

8-2007

PERFORMANCE OF VOIP SERVICES ON A DOCSIS NETWORK TARGETED BY A DENIAL OF SERVICE ATTACK

Benjamin Sangster

Clemson University, benjamin.sangster@gmail.com

Follow this and additional works at: https://tigerprints.clemson.edu/all_theses



Part of the [Computer Sciences Commons](#)

Recommended Citation

Sangster, Benjamin, "PERFORMANCE OF VOIP SERVICES ON A DOCSIS NETWORK TARGETED BY A DENIAL OF SERVICE ATTACK" (2007). *All Theses*. 182.

https://tigerprints.clemson.edu/all_theses/182

This Thesis is brought to you for free and open access by the Theses at TigerPrints. It has been accepted for inclusion in All Theses by an authorized administrator of TigerPrints. For more information, please contact kokeefe@clemson.edu.

PERFORMANCE OF VOIP SERVICES ON A DOCSIS NETWORK TARGETED BY
A DENIAL OF SERVICE ATTACK

A Thesis
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Computer Science

by
Captain Benjamin F. Sangster
August 2007

Accepted by:
James Martin, Committee Chair
Harold Grossman
Mark Smotherman

ABSTRACT

Over 48 million end users worldwide utilize cable modems as their means of accessing the Internet at high speeds. The United States accounts for 54% of those users. Networks which provide access via cable modems utilize Data Over Cable Service Interface Specifications (DOCSIS) as their means of network management. As availability to the Internet increases (especially at high speeds supported by broadband access), so does the opportunity for malicious activity against users utilizing the Internet. Denial-of-service (DoS) attacks are one form of malicious activity and one of the most common. In commonplace Ethernet-based wired networks, a DoS attack requires relatively high levels of computing and network resources to successfully deny service. In DOCSIS-based networks, high levels of computing and network resources aren't mandatory in order to sufficiently degrade a network segment, especially when the objective of the attack is to reduce the quality of Voice over Internet Protocol (VoIP) sessions. This phenomenon hinges on the Media Access Control layer protocol employed by DOCSIS used for managing access to the upstream transmission medium. Utilizing NS, a discrete event network simulator, we define and analyze a DoS attack that specifically targets DOCSIS-based networks. The attack consumes a small portion of the downstream bandwidth available over a cable network but can severely impact upstream performance. While the DoS attack can have any objective, we focus on an attack on best effort VoIP sessions. The implications of this phenomenon are widespread as end users looking for cost-saving voice telecommunications services migrate to best effort VoIP such as provided by Vonage. The contribution of this research is the formulation of

a DoS attack that exploits the relatively inefficient upstream channel in a DOCSIS system and analysis of the attack which explores the impact of the two attack parameters on VoIP performance. Those two attack parameters are the number of nodes attacked and the frequency at which each node is attacked.

DEDICATION

A special thanks to wife, Janine, and my son, Tucker. Without your love and support, this thesis would not have been possible.

ACKNOWLEDGMENTS

A special thanks to Brigadier General Dennis L. Via whose mentoring and help made it possible for this degree and future assignment. I am also grateful for Dr. Jim Martin, Dr. Harold Grossman, and Dr. Steve Stevenson who made my graduate experience a memorable one.

TABLE OF CONTENTS

	Page
TITLE PAGE	i
ABSTRACT	iii
DEDICATION	v
ACKNOWLEDGMENTS	vi
LIST OF FIGURES	ix
LIST OF TABLES	xi
CHAPTER	
1. INTRODUCTION	1
2. BACKGROUND	4
2.1 MAC Overview	4
2.2 Relevant Technologies and Protocols	4
2.3 DOCSIS	8
2.4 Voice over Internet Protocol	12
2.5 Denial-of-Service Attacks	13
2.6 Security	21
3. THE ATTACK DEFINED	25
4. METHODOLOGY	29
4.1 Overview	29
4.2 Simulator and Network Configuration	29
5. ANALYSIS	35
5.1 Baseline Results	35
5.2 Impact of Increasing N^A	38
5.3 Impact of Increasing and Decreasing R^A	48

Table of Contents (Continued)

