

Experimental Evaluation of DOCSIS 1.1 Upstream Performance

Radim Bartoš¹, Chaitanya K. Godsay^{1,2}, and Steve Fulton²

¹ Department of Computer Science

² InterOperability Laboratory

University of New Hampshire, Durham, NH 03824, USA

E-mail: rbartos@cs.unh.edu

ABSTRACT

Data-Over-Cable Service Interface Specification (DOCSIS) is one of the many last mile technologies intended to provide Internet access and packet-based services to the customer. DOCSIS uses the widely deployed hybrid fiber/coax (HFC) network as the physical link between multiple cable modems (CMs) and the cable modem termination system (CMTS).

This paper presents the upstream performance of DOCSIS 1.1 in the physical layer and MAC layer across various traffic patterns. The experiments are conducted on real devices and not simulators as used in all previous research. The use of real devices allows us to capture the complete complexity of the protocol, and gives us realistic results but it also limits our control over different parameters. The goal of the project is to study the impact of different parameters that can be controlled by the cable service provider and to compare different CMs with respect to upstream performance. The performance metrics used are upstream data rate and channel utilization. The results can be used by the cable operators to optimize their networks, by the CM and CMTS manufacturers to enhance their products and they may help in identifying protocol bottlenecks for upstream performance.

KEYWORDS

Data over cable, DOCSIS, upstream performance evaluation.

1 Introduction

Cable operators, in the early nineties, envisioned the growth of cable networks and were driven to explore possibilities for transmitting data from the residential user to the service provider. By providing this capability, packet-based services, such as high-speed Internet access, cheaper telephone connections, and video-conferencing could be deployed easily. This led to the formation of many research groups and Multimedia Cable Network System (MCNS), a collaboration of cable companies, was the first to come up with a specification. MCNS released the set of standards known as DOCSIS 1.0 (Data Over Cable Service Interface Specification) in March 1997. CableLabs, a non-profit research and development consortium, worked in collaboration with MCNS, and is now responsible for developing new specifications and product certification.

Specification [1] describes a DOCSIS network as a tree based network with the Cable Modem Termination System (CMTS) as the root of the tree and the Cable Modems (CMs) as the leaves of the tree. The CMTS is at the service provider facility and the CM is at the residential user home. The transmission of data from the CMTS to CM, termed as downstream, is a point to multipoint broadcast, whereas the transmission from the CM to CMTS, termed as upstream, is controlled by the CMTS and is multipoint to point TDMA. DOCSIS defines an asymmetric network in terms of upstream and downstream data rate, with downstream being substantially larger than the upstream. The residential user has the Customer Premise Equipment (CPE), such as computer, telephone, etc., connected to the CM. Data in upstream goes from the CM to the CMTS, which is then forwarded appropriately. Similarly, data in downstream, passes from the CMTS to the CM, which is forwarded to the CPE. Typically there are 1500 to 2000 CMs connected to a CMTS with distance between the CMTS and CM going up to 50 miles.

Let us briefly discuss a typical data transmission in upstream. Once the CM has registered with the CMTS, it is allowed to transmit data upstream. However, it can send data only when it is allowed by the CMTS to do so. Since upstream is multipoint to point TDMA arbitrated by the CMTS, the CMTS sends *bandwidth allocation MAPs*, simply termed as MAPs, at regular intervals. The MAP explicitly describes the time when a CM is allowed to transmit data upstream.

Transmitting data upstream is a three-step process. First the CM has to send a request for a data grant to the CMTS, then it has to wait to get a data grant from the CMTS (in the MAP) and then it must send the data (at time specified by the MAP). We call this process the RDS cycle (Request-Data grant wait-Send cycle).

