM receive all data sent over the RF channel. Further, it is possible to increase the provisioned service rates by modifying the configuration. To counter the theft of service, CableLabs extended the Baseline Privacy Interface (BPI) security service described in the DOCSIS 1.0 specification to Baseline Privacy Interface Plus (BPI+) in the DOCSIS 1.1/2.0 releases.

BPI+ addresses two areas of concern: securing the data as it travels across the network and preventing the theft of service. Both BPI and BPI+ require encryption of the frames essentially forming a virtual private network for all transmissions between the CMTS and the CM, in order to protect the customer’s data as it traverses the coaxial cable. The Data Encryption Standard cipher algorithm is specified to be used in cipher block chaining mode for encryption of both the upstream and downstream MAC frame’s packet data in both the BPI and BPI+ security services. Public key encryption is used by the CM to securely obtain the required keys from the CMTS. Each CM must contain a key pair for the purpose of obtaining these keys from the CMTS.

To prevent the theft of service BPI+ requires the use of secure modem authentication procedures to verify the legitimacy of a particular CM. CM’s download their firmware from the service provider each time they boot. BPI+ requires the CM to successfully boot only if the downloaded code file has a valid digital signature. When a CM makes an authorisation request to the CMTS it must provide a unique X.509 digital certificate. After receiving a properly signed X.509 certificate and verifying the 1024 bit key pair the CMTS will encrypt an authorisation key using the corresponding public key and send it to the CM. A trust chain is developed by using a three level certificate hierarchy. At the top level is the root certification authority (CA) which belongs to CableLabs. The root CA uses its certificate to sign a manufacturer’s CA certificate at the second level. The manufacturer CA certificates are then used to sign individual certificates for each CM produced by that particular manufacturer. This process ensures that a given CM is legitimate and that the keys for encrypting the user’s data are only distributed to trusted CM’s.

Although DOCSIS 1.1/2.0 specifies the use of these security procedures to protect both the service provider and the customer, like all security measures, the system’s

[2] The security vulnerability occurs when a subscriber configures his/her network with file or print sharing. There are many reports of how dangerous this can be for example http://cable-dsl.navasgroup.com/netbios.htm#Scour.

defence is jeopardised if they are not used. Prior to 2005, polls and press reports indicated that the majority of cable network operators had not enabled the security methods required by DOCSIS 1.1/2.0.

2.4 A Queueing Network Model (QNM)

This section presents a QNM of a general DOCSIS 1.1/2.0 HFC network. The main contributors to the delay experienced by data that arrives to the CM's and requires onward routing from the DOCSIS network are the buffering delay at the CM, contention delay via the CRA, scheduling delay at the CMTS (bandwidth requests or periodic grants) and transmission delay (of the bandwidth requests (when applicable) and data transmissions).

The QNM draws on modelling aspects from a QNM of part of a DOCSIS 1.1/2.0 HFC network given in [7] and an open QNM of a DOCSIS 1.1/2.0 HFC network in [8]. It models transmission in both the upstream and downstream channels.

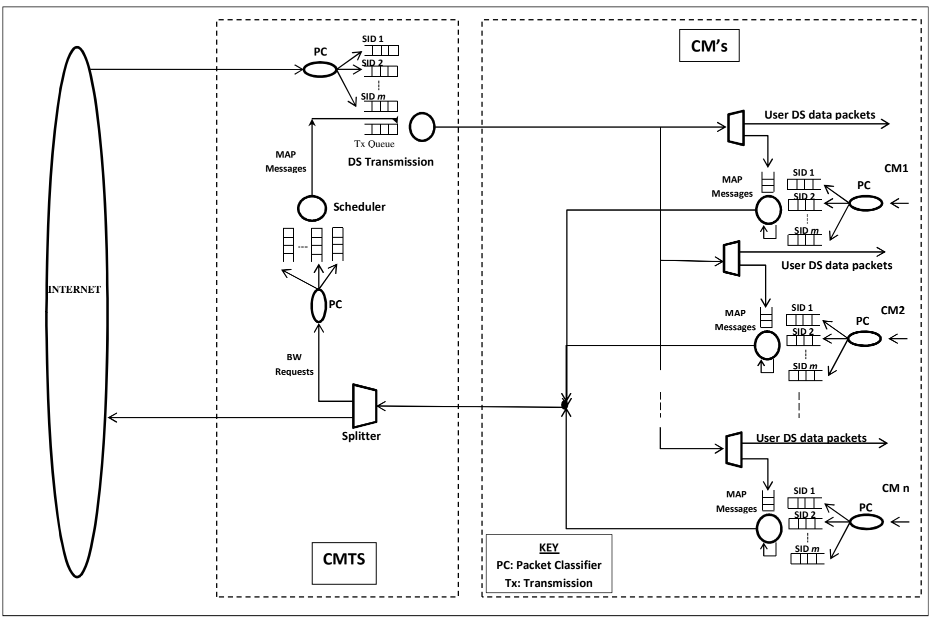

Fig. 3. QNM of a general DOCSIS 1.1/2.0 HFC network

The QNM illustrates that the bottleneck in a DOCSIS HFC network is upstream transmission due to the many-to-one access topology. In addition the upstream channels are restricted in their capacity to transport packets at high rates. This upstream packet rate limitation impacts both downstream and upstream throughput.

Most of this QNM (Fig. 3) is considered to be comprehensible in light of the overview of the protocol operation given in the previous sub-sections and the DOCSIS specification. The blocking at the CM schedulers represents contention and/or scheduling delay of data packets at the CM service station while they await an allocation of a transmission opportunity either in response to an (aperiodic) bandwidth request sent to the CMTS by the CM scheduler or via a pre-arranged periodic grant.

3 Performance Modelling of DOCSIS HFC Networks

This section summarises existing approaches and results of the performance modelling and evaluation of DOCSIS networks and centres on the following operational aspects among others: CRA's, upstream bandwidth/slot allocation algorithms, flow-priority scheduling disciplines and QoS provisioning. Finally an overview of research on the effect on DOCSIS network performance of running TCP applications is presented. Corresponding research surrounding HFC networks under the IEEE 802.14 protocol is included on some occasions due to its conceptual and operational similarity to DOCSIS in many respects and the implied opportunities for the cross-pollination of modelling approaches and their solutions between the two transmission technologies. Performance models of DAVIC/DVB HFC networks were not reviewed due to time limitations.

Similarities can be expressed as the existence of the following in both protocols: data encryption in both directions, CM ranging, time-slotted upstream channel, the use of random access methods for registering with and requesting bandwidth from the headend and employment of request-grant procedures for upstream bandwidth allocation. Piggybacking may be used by stations to make bandwidth requests. QoS support to differentiated flows by enabling the assignment of a subset of contention slots to a particular class of CM's and for reduced access delay both allow use of contention slots for short information transfers (immediate access). In addition both use the MPEG-2 format for downstream packet transmission. One of the major differences is in the variability of transmitted packet size and others include the implementation of the above algorithms for example the CRA, security and transport mechanisms and QoS support. The reader is directed to the following publications for detailed comparisons between these two protocols as well as the DAVIC/DVB standard and all their implementations: [9-11].

Overall, it was found that for DOCSIS-compliant HFC networks, performance was evaluated predominantly using simulation as opposed to analysis and little queue modelling has been carried out. This can be attributed to its inherently complex architecture and operation characterised by interdependence between several system-parameters. It was found that in all the cases of performance modelling i.e. simulation and analysis, simplifying assumptions were made to aid evaluation. Verification of models against real network data was starkly atypical.

3.1 Collision Resolution Algorithms (CRA's)

In order to control the performance experienced by packets arriving to the CM's due to the CRA, DOCSIS allows dynamic adjustment of the initial and maximum backoff range parameters. Therefore CRA performance models must essentially include these two parameters. If the initial backoff parameter is too low, then the frequency of collisions rises resulting in repeated attempts and hence greater delays whereas if either or both of the parameters are too large then minislots and thus bandwidth is wasted.

A great amount of analysis on the performance bounds experienced by contention requests under tBEB has been carried out in the literature and this has been helpfully summarised in [12]. Kwak et al carried out throughput and mean delay analysis of contention requests under a generalised exponential backoff CRA (where tBEB arose as a special case of the general version) [12]. The performance metrics were evaluated in terms of the above two backoff parameters for a fixed number of stations in a saturated network (i.e. one where each CM always has a packet to transmit) and where the collision resolution process was taken to be in equilibrium. Independence between successive collisions was assumed facilitating tractable mathematical analysis which was dominated by probabilistic arguments. The mean contention delay of requests in contention was taken as the mean total number of backoff time slots that a customer tarries while contending for transmission. Wang and Qiu [13] evaluated via simulation a proposed improvement to the tBEB algorithm that was claimed to offer improved delays to requests in contention.

The 802.14 CRA is based on a sophisticated n-ary tree splitting algorithm and analogous to tBEB it too has faced proposed extensions for example that by van den Broek et al [14]. Of the few queueing models encountered during the literature review on the performance modelling of HFC networks, one involved modelling two variants of a ternary tree-splitting contention algorithm using queueing models by Boxma et al [15]. The authors found that the first two moments of the contention delay in the free access contention tree algorithm are closely modelled by the sojourn time of a conventional finite-capacity machine repair network model with Random Order of Service (ROS)[3] whereas those of the blocked-access variant were found to match the sojourn time moments of the aforementioned network with a delay prior to the ROS queue. Noteworthy is the observation that unlike tBEB, the ternary-tree algorithm cannot be relied upon to preserve priority assignments when resolving collisions [16].

