Hybrid-Fibre Coax Network Throughput Characterization, with Implications for Network Provisioning

by
Blair Yoshida

A Thesis
Submitted to the Faculty of Graduate Studies
In Partial Fulfilment of the Requirements
For the Degree of:

MASTER OF SCIENCE

Department of Electrical and Computer Engineering
University of Manitoba
Winnipeg, Manitoba
Canada

THE UNIVERSITY OF MANITOBA

FACULTY OF GRADUATE STUDIES

*****

COPYRIGHT PERMISSION

Hybrid-Fibre Coax Network Throughput Characterization, with Implications for Network Provisioning

BY

Blair Yoshida

A Thesis/Practicum submitted to the Faculty of Graduate Studies of The University of

Manitoba in partial fulfillment of the requirement of the degree

Of

MASTER OF SCIENCE

"One in five of US home Net users (22 million) access the Web using a broadband, said NetRatings."

Half US Net access is via broadband
The Register, March 6, 2002
<http://www.theregister.co.uk>

"VoIP: At Last, a Killer App?"

John Barrett
Parks Associates
<http://www.parksassociates.com/>

Acknowledgements

I would like to take this opportunity to thank all those who have followed the progress of this endeavour for their patience. In addition, I would like to thank the staff at TRLabs for the use of their computing resources.

The support of my advisor Dr. R. McLeod, the department of Electrical and Computer Engineering, my family, and my friends in allowing the work to progress at my own pace is greatly appreciated.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

ADSL	Asymmetric Digital Subscriber Line
CBR	Constant Bit Rate
CM	Cable Modem
CMTS	Cable Modem Termination System
CSF13	Common Simulation Framework 13
DOCSIS	Data Over Cable Service Interface Specification
HFC	Hybrid Fibre-Coax
IE	Information Element
ITU	International Telecommunications Union
MAC	Media Access Control
MAP	Bandwidth Allocation Map
MCNS	Multimedia Cable Network System Partners Ltd.
PDU	Protocol Data Unit
p-FCFS	Prioritized First Come First Served
QoS	Quality of Service
SYNC	Time Synchronization
TCP/IP	Transmission Control Protocol / Internet Protocol
UCD	Upstream Channel Descriptor
VBR	Variable Bit Rate
VoIP	Voice Over Internet Protocol

1. Introduction

With the recent widespread acceptance and use of high speed internet services in the form of ADSL over twisted pair copper, and cable modems over coax, an infrastructure is being established which will allow much greater use of high-bandwidth applications. As a result, users may encounter less than optimal performance due to the increased competition over the available bandwidth.

It is currently predicted that levels of data traffic are expected to increase exponentially through the use of applications such as multimedia and video-conferencing [Mcp99]. Currently, traffic on the Internet is doubling about every five to ten months [Mcp99], with the result being that existing networks will be quickly approaching capacity.

With the increased use of broadband connections, services such as voice over IP (VoIP) introduce new ways of communicating which offer the potential to significantly reduce the cost of long distance voice communication but with the cost of increased network traffic. The potential reduction in long-distance voice communications costs alone may drive a change in network data traffic distributions and patterns. Currently, about 10 percent of all voice traffic is classified as VoIP, although less than 1 percent of those calls originate on a VoIP phone [CNET1].

In the case of cable modem networks, the increased traffic levels due to the increased number of subscribers, and the likelihood for increased traffic generated by individual users, have the potential to fully consume the available bandwidth. Increasing the bandwidth available in a cable modem network is non-trivial, and involves reducing the number of users accessing a given segment of the network. As a result, determining the most efficient use of the available resources is highly desirable.

With the introduction of services such as VoIP, which require some level of quality of service (QoS), service providers will likely offer tiered services with basic access and services at one price, and real-time (based on QoS) and other value-added services for a premium. Thus, in addition to determining the most efficient use of available bandwidth, there is the requirement for a solid understanding of the impact of the various services on the performance of the overall network in order to be able to maximize the overall throughput.

Unfortunately, the majority of the deployed cable modems are based upon a version of an industry standard that does not incorporate the provision of QoS [CaLa02]. Therefore, if it becomes necessary to implement a widespread replacement of existing equipment, the impact of new technologies or user behaviours on the network configuration need to be considered.

Sdralia et al. [SdSm99] have simulated early (non-QoS) cable modem environments to predict the upstream system throughput and mean access delay. However, with new user usage patterns, which include VoIP, etc., a number of questions need to be addressed: is there an optimal mix of voice and data traffic sources that can maximize throughput; and do the current standards in use have the flexibility to accommodate potential changes in user behaviour.

Though the scope of this research is limited to cable modem networks, the research is equally applicable to other networking situations, such as wireless, which use similar methods and protocols to control network traffic.

1.1 Report Structure

A discussion of background material will be introduced which will present a description of the operation of a cable modem network, and the protocols in place governing communication on that network.

The discussion will expand to describe the model and the methodology used to simulate the operation of a cable network. Followed by an analysis of the simulation results, with a focus on the impact of voice traffic on network throughput. Conclusions and recommendations will be discussed in closing.

2. The Broadband Network

This chapter will discuss the operation of a broadband network within the framework of a hybrid fibre-coax (HFC) environment using the DOCSIS (Data Over Cable Service Interface Specification) standard which forms cable television's standard configuration in North America.

2.1 The Cable Modem Environment

The cable modem network takes advantage of the existing cabling infrastructure in place for the distribution of cable television signals. Through the addition of a headend controller, located at the signal source or cable system control centre, and individual cable modems (CM), located at the subscribers' premises, a system that was originally designed for unidirectional signal flow can be used for bidirectional communications. Within the cable modem environment the headend controller is referred to as a cable modem termination system (CMTS).

The cable modem network can be described as a highly asymmetric environment where there is a latge amount of bandwidth available in the downstream direction (30 Mbps), i.e. from the provider to the user, and a relatively small amount of bandwidth (<5Mbps) available in the reverse direction. The upstream channel is shared by all users that are connected with a common cable. This common cable forms a local LAN segment, and like any shared resource, the upstream channel may encounter high rates of congestion as the penetration rate (or take rate) of subscribers goes up and/or the amount of traffic generated by the individual users increases.

Physically, the cable network uses a shared-medium, tree-and-branch architecture that carries an analog transmission as shown in Figure 1. The cable modem network traditionally has consisted of coaxial cable connecting a head-end to the individual consumers.

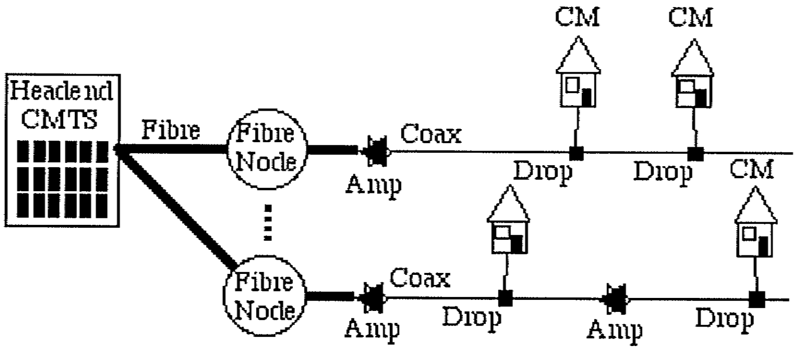

Figure 1 - Typical HFC Network

One of the reasons for the low upstream data rates is this tree-like architecture. Due to attenuation inherent in the cabling, there is a need to amplify signals in the network. The information travelling from the user to the headend needs to be combined and then amplified in order for it to be successfully received at the headend. The amplifiers for the return path back to the headend are receiving and amplifying the cumulative effect of all the noise on the individual branches of the network. The resulting effect is that the noise degrades the quality of the return signal and thus limits the data rates that can be achieved.

Due to the large distances traversed by the network (which can amount to tens of kilometres), and the resulting associated signal attenuation, the main coaxial trunks are normally replaced with high-reliability, low attenuation fibre optic cable. These trunks extend well into a community and are terminated at what is known as a fibre node. The fibre node provides a bridge between the fibre optic cable and the traditional coaxial cable, hence the hybrid fibre-coax nomenclature. As an added benefit, the use of HFC also increases the useable frequency range that can be transmitted.

From the fibre nodes, coaxial cable segments are run into the neighbourhoods and provide service to the individual subscribers' homes. Typically, from each fibre-node, coaxial cable is used to connect 500-2000 subscribers [Az97]. At the individual subscriber's location, a cable modem provides a bridge between the ethernet port of the subscriber's computer(s) and the HFC network. Through the cable modem, the local computer can make requests and be connected to machines to which the cable service provider is connected, i.e., the greater Internet.

The problem of bandwidth usage arises from the fact that all the houses connected to a common segment of coaxial cable share, and need to contest for access to, the available upstream bandwidth.

2.1.1 RF Spectrum

As was previously mentioned, the cable modem network uses the pre-existing infrastructure in place for the distribution of cable television signals. As a result, much of the potential signal spectrum is already occupied for the broadcast of television signals. In order to add the capability of carrying data signals, it required that the existing uses be taken into account and that data signalling not interfere with the pre-existing video signals.

In North America, the spectrum between 50 MHz and 550 MHz is reserved for NTSC (National Television System Committee) analogue television broadcasts where each television channel uses a 6 MHz band within this range. The NTSC standard forms the basis upon which television signals in North America are transmitted and displayed.

Prior to the introduction of peer-to-peer services, Internet use was primarily a "client server" environment where the individual user, acting as the client, would make a request for a large amount of data from a server. In this context, the amount of network traffic that would be supplied from the user can be considered only a small percentage of the total transaction. The traffic that would be directed to the user would make up the bulk of the network traffic. As a result, with network traffic highly asymmetric in nature the spectrum, which was allocated within the cable system for data communication, was divided into separate upstream and downstream frequency bands as shown in Figure 2. To deliver DOCSIS data services over a cable television network, one 6 MHz radio frequency (RF) channel in the 50 - 860 MHz spectrum range is typically allocated for downstream traffic to homes and another 6 MHz channel in the 40Mhz band (5 - 42 MHz) is used to carry upstream signals [DOCRFI1.1]. In this context the term channel is referring to a unit of the available spectrum which would normally be reserved for the broadcast of a television channel but is being used for data transmission.

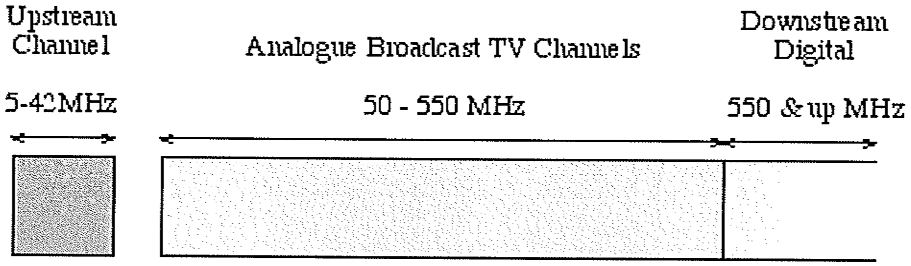

Figure 2 - Spectrum Allocation

The choice of the 5 - 42 MHz upstream band was a result of the fact that the bandwidth, being a valuable resource, in this particular frequency range is not normally occupied by any pre-existing video channels. This frequency range is not used for television signals due to problems associated with ingress noise and therefore was unsuitable for clear television transmission. Unfortunately the potential for ingress noise along with cascading amplifiers in the upstream direction resulting in noise funnelling which impacts the maximum data rate achievable for communication on the upstream channel. Due to the perceived low bandwidth requirements of the upstream channel in the original specification, this was not considered to be a problem.

2.1.2 The DOCSIS Specification

The DOCSIS standard is a cable modem system specification that was created by a limited partnership between a number of cable operators, which included TCI/AT&T, Time Warner, et.al., called the Multimedia Cable Network System Partners Ltd. (MCNS). DOCSIS was developed as an industry driven alternative to the **IEEE802.14 Cable TV Media Access Control and Physical Protocol** and provides a reference specification to facilitate interoperability amongst cable modems and cable modem termination systems. The 12 specifications which makeup DOCSIS 1.0 were first published in 1997. In March 1998, the International Telecommunications Union (ITU) accepted DOCSIS as a cable modem standard, called ITU J.112 [ITUW].

The DOCSIS 1.0 standard provides a scheme for transporting IP-based data packets over an HFC cable network. It provides a solution for supporting non-deterministic network behaviour such as web browsing, e-mail, etc.

In 1999 the cable television industry, through Cable Television Laboratories, Inc. (CableLabs), a nonprofit research and development consortium of cable television system operators, extended DOCSIS by issuing a second-generation specification called DOCSIS 1.1. The DOCSIS 1.1 specification, primarily aims at enhancing the limited QoS functionality of a DOCSIS 1.0 based cable access system. New media access control (MAC) messages have been defined for dynamic QoS signalling, and several new QoS parameter encodings have been defined in the existing MAC messages. A DOCSIS 1.1 CMTS can better support the requirements of delay/jitter sensitive traffic on a DOCSIS 1.1 CM. These improvements add support for IP telephony and other constant-bit-rate services and provide the bandwidth and latency guarantees required to offer tollquality voice, dedicated business-class data services and multimedia applications [CDN1].

It should be noted that, as of March 2002, most cable modems are only DOCSIS 1.0 compliant, and therefore not able to take advantage of the QoS provisions of DOCSIS 1.1. At this point in time, there are only 21 DOCSIS 1.1 compliant modems versus the 203 DOCSIS 1.0 compliant modems that have been certified by CableLabs [CaLa02], i.e. approximately 10%.

Ideally, from an economic point of view, the goal for a cable system operator is to oversubscribe, queue user requests, and fully utilize all the available bandwidth and infrastructure such as existing CMs.

2.1.3 Media Access

The operation of a DOCSIS compliant HFC network can be broken down into a number of different components. For some network operations the behaviour is distributed and for others the operation is extremely centralized. However, in general terms, the system operates by having the individual cable modems make requests for service from the headend control equipment (i.e. the CMTS), with the CMTS allocating the available bandwidth on both the upstream and downstream directions.

As a result of the division of the upstream and downstream spectrums, the CMs have no direct way to communicate with other CMs or the CMTS. CMs can only communicate with other CMs through the CMTS, and must contend for the CMTS' attention via the channel's contention mechanism.

Access to the downstream channel is completely under the control of the CMTS. The CMTS has direct control over the bandwidth that is allocated to the individual CMs. The MAC sublayer defines a single transmitter for each downstream channel (the CMTS) with all the CMs listening to all frames transmitted on the downstream channel upon which they are registered. Each individual CM accepts those frames which are addressed to itself.

The upstream channel is characterized by many transmitters (the CMs) and only one receiver (the CMTS). The upstream channel can be characterized of as a stream of "mini-slots," where a mini-slot is the unit of granularity for upstream transmission opportunities. Time in the upstream channel is slotted, providing for time division multiple access (TDMA) at regulated time "ticks" of 6.25 micro seconds. The relationship between the time ticks and the mini0-slots is shown in Table 1. The CMTS provides the time reference and controls the allowed usage for each mini-slot within the system.

Parameter	Value
Time Tick	6.25 microseconds
Ticks per Mini-slot	4
Microseconds per Mini-slot	25
Bytes per Mini-slot	16 (nominal assuming QPSK)
Symbols per Byte	4 (assuming QPSK)
Mini-slots per Second	40,000
Symbols per Second	2,560,000

Table 1 - Example of Relationship Between Ticks, Mini-Slots, etc.

The data format on the physical layer includes a variable-length modulated burst with precise timing beginning at boundaries spaced at integer multiples of 6.25 μsec apart (which is 16 symbols at the highest data rate[1]). Each burst supports a flexible modulation, symbol rate, preamble, randomization of the payload, and programmable FEC encoding. All of the upstream transmission parameters associated with burst transmission outputs from the CM are configurable by the CMTS.

The CMTS controls assignments on the upstream channel through the use of a global

---

[1] The upstream modulator is capable of both QPSK and 16QAM. Also both modulation schemes offer numerous symbol rates.

broadcast, on the downstream channel, of a bandwidth allocation map (MAP). The MAP mechanism provides a dynamic mix of contention- and reservation-based upstream transmit opportunities. The MAP operates by transmitting a message in which a predetermined mini-slot usage scheme is distributed to all the CMs. The MAP message can assign usage of a number of mini-slots to individual modems, and reserve others for bandwidth requests from the CMs using a contention mode scheme.

Allocation algorithms may vary the size of the MAPs over time to provide a balance of network utilization and latency under varying traffic loads. Neither the actual implementation, nor a description, of the allocation/scheduling algorithm is part of the DOCSIS specification. The implementation is left to the various hardware manufacturers in order to allow them to provide product differentiation.

2.1.3.1 Bandwidth Allocation Map (MAP)

The MAP is a MAC management message transmitted by the CMTS on the downstream channel. The MAP describes, for some time interval, how the upstream mini-slots must be utilized. The mini-slots within the MAP-defined time intervals may be used as follows: granted to individual CMs for transmissions; used by all CMs to contend for upstream bandwidth (Request/Contention); and others as an opportunity to perform maintenance of the link. The relationship of the MAP to the usage of the upstream channel is illustrated in Figure 3.

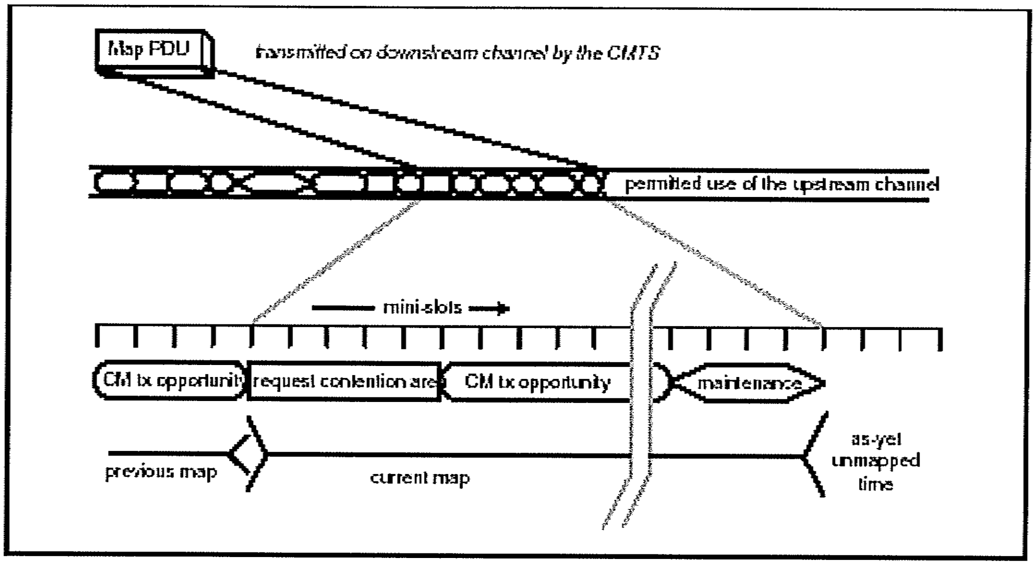

Figure 3 - Bandwidth allocation map (MAP). [DOCRFI1.1]

During the contention time slots, CMs attempt to place requests to the CMTS for the assignment of future time slots in which they may transmit their data.

Link maintenance provides opportunities for new stations to join the network. This allows the new stations to perform initial ranging to determine the maximum round-trip propagation delay and perform any necessary output power adjustments.

The MAP itself consists of a fixed-length header followed by a variable number of information elements (IEs) where each IE defines the allowed usage for a range of mini-slots. All CM stations must scan all the IEs, to determine how they may interact with the upstream channel.

The number of mini-slots described, and the usage of those mini-slots, may vary from MAP to MAP. At minimum, a MAP may describe a single mini-slot. This would be wasteful in both downstream bandwidth and in processing time within the CMs. At maximum, the time described in a MAP may stretch to tens of milliseconds. Such a MAP would provide poor upstream latency. Allocation algorithms may vary the size of the MAPs over time to provide a balance of network utilization and latency under varying traffic loads. At minimum, a MAP must contain two information elements (IEs): one to describe a time interval on the upstream channel, and a null IE to terminate the list. At a maximum, a MAP must be bounded by a limit of 240 IEs, where each IE is a 32-bit block.

The individual IEs can take the form of one of the following types:

- Request IE - which provides an upstream interval in which requests may be made for bandwidth for upstream data transmission. These may be broadcast, which provides an invitation for CMs to contend for access. Alternatively they may be unicast, which provides an invitation for a particular CM to request bandwidth.

- Short and Long Data Grant IEs - provides an opportunity for a CM to transmit one or more upstream PDUs.