When the CPE sends some data to the CM, the CM looks in the most recent MAP for the REQ or REQ/DATA region. It then creates a data grant request message indicating the grant size and transmits the request in the time specified for the REQ or REQ/DATA region in the MAP. These regions are subject to collisions as many CMs can try to send a data grant request message. If the data grant request message reaches the CMTS then it either sends a long data grant or short data grant to the CM in the following MAP. A long or a short data grant depends on how

much data the CM wants to send. The CMTS will send a data grant pending message in the following MAP, if the CMTS has received the data grant request from the CM but cannot allocate a data grant to it. The CM detects a collision when it does not get a short data grant, long data grant, or a data grant pending in the next MAP. In the case of a collision, the CM starts exponential backoff for collision resolution. The CM will then defer certain number of request opportunities before requesting again.

However, if the CM gets a short/long data grant successfully from the CMTS, then it extracts the time to send the data from the MAP and transmits the data to the CMTS at the time specified by the MAP. The CM thus has to go through one or more RDS cycles to transmit data upstream.

It can be seen that DOCSIS is a complex protocol that is difficult to model. This paper attempts to capture the full complexity of the protocol by measuring the performance of real devices.

2 Motivation and Goals

In the early nineties many research groups were formed to develop a specification for delivery of data in the last mile using the widely deployed cable networks. Organizations involved in this effort were Data Over Cable Service Interface Specification group (MCNS-DOCSIS), IEEE 802.14 working group, Society of Cable Telecommunications Engineers (SCTE), Digital Video Broadcasting (DVB), Digital Audio Video Council (DAVIC), and ATM Forums Residential Broadband Working Group (RBWG). The initial research in DOCSIS was on comparing different aspects of specifications developed by different groups, mainly IEEE 802.14 and MCNS-DOCSIS [2]. The effect of different upstream allocation and scheduling algorithms, MAP rates, MAP length, etc., were studied in [3, 4, 5]. Efforts to statistically predict the upstream requests and allocate data grants was presented in [6]. Recently there has been a study on performance evaluation of DOCSIS 1.1 using simulator [7]. All of the above research have employed simulators for experimentation with most of them using Opnet as the DOCSIS simulator.

The presented study has been carried out by using real devices with the goal to evaluate the impact of different upstream parameters, that can be controlled by the cable service provider, on upstream performance. Different parameters are considered from the PHY layer, MAC layer, and traffic patterns.

In the PHY layer the effect of different modulation formats and modulation rates is considered. In the MAC layer we consider performance enhancers, such as concatenation and piggybacking, that improve the upstream performance. However, there has been no research on the behavior of these enhancers. How much performance improvement do they provide? Is there a situation where using these enhancers might be detrimental? Are there bottlenecks in using a combination of these enhancers? This

paper presents the experimental evaluation of these aspects of the protocol.

Another aspect of the study is to understand the behavior of the DOCSIS network under different kinds of traffic loads. Devices are generally tested with a constant bit rate stream. However, a constant bit rate stream is hardly ever generated in the real world. Since Internet is the biggest application driving the cable modem market we decided to concentrate on generating Internet-like traffic. Some advanced testing devices now support traffic distributions for packet length. A traffic pattern that closely approximates realistic traffic was generated and transmitted through the DOCSIS network to aid in understanding the sensitivity of the DOCSIS network to different traffic patterns.

The change in performance due to the addition of CMs across different CM manufacturers is also analyzed. We hope that the results will be used by the vendors to enhance their products.

It should be noted that we do not test for the conformance to the DOCSIS 1.1 protocol. The conformance and certification testing of CMs and CMTSes is done by CableLabs only. We use certified and conformant CMs and CMTSes to get statistical results and analyze the impact of different protocol parameters.

3 Parameters considered in this study

We use upstream channel utilization and upstream data rate as the performance metrics. Latency, another measure of performance, is not considered in this paper, however, it is a subject of our current research. The parameters considered for evaluation of upstream performance are enlisted below.