Lin et al [17] compared via simulation the request access delay (RAD) and throughput of HFC networks operating under early versions of the 802.14 (draft 2) and DOCSIS protocols. RAD was defined as the time from arrival of a data customer to the CM to the receipt of CMTS acknowledgement of the request for bandwidth and thus it is a measure of the CRA efficiency. Fair comparison was achieved by examining the performance measures of the network operating under the fine-tuned parameter settings and the same minislot allocation strategy of the respective protocols. Throughput was found to be very similar and RAD was at most about 22% longer in the DOCSIS network with notable differences occurring in the load range of about 40 – 75%. For the rest of the load range, the delay was very similar. This difference in network performance during moderate load was attributed to the more efficient first transmission policy in 802.14 which reduces the numbers of requests contending the same minislot cluster.

3.2 Bandwidth/Slot Allocation

As mentioned previously in Section 2.4, the performance bottleneck in DOCSIS 1.1/2.0 HFC networks is the uplink and the upstream throughput is highly dependent

[3] They assert that indeed any work-conserving service discipline produces the same results as the ROS scheduling discipline as shown by a prior result employing the PS service discipline.

on the ratio of contention capacity to total upstream channel capacity [18].Too high a ratio and the transmission capacity is limited resulting in reduced upstream bandwidth utilisation. On the other hand, too low a ratio and the opportunities to gain transmission grants decrease resulting in longer delays at the CM's. Several mechanisms have been proposed to alleviate these limitations and to this end studies (both analytic and simulation) have also been carried out to determine the optimum ratio. It can be seen intuitively that allocating unused minislots for contention (corresponding to periods of light loading) and maintaining a minimum number of minislots for contention during periods of high load helps to reduced contention delay [2, 18].

Lambert et al [8, 19] modelled the upstream transmission (i.e. contention and reservation) in DOCSIS 1.1/2.0 HFC networks by an open QNM comprising two processor share (PS) queues in tandem and employed a decomposition technique to solve the network. It was found that a ratio of between 10%-15% yielded the least access delay for a wide range of inter-arrival time correlation levels in the arrival processes of data packets to the CM's. The work of Lambert et al is detailed in Section 4.2.

Cho et al [20] observed from simulation experimentation using the common simulation framework[4] version 13.0 (CSF v13) that the optimal ratio of contention to total upstream capacity, resulting in the highest throughput and least access delay was 0.15 for a MAP size of 2ms. It must be pointed out that though the assumed arrival process of customers to the CM's was not stated in [20], their result is accepted here because of support from Lambert et al's observed invariance of the ratio to correlation-levels in the arrival process [19] and thus it is thought invariance to the arrival process.

3.3 Flow-Priority Scheduling and Quality of Service (QoS) Provisioning

The scheduling of the DOCSIS-defined priority-services in networks is open to vendor-implementation. Several scheduling mechanisms have been designed and evaluated in the literature and feature priority queueing mechanisms such as the weighted fair queueing (WFQ) policy or its variants, Pre-emptive Resume (PR) priority scheduling and the earliest deadline first (EDF) policy among others. Existing works pertaining to the provision of QoS in DOCSIS HFC networks are summarised below.

It was found by Xiao and Bing [21] in a very sparsely-populated experimental DOCSIS network implemented using *Arris*™ CM's and CMTS that for CBR traffic flows to the CM's, the network performance varied with packet length in the following way: in both upstream and downstream directions the larger the packet, the greater the throughput and the less the delay and packet loss. This expected observation is attributable to the lower ratio of overhead to transmitted data for larger packets compared to that when smaller packets are being transmitted. Additionally, as expected, loss began to occur in both directions when the rate of sending of packets reached the corresponding link capacity and this coincided with levelling off of throughput and delay. It was found that the level of network performance experienced

[4] A DOCSIS simulation program created by OPNET Technologies Inc. in conjunction with CableLabs.

by customers to different CM's belonging to a particular priority-class was the same. For fixed packet lengths, all classes of CM's experienced the same level of network performance for increasing sending rates until they approached to within a close margin of a 'breakout' value, after which the throughputs levelled-off, loss rates ramped at different speeds and delay rose extremely quickly to different maximum values. Finally a simulation of the experimental DOCSIS network was carried out in OPNET with a standard WFQ CMTS scheduling mechanism where priorities were implemented by setting appropriate probabilities of being served. The simulation revealed that in the context of performance this scheduling discipline accurately modelled the behaviour shown by the *Arris*™ CMTS.

Hawa and Petr [22] proposed a preliminary new CMTS scheduling architecture to enable the five DOCSIS-defined QoS services satisfy the bandwidth and delay guarantees of CBR, VBR and BE traffic. The complex queueing station design had multiple buffers and grouped the five flows into three classes: Type 1 represented UGS data grants and unicast requests for rtPS and nrtPS flows, Type 2 represented flows with minimum bandwidth reservations and Type 3 related to flows with no bandwidth reservations. Type 1 flows were prioritised over Types 2 and 3 flows via a semi-pre-emptive mechanism (whereby customers were allowed to complete service when their deadlines preceded those of new arrivals to the queue). Flows within the latter two classes were differentiated through a priority WFQ system. Both random early detection (RED) and multi-priority RED were proposed for use in buffer management due to their support for TCP traffic. The authors stated that they were in the process of evaluating the scheduler's performance via simulation and analysis[5].

Zhenglin and Chongyang [23] modelled scheduling at the CMTS analytically under various simplifying assumptions including two traffic flows namely real-time CBR traffic and non-real-time data traffic under UGS and BE DOCSIS contention services respectively. The authors considered the real-time flow having higher priority over data and implemented this prioritisation using the PR scheduling discipline. The arrivals of the bandwidth requests of the two classes to the CMTS scheduler were assumed to be independent Poisson processes with different rates and the service times of these two classes were assumed to be independent and generally distributed with different means. Concatenation and immediate access were not modelled. The fixed contention slot allocation scheme was used with contention slots grouped at the end of each MAP. Mean performance metrics at the CMTS were derived via the celebrated P-K formula and stochastic and queueing theoretic arguments. Neither the analytic model nor formulae were verified against simulation or actual measurement. The experiments carried out showed, as expected, that the use of the PR scheduling discipline in this specific context enabled the CBR real-time traffic to meet its stringent time constraints obviously at the expense of non-real-time traffic whose requests could be timed-out. Finally it was shown that larger packets exhibited better bandwidth utilisation efficiency. This was attributed to the fact that larger packets use a relatively smaller physical overhead.

Lin and Lee [24] designed and evaluated an admission control policy at the CMTS to service time-sensitive flows. This QoS scheduling mechanism permitted flows

---

[5] At the time of publication of this tutorial, a later publication by Hawa and Petr presenting performance evaluation of their CMTS scheduling queueing architecture was not found after a brief search of the internet.

based on available bandwidth and delay guarantee provision using the tolerated jitter parameter. The EDF policy was employed at the CMTS scheduler and it was claimed via analysis and simulation that both delay and throughput were enhanced in this admission control policy compared to several existing scheduling schemes.

Unfortunately these performance gains come at the expense of rejection of flows whose QoS requirements cannot be fulfilled.

Droubi et al [25] proposed a CMTS architecture that provided bandwidth guarantees using the existing self-clocked weighted fair queueing (SCFQ) scheme. This scheme was chosen because it provides the least computation and implementation complexity among the WFQ algorithms. The SCFQ scheme can be implemented as a head of the line queue where customer priorities are their finish times calculated using the negotiated transmission rates thus incorporating bandwidth guarantees. In addition an implementation was proposed for providing the UGS service for delay-sensitive CBR traffic. The architectures were verified via simulation using the CSF package. Two sets of experiments were conducted over a sparsely populated network of 15 CM's characterised by Poisson then Markov-modulated Poisson process (MMPP) arrivals respectively.

Bushmitch et al [26] proposed a new upstream service flow scheduling service, UGS with piggybacking and showed via simulation, using real video traces, that it improved both the overall upstream bandwidth utilisation and delay experienced by real-time upstream VBR video packets when compared to the existing UGS (low delay, CBR allocation) and rtPS (good bandwidth utilisation for both CBR and VBR but higher delay) service flow provisioning. This came at the expense of more complex implementation and the degraded QoS experience of lower priority SID flows e.g. BE flows. It must be noted that in DOCSIS 1.1/2.0 piggybacking is not permitted with UGS nor are any other contention mechanisms and therefore the aim of this proposal was to highlight possible areas of improvement to the DOCSIS 1.1/2.0 specification. The application of the proposed scheduling service assumed that the real-time VBR traffic had been 'smoothed' to reduce burstiness. The authors referred to works which state that compressed digital video and other types of video streams are long range dependent exhibiting burstiness over multiple time scales. Several 'smoothing' techniques of video streams were described, most of which result in video streams comprising a significant CBR component and an additional bursty component which cannot be avoided. It is this CBR component that was supported by the UGS part of the scheduling discipline and the piggyback requests accommodated the extra bandwidth required for the bursts, while maintaining better delay constraints than when using rtPS.