	Page
6. CONCLUSION.....	63
APPENDIX.....	64
REFERENCES	65

LIST OF FIGURES

Figure		Page
2.1	DOCSIS Protocol Stack.....	9
2.2	DOCSIS Bandwidth Allocation Map.....	11
2.3	Deauthentication Attack.....	16
2.4	SIFS/DIFS Attack.....	20
3.1	Attack Defined.....	27
4.1	Network Map.....	31
5.1	BASELINE: Mean Opinion Score.....	35
5.2	BASELINE: US & DS Utilization.....	36
5.3	BASELINE: Piggybacking - Contention - Concatenation.....	36
5.4	ATTACK: Aggregate Downstream Bandwidth.....	39
5.5	ATTACK: Aggregate Upstream Bandwidth.....	39
5.6.1	Mean Opinion Score.....	41
5.6.2	ATTACK: Collision Rate.....	42
5.6.3	VOIP: Collision Rate.....	42
5.7	ATTACK: DS & US Utilization.....	44
5.8	ATTACK: Piggybacking - Contention - Concatenation.....	45
5.8.1	VOIP: Piggybacking - Contention - Concatenation.....	46
5.9	ATTACK: Piggybacking.....	46
5.10	ATTACK: Collision Rate Ratio.....	47

List of Figures (Continued)

Figure		Page
5.11	ATTACK: Concatenation	47
5.11.1	ATTACK: Mean Opinion Score	49
5.11.2	ATTACK: Mean Opinion Score	49
5.12	ATTACK: Aggregate Downstream Packet Bandwidth	50
5.13	ATTACK: Aggregate Downstream Packet Bandwidth	51
5.14	ATTACK: Aggregate Upstream Packet Bandwidth	52
5.15	ATTACK: Aggregate Upstream Packet Bandwidth	52
5.16	Downstream Utilization	53
5.17	Downstream Utilization	54
5.18	Upstream Utilization	54
5.19	Upstream Utilization	55
5.20	ATTACK: Collision Rate	56
5.21	ATTACK: Collision Rate	56
5.22	ATTACK: Contention Request Ratio	57
5.23	ATTACK: Contention Request Ratio	58
5.24	ATTACK: Piggybacking	59
5.25	ATTACK: Piggybacking	59
5.26	ATTACK: Concatenation	60
5.27	ATTACK: Concatenation	60

LIST OF TABLES

Table		Page
5.1	Baseline Configuration	37
5.2	Piggybacking.....	61
5.3	Piggybacking.....	61
5.4	Contention Request Ratio	61
5.5	Contention Request Ratio	62
5.6	Concatenation	62
5.7	Concatenation	62

INTRODUCTION

In the mid-1990's, the cable industry launched an effort to create a set of Data Over Cable Service Interface Specifications (DOCSIS). The goal was to provide a set of standards in which cable modems and associated hardware could be engineered and manufactured by various companies while maintaining interoperability among one another. In March 1997, the main specification work for DOCSIS 1.0 was completed.

Since the inception of DOCSIS 1.0, the cable broadband access market has witnessed unprecedented growth. In February 2007, U. S. broadband penetration reached 80.16% among active Internet users [“U. S. Broadband Penetration Breaks 80% Among Active Internet Users,” 2007]. Of those users, 41% utilize cable modems as their means of obtaining high-speed Internet access [“DSL overtakes Cable in the U. S.,” 2006]. As the availability and usage of high-speed broadband access increases, so does the demand for broadband applications, such as Voice over Internet Protocol, also increases. VoIP usage is projected to reach 12.1 million subscribers by 2009 [Meckler, 2004]. While the majority of telephony service provided by cable service providers generally uses the DOCSIS QoS mechanisms (and is therefore isolated from a DoS attack), a growing amount of best effort VoIP is also utilized. The driver for best effort telephony from companies such as Vonage is cost.

With the number of households that utilize broadband access and VoIP services reaching such remarkable levels, the opportunity for malicious activity against those households, at a minimum, increases at the same rate. Malicious activity occurs in many different forms. In this study, we focus on malicious activity which inhibits authorized

users from utilizing network resources and services. This form of malicious activity is referred to as a denial-of-service (DoS) attack. In wired networks such as switched Ethernet, a DoS attack would need a node or group of nodes capable of producing sufficient levels of network traffic to saturate the network and successfully deny service to users of that network segment.