Modulation Formats and Modulation Rates: DOCSIS

1.1 allows two modulation formats for upstream transmission, QPSK (2 bits/symbol) and 16QAM (4 bits/symbol). DOCSIS 1.1 also supports five modulation rates for upstream transmission, 160 ksym/s, 320 ksym/s, 640 ksym/s, 1280 ksym/s, and 2560 ksym/s which correspond to the channel widths of 200 kHz, 400 kHz, 800 kHz, 1600 kHz, and 3200 kHz respectively. The product of modulation rate and modulation formats gives us the theoretical maximum upstream data rate for each combination of modulation rate and modulation format. We thus have 0.32 Mbps, 0.64 Mbps, 1.28 Mbps, 2.56 Mbps, and 5.12 Mbps as the theoretical maximum data rate for QPSK and 0.64 Mbps, 1.28 Mbps, 2.56 Mbps, 5.12 Mbps, and 10.24 Mbps for theoretical maximum data rate for 16QAM. The product of modulation rate and modulation format is also referred to as the *channel rate*.

Concatenation: The CM can send a concatenated burst of packets instead of small packets if allowed by the configuration file and the CMTS. The configuration file provides operational parameters to the CM when it

registers with the CMTS. We use the configuration file to control the performance enhancers (on/off). The improvement in upstream performance by using concatenation, in terms of maximum data rate achieved without dropping any packets, is studied.

Piggybacking: If the CM wants to send data upstream, it has to request a data grant in the REQ or REQ/DATA region. Alternatively, it can request a data grant by piggybacking a request with the data packet being sent. The CM does so by adding an extended header to the data packet being sent. Piggybacking is enabled by using the configuration file, and should be allowed by the CMTS.

Traffic profiles: We conduct experiments on Constant Packet Length-Constant Bit Rate (CPL-CBR) traffic, as it provides useful insights into the behavior of CMs. However, since the biggest application driving the cable modem market is high-speed Internet connectivity, we study the behavior of a DOCSIS network with the source, transmitting Internet-like traffic. We define Distributed Packet Length-Constant Bit Rate (DPL-CBR) as traffic that would transmit different packet sizes at the micro level but maintain a constant bit rate at the macro level. Several advanced traffic generators now support quad-modal packet length distribution (four Gaussian distributions superimposed, each with adjustable mean and half-point width). The above feature is used to compare, understand and analyze the behavior of the DOCSIS network under realistic traffic loads.

Number of CMs: Since the upstream performance is limited by the RDS cycle, having more active CMs on the network with the potential of transmitting data upstream increases the probability of collisions. Since the number of RDS cycles increase on collision, we study the effect of upstream performance for different numbers of CMs on the network.

CM chipset manufacturers: We conduct experiments on devices based on different chipsets. This provides important insights into the sensitivity of different CMs to different performance parameters. These results can also be used by the CM manufacturers to enhance their devices.

4 Experimental setup

The experiments presented in this paper were conducted on the testbed shown in Figure 1. The test network consists of a CMTS connected via coaxial cable plant to one or more CMs (the number of CMs varies with the experiment as described in the next section). The upstream traffic is generated by a traffic generator connected to the CM over an Ethernet network. The output of the CMTS is routed over another Ethernet segment to the traffic analyzer. Two traffic

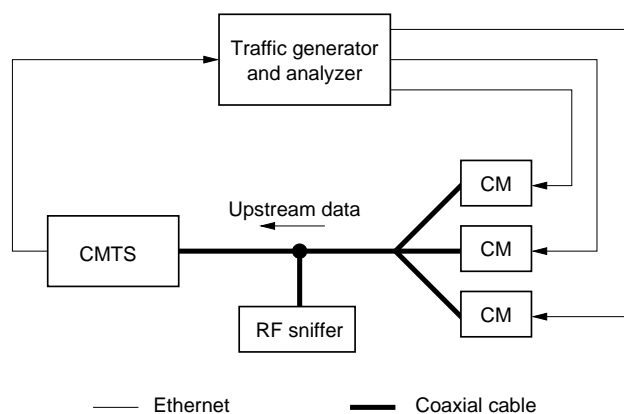


Figure 1. Experimental setup.

generators/analyzers were used in the experiments: Smart-Bits 600 chassis with LAN 3101A cards for CPL-CBR experiments and Adtech AX4000 chassis with 10/100BaseT Ethernet interfaces for DPL-CBR experiments¹.