Golmie et al [16] on the other hand proposed facilitating differentiated service by a three-pronged approach. Firstly priorities were assigned to contention slots such that a flow with a new request waited for a group of contention slots with its own priority in order to transmit and did so with probability one within this range of slots. The second scheme involved customising the tBEB backoff range offering such that for high priority requests the maximum backoff parameter was set to the number of contention slots reserved for that priority. Thus high priority customers retransmitted in the assigned contention slots in the next available MAP with probability one and this way delays were minimised. Finally the ratio of data to contention slots was adjusted dynamically according to an algorithm that was a slight modification of an existing one. The authors evaluated via simulation a DOCSIS network comprising up to 200 CM's servicing Poisson arrivals and it was clearly observed that indeed differentiated service was provided in terms of access delays to the different classes of customers. This occurred however at the expense of the delays experienced by the lower priority classes and extra processing at the CMTS and CM's.

Sdralia et al [27] evaluated via simulation using CSF v12, a priority-FCFS scheduling mechanism at the CMTS in the interest of providing reference statistics against which the performance of the CMTS under other scheduling mechanisms could be compared. Here requests for bandwidth that arrived at the CMTS were queued in their respective priority buffers. The authors simulated a network comprising 200 CM's and eight priorities while ignoring concatenation. It was found that the maximum upstream throughput efficiency was about 77% with larger packet sizes of 1.5 kB and only 61% for smaller packet sizes of 100 bytes. These conservative maxima are due to MAC and physical layer overheads, unused capacity and the MAP structure. The authors also found that small packet sizes exhibited high access delay which could be reduced with concatenation. They asserted that large packet sizes make more efficient use of bandwidth but under saturation even large packets suffer and thus justified the inclusion of prioritisation.

A proposal for network-wide QoS provisioning via a QoS management device connected to the CMTS on the one hand and to proposed QoS controllers in the CM's on the other was made by Adjih et al [28]. The QoS management device was designed to fulfil QoS levels to requesting subscribers by logging the network usage statistics and when bandwidth is limited, negotiating with the QoS controllers more suitable traffic demands. In this case the authors proposed a network of adaptable CM's sensitive to bandwidth availability.

This QoS support however, comes at the expense of implementation complexity and increased network traffic due to the additional management packets traversing the network.

Lin et al [17] compared via simulation the performance of a DOCSIS HFC network under three upstream scheduling disciplines: shortest job first (SJF), longest job first (LJF) and modified-SJF. Here the size of job (i.e. short or long) refers to the amount of bandwidth requested by the CM. The SJF discipline showed poorer RAD but lower data transfer delay (DTD), a measure of the efficiency of the upstream transmission scheduling algorithm and defined as the time between receipt of bandwidth request at the CMTS and subsequent receipt of full data packets there. The larger RAD was attributed to the shorter DTD that results in larger proportions of time that the CM is empty and consequently a larger proportion of arrivals to the CM having to contend for channel transmission via the CRA. The modified-SJF was sought to help to alleviate this limitation by splitting data grants allocated to a single CM into smaller sizes and distributing these across several minislots. The network running under the modified-SJF scheduling discipline exhibited the most balanced performance of the three disciplines.

3.4 TCP Applications over DOCSIS HFC Networks

The intended use of DOCSIS cable networks for IP transmission necessitates the study of the behaviour of TCP traffic over the DOCSIS network as this transmission service forms a major proportion of Internet traffic. Further, the behaviour of TCP over asymmetric paths in other infrastructures such as wireless systems [29-32] has implications on its effect on the performance of DOCSIS networks.

A network exhibits bandwidth asymmetry when running TCP applications if achieved throughput is not solely a function of the link and traffic characteristics of the forward direction but in fact depends on the impact of transmission in the reverse direction too. Most of the prior work focused on highly asymmetric paths with respect to bandwidth where the normalised asymmetry level (the ratio of raw bandwidths to the ratio of packet sizes in both directions) typically would be of the order of 2-4 [29]. In DOCSIS HFC networks the upstream channel exhibits packet rate asymmetry due to low upstream packet rates with respect to downstream capacity. However the problem symptoms are similar. Various methods have been proposed to alleviate the TCP over asymmetric path problems including header compression and modified upstream queue policies (drop-from-front, TCP acknowledgement prioritisation, TCP acknowledgement filtering). Some of these ideas can be applied to DOCSIS networks. For example, a CM that supports TCP acknowledgement filtering could drop 'redundant' TCP acknowledgements that are queued. While this would increase the TCP acknowledgement rate, it would also increase the level of TCP acknowledgement compression. TCP acknowledgement reconstruction could be implemented in the CMTS to prevent the increased level of TCP acknowledgement compression from affecting network performance.

Liao and Ju [33] designed and evaluated two novel mechanisms to improve TCP transmission that is downstream-heavy via an ns-2 simulation of a small DOCSIS network comprising 30 CM's. In the first, bandwidth requests were sent faster in order to reduce the asymmetry ratio thereby helping to reduce upstream access delay while maintaining TCP downstream data transmission rates. In the mechanism 'piggybacked' bandwidth requests were sent in reserved unicast minislots at the front of the MAP if the data grant was at the backend of its MAP and the new transmission cycle had not yet started. If the new packet arrived before the start of the data grant but after the start of the reserved minislot, the CM would send the request via piggybacking in the normal way.

Naturally, this can be seen to occur at the expense of additional contention delay for new stations attempting to begin transmitting and additional implementation complexity.

In the second mechanism, upstream TCP acknowledgements were prioritised by reducing the sending rates of the larger data packets compared to those of the (smaller) TCP acknowledgements. This had an adverse effect on upstream TCP transfer latency though not significantly, it was claimed, and additional implementation complexity at the CMTS (only).

Elloumi et al [34] found that TCP throughput over an 802.14 network was low primarily due to TCP acknowledgement compression. The authors proposed two solutions: one involving piggybacking and a second involving TCP rate smoothing by controlling the TCP acknowledgement spacing. It was found that piggybacking helped reduce the burstiness associated with the TCP acknowledgement stream in certain situations. However it was limited in its ability to effectively match offered load over a range of operating conditions. The authors' second solution was to control the TCP sending rate by measuring the available bandwidth and calculating an appropriate TCP acknowledgement rate and allowing the CM to request a periodic grant that would provide sufficient upstream bandwidth to meet the required TCP acknowledgement rate.

Cohen and Ramanathan observed that an HFC network presents difficulties for TCP due to the asymmetry between upstream and downstream bandwidth's and due to high loss rates (the authors assumed channel loss rates as high as 10-50%) [35]. Because of the known problems associated with TCP/Reno in these environments [36-38], the authors proposed a 'faster than fast' retransmit operation where a TCP sender assumes that a packet is dropped when the first duplicate TCP acknowledgement is received (rather than the usual triple duplicate TCP acknowledgement indication).

4 Case Studies

In this section, two performance modelling studies of DOCSIS 1.1/2.0 HFC networks, based on open QNM's and evaluated via simulation and analysis respectively are detailed.

The first study evaluates the performance of a DOCSIS 1.1/2.0 HFC network via and '*ns*' simulation [2]. In light of the previously observed impact of DOCSIS network configurations on performance [39, 40], two sets of experiments were conducted to investigate the impact of different network configurations and those of different upstream bandwidth allocation strategies on the network performance when carrying TCP/IP traffic. Parameter values were discovered that showed a marked improvement in the network performance characterised by almost perfect downstream utilisation, significantly reduced access delay and lower collision rates and web response times (WRT's).

The second performance model is a high level abstraction of a DOCSIS 1.1/2.0 HFC network represented by an open QNM with blocking [8]. This was solved approximately by decomposing the QNM into two dependent sub-models: a closed QNM and a group of single-server queues. An optimum range of ratios of contention channel capacity to entire uplink channel capacity that minimised the mean time for packets to exit the cable network (equal to the mean time to access the wide area network/Internet) was derived.

4.1 Simulation

The simulation modelled the behaviour of the DOCSIS 1.1/2.0 MAC and physical layers as defined in [1] over a cable network which is illustrated in Fig. 4. A detailed discussion of the validation of the model is presented in [41]. The implementation of the simulation network model and associated web traffic models were based on the "flexbell" model with user-session variables characterised by heavy-tailed distributions with infinite variance as defined in [42]. These user-session variables include inter-arrival times of web pages, number of objects per webpage and the size of objects and so on. Withstanding the challenges of simulating real networks with web traffic characterised by not only self-similar (mono-fractal) but also multi-fractal properties at small time-scales, the flexbell model with session variables satisfying different Pareto distributions has been found to provide reasonable estimates to real

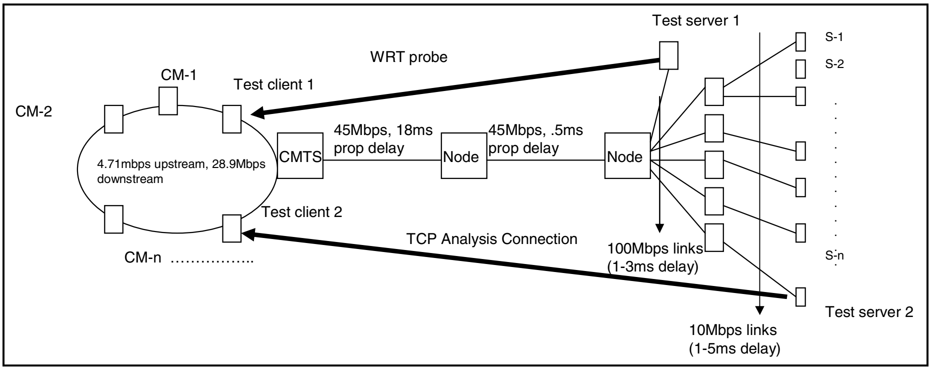

Fig. 4. Simulation network model

network behaviour, exhibiting self-similar properties despite not convincingly bearing multi-fractal scaling [42]. This makes such a model an attractive basis for use in studying network transmission technologies with current multimedia traffic profiles.