- Request/Data IE - which provides an upstream interval in which requests for bandwidth or short data packets may be transmitted and is a variation of the Request IE.

- Initial Maintenance IE - provides an interval in which new stations may join the network.
- Station Maintenance IE - provides an interval in which stations are expected to perform some aspect of routine network maintenance, such as ranging or power adjustment.
- Data Acknowledge IE - acknowledges that a data PDU was received.
- Expansion IE provides for extensibility,
- Null IE terminates all actual allocations in the IE list.

During periods defined by the Request and Request/Data IEs, the data transmitted by the CMs is at risk of encountering collisions that must be detected by the CMs. Collision detection is performed by parsing subsequent MAPs for an acknowledgment from CMTS for transmission by the CM. The lack of a return acknowledgement would imply a collision occured and that the CMTS did not receive the request during the previous contention period. The operation of the system is in some ways similar to the operation of an IEEE802.11 wireless ethernet network, in effect using a CSMA/CA scheme.

The method of contention resolution is based on a truncated binary exponential back-off scheme, with the initial back-off window and the maximum back-off window controlled by the CMTS. The values are specified as part of the MAP MAC message and represents a power-of-two value.

When a CM has information to send, and wants to enter the contention resolution process, it sets its internal back-off window equal to the Data Backoff Start defined in the MAP currently in effect. The CM randomly selects a number within its back-off window. This random value indicates the number of contention transmit opportunities that the CM defers before transmitting. A CM only considers contention transmit opportunities for which this transmission would have been eligible. These are defined by either Request IEs or Request/Data IEs in the MAP. After a contention transmission, the CM waits for a Data Grant, or Data Acknowledge, in a subsequent MAP. Once either is received, the contention resolution is complete. The CM determines that the contention transmission was lost when it cannot find a subsequent MAP with a Data Grant or Data Acknowledge for it and with an Ack time more recent than the time of transmission.

As previously indicated, the upstream and the downstream communication channels occupy different portions of the radio frequency spectrum. As a result, access to the two channels is handled by two separate mechanisms, however in both cases the CMTS manages channel access. It is the function of the CMTS to manage the available bandwidth, and sort the incoming requests in order to provide fair access to the requesting CMs through the use of appropriate scheduling techniques. As a result, all of the subscribers' cable modems on a particular LAN segment will be competing to get the attention of the CMTS in order to obtain access to the channel.

DOCSIS, like IEEE 802.14, leaves the implementation of any algorithms for computing bandwidth allocation and scheduling to the individual CMTS vendors. The single algorithm implemented within the OPNET common simulation framework model is a prioritized, first come first serve scheme (p-FCFS).

2.1.4 Quality of Service

The introduction of QoS in DOCSIS 1.1 is accomplished through the use of a policy type mechanism referred to as a service flow. These pre-configured service flows provide a mechanism for upstream and downstream QoS management and are central to the operation of the MAC protocol.

The basic model is that the classifiers within the CM, the CMTS, or both, associate packets into exactly one Service Flow. The Service Flow encodings provide the QoS parameters for treatment of those packets on the RF interface.

A service flow defines a particular unidirectional mapping between a CM and the CMTS which includes some associated QoS. The associated QoS is defined by a number of parameters such as latency, jitter, and throughput assurances.

The CM and CMTS provide this QoS by shaping, policing, and prioritizing traffic according to some predefined QoS parameter set. The CMTS can assign one or more service flows to each CM corresponding to the level of service required by the individual CMs. This mapping can indicate the characteristics of the service that will be provided via these service flow, each of which can be used to uniquely specify, for example, CBR and VBR service flows. This mapping can be negotiated between the CMTS and the CM during CM registration or via dynamic service establishment.

In a basic CM implementation, two service flows (one upstream, one downstream) are used. However, service flows allow for the development of multiple increasingly structured service classes tied to specific traffic types while supporting interoperability with more basic modems. It is possible that certain service flows can be configured to have a maximum packet size or be restricted to small fixed size. As a result it enables the separation of differentiated traffic. For the purpose of this study, multiple service flows are used to provide separate flows for voice traffic and data traffic.

Within the DOCSIS specification, when a unique service flow ID is being *used,* there is an associated service ID (SID) assigned. Upstream bandwidth is allocated to SIDs, and hence to the CMs, by the CMTS. The SIDs provide the mechanism by which upstream QoS is implemented.

DOCSIS defines six QoS services: unsolicited grant service (UGS); unsolicited grant service with activity detection (UGS-AD); real-time polling service (rtPS); nonreal-time polling service (nrtPS); best effort (BE) service; and committed information rate (CIR) service [Yin98]. Each service can be tailored to a specific type of data flow with UGS as an example of an upstream flow scheduling service type that can be used for mapping constant bit rate (CBR) traffic onto the service flows.

For the purposes of this study, the focus will be on the mix of voice and data sources. As a result, only two service flow types will be used within the simulations performed and are described in further detail below.

2.1.4.1 Unsolicited Grant Service (UGS)

In this study, UGS service is provided to those modems generating voice traffic. The UGS is designed to support real-time service flows that generate fixed-size data packets on a periodic basis, such as Voice over IP. This service, offered by the CMTS, provides fixed-size grants on a real-time periodic basis to a service flow. This eliminates the overhead and latency of a CM needing to make requests via the contention mechanism and assures that grants will be available to meet the flow's real-time needs.

2.1.4.2 Best Effort Service

The Best Effort (BE) service provides efficient service to best effort traffic. This service can be assigned when a CM makes a bandwidth request via the contention request opportunities. However, unlike UGS service, there is no periodicity to the issuing of BE grants, i.e. grants will be issued via the MAP only if bandwidth is available. BE service was provided to those modems in the study generating data traffic.

3. Analysis/Methodology

The approach taken to analyze the network performance was via modelling the network in software, and simulating the operation of the various components under varying scenarios. Data was collected at numerous points in the network for every scenario. Once all scenarios were simulated, the individual parameters were collated to reflect the impact of various operating conditions and are available in the appendices.

3.1 OPNET

The analysis of the HFC network was performed by modelling, and then simulating the operation of an HFC network using OPNET Modeler. OPNET Modeler is a software environment produced by OPNET Technologies/MIL3 that allows the design and study of communication networks, devices, protocols, and applications.

The model used to simulate the HFC network was the DOCSIS Common Simulation Framework 13 (CSF13). The CSF13 model was developed for CableLabs and is not a standard OPNET simulation model. The CSF13 model is based upon the initial model, called the Common Simulation Framework (CSF), which was developed by MIL3 to allow different companies investing in cable modem MAC protocol development to have a common ground for comparing one protocol implementation to another [OpDoc]. The simulations were performed using the CSF13 model with OPNET Modeler release 7.0.B PL8 for Solaris.

The use of the CSF13 model provided a means to analyze and develop specific protocols in the OPNET development environment using pre-existing traffic sources and sinks for communicating data, video, and voice across the cable network. Having identical traffic characteristics generated over different MAC implementations provided an easy mechanism to observe how certain implementations, or portions of implementations, would perform in relation to others.

Though the CSF13 is the latest version of the OPNET DOCSIS model, it does not completely model the DOCSIS1.1 environment. There are a number of features, such as the full OSI protocol stack (i.e. TCP/IP), which are missing from the model. However, the basics of:

- Upstream and downstream communication;
- Packet size conformance;
- Upstream bandwidth allocation;
- Best effort, real-time polling, and unsolicited grant;
- Ranging [2];

---

[2] Ranging is performed to determine the distance between the CM and the CMTS. Since all data to the CMTS must arrive during predefined time slots, the distance needs to be known in order to calculate propagation delays.

- A restricted QoS model in which QoS has been implemented to a limited extent in the current model. That is, each CM is capable of supporting a single QoS and priority.

are provided and allow for a meaningful analysis to be performed. The upstream bandwidth allocation and the QoS features represent the primary elements of interest for the purposes of this study.

The limited nature of the CSF13 model was due to the fact that the intention behind the model was to provide an environment where bandwidth allocation on the upstream channel can be observed [OpDoc]. Since bandwidth allocation algorithms do not form part of the DOCSIS specification and are left to manufacturers to implement, only a single scheduling algorithm for use by the CMTS is provided [OpDoc] as part of the model. Programming interfaces are provided for the manufacturers using the framework to provide their own scheduling code.

3.1.1 Network Model

The network model used in this study consists of a number of elements as shown in Figure 4. The model consists of a CMTS (in the centre) which is connected to a variable number of CMs (circles to the centre right) via a downstream bus (top, heavy line) representing the downstream channel, and an upstream bus representing the upstream channel (bottom, heavy line). Also connected to the CMTS are: a statistic collection unit, which gathers the various parameters of interest during the simulation and stores them to a file; and a network cloud that can source or sink data (though for the purposes of these experiments was not used).

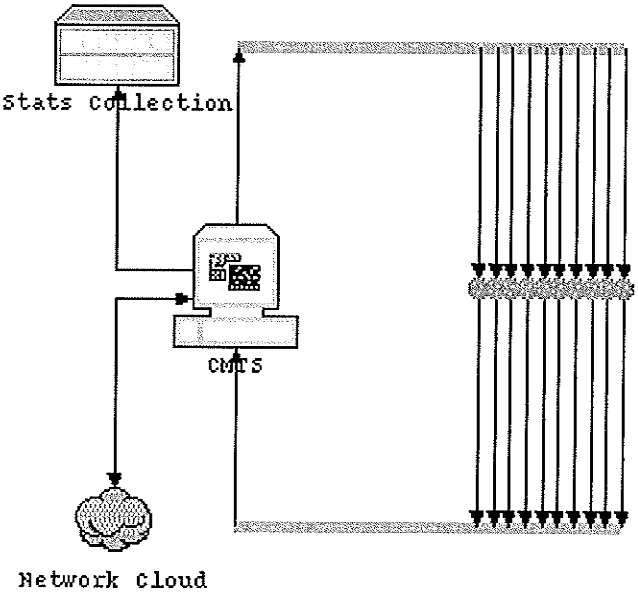

Figure 4 - Network Model

3.1.1.1 Cable Modem Termination System

The main functions of the CMTS is to allocate bandwidth on the upstream channel and provide appropriate priority queuing on the downstream channel. The downstream channels function as point to point connections and the upstream channel(s) are shared.

The CMTS model takes traffic sent on the upstream channel from a CM and determines whether that traffic is a request for bandwidth or application related traffic. Under the condition it is a request for bandwidth, the model is responsible for invoking the appropriate scheduling algorithms for allocating the requested bandwidth. The default, and only implemented, scheduling algorithm provided in the CSF 13 model is a p-FCFS scheme (§2.1.3 Media Access.) If the information is application traffic, the model will forward the traffic to the higher layer module.

If the information the CMTS receives is from a higher layer module, the model will process the information and schedule it to be sent to the destination CM. At the appropriate time, the information will be sent to the individual CM on the downstream channel.

3.1.1.2 Cable Modem

The CM embodies the client side of the DOCSIS protocol and maintains a single queue for any attached device(s) (eg. a local PC.) The role of the CM is to provide customer connectivity to the HFC (hybrid fiber-coaxial) network. Within the simulation environment, its principal function is to transmit packets from the subscribers' locations to the headend across the cable network on the upstream channel. In the OpNET model the CM generates traffic to be sent on the upstream channel and receives traffic generated from the CMTS on the downstream channel. It receives MAC management messages such as UCDs, SYNCs, MAPs, etc. The CM model included in the CSF13 framework can simulate both data and voice communications.

In terms of data sources, there are a number of types available including: a voice application traffic source; two customizable application traffic sources; a constant bit rate source; and an MPEG application traffic source. However to maintain a reasonable scope to the project, only the first two types of sources were used.

3.1.1.2.1 Data Sources

The data source model used in CSF13 is the simplest of traffic sources. It is represented mainly by a single state FSM that generates application packets based on an exponentially distributed inter-arrival time and a packet size distribution specified via external probability density functions. Figure 5 provides an example of the probability distribution that governs the packet size generated by the data sources.

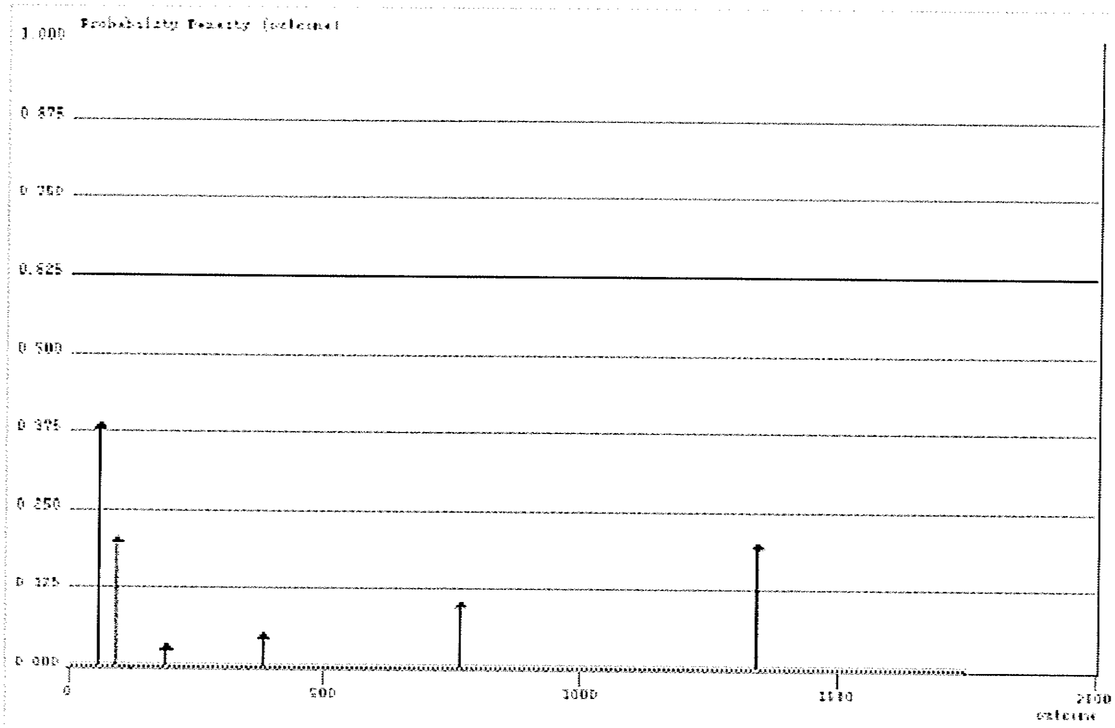

Figure 5 - Packet Size Probability Distribution Function

3.1.1.2.2 Voice Sources

The voice source model controls the aspects of a voice session. This mainly consists of setting up the voice session and alternating between high and low speech activity levels (talk and silence). Activity levels are sent to an encoder for modulation of the speech into voice samples.

The process model FSM invokes a calling process that will perform the call setup to the destination. In the model the call setup is not explicitly implemented, but rather a delay is modelled to represent the length of time required to establish the connection. Probability distribution functions are used to characterize the different speech activity levels. A simple on-off type process model is used.

3.1.1.3 Data Busses

The interconnections of the CMTS and the CMs are accomplished using the bus link model supported in OPNET. It provides a mechanism to model transmissions on a bus-based shared medium.

3.1.1.4 Stats Collection

The function of the stats collection node is to receive application traffic forwarded from the CMTS, parse the traffic based on its type, and send it to the appropriate sink module. Each sink module is responsible for recording related statistics, destroying the packet, and freeing up memory.

3.1.1.5 Network Cloud

The network cloud, though shown in the network model (Figure 4), plays no role in the current model. Normally it would offer a means to communication to/from the cable network from/to an external network.

3.2 Test Scenarios

Initially, the throughput of the upstream channel was simulated under "as-intended" conditions, i.e. assuming that the users are operating in an asymmetric environment where the upstream traffic is merely generating requests for data to be sent on the downstream channel. Under these conditions, network simulations were performed to determine the number of CMs required to fully saturate the capacity of the upstream channel.

In addition to determining the number of modems which would be involved in each simulation run it was necessary to determine the length of time which would need to be simulated in order to obtain reasonable output. Since the CSF model was only designed to operate under steady state conditions [OpDoc], it was necessary to determine the point at which steady state operation was attained.

Figure 6 provides a representative sample of the output generated by a single simulation run. As can be seen on both of the traces, data generated within approximately the first 30 seconds of simulation appear noticeably different in character to the remainder of the output. In addition, after this initial period the output data, though not constant, remains bounded within a relatively small range of values.

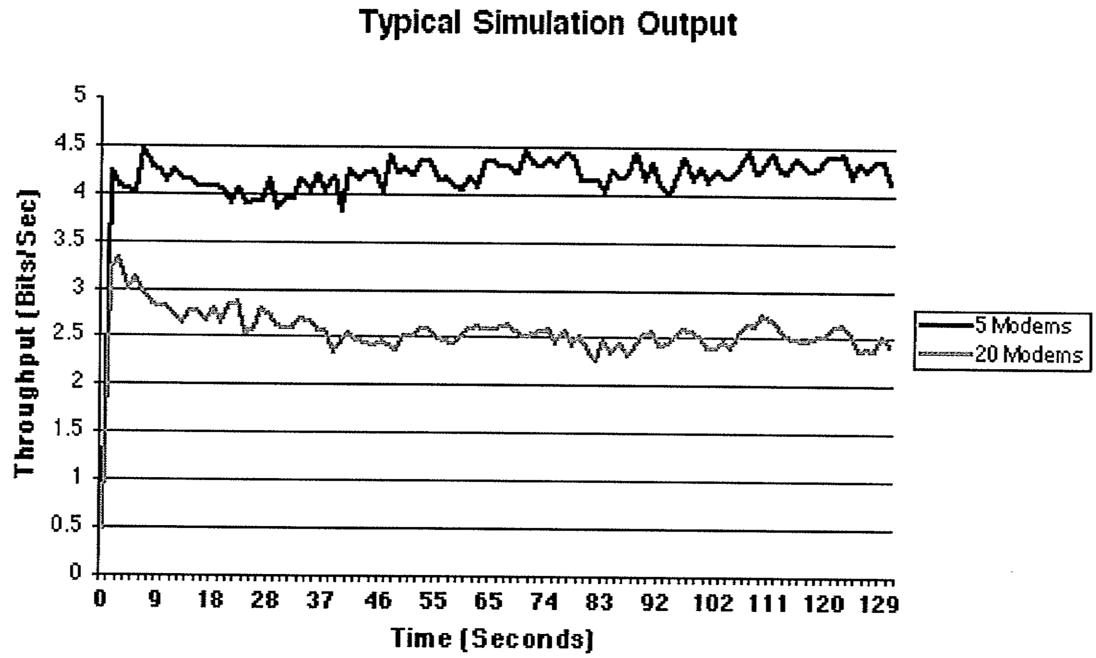

Figure 6 - Typical Output From Simulation

As a result, it was decided to limit the length of simulated network time to a period of 130 seconds, which would provide approximately 100 seconds of useable data. The primary reason for limiting the simulation time was in order to keep the amount of CPU time required for each simulation to a reasonable level as there was approximately a 15:1 ratio of CPU time to simulation time.

In order to fully assess the effect of mixing voice and data traffic, a large number of scenarios were simulated. These simulations included scenarios involving the following:

a) Data sources only - the number of sources was increased linearly to a maximum number which was determined when the upstream channel reached saturation.

b) Voice sources only - the number of sources was increased linearly to a maximum number which was determined when the channel reached saturation.

c) A combination of data and voice sources, increasing linearly the numbers of each type, up to the maximum as determined in parts a) and b).

In total, 169 different scenarios were simulated.

3.2.1 Measured Parameters

A select number of data parameters were collected which could provide an indication of the operation and performance of the system. Overall throughput, the amount of time provided to all the modems for transmission, and amount of time available for BE service to gain access to the upstream channel were selected to provide an overview of the upstream channel. The following three parameters were selected to explore these characteristics:

- The throughput on the upstream bus;
- The size of the bandwidth allocation map;
- The amount of time devoted to contention.

A full list of available parameters is provided in Appendix D

4. Results

4.1 Throughput

As the upstream channel occupies a standard 6MHz television channel, where the analog RF is 4.2MHz vestigial sideband (+ 25KHz FM aural carrier for a channel of 6 Mhz), it should theoretical provide a throughput of 4.2 Mbps.

The initial scenarios utilizing only non-voice data sources generated sufficient traffic with 60 data modems to fully utilize the available upstream bandwidth. When compared to the typical 500-2000 subscribers on a LAN segment, this would imply that only between 3%-12% of all subscribers may be active at any given time. However, once the voice sources were introduced the impact on upstream throughput was immediately apparent by the reduction in the number of users required to consume the available bandwidth.