Sigtek ST-260B DOCSIS 1.0/1.1 RF sniffer/traffic analyzer was used to ascertain that the traffic generated by the CMs adhered to the set parameters. This helped us eliminate several CMs that incorrectly implemented certain aspects of the protocol (e.g., piggybacking) and led to the final decision to limit the study to only CableLabs certified CMs.

Two sets of CMs, each consisting of identical devices, were used in the presented study. One set consisted of two CMs based on Broadcom BCM3300 QAMLink chipsets. The CMs in the second set were based on Texas Instruments TNETC4040 chipsets.

We use the above test setup to find the maximum data rate (throughput) at which the CM can transmit upstream without packet loss². A script was used to set the CM parameters, to control the traffic generator, and to process the results from the analyzer. A binary search algorithm was used to find the maximum throughput of the modem. The search algorithm starts with 0 as the minimum data rate and theoretical maximum data rate as the maximum. It then averages the minimum and maximum to find the current data rate to transmit. If the transmission succeeds then the current is made the minimum and the process continues. If the current data rate transmission fails then the current is made the maximum and the process of averaging the minimum and maximum, followed by transmission and update continues. This process will terminate only when the difference between the maximum and minimum is within the tolerance specified.

The test setup is such that there is no injected physical noise or interference, only a few feet of cable and all

¹Several experiments were conducted on both traffic generators/analyzers to verify that they yield the same results.

²There is small unavoidable packet loss in the experiments with multiple modems. This loss occurs regardless of the offered data traffic rate and is a result of collisions at the moment the traffic streams are started. The script was augmented to ignore this loss.




Figure 2. Experiment 1 – Maximum data rate (single Broadcom-based CM, modulation format 16QAM, piggybacking and concatenation on).




Figure 3. Experiment 1 – Channel utilization (single Broadcom-based CM, modulation format 16QAM, piggybacking and concatenation on).

the devices close to each other. This is almost an ideal condition. In the deployed cable networks the distances are more and there is significantly more noise and interference. However, experimental setup for generating and controlling such an environment was not available for the experiments. We have thus selected only those parameters that are not affected, or minimally affected by distance and interference.

5 Experimental evaluation

This section presents a selection of the results that were obtained in the study. A full set of results can be found at <http://www.cs.unh.edu/cnrg/cgodsay/DOCSIS-study>.

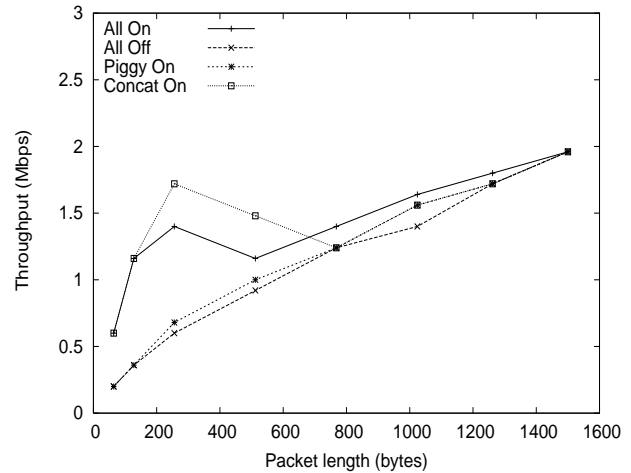


Figure 4. Experiment 2 – Performance enhancers (single Broadcom-based CM, channel rate 5.12 Mbps, modulation format 16QAM).