The "flexbell" topology represents numerous clients (which in the case of DOCSIS networks are the CM's) at one end of the network connected via a single bottleneck link to numerous sets of servers at the other. Transmission to a particular set of servers is through a single node as illustrated at the right-hand side of the simulation network model in Fig. 4 [42].

The simulation modelled the CM contention process, TDMA upstream bandwidth allocation and upstream and downstream packet transmission with all the nodes in the simulation network model (Fig. 4) acting as delay-stations, modelled as finite-capacity queues. The maximum size of each queue was a simulation parameter.

All experiments involved a variable number of CM's (i.e. CM-1 through CM-n in the simulation network model) that interacted with a set of servers (S-1 through S-n). The RTT from the CM's to the servers was randomly selected in the range between 42 – 54 ms.

Downstream web traffic was simulated via a four-dimensional traffic model where each constituent component (i.e. each dimension) modelled a different user-session variable satisfying a heavy-tailed infinite variance distribution. These variables were simulated using different Pareto distributions whose parameters values are given below. In addition to downstream web traffic, 5% of the CM's were configured to generate downstream low speed UDP streaming traffic (i.e. a 56Kbps audio stream), 2% of the CM's downstream high speed UDP streaming traffic (i.e. a 300Kbps video stream) and 5% of the CM's to generate downstream P2P traffic. The P2P model (based on [43]) incorporated an exponential on/off TCP traffic generator that periodically downloaded on average 4MB of data with an average idle time of 5s between each download.

The limitations of the simulation were as follows: i) CM's were confined to a single default BE service flow and a single UGS or rtPS flow; ii) the model was limited to one upstream channel for each downstream channel; iii) the model did not support dynamic service provisioning; iv) physical layer impairments were not modelled; v) the model assumed that the CMTS and the CM clocks were synchronised.

The model accounted for MAC and physical layer overhead including forward error correcting (FEC) data in both the upstream and downstream directions. For the simulations an FEC overhead of 4.7% (8% in the upstream direction) was assumed and this was modelled by reducing channel capacity accordingly[6]. The downstream and upstream channels supported an optional service rate. Service rates were implemented using token buckets where the rate and maximum token bucket size were simulation parameters.

Traffic arriving at either the CMTS or the CM for transmission over the DOCSIS 1.1/2.0 HFC network was mapped to an existing SID and treated based on the profile. In this DOCSIS HFC network model, when a CM session began, it registered itself with the CMTS which established the default upstream and downstream SID. A CM had an upstream FIFO queue for each SID. In the downstream direction there were per SID queues as well as a single transmission queue. Queuing occurred at the SID queue only if downstream rate control was enabled. All downstream queues were FIFO with the exception that MAP messages were inserted at the head of the transmission queue.

The scheduler had a configured MAP time (through a MAP_TIME parameter) which was the amount of time covered in a MAP message. The MAP_FREQUENCY parameter specified how often the CMTS sent a MAP message. Usually these two parameters were set between 1 – 10 ms. The scheduling algorithm supported dynamic MAP times through the use of a MAP_LOOKAHEAD parameter which specified the maximum MAP time the scheduler could 'look ahead'. If this parameter was 0, MAP messages were limited to MAP_TIME amount of time in the future. If set to 255 the scheduler could allocate up to 255 slots in the future. This was only used on BE traffic and only if there were no conflicting periodic UGS or rtPS allocations.

The grant allocation algorithm (i.e. the scheduling algorithm) modelled requests as jobs of a non-pre-emptive soft real-time system [44]. The system could hold two types of jobs: periodic and aperiodic. Periodic jobs resulted in UGS periodic data grants and rtPS periodic unicast request grants. Aperiodic jobs were in response to rtPS and BE requests for upstream bandwidth. Every job had a release time, a deadline and a period. The release-time denoted the time after which the job could be processed. The deadline denoted the time before which the job had to have been processed. For periodic jobs, the period was used to determine the next release time of the job.

The scheduler maintained four queues of jobs where a lower number queue had priority over a higher number queue. The first and second queues contained UGS and rtPS periodic jobs respectively both operating under the EDF policy. UGS jobs were unsolicited grants and rtPS jobs were unsolicited polls to CM's for bandwidth requests. The third queue contained all the bandwidth requests that were in response to previous unicast request grants. Similarly, the fourth queue contained the bandwidth requests that arrived successfully from the contention request process with the latter two queues serviced according to the FIFO discipline. The CMTS processed jobs from the four queues in strict priority order with no pre-emption.

---

[6] To account for FEC overhead the upstream channel capacity was reduced by 8%. This approximation was suggested by CISCO Systems Inc. (www.cisco.com). The DOCSIS 1.1/2.0 framing overhead adds an additional 30 bytes to an IP packet. A system tick of 6.25 μs and an effective channel capacity of 4.71Mbps lead to 18 bytes of data per slot for a total of 85 slots required for a 1500 byte IP packet.

The parameters associated with a UGS service flow included the grant size, the grant interval and the maximum tolerated jitter. When a CM registered a UGS flow with the CMTS, the CMTS released a periodic job in the system with release time set to the current time and the deadline set to the release time plus maximum tolerated jitter. Finally, the period was set to the grant interval. After every period, a new instance of the job was released.

The same algorithm was used for rtPS except that the maximum poll jitter was used to determine the deadline. Requests for bandwidth allocations from BE contention or from rtPS polling were treated as aperiodic jobs. Periodic jobs with the earliest deadline were serviced first. Remaining bandwidth was then allocated to aperiodic jobs. The scheduler had an additional parameter PROPORTION that was used to establish a relative priority between rtPS allocations and BE allocations.

The two sets of simulation experiments (I and II) were based on the network depicted in Fig. 4 and the respective simulation delay model in Fig. 5 below. The second set differed from the first set in several significant ways: i) the scheduler allocated unused slots for contention requests; ii) the number of IP packets allowed in a concatenated frame was no longer limited to two; iii) the buffer size at the CMTS downstream queue was increased from 50 to 300 packets; iv) the number of system ticks per slot was increased from 4 to 5 which decreased the number of slots per map from 80 to 64.

The underlying simulation delay model is illustrated below in Fig. 5.

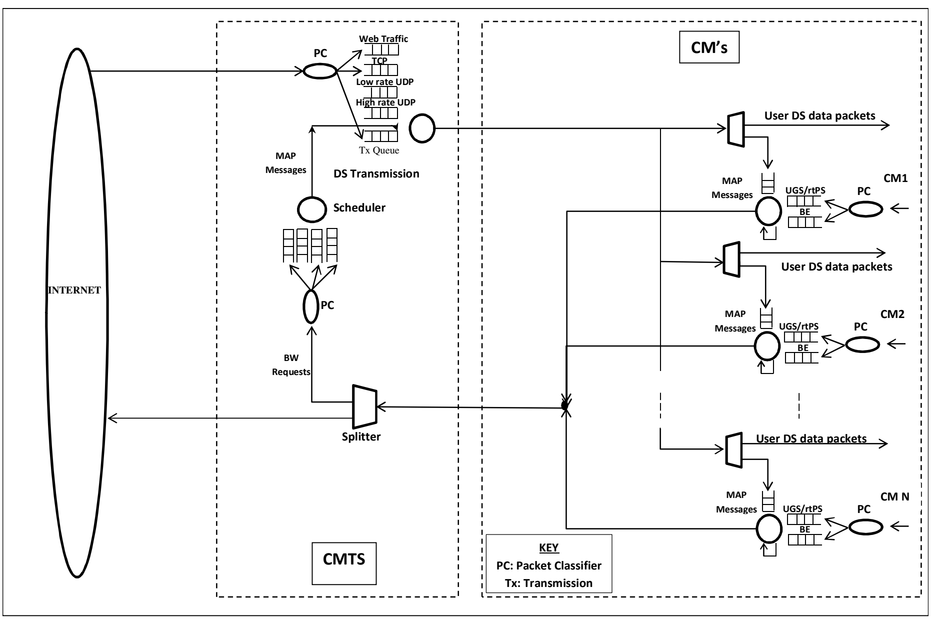

Fig. 5. Simulation delay model

The simulation model and downstream traffic parameter-values are given, respectively, in Table 1 and Table 2 below. Both sets of experiments were conducted over a range of settings of MAP_TIME and for a given MAP_TIME setting the number of CM's was varied from 100 to 500[7]. This was carried out over six MAP_TIME settings ranging from .001 to .01 s. The default MAP time setting was 2 ms (80 minislots per MAP).