In this thesis we show that the media access control (MAC) layer protocol used in DOCSIS cable systems make it possible for a DoS attack to successfully degrade a network without requiring large amounts of malicious network traffic. This is especially true when the objective of the attack is to reduce the quality of Voice over Internet Protocol (VoIP) sessions. Utilizing NS, a discrete event network simulator, we define and analyze a DoS attack that specifically targets DOCSIS-based networks. The attack consumes a small portion of the downstream bandwidth available over a cable network but can severely impact upstream performance. This non-intuitive result is possible in moderately congested DOCSIS networks. The attack “chokes” subscribers upstream bandwidth by consuming upstream contention request slots.

The objective of the DoS attack is to reduce the quality of latency sensitive applications such as VoIP. To effectively achieve the attack objectives, a specific number of nodes targeted by the DoS attack at a given intensity (or attack rate) will result in network performance that restricts VoIP service on a simulated DOCSIS network segment. The implications of this phenomenon are widespread as end users looking for cost-saving voice telecommunications services migrate to VoIP.

In the research presented in this thesis, we define a DoS attack that the following properties:

1. An attacking node located outside of the DOCSIS cable network requires a small amount of downstream bandwidth.
2. The attack has an optimal point that minimizes the downstream bandwidth consumed but maximizes the damage to the target network. Beyond this point, if the rate of attack packets is increased, performance in the DOCSIS network might actually improve as the attacked nodes might take advantage of piggybacking or concatenation as the network becomes more congested.

BACKGROUND

2.1 MAC Overview

The media access control layer is a sub layer of the data link layer specified by the Open Systems Interconnection Reference Model (OSI Reference Model). This sub layer provides addressing and channel access control mechanisms which enable multiple nodes on a network to communicate. MAC protocols are the foundation for network architectures and significantly effect the performance of higher level protocols such as File Transfer Protocol (FTP), Hyper Text Transfer Protocol (HTTP), Transfer Control Protocol (TCP), and Internet Protocol (IP) [Peyravi, 1999].

2.2 Relevant Technologies and Protocols

2.2.1 ALOHA

ALOHA, also known as ALOHAnet, was a technology developed at the University of Hawaii in 1970. Its purpose was to connect various campuses of the University spread across the physically separate Hawaiian islands creating a network capable of sharing information. The original implementation utilized a hub-star configuration. The hub broadcast packets on the outbound channels to the client stations. The client stations transmitted their data on the inbound channel to the hub. The hub would then retransmit the data it successfully received. Client stations listened to see if their transmission was successful. If it was not, the client station waited a short period and attempted to retransmit. This mechanism addressed the issue of two client stations transmitting at the same time resulting in a collision and subsequent corrupted data.

ALOHA is important because, just like Switched Ethernet discussed in section 2.2.2, it utilized a shared transmission medium [ALOHAnet, 2007].

Several versions of ALOHA have evolved since its inception in 1970. Two versions important to this thesis are Reservation Aloha (R-Aloha) and Aloha Reservation (Aloha-R). R-Aloha is the simple form of reservation protocols that is based on distributed contention. Stations transmit in time slots with successful transmission resulting in implicit reservation of future time slots corresponding to the slot successful transmission occurred. Time slots remain reserved for the same station as long as data remains to be sent. Initial access to the transmission medium is random. Aloha-R is a distributed contention-oriented reservation protocol that utilizes an explicit reservation mechanism. An Aloha-R based frame is divided into equal length time slots. One of the time slots is further divided into mini-slots which are used by stations to request reserved data slots [Peyravi, 1999].

2.2.2 802.3 (Ethernet)

Ethernet, also known as IEEE 802.3, is a network technology that enables multiple stations to communicate over a shared, wired transmission medium. Original implementations of Ethernet utilized coaxial cable transmitting at speeds of 3 Mbps. As Ethernet has advanced over the years, twisted-pair and fiber optic cable have replaced coaxial cable and transmission speeds have eclipsed 1 Gbps. One aspect of Ethernet that has remained consistent is its frame format which has enabled the wide range of Ethernet implementations to communicate amongst each other.