Experiment 1: Channel rate

The purpose of this experiment was to study the impact of varying channel rates by changing the channel widths and keeping the modulation format the same. Figure 2 shows the throughput performance of a network with one modem. It can be seen that as the channel rate increases the modem throughput increases. Figure 3 displays the results of the same experiment as channel utilization percentages (ratio of observed data rate and channel rate). The channel utilization is well below the theoretical maximum data rate and decreases as the channel rate increases.

Experiment 2: Performance enhancers

Figure 4 shows the results of experiments that evaluate the impact of performance enhancers, piggybacking and concatenation. It can be seen that concatenation significantly improves the throughput for small packet lengths. In our experiments, piggybacking did not have a significant impact on the performance. We have conducted the same experiment for all possible channel rates and obtained similar results.

Experiment 3: Modulation format

This experiment evaluated the impact of the two available modulation formats on the performance (see Figure 5). As outlined earlier, the experiment setup did not allow injection of physical layer impairments that would truly test the benefits of each modulation format. Instead, the experiment concentrated on the protocol-level aspects. The most pronounced difference in performance was observed for channel rate 2.56 Mbps where QPSK clearly outperformed 16QAM. The results were mixed for the remaining channel rates.




Figure 5. Experiment 3 – Modulation format (single Broadcom-based CM, piggybacking and concatenation on).

Experiment 4: CM chipset

Modems based on two different chipsets were available for the experiments (2 based on Broadcom BCM3300 QAM-Link and 3 based on Texas Instruments TNETC4040). Figure 6 compares the results for single modem experiments. The performance is comparable for lower channel rates while for higher channel rates the performance of the TI-based CM drops dramatically when packet length exceeds 780 bytes. This behavior was observed consistently over a wider range of parameters than that shown in the figure. Since we can only observe the external behavior of a CM, we are so far unable to determine the cause of the performance drop.

Experiment 5: Number of modems

This experiment evaluated the impact multiple simultaneously transmitting CMs. Figures 7 and 8 show the performance for modems based on both chipsets. In both cases, the per-CM throughput remains roughly the same even if two modems transmit simultaneously. This is not surprising given the results of the previous experiments where sole CMs were unable to achieve channel utilization better than a fraction of the theoretical maximum data rate. In the case of three modems, the carrying capacity of the network is reached and the per-CM throughput decreases.

Experiment 6: Internet traffic mix

In this experiment, the network was subjected to traffic with packet length distribution that approximates packet length distribution of real Internet traffic [8, 9]. We have used quad-modal distribution that superimposes four Gaussian distributions with parameters shown in Table 1. Table 2 gives the results for the four possible combinations of performance enhancers and both models of CMs. By com-




Figure 6. Experiment 4 – Comparison of CMs based on different chipsets (single CM, channel rate 5.12 Mbps, modulation format 16QAM, piggybacking and concatenation on).

Distribution number	1	2	3	4
Mean in bytes	46	300	576	1500
Half-point width	0.1	40	0.1	0.1
Weight	50%	20%	15%	15%

Table 1. Internet traffic mix parameters.

paring the outcomes of this experiment with the results obtained in Experiment 2 for fixed-size packets, it can be concluded that the packet distribution does not significantly affect the CM performance. The values obtained can be compared to Figure 4 for the Broadcom-based CM and we can conclude that the behavior was approximately equivalent to passing 500-byte packets with CPL-CBR.

6 Conclusions and Future Work

DOCSIS is a complex protocol with complex interactions. This makes it difficult to provide generalized data throughput projections based solely on channel capacity and offered data rate. However, we have analyzed different parameters involved in each layer of DOCSIS, namely PHY and MAC layers across different traffic profiles, with focus on upstream channel utilization and upstream data rate.

We have observed that one CM is unable to utilize all the available bandwidth even in nearly ideal conditions. The throughput of 16QAM was not twice QPSK and as channel rate increased, channel utilization decreased. Concatenation and piggybacking helped significantly for packet sizes below 800 bytes. The comparison between different CM manufacturers becomes more pronounced as the channel rate increases. Running the tests for more number of CMs gave expected results. The per CM data rate remained constant until the carrying capacity of the network was reached after which the per CM data rate dropped




Figure 7. Experiment 5 – Per-modem throughput for one and two Broadcom-based CMs (modulation format 16QAM, piggybacking and concatenation on).