Table 1. Simulation model parameter settings

Parameter	Value
upstream bandwidth	5.12 Mbps
downstream bandwidth	30.34 Mbps
Preamble	80 bits
Ticks per minislot	4
Fragmentation	OFF
Concatenation	ON
MAP_LOOKAHEAD	255 slots
Backoff Start	8 slots
Backoff stop	128 slots
Contention slots	12
Management slots	3
Simulation time	1000 s

Table 2. Downstream web traffic model parameter values

Traffic component	Pareto Mean	Pareto Shape Parameter
Inter-page interval	10	2
Objects per page	3	1.5
Inter-object interval	0.5	1.5
Object size	12 (segments)	1.2

For each experiment the following statistics were obtained:

Collision rate. Each time a CM detected a collision it incremented a counter. The collision rate was the ratio of the number of collisions to the total number of upstream packet transmissions attempted.

Downstream and upstream channel utilisation. At the end of a run, the CMTS computed the ratio of the total bandwidth consumed to the configured raw channel bandwidth. The utilisation value reflects the MAC and physical layer overhead including FEC bits.

Average upstream access delay. All CM's kept track of the delay from when an IP packet arrived at the CM in the upstream direction until it got transmitted. This statistic is the mean of all of the samples.

[7] Many providers provision a downstream RF channel by assigning 2000 households per channel which made this range of active CM's reasonable.

Web response time (WRT). A simple TCP client server application was run between Test Client 1 and the Test Server 1. Test Server 1 periodically sent 20KB of data to Test Client 1. With each iteration the client obtained a response time sample. The iteration delay was set to 2 s. At the end of the test, the mean of the response times was computed. The mean WRT was linked to end user perceived quality by using a very coarse rule of thumb which proposes that end users are bothered by lengthy download times characterised by WRT > 1s. This value was not advocated to be an accurate measure of end user quality of experience but rather it was used to simply provide a convenient network performance reference.

Experiment Set I. When the dominant application is web browsing (which uses the TCP service of TCP/IP's Transport Layer) the majority of data travels in the downstream direction. However, for certain configurations, the system can become packet rate bound in the upstream direction which can limit downstream throughput due to a reduced TCP acknowledgement rate. For the first set of experiments, piggybacking and concatenation were enabled however the maximum number of packets that could be concatenated into a single upstream transmission was limited to two.

Fig. 6 shows that the collision rates got extremely high as the number of active CM's increased. When only 100 users were active, the collision rate was about 50%. As the load increased, the collision rate approached 90-100% depending on the MAP_TIME setting. The behaviour of the system was influenced by several MAC protocol parameters. First, the number of contention slots assigned per map (i.e. the CONTENTION_SLOTS) directly impacted the collision rates at high loads. This set of experiments used a fixed number of contention slots, 12 per MAP which, as illustrated in Fig. 6, was insufficient at high loads. The set of curves in Fig. 6 illustrate the collision rate at different MAP_TIME settings. The collision rate was roughly 10 percent higher for the largest MAP_TIME than for the smallest MAP_TIME. This was a direct result of the MAP allocation algorithm which allocated a fixed number of contention slots each map time. If the scheduler's behaviour was altered so as to assign all unused data slots for contention, the collision rate would have been significantly lower. As the MAP_TIME was increased the bandwidth allocated for contention requests was effectively reduced.

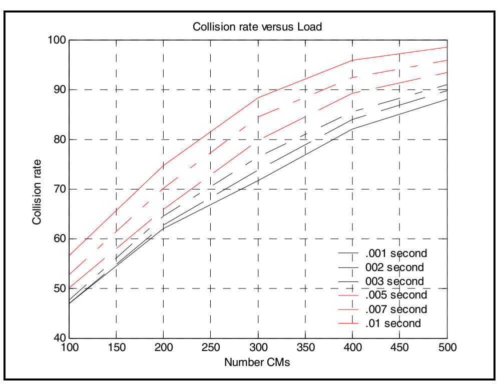

Fig. 6. Upstream collision rates as the number of CM's increase. (Experiment set I).

Fig. 7 a and b plot the channel utilisation as the load (i.e. number of active CM's) was increased. The downstream utilisation reached a maximum of about 64% with a MAP_TIME setting of .001 s. In this case, 12 contention slots per MAP were sufficient. For smaller MAP_TIME values, the downstream utilisation ramped up to its maximum value and then decreased at varying rates as the load was increased. As the collision rate increased, downstream TCP connection throughput decreased. Larger MAP_TIME values resulted in fewer contention-slot allocations leading to higher collision rates and reduced downstream utilization.

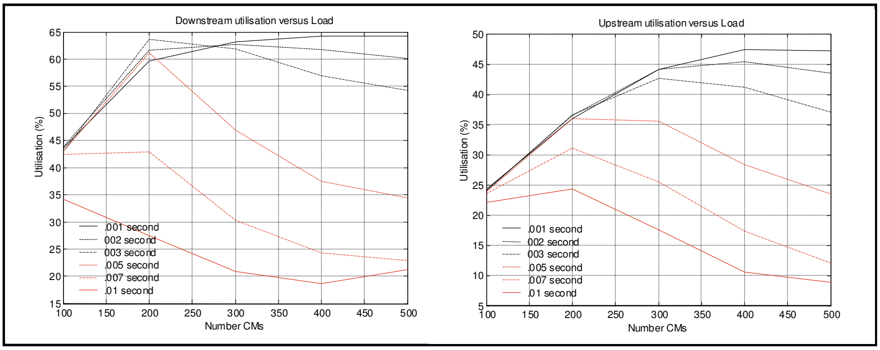

Fig. 7a. Downstream channel utilisation. (Experiment set I).

Fig. 7b. Upstream channel utilisation

Further illustrating this behaviour, Fig. 8 a shows that the average upstream access delay became very large at high loads when configured with large MAP_TIME settings. Even for lower MAP_TIME values, the access delay was significant. For a MAP_TIME of .002 s, the access delay exceeded .5s at the highest load level. To assess the end-to-end cable network performance WRT's were monitored. Using the rule of thumb described earlier, Fig. 8 b indicates that for MAP_TIME settings less than .005, up to 300 active users were accommodated before performance became bothersome. This result is clearly not generally applicable as it depends on the specific choice of simulation parameters.

Rather than making the full channel capacity available to subscribers, MSO's typically offer different service plans where each plan is defined by a service rate. For example, Charter communications offers a 3Mbps downstream rate and 512Kbps upstream rate [45]. While reduced service rates prevent customers from consuming more than their fair share of bandwidth at the expense of other customers, they offer little benefit when the network becomes congested. Fig. 9 a and b illustrate the results of an experiment that was identical to the web congestion scenario of Fig. 8 a and b except that CM's were restricted to a 2Mbps downstream service rate. Fig. 9 a shows the average upstream access delay was almost identical to that observed in the scenario without rate control. The WRT results shown in Fig. 9 b further suggest that a 2Mbps downstream service rate limit was of little use.

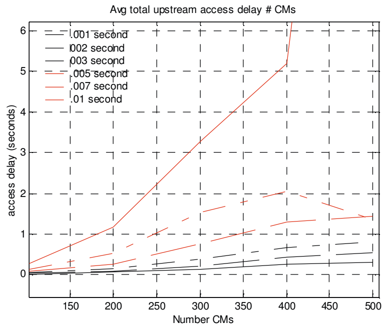

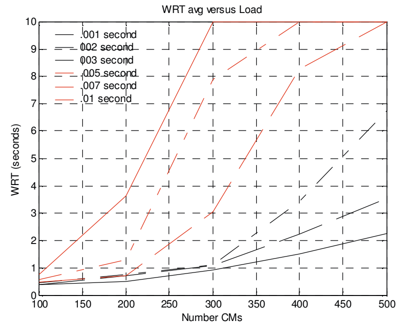

Fig. 8a. Upstream access delay (no rate control) (Experiment set I).

Fig. 8b. WRT

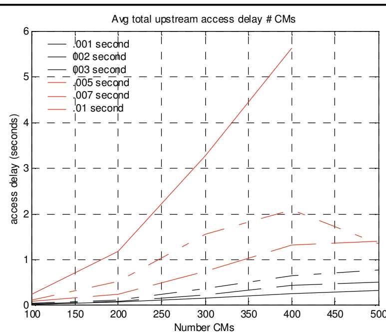

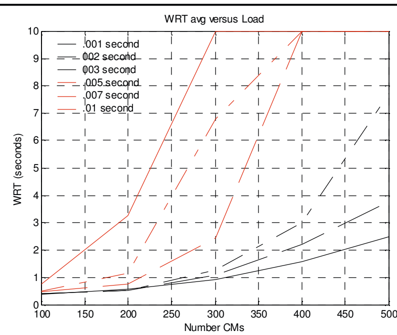

Fig. 9a. Upstream access delay (with rate control) (Experiment set I).

Fig. 9b. WRT

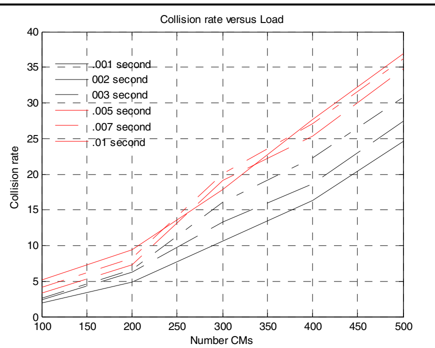

Fig. 10. Upstream collision rates. (Experiment set II).