The throughput results, as shown in Figure 7 and Appendix A, clearly illustrate that the inclusion of voice sources severely reduces the maximum throughput of the network. The maximum throughput when using only data sources yields a maximum throughput approximately equal to the theoretical maximum of 4.2 Mbps. However, once voice sources are introduced, the maximum throughput begins to drop, with the worst case (50+ data sources and 15 voice sources) only achieving approximately 50% of the theoretical maximum. A scenario of voice sources alone resulted in a maximum throughput of 2/3 of the theoretical maximum. An important observation is that for certain combinations of voice and data sources a throughput greater than the worst case values can be achieved.

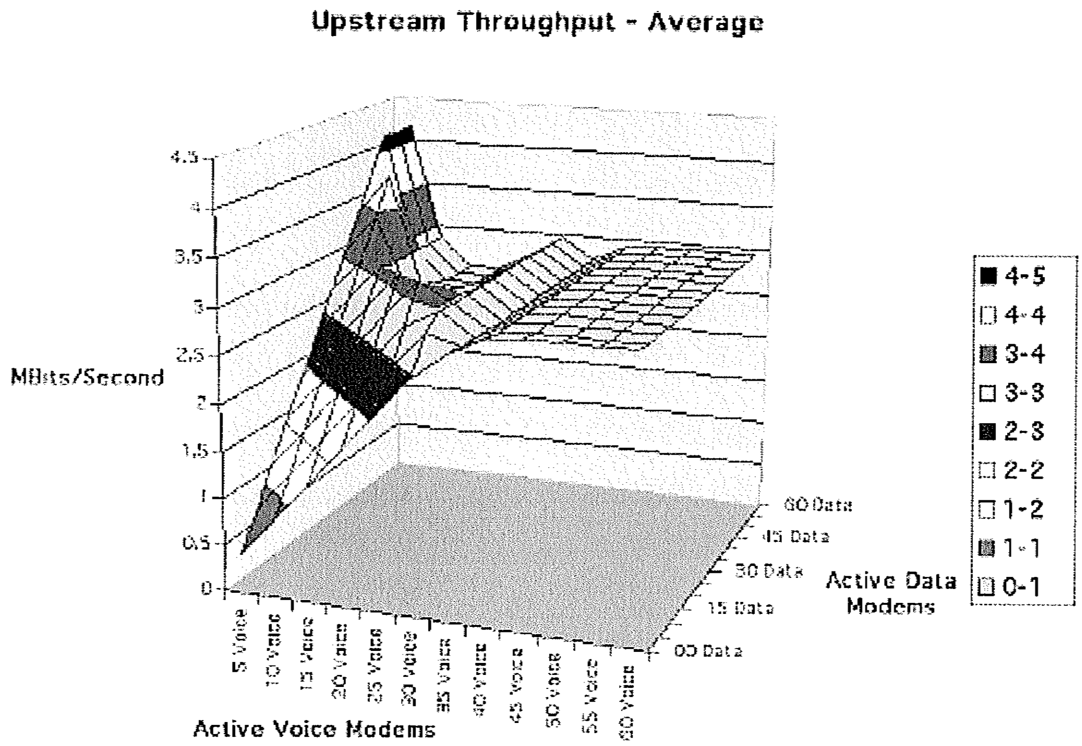

Figure 7 - Upstream Throughput for Voice & Data Mix (3-D)

Figures 8 and 9 show the same data as Figure ? but as viewed from above, and emphasizes the range of voice and data source combinations which will yield the greatest throughput. Figure 9 highlights the specific range of maximum throughput.

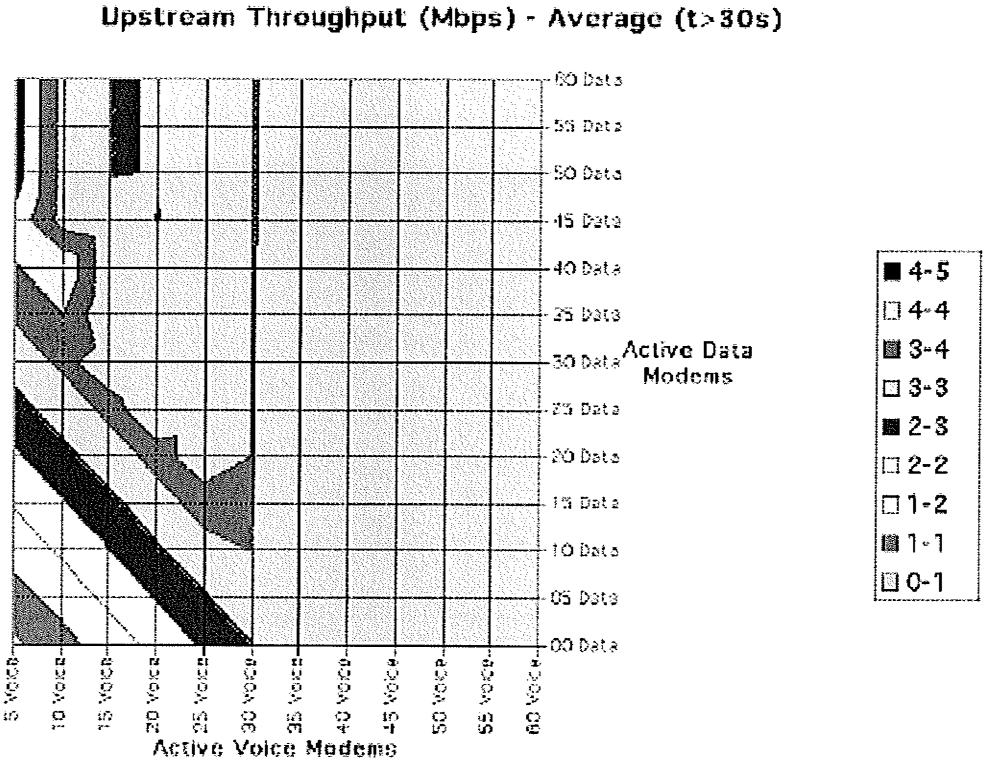

Figure 8 - Upstream Throughput Contour Map (2-D)

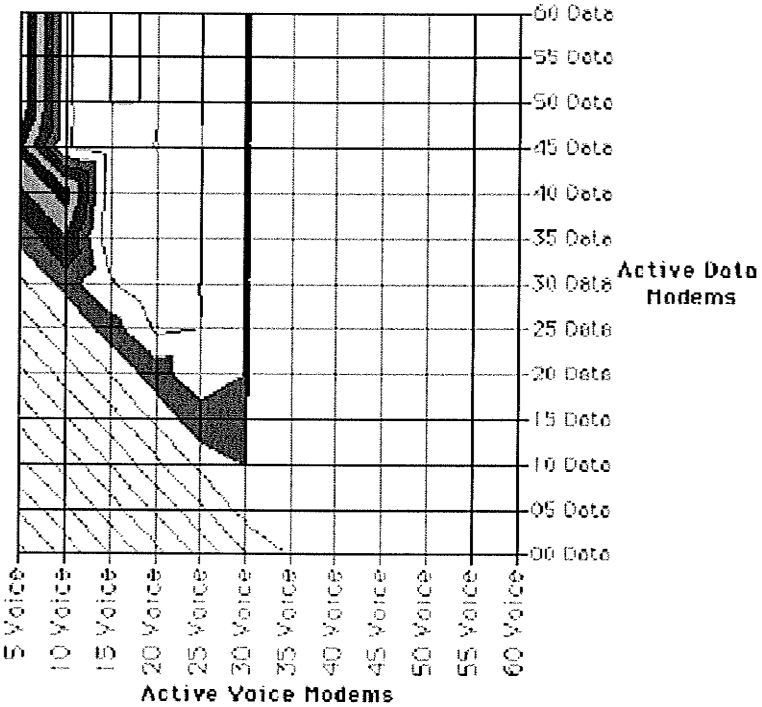

Figure 9 - Region of Maximum Throughput (2-D)

These results would indicate that it is necessary to closely monitor the number of VoIP users on a particular LAN segment. If traffic levels are approaching the channel capacity, it may be necessary to limit the number of voice users in order to maintain throughput at levels at or near the full potential of the upstream channel.

4.2 Bandwidth Allocation MAP and Contention Slot Time

In order to further investigate the drop in throughput due to the inclusion of voice sources, the characteristics of the MAP were isolated. As the MAP dictates how the upstream resources are utilized, the MAP was examined in terms of the number of IEs which it contained, and the amount of time dedicated to contention on the channel (Appendix B and Appenix C.)

The number of IEs within each MAP should provide an indication of number of grants being provided by the CMTS to the CMs, therefore providing an indication of the number of users which were able to access the upstream channel. Unfortunately this did not turn out to be the case. It turned out that there was not a direct correlation between the IEs in the MAP and the data throughput on the upstream channel.

Figures 10 and 11 illustrate how the number of IEs contained within the MAP changed with respect to the number of active sources. Correlating this data to the the throughput data of Figure 7, it is interesting to note the dramatic change in the number of IEs when upstream channel is at capacity and there is a mix of voice and data sources. This can clearly be seen in Figure 11. In addition, the number of voice sources has a much greater impact than the number of data sources. This can be seen by comparing the gradient along the voice modem axis relative to the data modem axis.

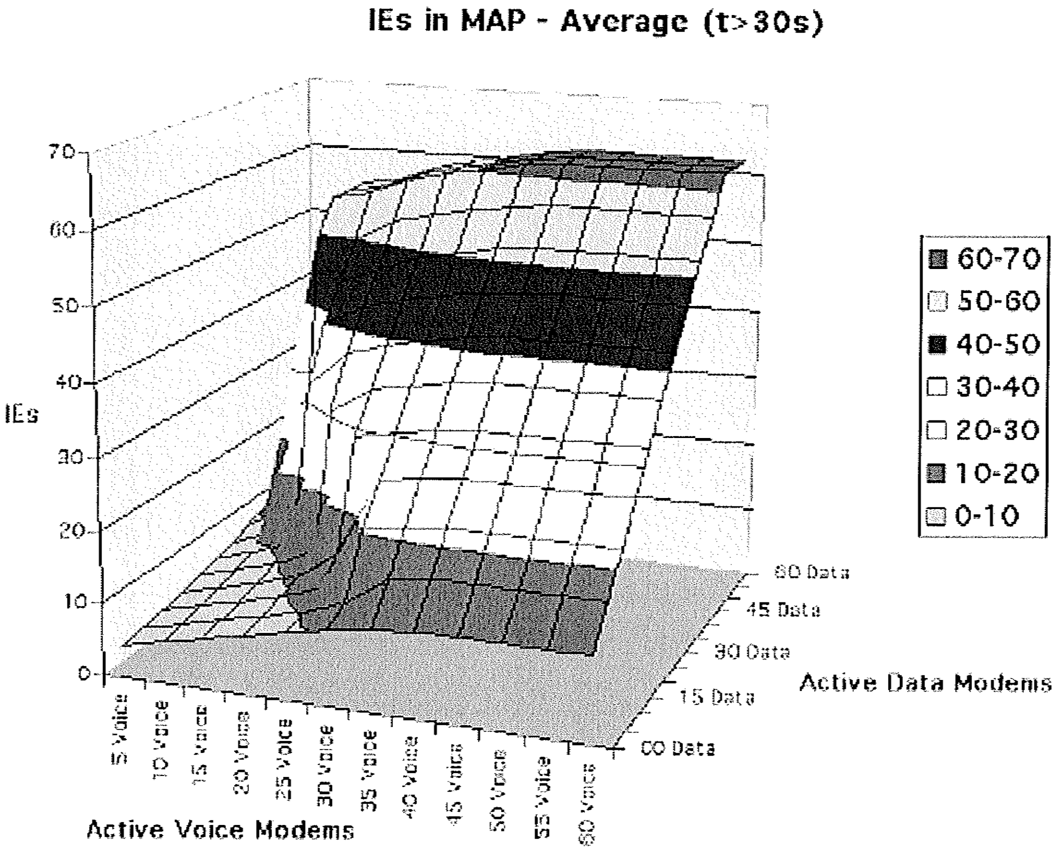

Figure 10 - Average Number of IEs in MAP

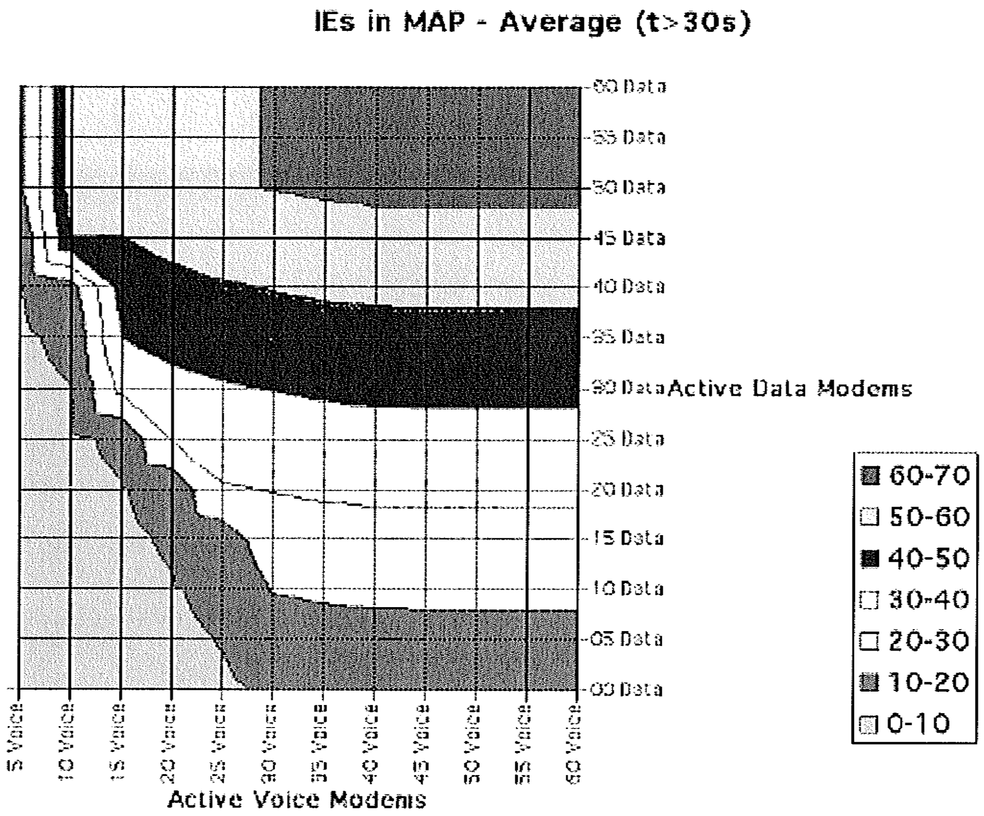

Figure 11 - Average Number IEs in MAP

As can be seen, as the number of modems increases, regardless of the modem type, the number of IEs within the MAP increases up to a maximum of over 60. By looking at the effect of the number voice modems, regardless of the number of data sources, once the number of voice sources reaches approximately 35 the number of IEs stops increasing. This correlates with the plot in Figure 7 where the upstream throughput stops increasing at approximately the same point. However, the number of IEs in the MAP continues to increases with the total number of sources (both voice and data) even though throughput does not increase. The number of IEs reaches maximum values when there are 35 or more voice modems and 45 or more data modems. This would indicate that the additional modems have some impact on the allocation of the upstream channel, but are not able to put data on the channel and contribute to the upstream throughput.

The exact nature of the MAP during the period when there was an increase in IEs without a corresponding increase in upstream throughput was not determined. However, the lack of any correlation between IEs and throughput would imply that the additional IEs were not likely to be Data Grant IEs. It is possible that the additional IEs are Data Grants of zero length implying a grant pending condition. This would occur if a CM successfully used the contention time to make a request but the CMTS was unable to allocate any resources on the upstream channel.

As part of the MAP, contention time is allocated. Figures 12 and 13 illustrates the impact on average contention time within each MAP.

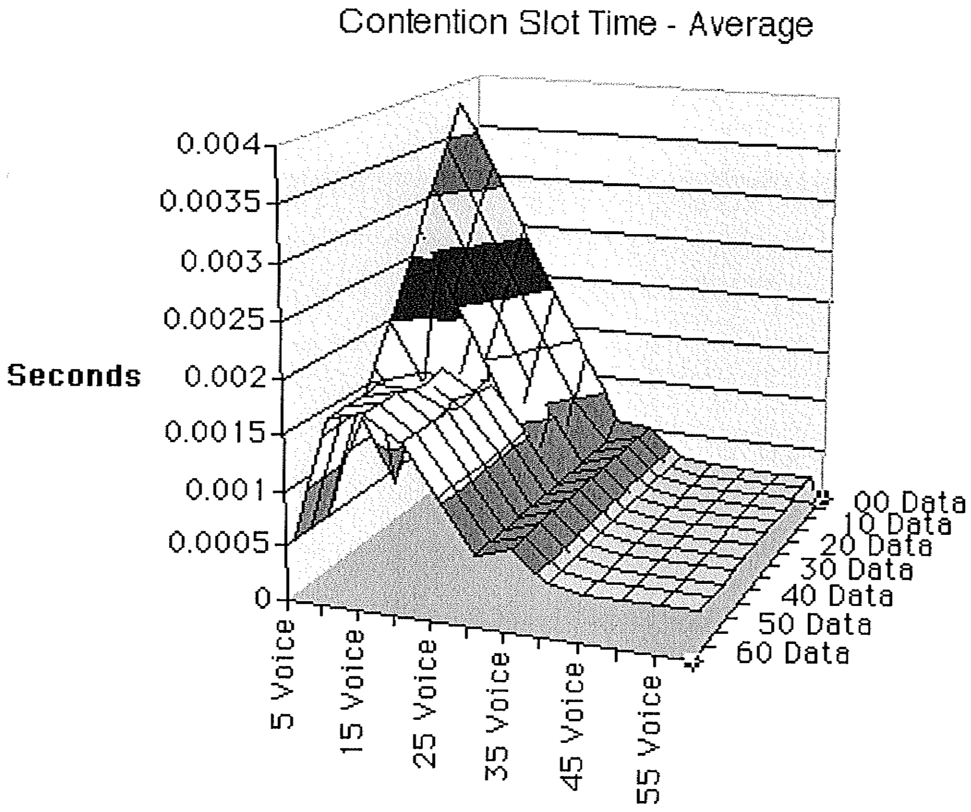

Figure 12 - MAP Contention Time

The reduced amount of contention time under a large number of voice and/or data modems would indicate that the additional IEs in the MAP are likely not Request IEs.

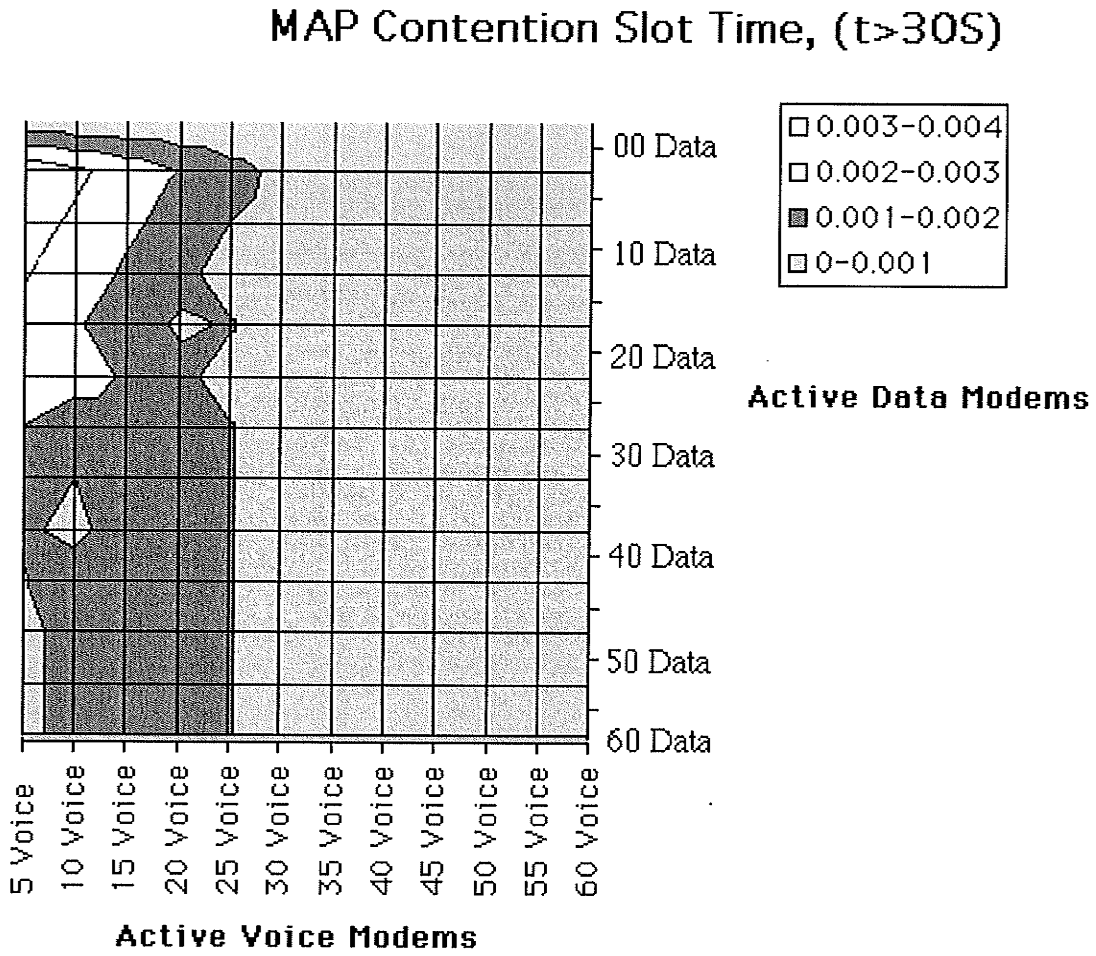

Figure 13 - MAP Contention Time

As can be seen in both Figure 12 and Figure 13, the amount of time devoted to contention decreases with the number of data modems. This is as would be expected as the additional modems generate more traffic the amount of available bandwidth would decrease and therefore the amount of time provided for contention is decreased. From Figure 12, the minimum amount of contention time plateaus at a level of approximately

0.5 microseconds.