Figure 8. Experiment 5 – Per-modem throughput for one, two, and three TI-based CMs (modulation format 16QAM, piggybacking and concatenation on).

as we added more CMs. Passing distributed packet length traffic did not significantly effect the CM performance.

The main contribution of this project is the experimental evaluation of upstream performance of DOCSIS 1.1 networks. There have been numerous studies on the subject that utilized either analytical models or simulators. This is the first study that attempts to evaluate the performance of real devices.

Our future work will focus on latency experiments and on finding additional measures to compare CM performance. The percentage of MAP opportunities used by a CM as a comparison measure seems promising. Experiments will also be run on different CM and CMTS vendors with more number of CMs on the network.

We are working on a method to predict the upstream performance, to be used by the cable service providers. By using this method the cable service providers can study var-

Piggybacking	off	on	off	on
Concatenation	off	off	on	on
Broadcom-based CM	1.00	1.08	1.48	1.48
TI-based CM	0.84	0.84	0.92	1.00

Table 2. Experiment 6 - Throughput in Mbps under Internet traffic mix (single Broadcom-based CM, channel rate 5.12 Mbps, modulation format 16QAM).

ious traffic parameters, such as packet length distribution, of their networks and be able to predict the number of CMs it can support for a particular load or vice versa.

The experiments will be extended to study the upstream performance of DOCSIS 2.0, to compare the performance improvement provided by DOCSIS 2.0 over DOCSIS 1.1.

References

- [1] Cable Television Laboratories, Inc., “Data-Over-Cable Service Interface Specifications-Radio Frequency Interface Specification,” SP-RFIV1.1-I06-001215, Version 1.1, Dec. 2000.
- [2] N. Golmie, F. Mouveaux, and D. Su, “A Comparison of MAC Protocols for Hybrid Fiber/Coax Networks: IEEE 802.14 vs. MCNS,” Proc. IEEE ICC’99, Jun. 1999.
- [3] Y. D. Lin, C. Y. Huang and W. M. Yin, “Allocation and Scheduling Algorithms for IEEE 802.14 and MCNS in Hybrid Fiber Coaxial Networks,” IEEE Transactions on Broadcasting, vol. 44, no. 4, pp. 427-35, Dec. 1998.
- [4] V. Sdralia, C. Smythe, P. Tzerefos, and S. Cvetkovic, “Performance Characterization of the MCNS DOCSIS 1.0 CATV Protocol with Prioritized First Come First Serve Scheduling,” IEEE Trans. on Broadcasting, vol. 45, no. 2, pp. 196-205, June 1999.
- [5] Sung-Hyun Cho, Jae-Hyun Kim and Sung-Han Park, “Performance Evaluation of the DOCSIS 1.1 MAC Protocol According to the Structure of a MAP Message,” <http://viplab.hanyang.ac.kr/pdf/5-20.pdf>
- [6] F. Abi-Nassif, W. C. Lee, and I. Stavrakakis, “Offered Load Estimation in a Multimedia Cable Network System,” Proc. IEEE ICC’99, June 1999.
- [7] G. Chandrasekaran, M. Hawa and D. W. Petr, “Preliminary Performance Evaluation of QoS in DOCSIS 1.1,” Technical Report, Information and Telecommunication Technology Center, University of Kansas, ITTC-FY2003-TR-22736-01, Jan. 2003.
- [8] S. McCreary and K. C. Claffy, “Trends in Wide Area IP Traffic Patterns,” Tech. Rep., CAIDA, Feb. 2000.
- [9] K. Thompson, G. J. Miller, and R. Wilder, “Wide-Area Internet Traffic Patterns and Characteristics,” IEEE Network, pp. 10-23, Nov. 1997.