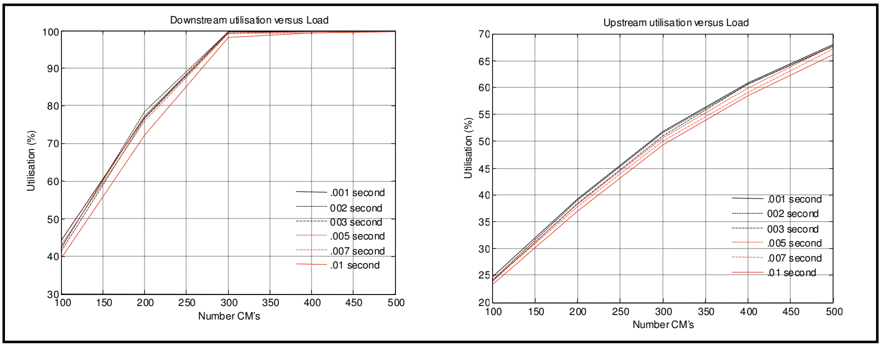

Fig. 11a. Downstream channel utilisation (Experiment set II)

Fig. 11b. Upstream channel utilisation

Experiment Set II. In the second set of experiments, the change that had the most impact was the increased bandwidth allocated for upstream contention requests. Fig. 10 shows that the collision rate ranged from 2% - 37% compared to 50% - 100% for set I. Collision rates were lowest for the runs with smaller MAP times. As the offered load to the system increased, the number of unused slots became smaller consequently reducing the number of contention slots. Hence the proportion of bandwidth allocated for contention slots was greater for small MAP times.

Fig. 11 a and b show that the utilisations in both directions were higher with a marked increase in the downstream utilisation and both were not affected by MAP times, unlike the first set of experiments. The invariance of upstream utilisation to MAP size was attributed to the profitable use of piggybacking by the runs with larger MAP times thus countering the adverse impact of their larger collision rates. The upstream rates of TCP acknowledgement packets in turn correspondingly affect the downstream rates, explaining the invariance to MAP times in the downstream direction.

The increased upstream utilisation was attributed to the increase in number of packets permitted to be concatenated and the greater number of transmission grants as a consequence of more slots available for bandwidth requests. Higher upstream TCP acknowledgement rate in turn has a positive effect on the downstream utilisation as does the larger downstream buffer. The increased downstream efficiency in turn leads to greater upstream utilisation.

Fig. 12 illustrates that from 40% - 90% of all packets sent upstream used a piggyback bandwidth request depending on the MAP size used. The runs with large MAP times were able to take advantage of piggybacking more than the runs with small MAP times because there was more time for packets to accumulate while waiting for a data grant.

The experiments were repeated with the concatenation limit relaxed and similar results were obtained with the exception that extreme levels of TCP acknowledgement-packet compression occurred. It has been shown that TCP acknowledgement compression leads to higher loss rates and that it makes it difficult for protocols that estimate bottleneck bandwidth's or that monitor packet delays to operate correctly [46-48]. Also, concatenation significantly increases access delay experienced by packets at other CM's therefore it is avoided by MSO's. However since all nodes in the network model were configured with adequate buffers in the simulation, network performance was not adversely impacted by the bursty traffic profiles caused by the TCP acknowledgement compression.

Piggybacking is less effective than concatenation for primarily downstream TCP traffic. It tends to be more advantageous for scenarios involving constant upstream traffic or backlogged upstream flows as it ensures that one packet is transmitted every cycle.

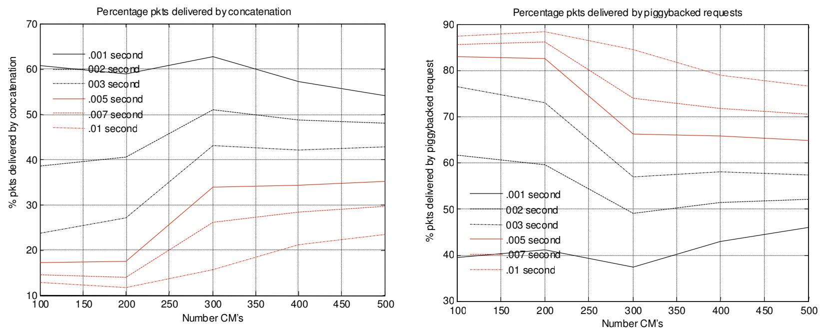

a. Percentage of concatenation b. Percentage of piggybacking

Fig. 12. Proportions of packets delivered by concatenation or due to piggybacked requests. (Experiment set II).

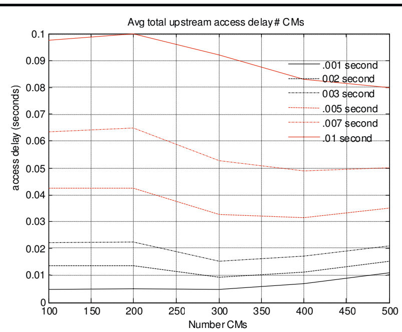

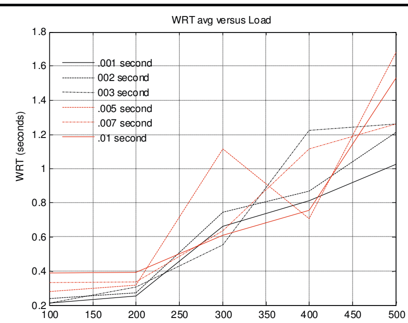

Fig. 13a. Upstream access delay (Experiment set II)

Fig. 13b. WRT

The upstream access delay was an order of magnitude lower than the first set of experiments, attributable to the lower collision rate which was in turn due to more bandwidth available for contention. The WRT too, was around one order of magnitude lower because of less frequent TCP retransmissions due to reduced packet-loss, which in turn was a direct consequence of the larger downstream buffer (Fig. 13).

The network performance improvements gained by transitioning from the parameter-settings of experiment set I to II were dramatic and they can be summarised as follows:

- As expected, the collision rate decreased dramatically due primarily to higher levels of bandwidth allocated for contention.
- The utilisation in both directions was higher with a marked increase in the downstream utilisation and both were not affected by MAP times, unlike the first set of experiments due to concatenation restriction relaxation and the dynamic contention bandwidth allocation.
- The access delay is more than an order of magnitude lower because of the reduced collision rate.
- The WRT metric was also around one order of magnitude lower because of fewer TCP retransmissions due to lower packet loss.

These performance gains can be seen to be achieved at the expense of increased implementation complexity and at the cost of providing greater downstream buffer capacity.

4.2 Analysis

In this section an overview is given of the optimisation of the upstream bandwidth allocation in DOCSIS 1.1/2.0 HFC networks via an open QNM and its approximate analytic solution [8]. Lambert et al [8] abstracted the upstream contention and transmission processes as an open QNM and derived the optimum ratio of contention to total uplink channel capacity that minimised the sum of the mean access and mean transmission times. The authors also assessed the effect of varying several system parameters on this ratio, for example the number of CM's, initial backoff parameter and so on.

The open QNM operates with blocking and comprises a group of single server finite-capacity queues (modelling the CM's) connected to the first of two single-server finite-capacity PS queues in tandem. The first PS queue (referred to as the contention queue) models the contention delay and the second (referred to as the reservation queue) models the transmission time of the data packets. The latter queue models the complement of the contention channel capacity, called the reservation channel and in this case the PS service discipline is chosen for its ability to appropriately capture the DOCSIS MAC design principle of distributing the transmission capacity fairly among active stations. The service rate of the contention queue was taken to be the saturation throughput of tBEB at equilibrium according to established modelling practice within the 802.11 environment and this was derived analytically.

The following network diagram (Fig. 14) has been inferred from the open QNM's textual description in [8]:

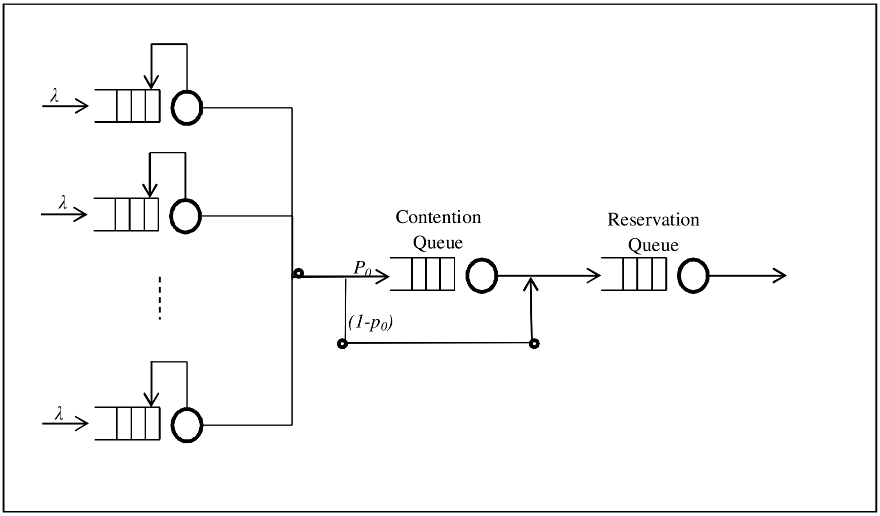

Fig. 14. DOCSIS HFC network abstracted open QNM

The QNM is constrained by the condition that there is only one customer per CM in either the contention or reservation queues at any one time therefore a customer at the head of the waiting line in a CM queue is blocked until its previous counterpart leaves the open network.