Figure 12 and Figure 13 also show that the decrease in contention time with increased voice modems. This is clearly visible in Figure 13 if the gradient for the voice axis is compared to the data axis. The decrease occurs more quickly for voice sources and, from Figure 12, falls to a level which is approximately half of the data modem value.

The reduced amount of time for contention time, coupled with the higher number of modems which would be contesting for that time, would likely result in an increase in the number of collisions which would occur. Interestingly, the contention time is never completely eliminated regardless of the amount of traffic being generated.

It is possible that with the higher traffic levels the contention slots are in a constant state of collision which would be basically wasting those time slots and thus contribute to the reduction in upstream throughput.

5. Conclusions

Based upon the simulation results, it would appear that the typical 500-2000 subscribers [Az97] per LAN segment represents a gross oversubscription of each LAN segment when providing voice service in addition to data service over an HFC network.

Results showed that if as few as 35 CMs were generating voice traffic, they were able to consume all the available upstream bandwidth. Under the worst-case conditions, where 2000 subscribers are sharing the upstream channel, this represent only 1.75% of the CMs on the segment. Therefore it may not be reasonable to simply upgrade CMs without reconsidering network provisioning.

The number of voice sources also affects the efficiency of the upstream channel. It was shown that under high demand, the upstream channel is unable to carry the same data rates that can be achieved when carrying data sources only. Though this effect was clearly seen in the throughput results, it may not be tied solely to the addition of voice sources. By analyzing the characteristics of the upstream bandwidth allocation mechanism, the reduction in overall throughput could be a result of the p-FCFS scheduling algorithm implementation used within the OPNET model. This was illustrated by comparing the the amount of information contained within the MAP and the throughput.

Since this work did not exhaustively analyze the MAP mechanism additional research should be done to understand the characteristics of the bandwidth allocation map and their impact on throughput. However, limitations of CSF13 model may impact the ability to determine work-arounds for problems in upstream bandwidth allocation. The limitations of CSF13 model include: only providing an open-loop model of the HFC environment (cannot simulate TCP/IP); and that the model only operates in steady state conditions.

Also, additional work should be done to determine the characteristics of actual user behaviour on the network. It may be possible, since at the moment there is not a universal way to connect VoIP traffic to the switched telephone network, that VoIP is not in widespread use by residential cable modem users. However it would appear that if as few as 35 subscribers using VoIP applications can adversely impact the performance of the network then the provisioning practices currently in place (i.e. 500-2000 users/LAN segment) may not be adequate to provide reasonable VoIP performance.

6 References

[Az97] Azzam, A., High Speed Cable Modems, McGraw-Hill, 1997

[CaLa02] Total of 21 1.1 Modems Certified to date - CableLabs® Certifies More DOCSIS™ 1.1 Modems,
<http://www.cablelabs.com/news/pr/2002/2002_03_28.html>, August 2002

[CDN1] Cable Datacom News, Cable Modem Standards and Specifications,
<http://www.cabledatacomnews.com/cmic/cmic3.html>, July 2002

[CNET1] Charny, B. and Hu, J., Time Warner Cable reaches VoIP deals, CNET News.com, Dec 2003
<http://news.com.com/2100-7352_3-5116936.html?tag=nefd_top>

[DOCRFI1.1] Data-Over-Cable Service Interface Specifications SP-RFIv1.1-I05-000714 Radio Frequency Interface Specification, 2000 Cable Television Laboratories, Inc., 2000

[ITUW] ITU web site
<http://www.itu.int>, July 2002

[Mcp99] McPhillips, E., The Factors Affecting the Growth of VoIP,
Extended Enterprise Laboratory, HP Laboratories Bristol, HPL-IRI-1999-003, May 1999

[OpDoc] OPNET DOCSIS Model Documentation, MIL 3, Inc., 1999

[SdSm99] V. Sdralia, C. Smythe, P. Tzerefos, and S. Cvetkovic, Performance Characterization of the MCNS DOCSIS 1.0 CATV Protocol with Prioritized First Come First Served Scheduling, IEEE Trans. Broadcasting, vol. 45, no. 2, 1999, pp. 196–205.

[Yin98] Ying-Dar Lin, Wei-Ming Yin, and Chen-Yu Huang, An Investigation Into HFC MAC Protocols: Mechanisms, Implementation, And Research Issues, IEEE Communications Magazine, September 1998

Appendix A

Upstream Throughput

The graphs presented in body of this document described only the average values for the parameters under discussion. For completeness, additional statistics describing the data are provided here. The additional information includes, the minimum, maximum, as well as the standard deviation of the raw data.

Upstream Throughput - Average (t>30s)

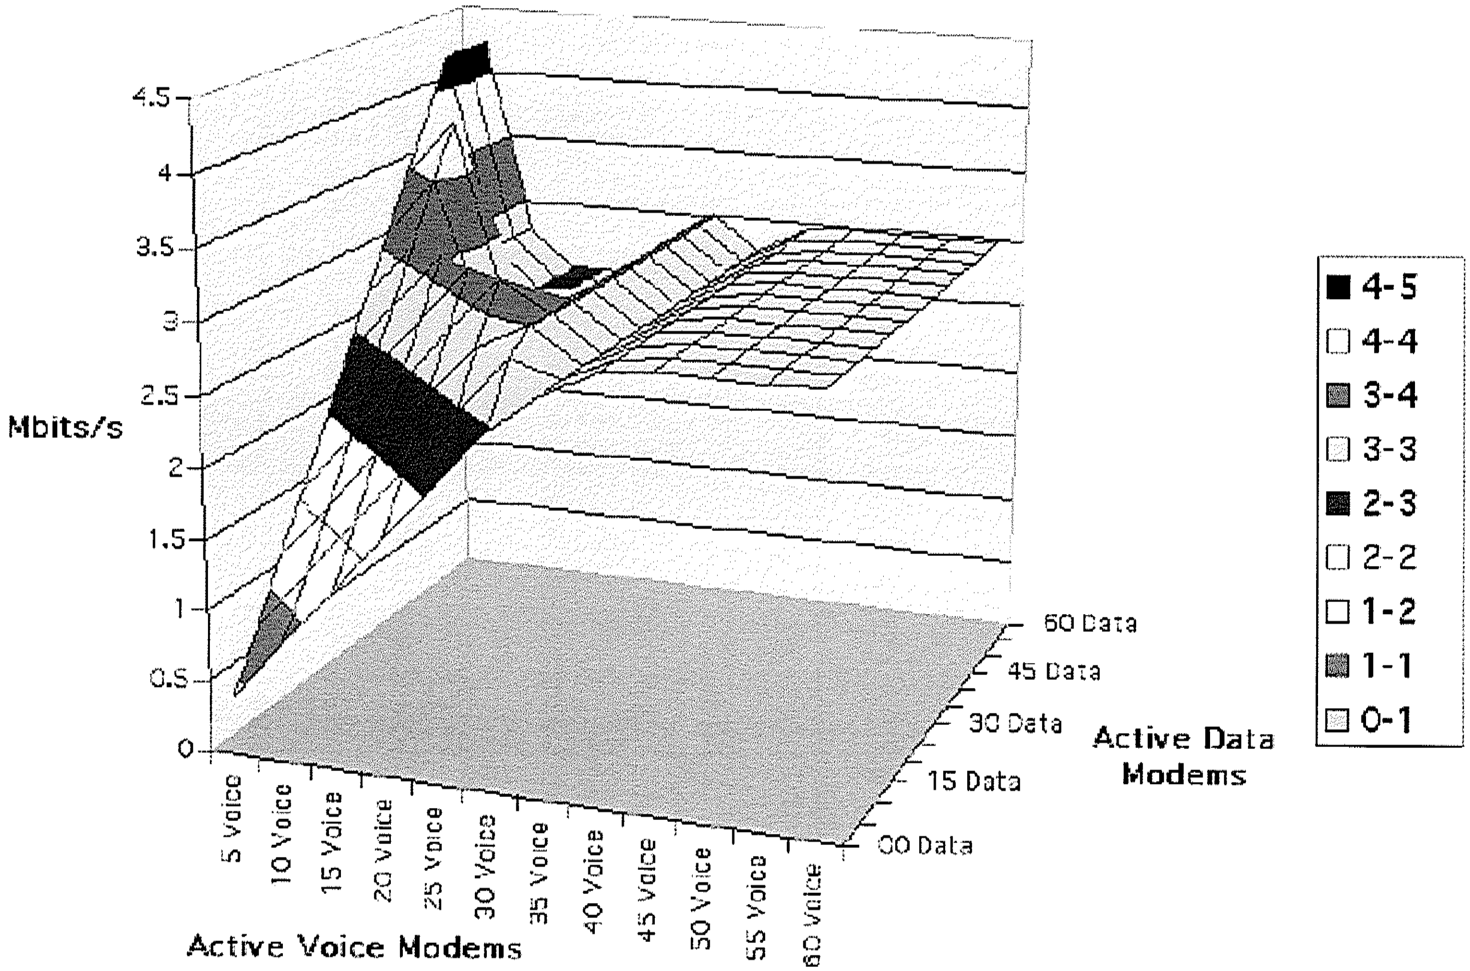

Upstream Throughput (Mbps) - Average (t>30s)

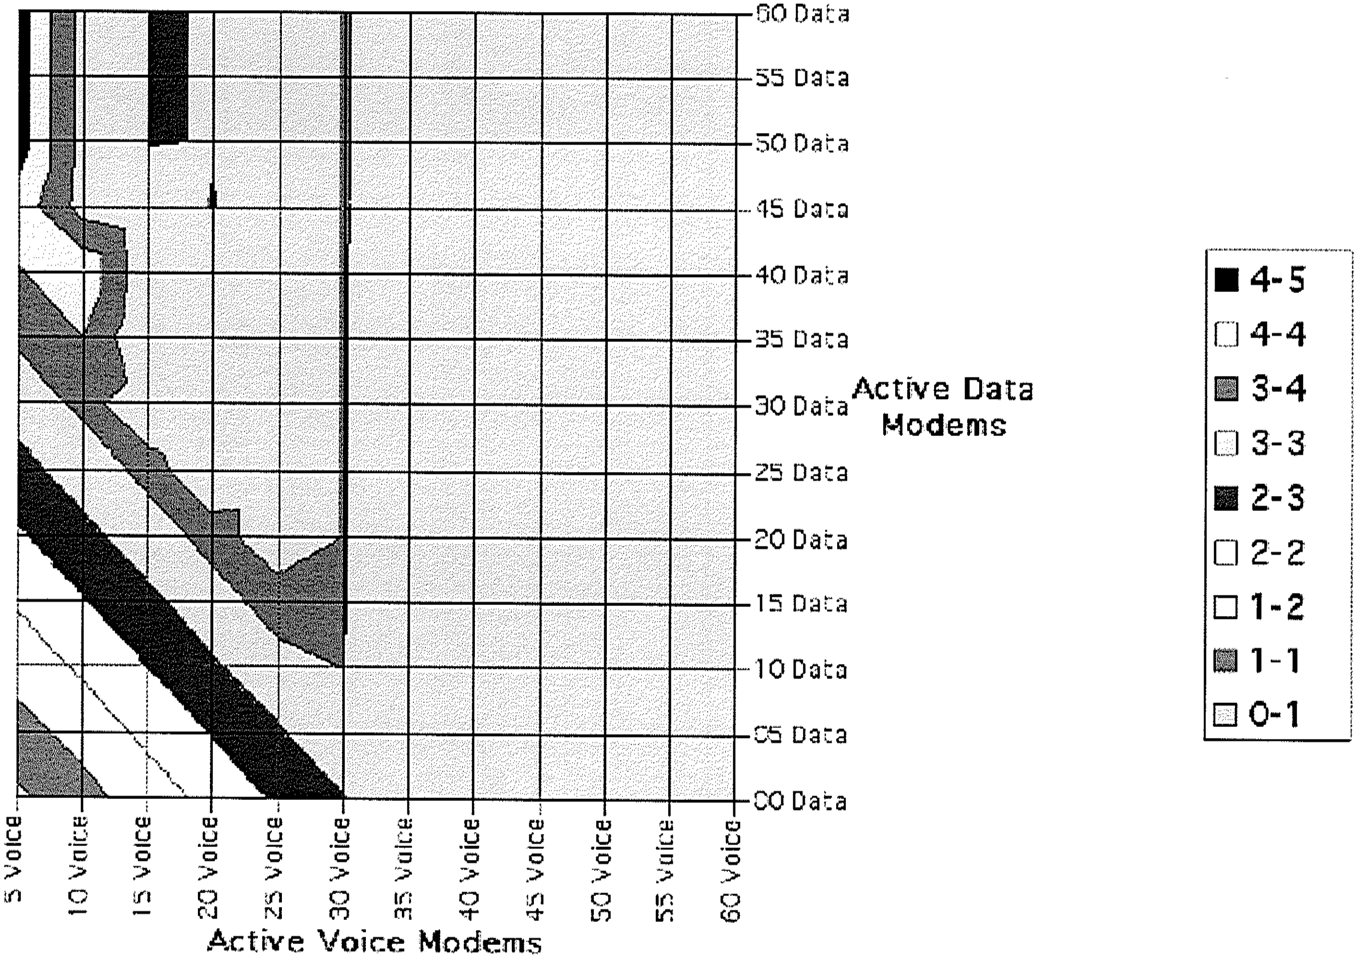

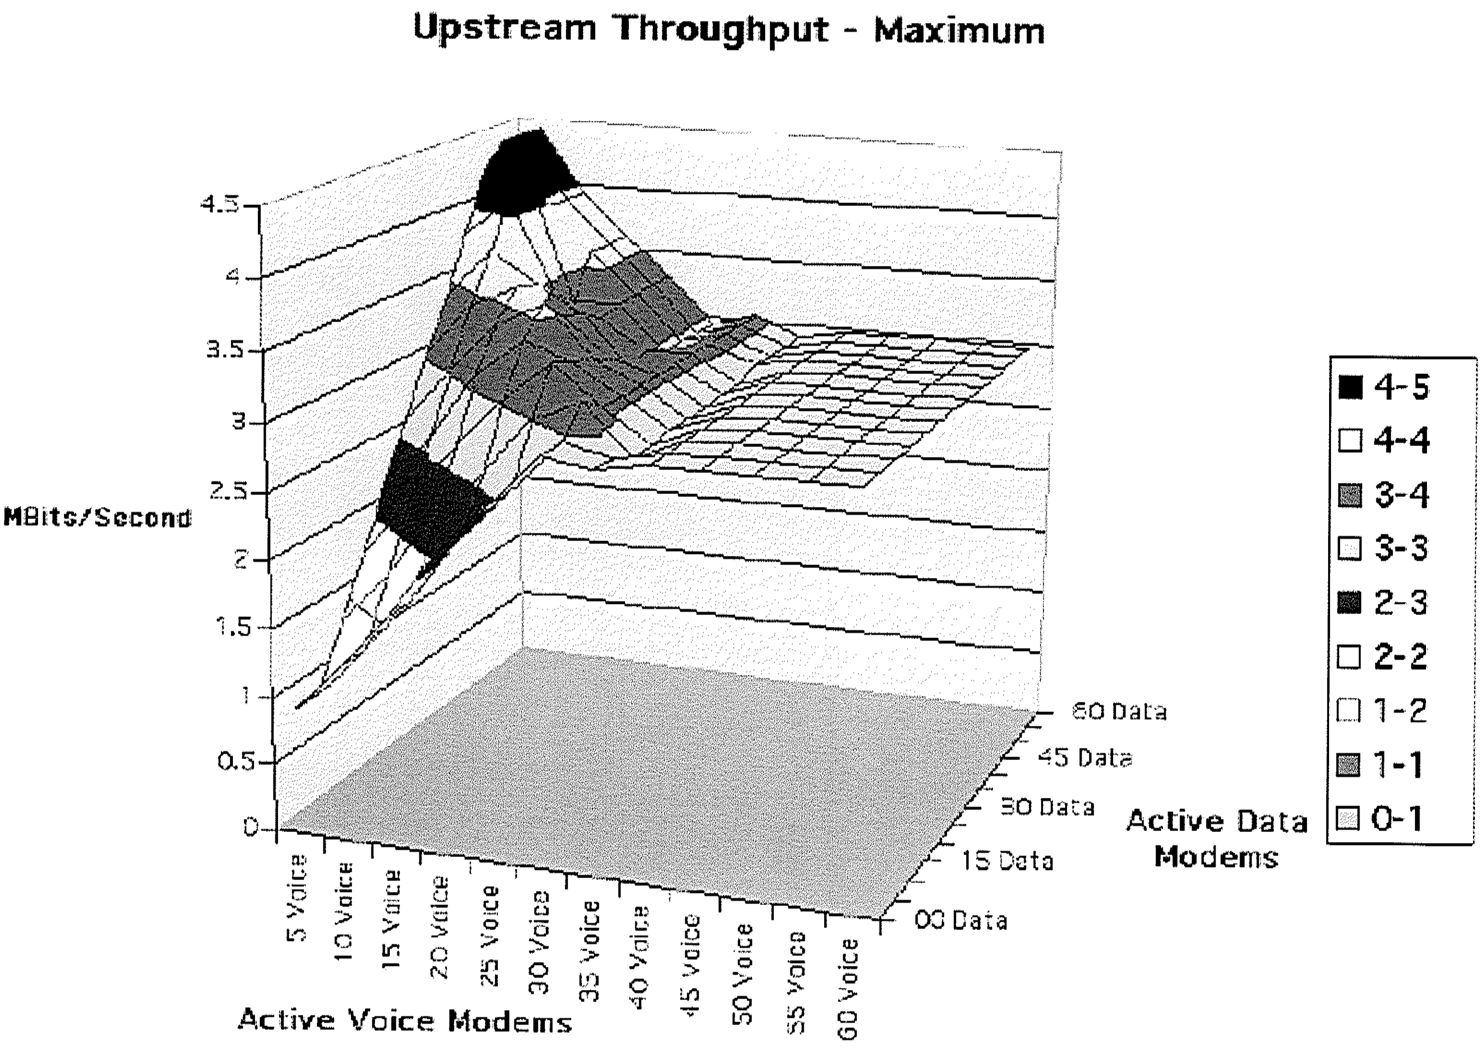
Upstream Throughput - Maximum
MBits/Second
4.5
4
3.5
3
2.5
2
1.5
1
0.5
0
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
Active Voice Modems
60 Data
45 Data
30 Data
15 Data
00 Data
Active Data Modems
4-5
4-4
3-4
3-3
2-3
2-2
1-2
1-1
0-1