The arrival process to the network i.e. to the CM's was modelled as Poisson with rate λ packets/ms. It was claimed in [8] that Poisson arrivals to the CM queues result in a pessimistic prediction for the amount of contention channel needed because such arrivals operating with piggybacking forfeit a significant amount of performance gain experienced by bursty arrivals sending bandwidth requests via piggybacking. All arriving packets to nonempty CM's rely on piggybacking to send their bandwidth request and thus they bypass the contention process.

The solution involved decomposing this open QNM into two interdependent sub-models (a closed QNM and a collection of independent single-server queues) and then evaluating these sub-models in repeated succession of each other whereby certain input values required for the solution of one of the sub-models was taken from the most recent solution of the other. Thus the intermediate solutions converge and the iterative solution process stopped when the desired level of accuracy is achieved.

Sub-Model I. The closed QNM was obtained by removing the buffers (waiting rooms) of the CM queues and forming an unbroken loop as shown in Fig. 15 below. The packet length distribution is irrelevant as the performance measures of such a network are insensitive to the service-time distribution of its constituent PS service centres. The mean residence times at the contention and reservation queues $E[R_{Cont}]$ and $E[R_{Res}]$ respectively are required for solving the second subsystem and they can be calculated using any algorithm used to solve closed QNM's with a fixed level of multiprogramming. Here p_0 is the probability of having an empty CM queue.

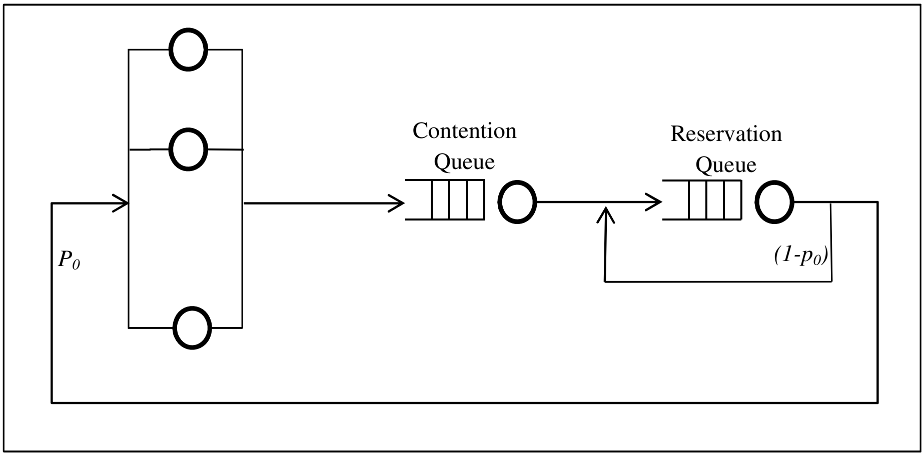

Fig. 15. Sub-model I: closed QNM

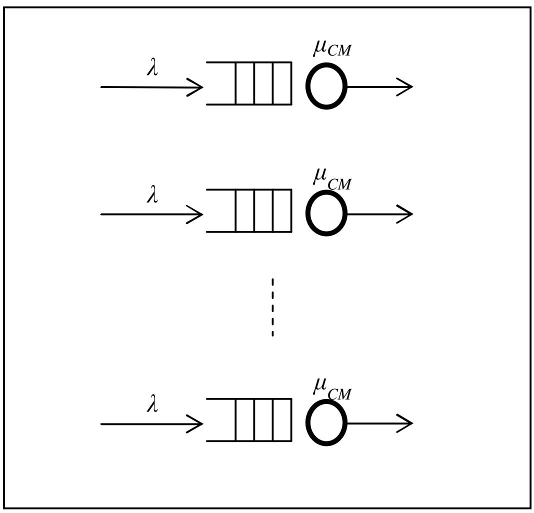

Fig. 16. Sub-model II: Collection of single-server queues

Sub-Model II. The independent single-server queues comprising this sub-model are standard M/M/1/N queues operating under the FCFS scheduling discipline and they represent the CM's (Fig. 16**.**). In a later publication in order to assess the impact of a correlated arrival process on the optimal fraction, the authors employed a multiple class Markovian arrival process with marked transitions (the MMAP process). The MMAP arrival process is a non-renewal point-process capable of modelling correlation with analytic tractability and it generalises a large group of inter-arrival time distributions for example those characterising the MMPP, the phase-type renewal process and their respective superpositions [19].

The mean service time of each queue, $1/\mu_{CM}$ is the mean time spent by the customer contending for access and transmitting its data and it is calculated as follows:

$$\frac{1}{\mu_{CM}} = E[R_{Cont}]p_0 + E[R_{Res}] \,.$$

Subsequently p_0 is updated via the steady-state probability distribution for an M/M/1/N queue and this new value of p_0 is fed into the evaluation of the first subsystem. This iterative process continues till the desired accuracy of p_0 is achieved.

It was found that the analytic results of the saturation throughput and mean access and transmission times of the network compared reasonably with the simulation of the original open QNM with increasing numbers of CM's and arrival rates as well as increasing ratio of contention to total upstream channel capacity. This implies that the decomposition technique used provides a good means for solving complex open QNM's of this kind.

It was deduced that assigning 10 – 15% of the upstream minislots for contention yields near optimum results. Interestingly an identical range was discovered as optimal in the advanced system with a wide range of levels of correlation in the arrival process. In other words the fraction of contention to total upstream channel capacity was found to be invariant to the level of correlation in the arrival process at the CM's [19]. The exact value (within this range) depends on the specific system-parameter values such as data load level, minimum contention window and number of CM's among others.

It would be interesting to discern the level of accuracy of the open QNM by say comparing its simulation against actual network measurements or another independently constructed simulation.

5 Conclusions and Future Work

The DOCSIS 1.1/2.0 protocol over HFC cable networks constitute a complex system with many interdependent parameters, the intricacies of which are further heightened by the presence of bursty and/or self-similar input traffic flows characteristic of current internet traffic. This has often necessitated the performance evaluation of DOCSIS 1.1/2.0 HFC networks via simulation rather than analytic modelling especially when the nature and extent of the interdependence among several network characteristics is being studied. On the other hand analytic methodologies provide a cost-effective means to derive optimal (or optimal narrow ranges of) parameter-estimates for dimensioning a limited number of operational aspects of DOCSIS HFC networks. This tutorial has shown how both simulation and analytic approaches can successfully be used to optimise the performance of DOCSIS 1.1/2.0 HFC network configurations. Moreover the respective tradeoffs encountered as a consequence of performance improvements to DOCSIS networks were identified.

The DOCSIS protocol continues to evolve and cable network equipment implementing the latest DOCSIS 3.0 standard is now being deployed. This latest version supports Internet Protocol version 6 and achieves much higher service rates in both the upstream and downstream directions as multiple channels can be 'bonded' together and thus deliver more packets simultaneously [49]. While channel bonding can greatly increase raw capacity, limitations may still be experienced say in downstream DOCSIS network throughput of TCP traffic, which is directly affected by the rate at which TCP acknowledgements can be transported upstream by the cable modems. DOCSIS 3.0 does take steps to reduce this bottleneck by for example allowing individual cable modems to have multiple requests outstanding at any given time [49].

It is expected that many of the performance impacting behaviours observed in the former releases of DOCSIS over HFC networks will also broadly exist under DOCSIS 3.0 [49] and future extensions. In this context the performance models and their quantitative analyses reviewed in this tutorial could also be used, with appropriate enhancements to evaluate and predict the performance of new and future releases of DOCSIS protocols over HFC networks.

References

1. Cable Television Laboratories Inc.: Radio Frequency Interface Specification, Data Over Cable Service Interface Specifications DOCSIS 2.0, `http://www.cablelabs.com/cablemodem/specifications/specifications20.html`
2. Shah, N., et al.: A tutorial on DOCSIS: protocol and performance models. In: Third International Working Conference on the Performance Modelling and Evaluation of Heterogeneous Networks, Ilkley, UK (2005)
3. Shirali, C., Shahar, M., Doucet, K.: High-bandwidth interface for multimedia communications over fixed wireless systems. IEEE Multimedia 8(3), 87–95 (2001)
4. Elfeitori, A.A., Alnuweiri, H.: Network architecture and medium access control for deploying third generation (3G) wireless systems over CATV networks: Research Articles. Wirel. Commun. Mob. Comput. 5(2), 139–152 (2005)
5. Tanenbaum, A.S.: Computer Networks, 4th edn. Prentice Hall PTR, Englewood Cliffs (2003)
6. Abramson, N.: THE ALOHA SYSTEM: another alternative for computer communications. In: Proceedings of the Fall Joint Computer Conference, November 17-19, pp. 281–285. ACM, Houston (1970)
7. Laubach, M.E., Farber, D.J., Dukes, S.D.: Delivering Internet Connections over Cable: Breaking the Access Barrier. John Wiley, New York (2001)
8. Lambert, J., Van Houdt, B., Blondia, C.: Dimensioning the Contention Channel of DOCSIS Cable Modem Networks. In: Boutaba, R., Almeroth, K.C., Puigjaner, R., Shen, S., Black, J.P. (eds.) NETWORKING 2005. LNCS, vol. 3462, pp. 342–357. Springer, Heidelberg (2005)
9. Lin, Y.-D., Yin, W.-M., Huang, C.-Y.: An investigation into HFC MAC protocols: Mechanisms, implementation, and research issues. IEEE Communications Surveys & Tutorials 3(3), 2–13 (2000)
10. Ali, M.T., et al.: Performance evaluation of candidate MAC protocols for LMCS/LMDS networks. IEEE Journal on Selected Areas in Communications 18(7), 1261–1270 (2000)
11. Perkins, S., Gatherer, A.: Two-way broadband CATV-HFC networks: state-of-the-art and future trends. Computer Networks 31(4), 313–326 (1999)
12. Kwak, B.-J., Song, N.-O., Miller, L.E.: Performance analysis of exponential backoff. IEEE/ACM Trans. Netw. 13(2), 343–355 (2005)
13. Wang, B., Qiu, K.: Study of Collision Resolution Algorithms of DOCSIS. Journal of University of Electronic Science and Technology of China 32(33), 293–295 (2003)

14. van den Broek, M.X., et al.: A novel mechanism for contention resolution in HFC networks. In: Twenty-Second Annual Joint Conference of the IEEE Computer and Communications, INFOCOM 2003, IEEE Societies, Los Alamitos (2003)
15. Boxma, O., Denteneer, D., Resing, J.: Delay models for contention trees in closed populations. Performance Evaluation 53(3-4), 169–185 (2003)
16. Golmie, N., Mouveaux, F., Su, D.H.: Differentiated services over cable networks. In: Global Telecommunications Conference, GLOBECOM 1999 (1999)
17. Lin, Y.-D., Huang, C.-Y., Yin, W.-M.: Allocation and scheduling algorithms for IEEE 802.14 and MCNS in hybrid fiber coaxial networks. IEEE Transactions on Broadcasting 44(4), 427–435 (1998)
18. Chu, K.-C., et al.: A novel mechanism for providing service differentiation over CATV network. Computer Communications 25(13), 1214–1229 (2002)
19. Lambert, J., Van Houdt, B., Blondia, C.: Queues in DOCSIS cable modem networks. Computers & Operations Research 35(8), 2482–2496 (2008)
20. Cho, S.-H., Kim, J.-H., Park, S.-H.: Performance evaluation of the DOCSIS 1.1 MAC protocol according to the structure of a MAP message. In: IEEE International Conference on Communications, ICC 2001 (2001)
21. Xiao, C., Bing, B.: Measured QoS performance of the DOCSIS hybrid-fiber coax cable network. In: The 13th IEEE Workshop on Local and Metropolitan Area Networks, LANMAN 2004 (2004)
22. Hawa, M., Petr, D.W.: Quality of service scheduling in cable and broadband wireless access systems. In: Tenth IEEE International Workshop on Quality of Service (2002)
23. Zhenglin, L., Chongyang, X.: An Analytical Model for the Performance of the DOCSIS CATV Network. The Computer Journal 45(3), 278–284 (2002)
24. Lin, J.-T., Lee, W.-T.: Bandwidth admission control mechanism for supporting QoS over DOCSIS 1.1 HFC networks. In: 10th IEEE International Conference on Networks, ICON 2002 (2002)
25. Droubi, M., Idirene, N., Chen, C.: Dynamic bandwidth allocation for the HFC DOCSIS MAC protocol. In: Proceedings of the Ninth International Conference on Computer Communications and Networks (2000)
26. Bushmitch, D., et al.: Supporting MPEG video transport on DOCSIS-compliant cable networks. IEEE Journal on Selected Areas in Communications 18(9), 1581–1596 (2000)
27. Sdralia, V., et al.: Performance characterisation of the MCNS DOCSIS 1.0 CATV protocol with prioritised first come first served scheduling. IEEE Transactions on Broadcasting 45(2), 196–205 (1999)
28. Adjih, C., et al.: An Architecture for IP Quality of Service Provisioning Over CATV Networks. In: Roger, J.-Y., Stanford-Smith, B., Kidd, P.T. (eds.) Business and Work in the Information Society: New Technologies and Applications, pp. 368–374. IOS Press, Ohmsha (1999)
29. Balakrishnan, H., Padmanabhan, V.N., Katz, R.H.: The effects of asymmetry on TCP performance. Mobile Networks and Applications 4(3), 219–241 (1999)
30. Lakshman, T.V., Madhow, U., Suter, B.: Window-based error recovery and flow control with a slow acknowledgement channel: a study of TCP/IP performance. In: Proceedings of Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE INFOCOM 1997 (1997)
31. Jacobson, V.: Compressing TCP/IP headers for low-speed serial links (1990)
32. Kalampoukas, L., Varma, A., Ramakrishnan, K.K.: Improving TCP throughput over two-way asymmetric links: analysis and solutions. In: Proceedings of the 1998 ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems, pp. 78–89. ACM, Madison (1998)

33. Liao, W., Ju, H.-J.: Adaptive slot allocation in DOCSIS-based CATV networks. IEEE Transactions on Multimedia 6(3), 479–488 (2004)
34. Elloumi, O., et al.: A simulation-based study of TCP dynamics over HFC networks. Computer Networks 32(3), 307–323 (2000)
35. Cohen, R., Ramanathan, S.: TCP for high performance in hybrid fiber coaxial broad-band access networks. IEEE/ACM Trans. Netw. 6(1), 15–29 (1998)
36. Elloumi, O., Afifi, H., Hamdi, M.: Improving congestion avoidance algorithms for asymmetric networks. In: 1997 IEEE International Conference on Communications, ICC 1997, Towards the Knowledge Millennium, Montreal (1997)
37. Fall, K., Floyd, S.: Simulation-based comparisons of Tahoe, Reno and SACK TCP. SIGCOMM Comput. Commun. Rev. 26(3), 5–21 (1996)
38. Hoe, J.C.: Improving the start-up behavior of a congestion control scheme for TCP. In: Conference Proceedings on Applications, Technologies, Architectures, and Protocols for Computer Communications, pp. 270–280. ACM, Palo Alto (1996)
39. Martin, J., Shrivastav, N.: Modeling the DOCSIS 1.1/2.0 MAC protocol. In: Proceedingsof the 12th International Conference on Computer Communications and Networks, ICCCN 2003 (2003)
40. Martin, J.: The Interaction Between the DOCSIS 1.1/2.0 MAC Protocol and TCP Application Performance. In: Second International Working Conference on the Performance modelling and Evaluation of Heterogeneous Networks, Ilkley, UK (2004)
41. Martin, J., Westall, J.: A Simulation Model of the DOCSIS Protocol. SIMULATION 83(2), 139–155 (2007)
42. Feldmann, A., et al.: Dynamics of IP traffic: a study of the role of variability and the impact of control. SIGCOMM Comput. Commun. Rev. 29(4), 301–313 (1999)
43. Saroiu, S., Gummadi, P., Gribble, S.: A Measurement Study of Peer-to-Peer File Sharing Systems. In: Multimedia Computing and Networking. San Jose, CA, USA (2002)
44. Stankovic, J., et al.: Deadline Scheduling for Real-time Systems - EDF and Related Algorithms. Kluwer Academic Publishers, Dordrecht (1998)
45. Charter Communications, `http://www.charter.com/`
46. Balakrishnan, H., et al.: TCP behavior of a busy Internet server: analysis and improvements. In: Proceedings of the Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM 1998, IEEE, Los Alamitos (1998)
47. Paxson, V.: Measurements and Analysis of End-to-End Internet Dynamics. In: Computer Science Division. University of California, Berkeley (1997)
48. Martin, J., Nilsson, A., Injong, R.: Delay-based congestion avoidance for TCP. IEEE/ACM Transactions on Networking 11(3), 356–369 (2003)
49. Cable Television Laboratories Inc.: MAC and Upper Layer Protocols Interface Specification, Data-Over-Cable Service Interface Specifications DOCSIS 3.0, `http://www.cablelabs.com/cablemodem/specifications/specifications30.html`

Appendix I: List of Acronyms

ABR	Available bit rate
ATM	Asynchronous transfer mode
BE	Best effort
BPI	Baseline privacy interface
BPI+	Baseline privacy interface plus
CA	Certification authority
CATV	Community antenna television
CBR	Constant bit rate
CM	Cable modem
CMTS	Cable modem termination system
CRA	Collision resolution algorithm
CSF vX.Y	Common simulation framework version X.Y
DAVIC	Digital Audio Visual Council
DOCSIS	Data-over-cable service interface specification
DSL	Digital subscriber line
DTD	Data transfer delay
DVB-RCC	Digital video broadcast return channel for cable
EDF	Earliest deadline first
FEC	Forward error correcting
HFC	Hybrid fibre coax
LJF	Longest job first
MAC	Media Access Control
MCNS	Multimedia cable network system
MMAP	Markovian arrival process with marked transitions
MMPP	Markov-modulated Poisson process
MSO	Multiple system operator
nrtPS	Non-real-time polling service
PR	Pre-emptive resume
PS	Processor share
QNM	Queueing network model
QoS	Quality of Service
RAD	Request access delay
RED	Random early detection
RF	Radio frequency
ROS	Random order of service
rtPS	Real-time polling service
SCFQ	Self-clocked weighted fair queueing
SID	Service flow ID
SJF	Shortest job first
tBEB	Truncated binary exponential backoff
TCP	Transmission Control Protocol
TDMA	Time division multiple access
UBR	Unspecified bit rate
UGS	Unsolicited grant service
UGS-AD	Unsolicited grant service with activity detection
VBR	Variable bit rate
VoIP	Voice over IP
WFQ	Weighted fair queueing
WRT	Web response time