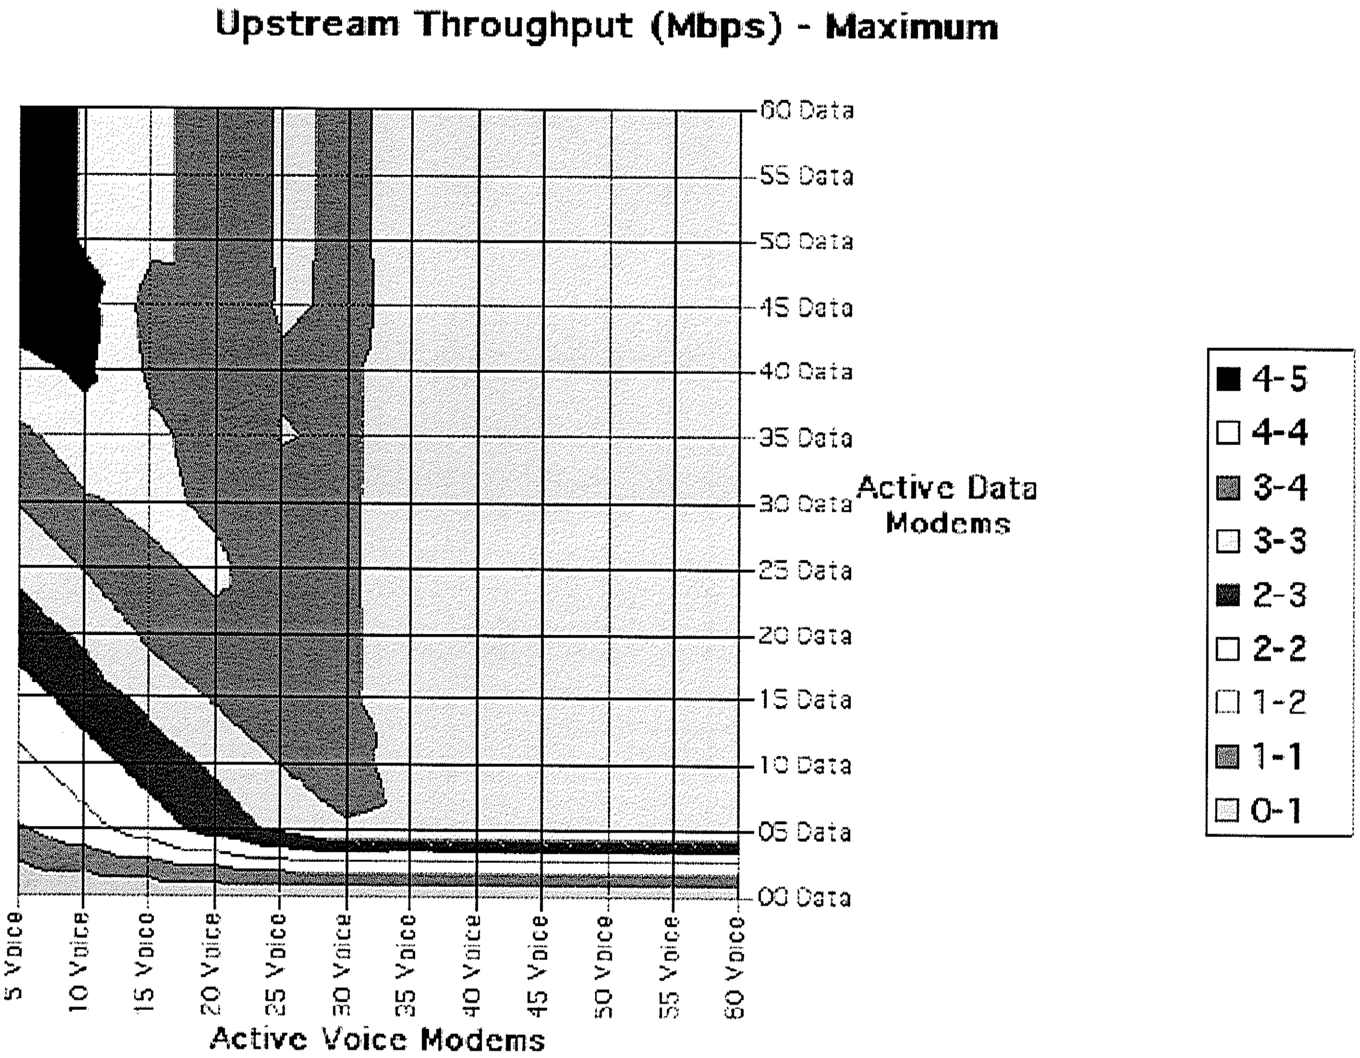
Upstream Throughput (Mbps) - Maximum
60 Data
55 Data
50 Data
45 Data
40 Data
35 Data
30 Data
25 Data
20 Data
15 Data
10 Data
05 Data
00 Data
Active Data Modems
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
Active Voice Modems
4-5
4-4
3-4
3-3
2-3
2-2
1-2
1-1
0-1

Upstream Throughput (Mbps) - Minimum

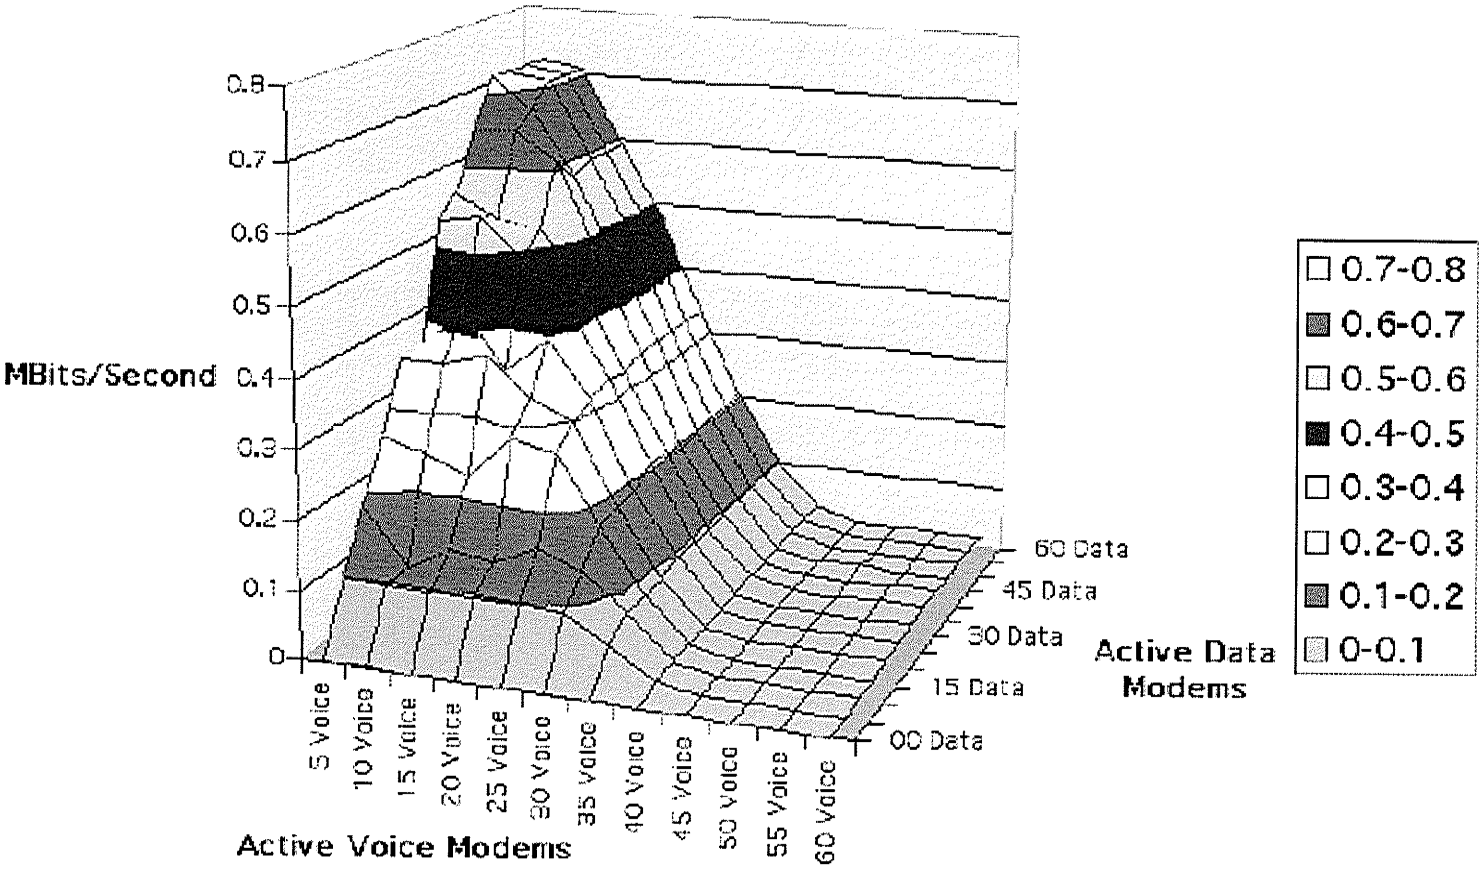

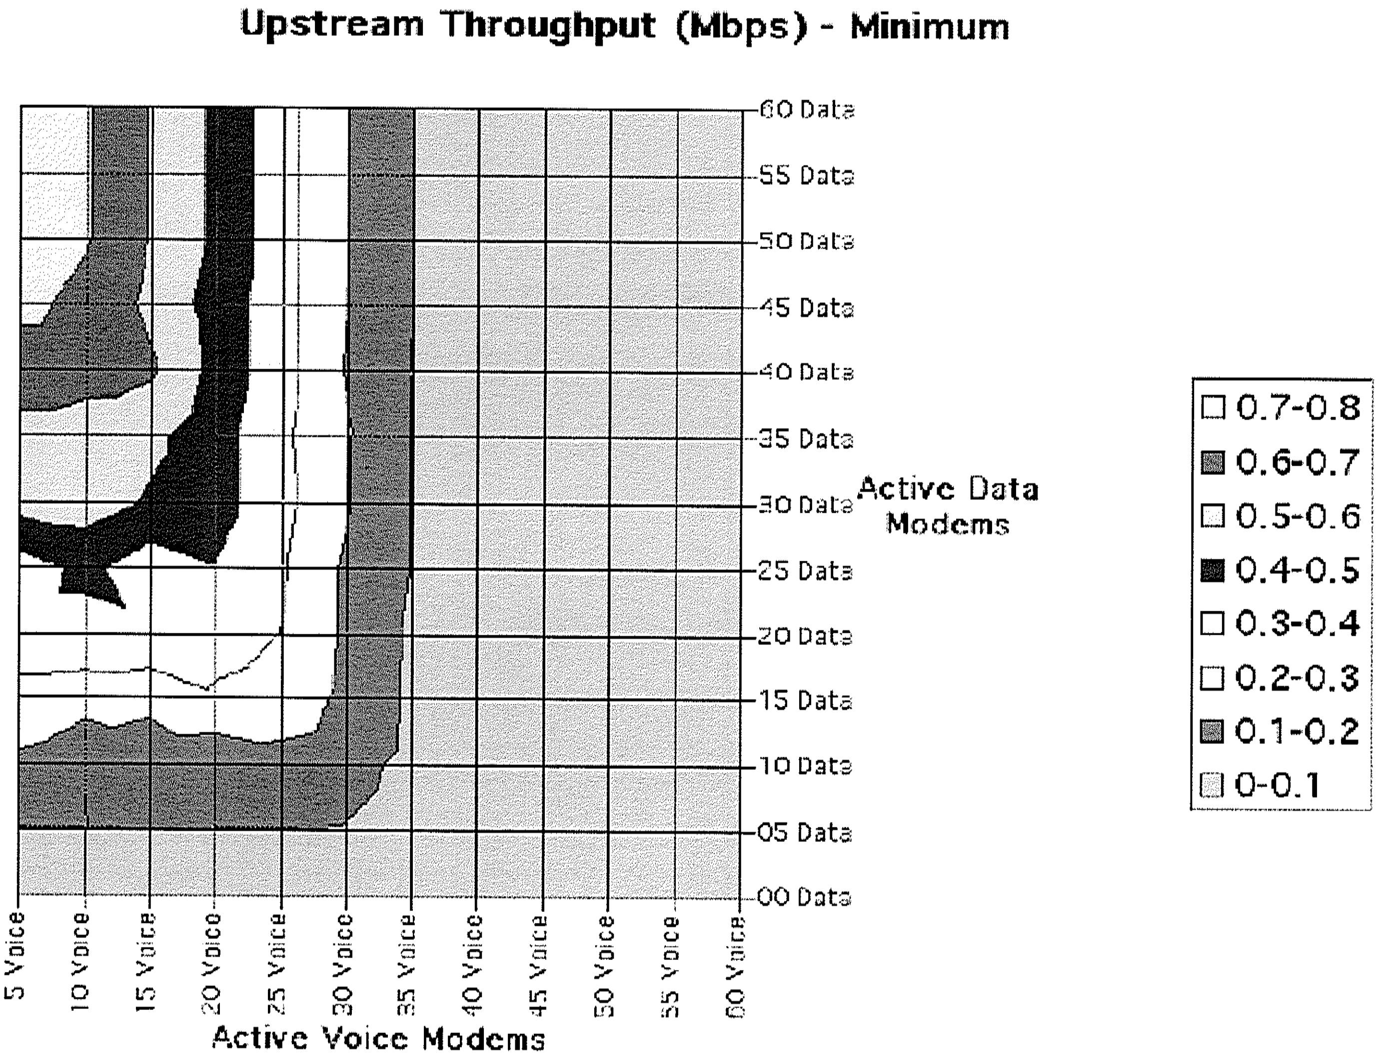
Upstream Throughput (Mbps) - Minimum
60 Data
55 Data
50 Data
45 Data
40 Data
35 Data
30 Data
25 Data
20 Data
15 Data
10 Data
05 Data
00 Data
Active Data Modems
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
Active Voice Modems
☐ 0.7-0.8
▨ 0.6-0.7
☐ 0.5-0.6
■ 0.4-0.5
☐ 0.3-0.4
☐ 0.2-0.3
▨ 0.1-0.2
☐ 0-0.1

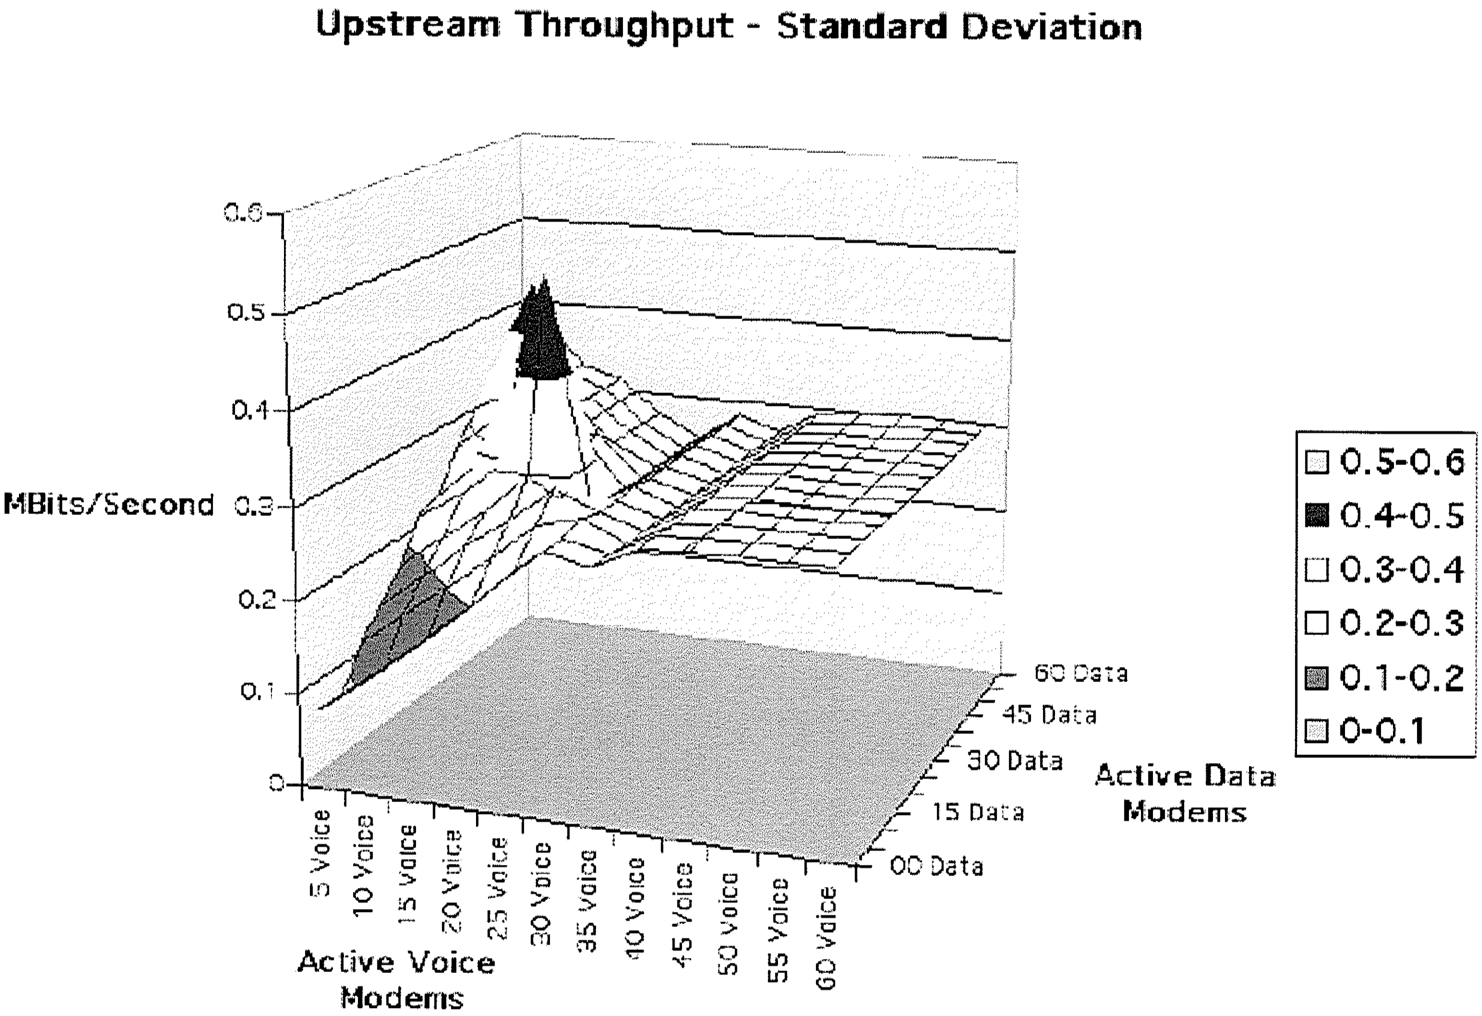
Upstream Throughput - Standard Deviation
MBits/Second
0.6
0.5
0.4
0.3
0.2
0.1
0
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
Active Voice Modems
00 Data
15 Data
30 Data
45 Data
60 Data
Active Data Modems
□ 0.5-0.6
■ 0.4-0.5
□ 0.3-0.4
□ 0.2-0.3
▨ 0.1-0.2
□ 0-0.1

Upstream Throughput (Mbps) - Standard Deviation

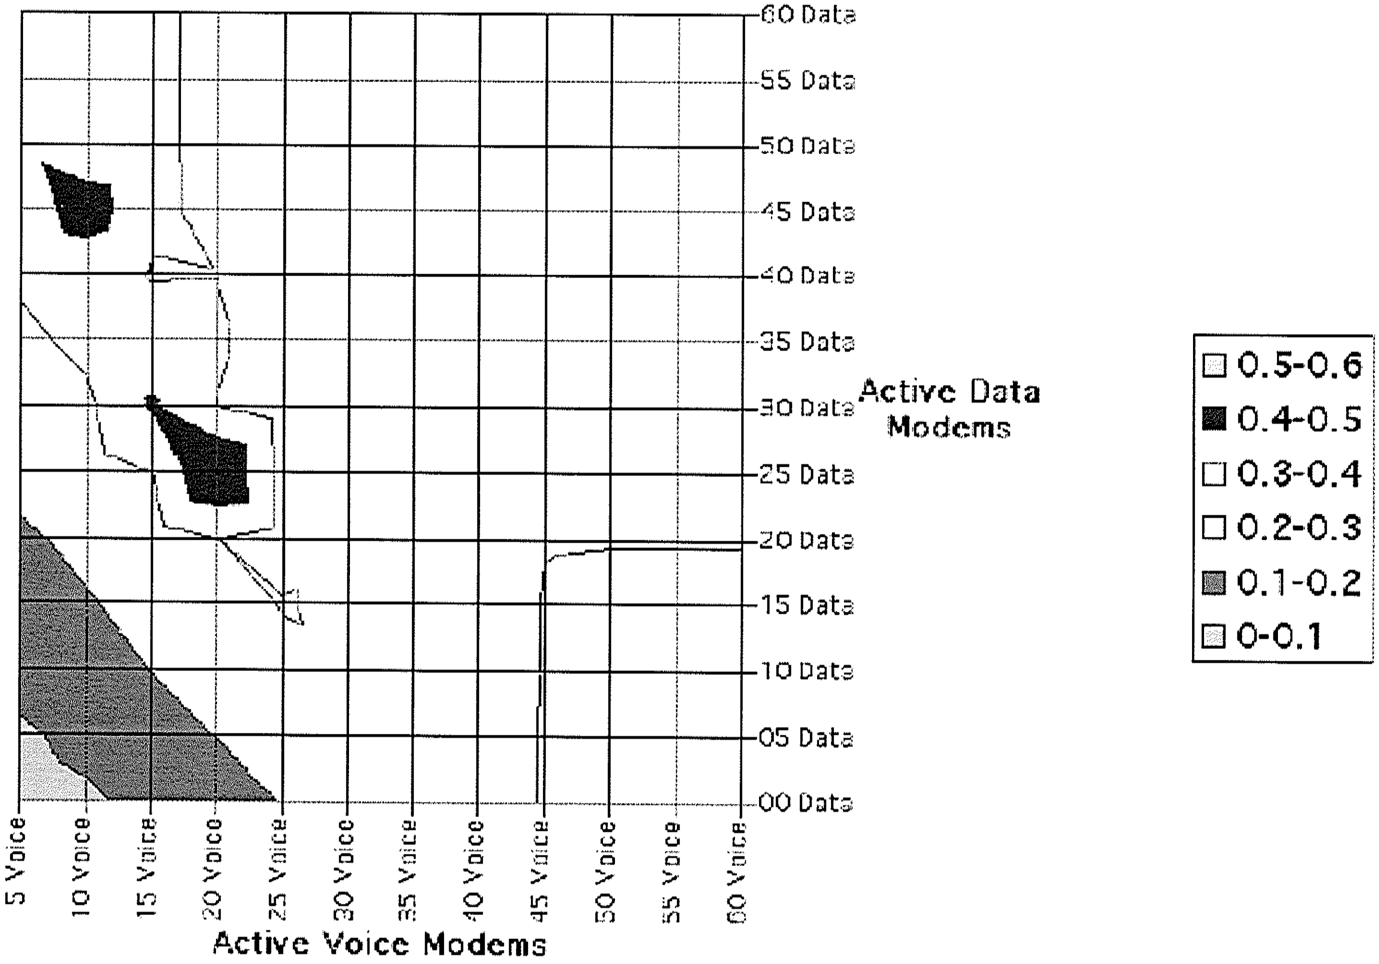

Appendix B

MAP IEs

The graphs presented in body of this document described only the average values for the parameters under discussion. For completeness, additional statistics describing the data are provided here. The additional information includes, the minimum, maximum, as well as the standard deviation of the raw data.

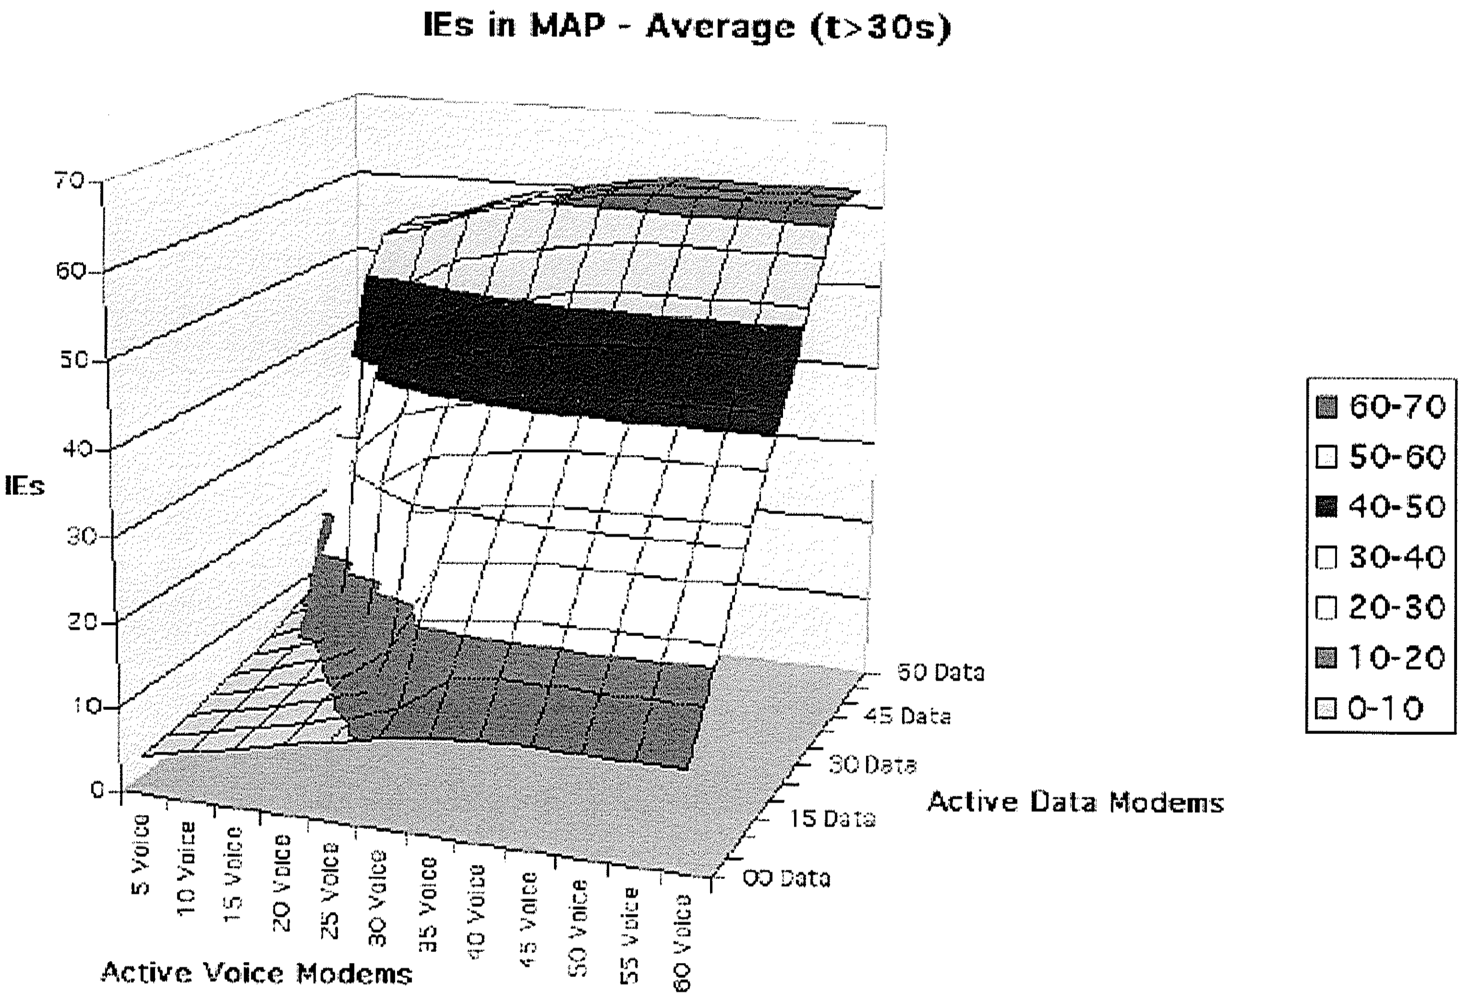
IEs in MAP - Average (t>30s)
IEs
70
60
50
40
30
20
10
0
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
Active Voice Modems
00 Data
15 Data
30 Data
45 Data
50 Data
Active Data Modems
▨ 60-70
☐ 50-60
■ 40-50
☐ 30-40
☐ 20-30
▨ 10-20
☐ 0-10

IEs in MAP - Average (t>30s)

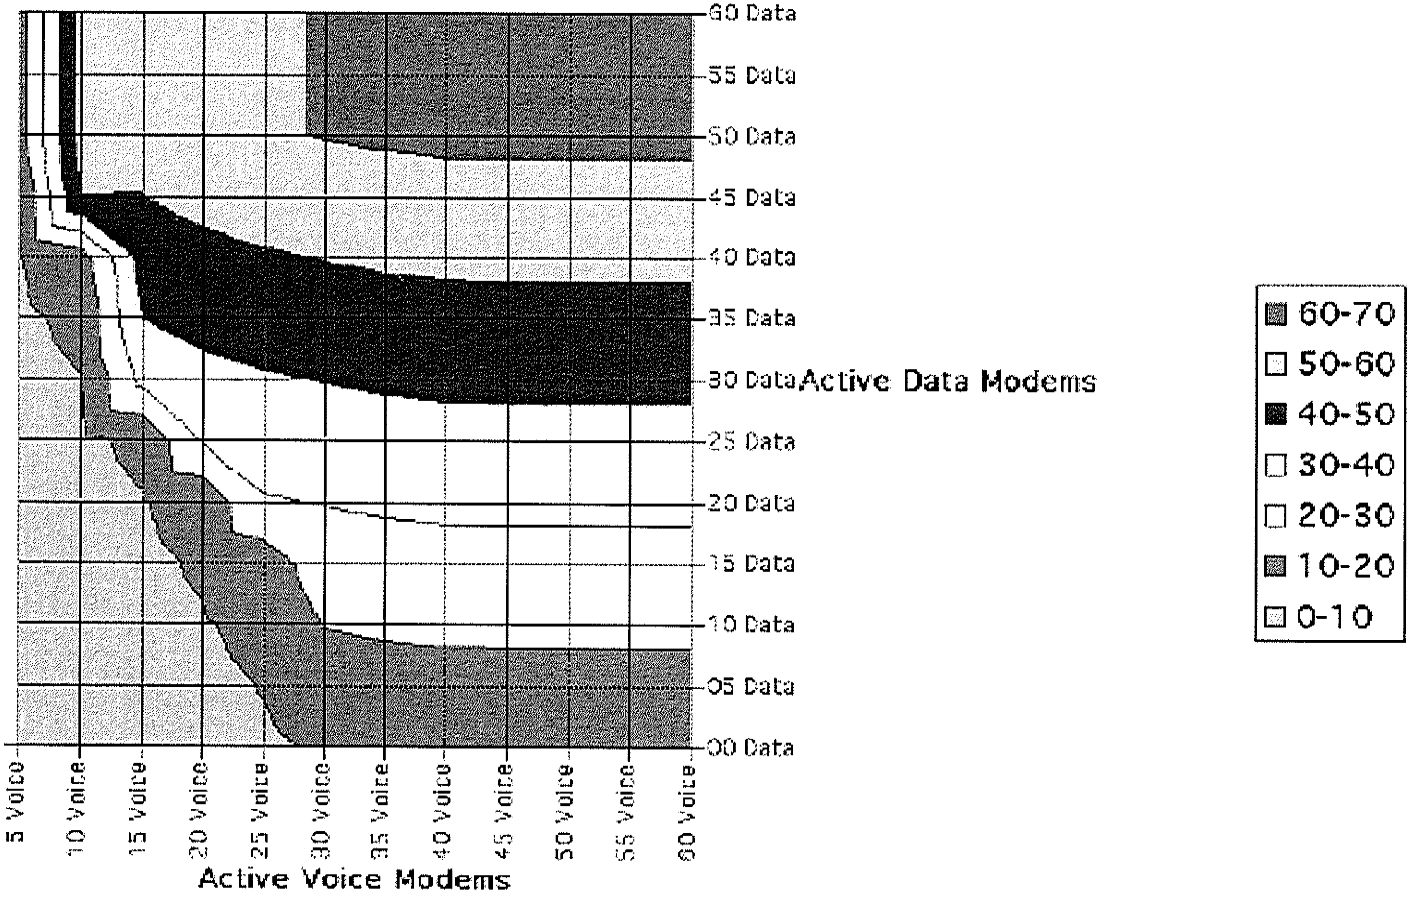

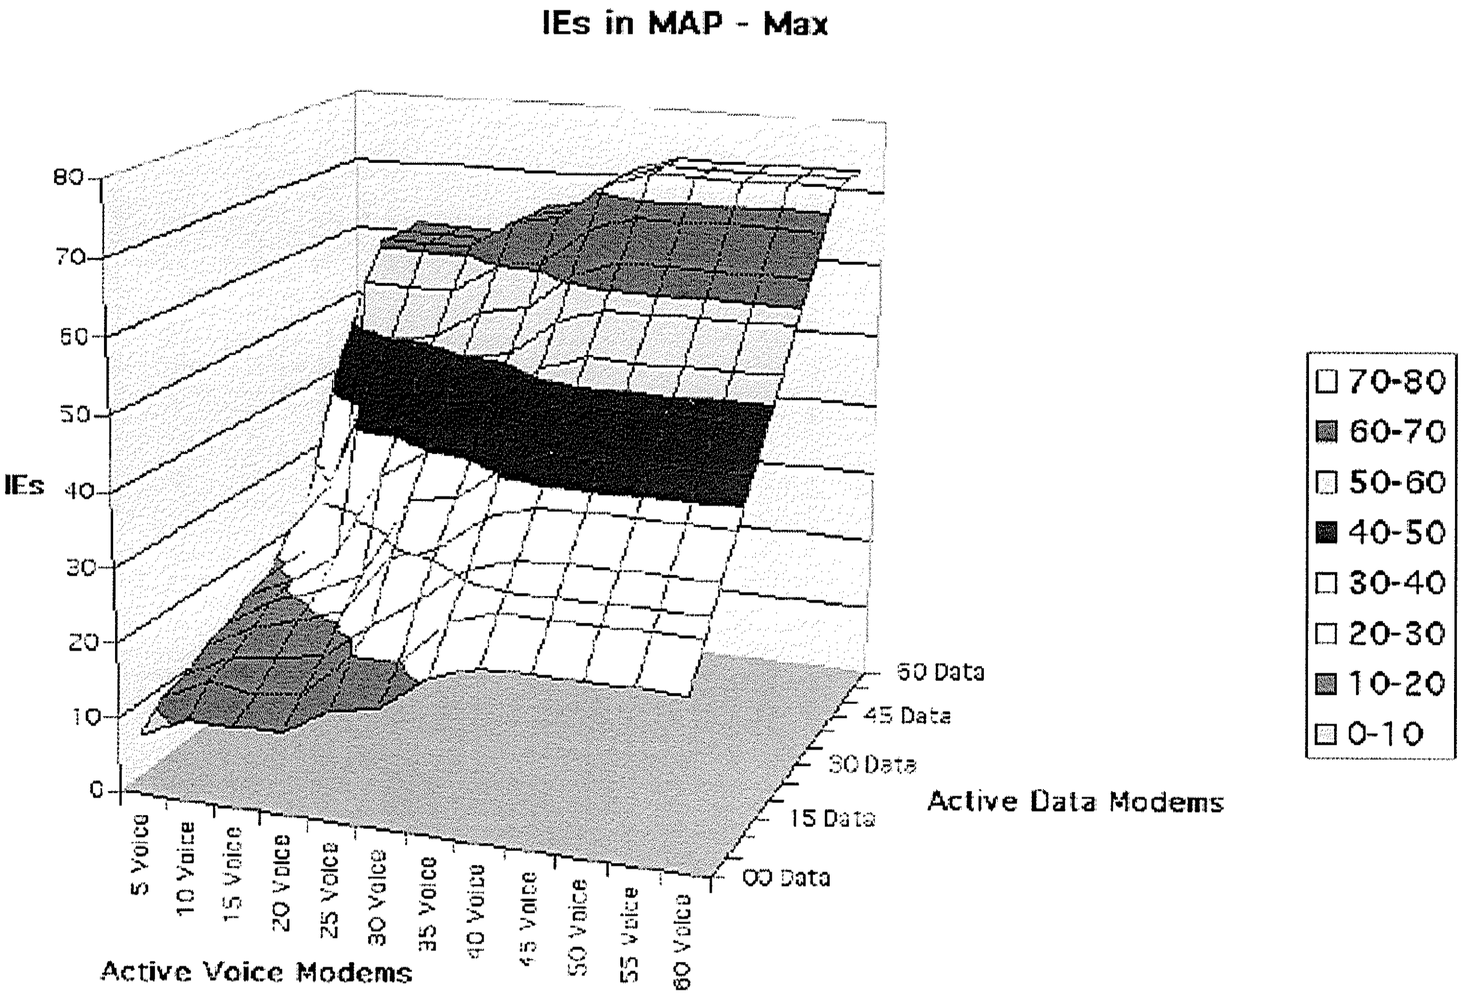
IEs in MAP - Max
IEs
80
70
60
50
40
30
20
10
0
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
Active Voice Modems
00 Data
15 Data
30 Data
45 Data
60 Data
Active Data Modems
☐ 70-80
▨ 60-70
☐ 50-60
■ 40-50
☐ 30-40
☐ 20-30
▨ 10-20
☐ 0-10

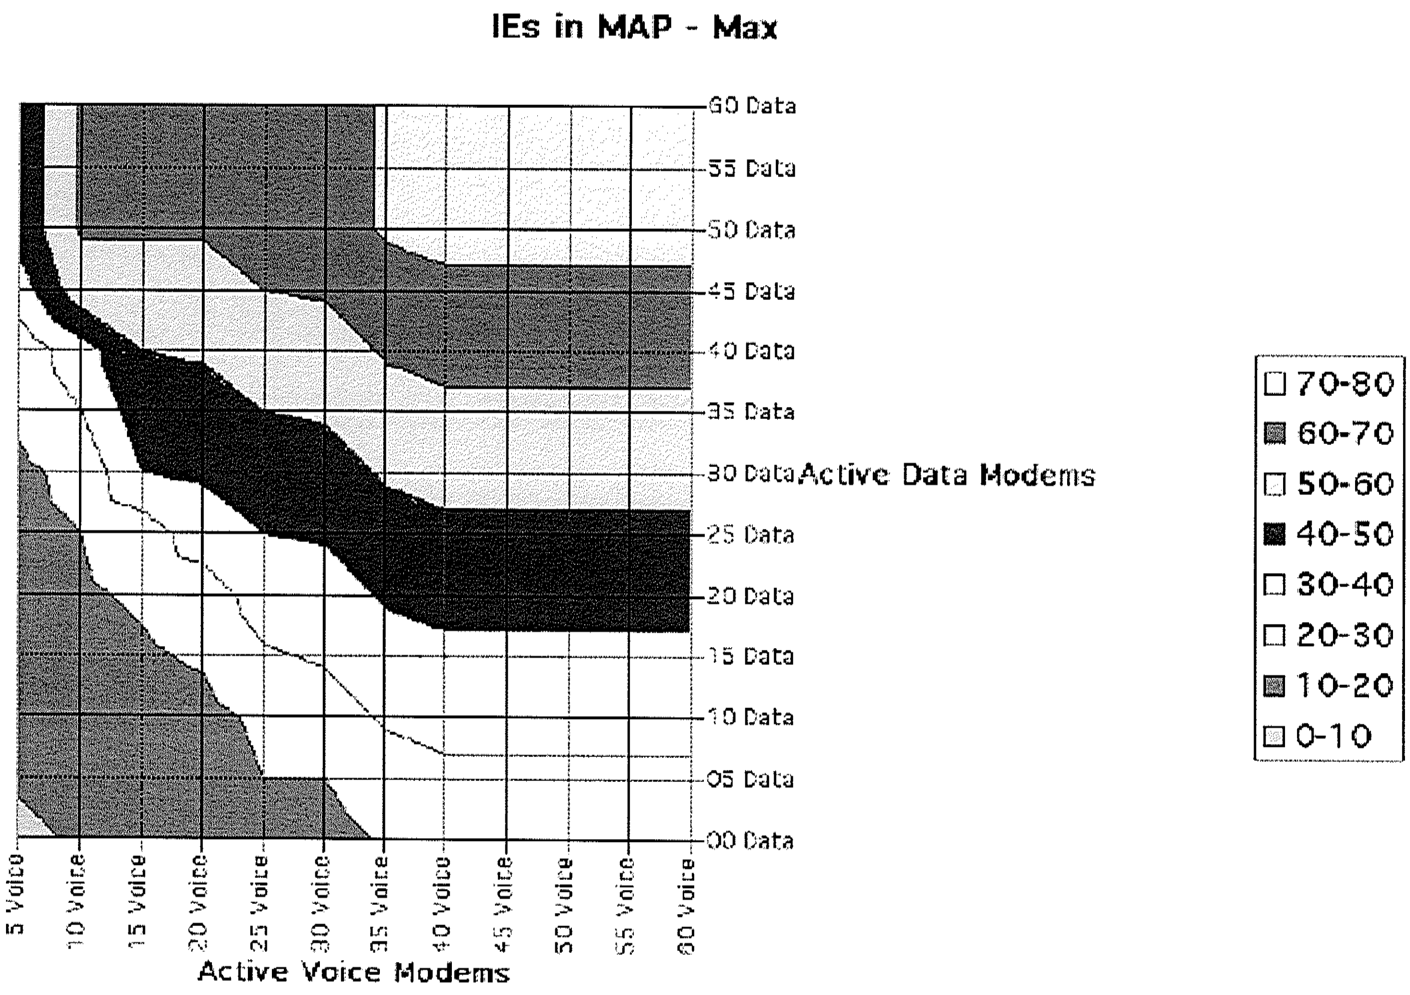
IEs in MAP - Max
60 Data
55 Data
50 Data
45 Data
40 Data
35 Data
30 Data
25 Data
20 Data
15 Data
10 Data
05 Data
00 Data
Active Data Modems
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
Active Voice Modems
□ 70-80
▨ 60-70
□ 50-60
■ 40-50
□ 30-40
□ 20-30
▨ 10-20
□ 0-10

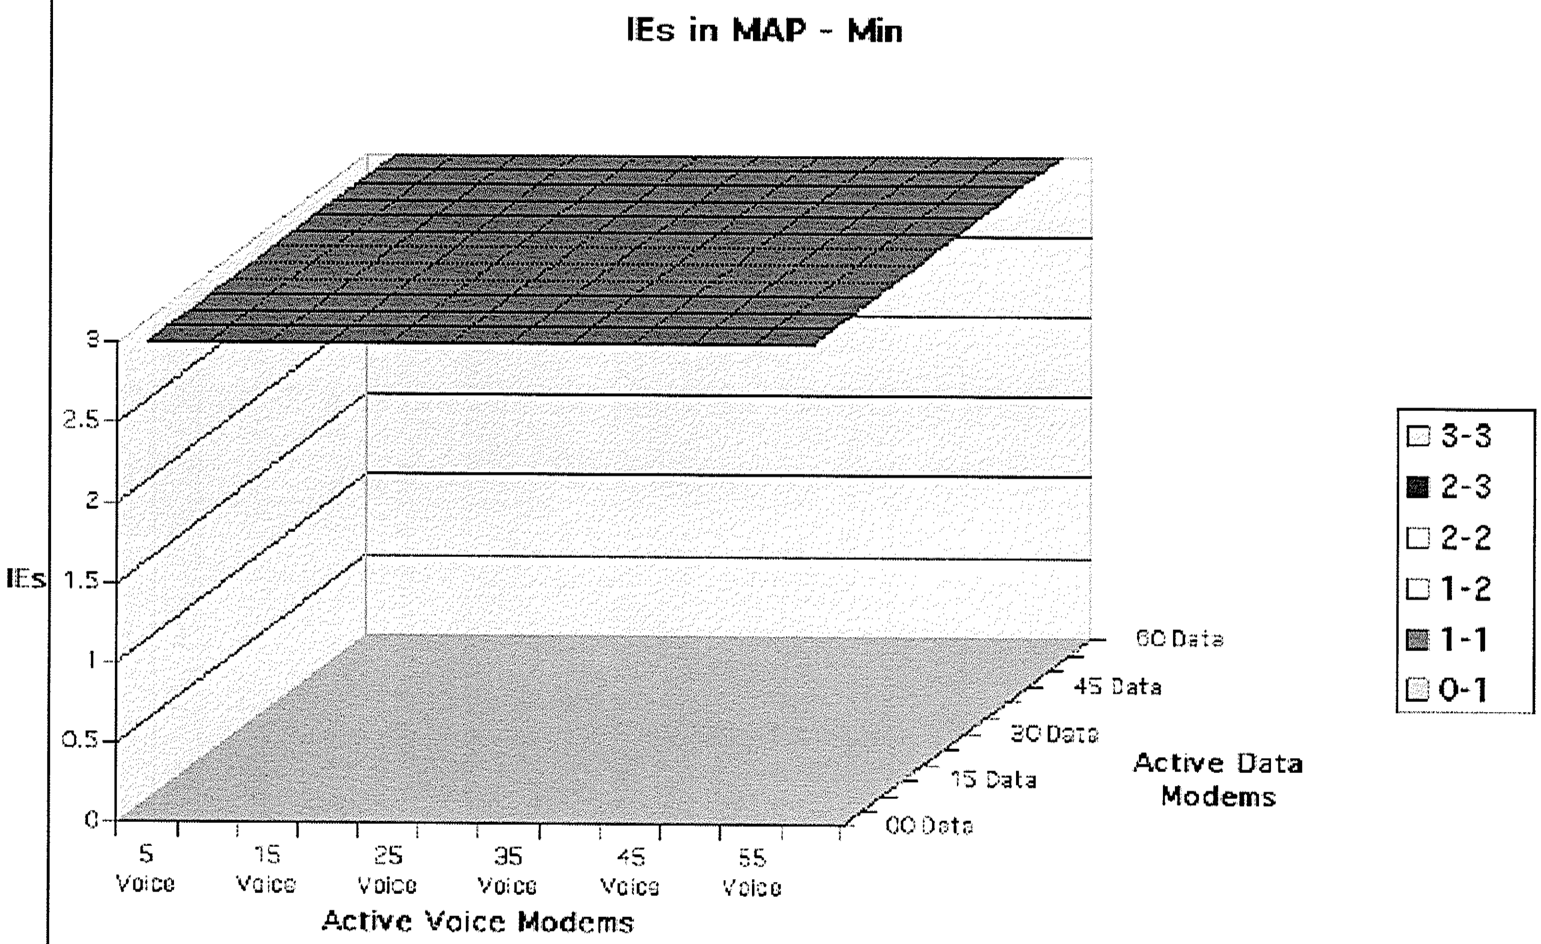
IEs in MAP - Min
IEs
3
2.5
2
1.5
1
0.5
0
5 Voice
15 Voice
25 Voice
35 Voice
45 Voice
55 Voice
Active Voice Modems
00 Data
15 Data
30 Data
45 Data
60 Data
Active Data Modems
☐ 3-3
■ 2-3
☐ 2-2
☐ 1-2
▨ 1-1
☐ 0-1

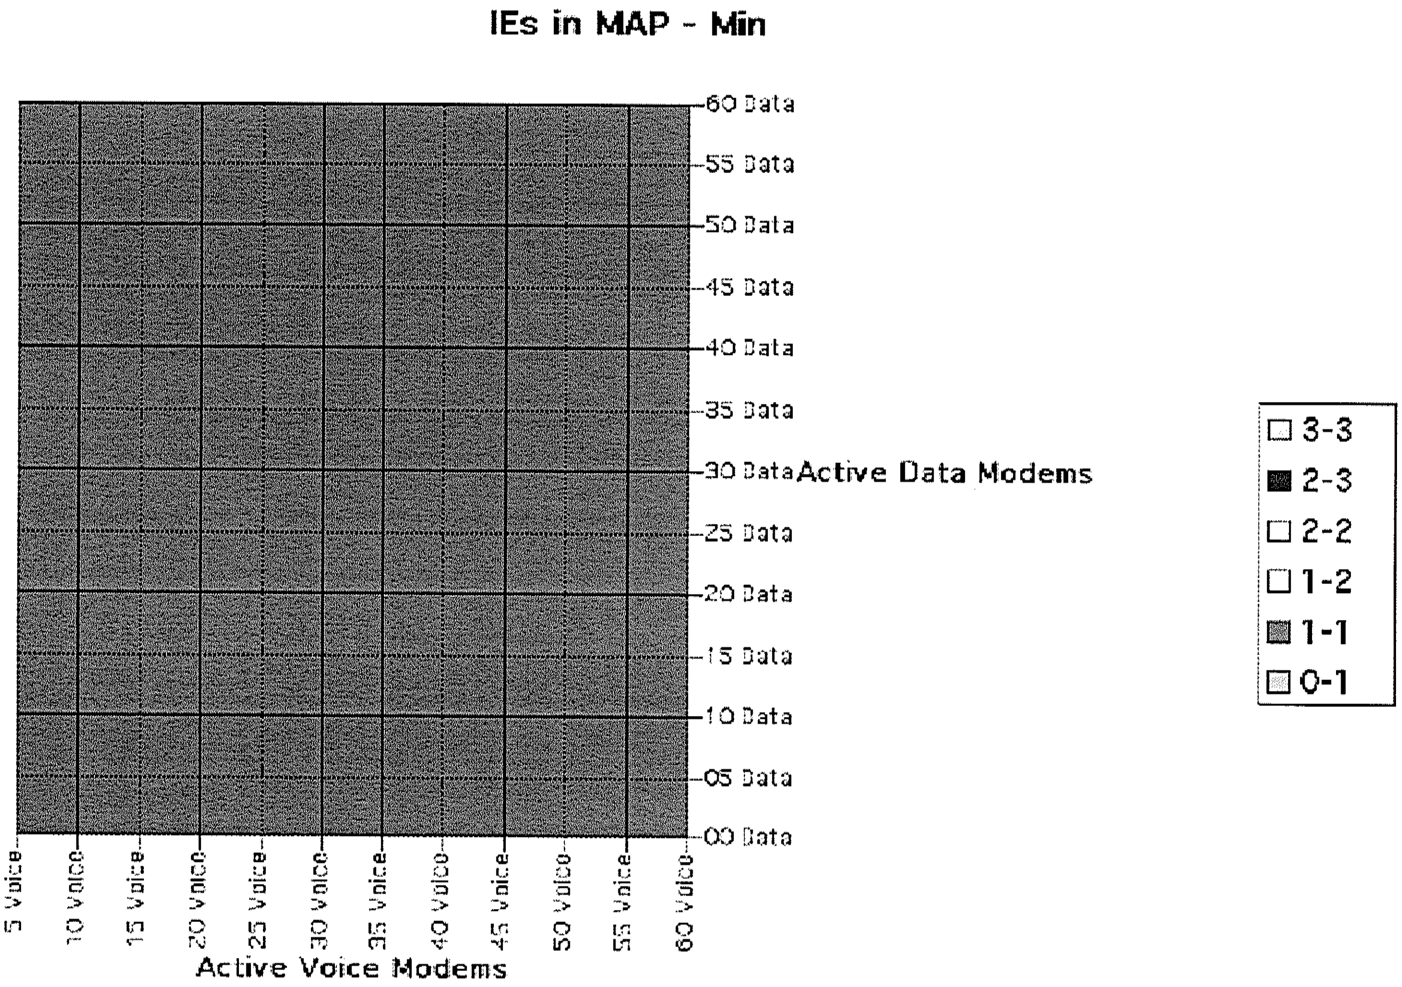
IEs in MAP - Min
60 Data
55 Data
50 Data
45 Data
40 Data
35 Data
30 Data
25 Data
20 Data
15 Data
10 Data
05 Data
00 Data
Active Data Modems
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
Active Voice Modems
3-3
2-3
2-2
1-2
1-1
0-1

IEs in MAP - Std Dev

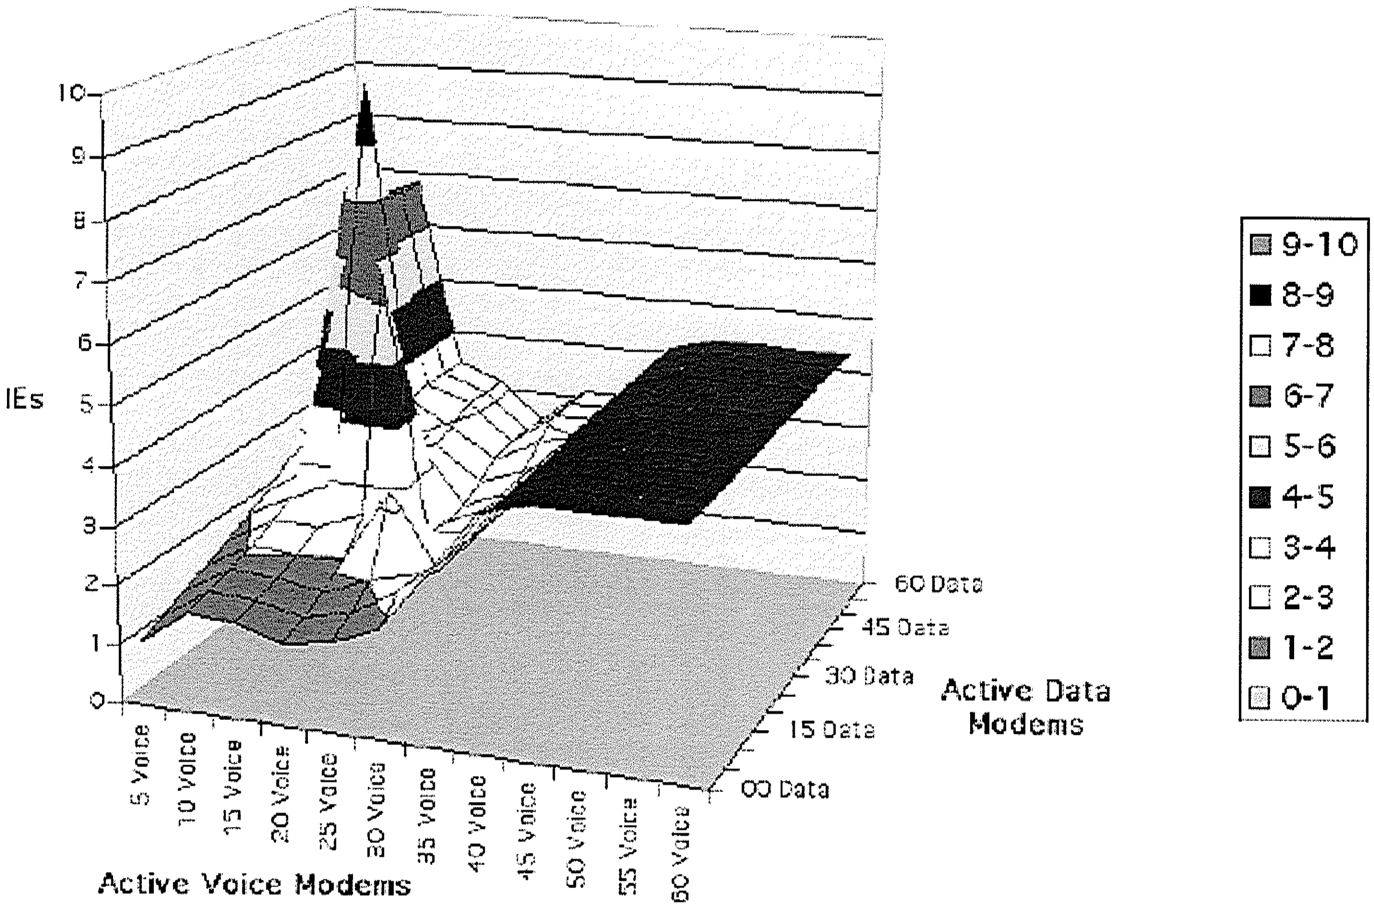

IE in MAP - Std Dev

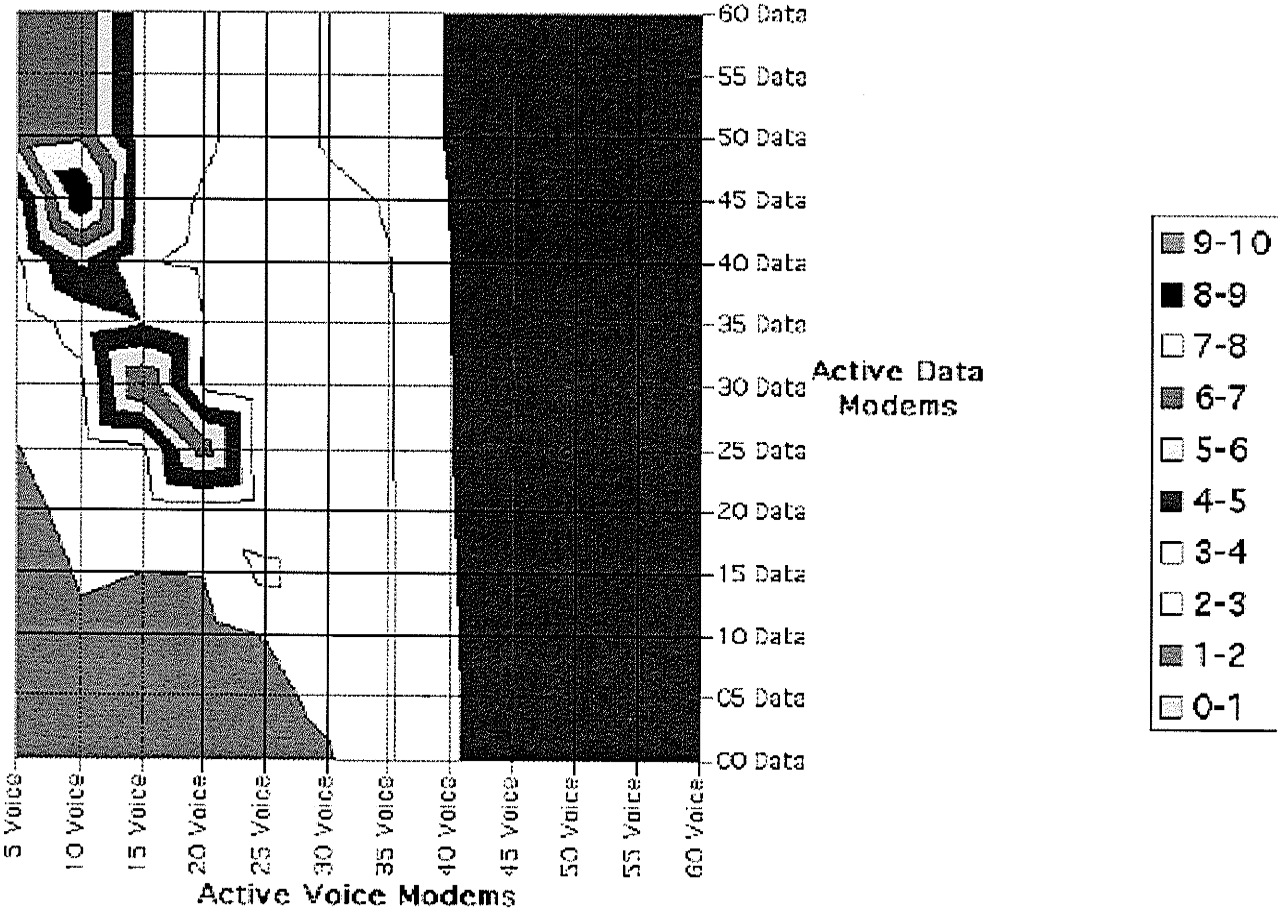

Appendix C

Contention Slot Time

The graphs presented in body of this document described only the average values for the parameters under discussion. For completeness, additional statistics describing the data are provided here. The additional information includes, the minimum, maximum, as well as the standard deviation of the raw data.

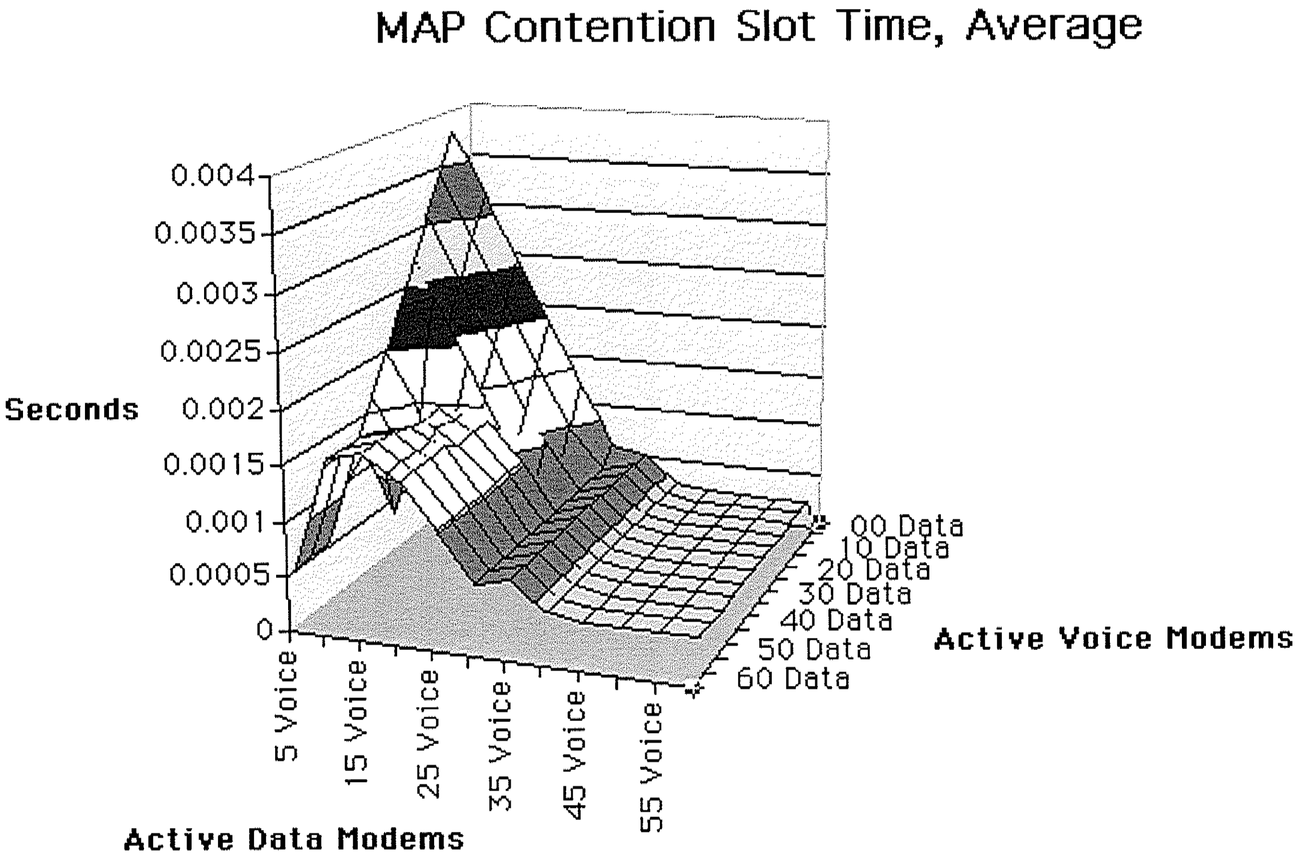
MAP Contention Slot Time, Average
Seconds
0.004
0.0035
0.003
0.0025
0.002
0.0015
0.001
0.0005
0
5 Voice
15 Voice
25 Voice
35 Voice
45 Voice
55 Voice
Active Data Modems
00 Data
10 Data
20 Data
30 Data
40 Data
50 Data
60 Data
Active Voice Modems

MAP Contention Slot Time, (t>30S)

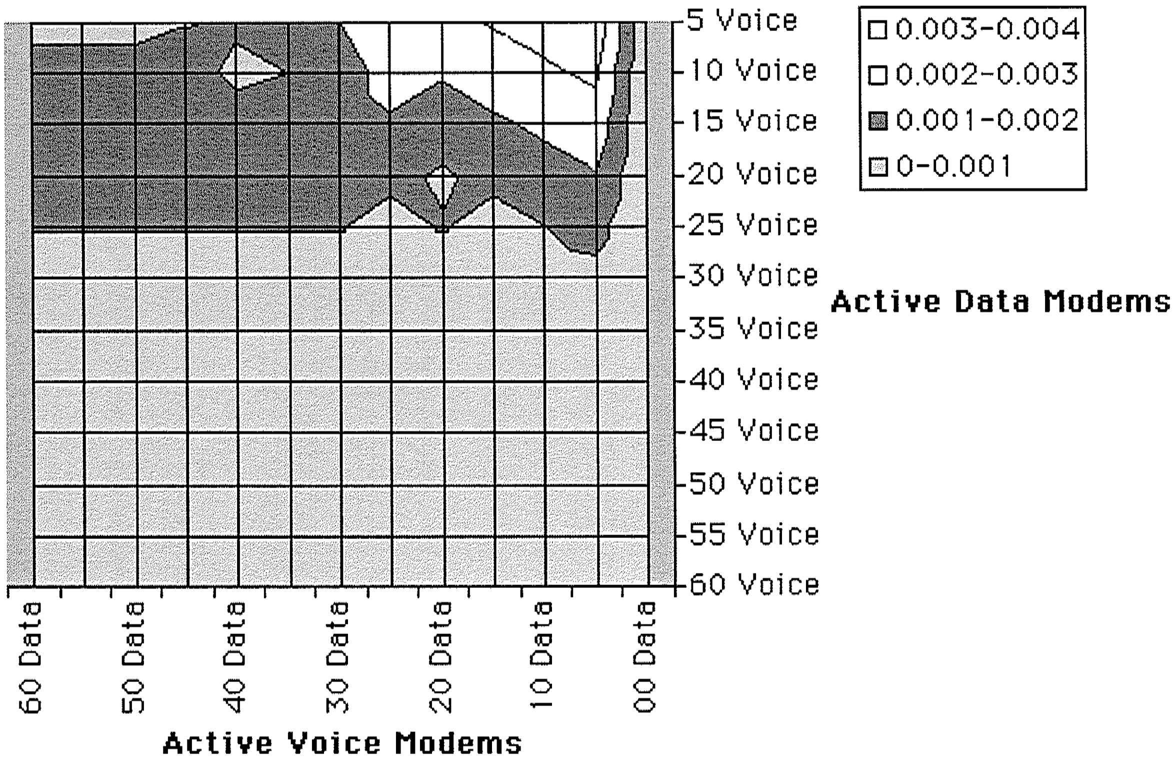

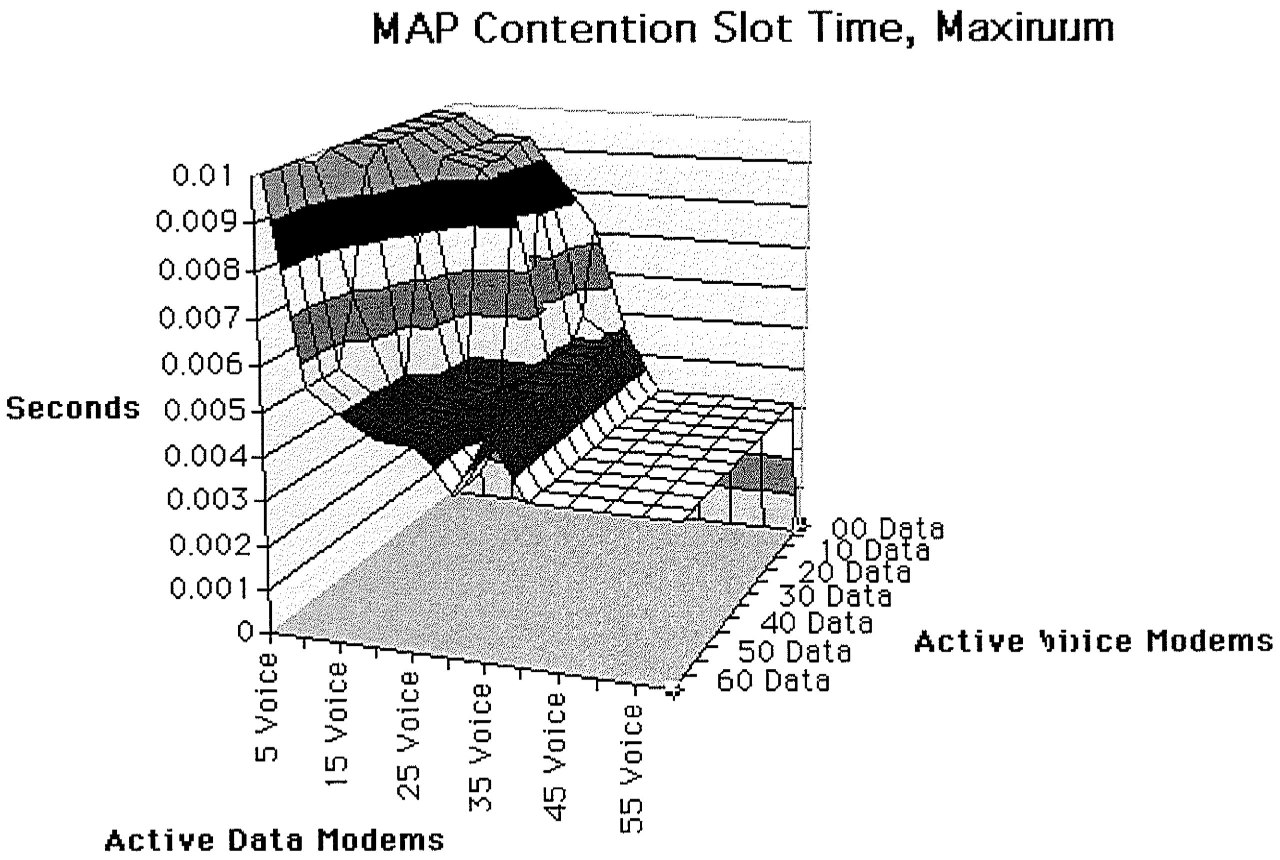
MAP Contention Slot Time, Maximum
Seconds
0.01
0.009
0.008
0.007
0.006
0.005
0.004
0.003
0.002
0.001
0
5 Voice
15 Voice
25 Voice
35 Voice
45 Voice
55 Voice
Active Data Modems
00 Data
10 Data
20 Data
30 Data
40 Data
50 Data
60 Data
Active Voice Modems

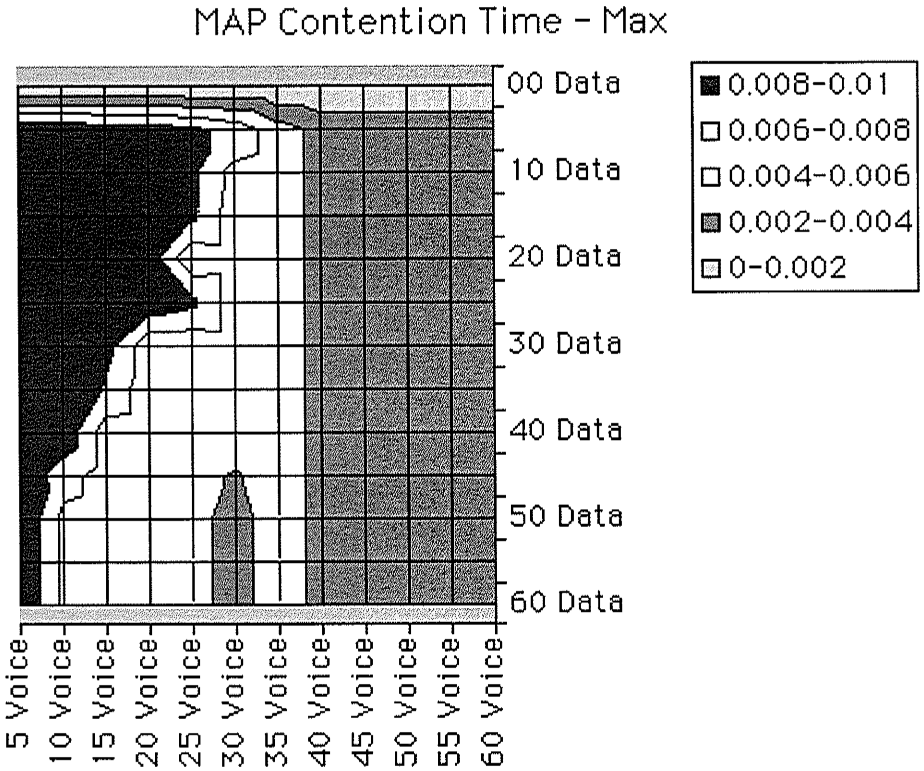
MAP Contention Time - Max
00 Data
10 Data
20 Data
30 Data
40 Data
50 Data
60 Data
5 Voice
10 Voice
15 Voice
20 Voice
25 Voice
30 Voice
35 Voice
40 Voice
45 Voice
50 Voice
55 Voice
60 Voice
■ 0.008-0.01
□ 0.006-0.008
□ 0.004-0.006
▨ 0.002-0.004
▨ 0-0.002

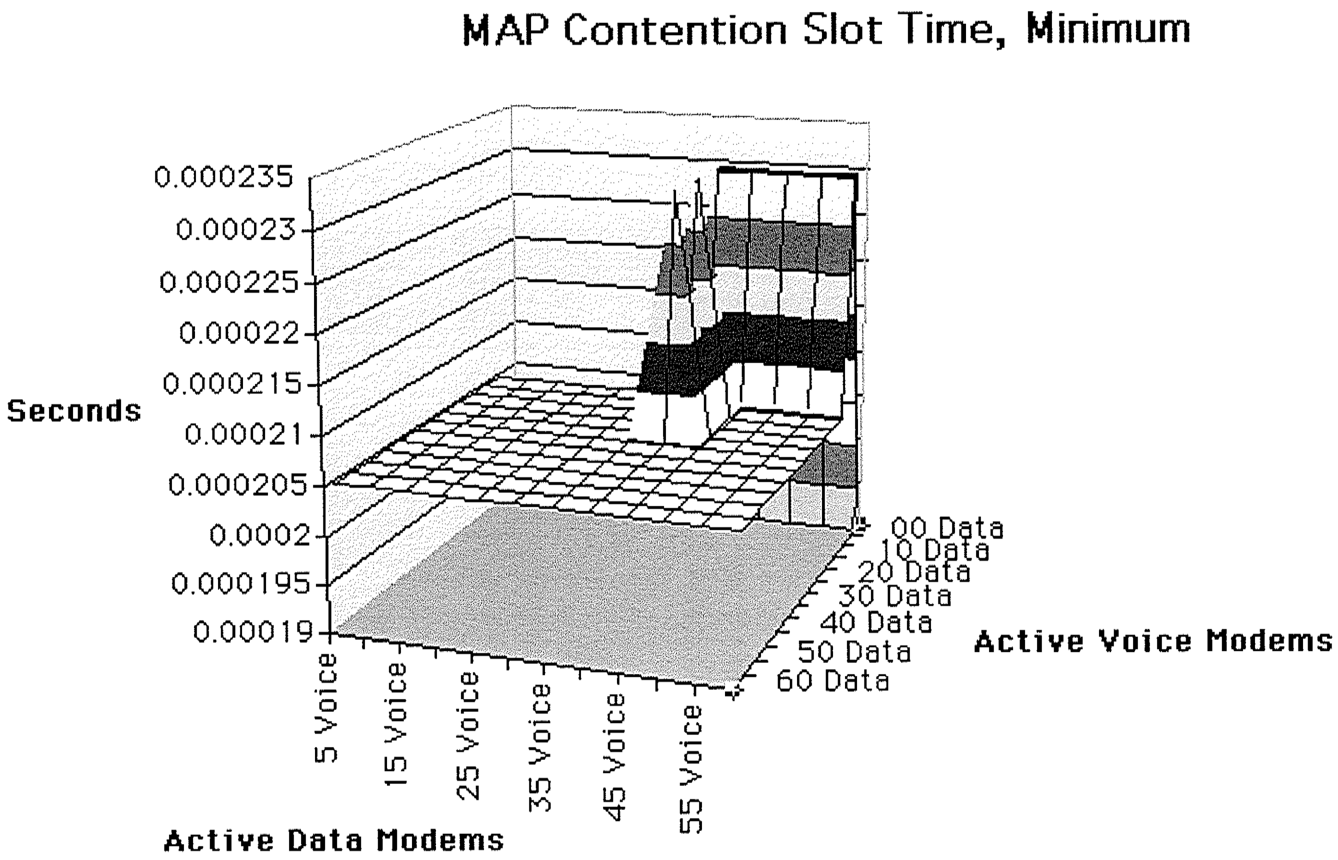
MAP Contention Slot Time, Minimum
Seconds
0.000235
0.00023
0.000225
0.00022
0.000215
0.00021
0.000205
0.0002
0.000195
0.00019
5 Voice
15 Voice
25 Voice
35 Voice
45 Voice
55 Voice
Active Data Modems
00 Data
10 Data
20 Data
30 Data
40 Data
50 Data
60 Data
Active Voice Modems

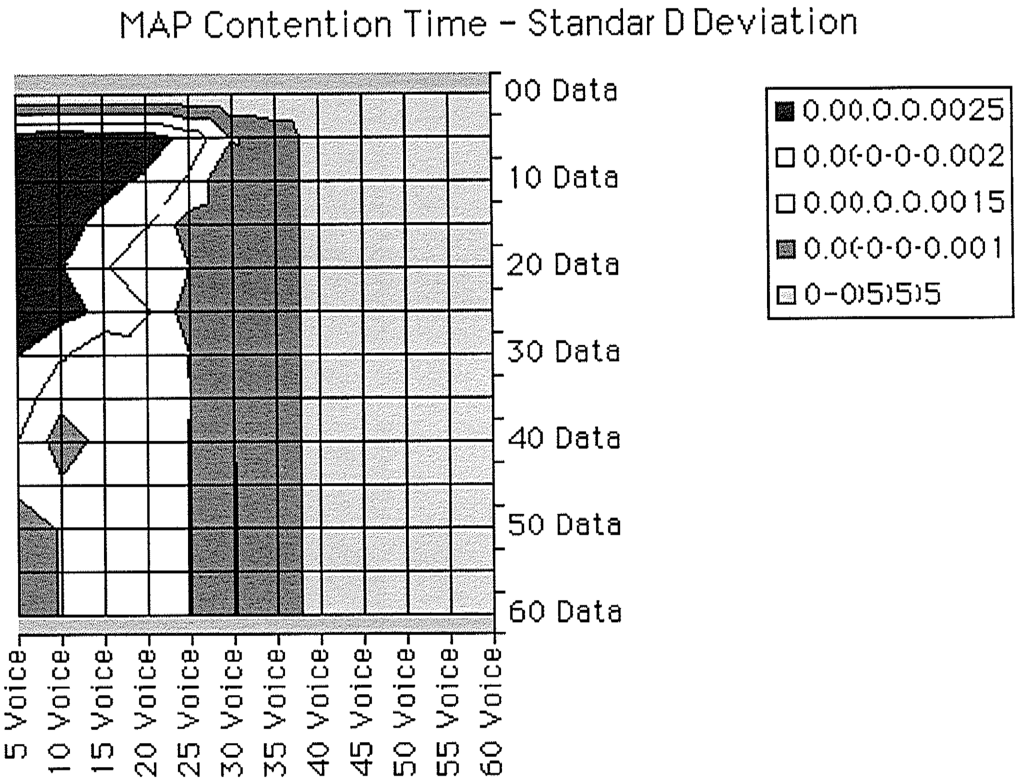

MAP Contention Time - Standar D Deviation

Appendix D

Available CSF13 Parameters

DOCSIS Global Statistics

Access Delay for All Downstream Sources

Access Delay for All Sources

Accxess Delay for Fragmentation-disabled Modems

Access Delay for Fragmentation-enabled Modems

Access Delay for Optional MPEG Source

Access Delay for Optional ON-OFF Source

Access Delay for Optional WWW Source

Access Delay for Priority 0 CMs

Access Delay for Priority 1 CMs

Access Delay for Priority 2 CMs

Access Delay for Priority 3 CMs

Access Delay for Priority 4 CMs

Access Delay for Priority 5 CMs

Access Delay for Priority 6 CMs

Access Delay for Priority 7 CMs

Access Delay for Source Type 1

Access Delay for Source Type 2

Access Delay for Source Type 3

Access Delay for Voice Sources

Global_defer_stat

Map Size

Request Indication

Throughput for All Sources

Throughput for All Downstream Sources

Throughput for Optional MPEG Source

Throughput for Optional ONOFF Source

Throughput for Optional WWW Source

Throughput for Priority 0 CMs

Throughput for Priority 1 CMs

Throughput for Priority 2 CMs

Throughput for Priority 3 CMs

Throughput for Priority 4 CMs

Throughput for Priority 4 CMs

Throughput for Priority 5 CMs

Throughput for Priority 6 CMs

Throughput for Priority 4 CMs

Throughput for Priority 5 CMs

Throughput for Priority 6 CMs

Throughput for Priority 7 CMs

Throughput for Source Type 1

Throughput for Source Type 2

Throughput for Source Type 3

Throughput for Voice Sources

window_size_stat

DOCSIS Local Statistics

Contention Slot Time in Map

Encoder Average Bit Rate

Encoder Average Packet Rate

Encoder Instantaneous Bits Sent

Encoder Instantaneous Bit Rate

Encoder Instantaneous Packets Sent

Encoder Total Bits Sent

Encoder Total Packets Sent

Number of Data Ack IEs in MAP

Number of Grant IEs in MAP

Number of IEs in MAP

Number of Pending IEs in MAP

Number of Request IEs in MAP

Number of Voice Sessions

Packetizer Number of Voice Frames Per Packet

Packetizer Packet Size

Packetizer Packing Latency

Request Queue Length

Request Subqueue Length

Reserved Slot Time in MAP

Time Described in MAP

Time Between MAPs

Time Voice Source Active

Upstream Channel Collision Status

Voice Duration of Last Session

Voice Speech Level[3]

---

[3] It's the end of the world as we know it and I feel fine – R.E.M./Document