Ethernet's distinctive characteristic which enables multiple stations to communicate on the same physical transmission medium is known as Carrier Sense Multiple Access with Collision Detection (CSMA-CD). When a station on an Ethernet network needs to transmit, it follows the following algorithm:

1. Ethernet frame ready to transmit.
2. Is the transmission medium idle?
3. If yes, begin transmitting.
4. If no, wait for the transmission medium to become idle and then wait the inter frame gap (varies among implementations).
5. Continue to monitor transmission medium to determine if collision occurs.
6. No collision, end successful transmission.
7. Collision occurs, implement #4 (collision detection procedure).
8. Collision Detection Procedure:
9. Continue transmitting current transmission to enable all stations the opportunity to detect collision.
10. Has the maximum number of transmission attempts been reached?
11. If yes, abort transmission.
12. If no, determine random back-off interval and wait that amount of time before retransmitting.
13. Return to #1 and attempt to retransmit.

Another characteristic of Ethernet which is important to our study is the use of 48-bit addresses to uniquely identify stations on an Ethernet network. This unique addressing enables stations to identify both the source and destination of packets transmitted. Having the source and destination of each packet identified provides Ethernet networks the capability to more efficiently route packets to the specific destination instead of each station on the network checking the packets. Ethernet networks can also use this addressing scheme to provide an additional layer of security to

networks by preventing certain address ranges from injecting traffic into a network segment or identifying a specific range of addresses that may be attempting to maliciously effect a network.

2.2.3 802.11 (Wireless)

Wireless networks, also known by the protocol they are based on 802.11, provide network users the ease of mobility without the hassles of wires and the physical limitation of wires. One could say that 802.11 networks give users the mobility and flexibility that wired networks inhibits. The 802.11 wireless network standard accomplishes this by utilizing radio broadcasts operating in the industrial, scientific, and medical (ISM) bands of the radio spectrum. Specifically, the 2.4-GHz ISM band and the 5-GHz band are utilized by the 802.11 standards. Within those bands, government regulations constrain the power that can be emitted by 802.11 technologies that utilize those radio bands.

802.11 networks are comprised of four primary physical components. Those four components are stations, access points (AP), wireless medium, and distribution systems. Stations are computing devices that enable users to transfer data between one another via wireless network interfaces. Devices called “access points” perform wireless-to-wire bridging functions which convert frames on an 802.11 network to another type of frame for delivery to the rest of the world. AP's perform a number of other functions, but bridging is considered to be the most important.

In order to move data from station-to-station on an 802.11 network, the standard uses a wireless medium. Several physical layers are defined for the wireless medium. Two radio frequency (RF) layers and one infrared (IR) layer were initially defined with

the RF layers experiencing wider use. When several AP's are joined together to form one large coverage area, each AP must communicate with the other AP's in that coverage area to track the movement of stations from one AP to the next. The distribution system is the logical portion of the 802.11 network that forwards the data/frames from the sending station to the receiving station. No specific technology is defined by the 802.11 standard for use in a distribution system. In most commercial uses, some form of bridging engine along with the distribution system medium is utilized to transmit data/frames between AP's. The most common term for this part of the network is the backbone network. The most common technology utilized as the backbone network in 802.11 distribution systems is the Ethernet technology.

2.3 DOCSIS

2.3.1 DOCSIS History and Overview

The cable systems that DOCSIS was created for consisted of a head end, transmission medium, cable modem termination system (CMTS), and cable modems (CM). The head end was where bidirectional frequency division multiplexed (FDM) signals originated. Those multiplexed signals would then travel over coaxial cable to cable modem termination systems. Eventually, coaxial cable was replaced by fiber optic cable between the head end and the cable modem termination systems. Once the FDM signal reached the cable modem termination system, it was passed onto a bidirectional bus architecture network capable of supporting multiple cable modems. From the cable modem termination system to the cable modems, coaxial cable was used as the transmission medium. Upstream data (from cable modems to cable modem termination

system) utilized higher frequencies for transmission while downstream data (from cable modem termination system to cable modems) utilized lower frequencies.




Figure 2.1. The DOCSIS cable modem protocol stack. The physical layer is where modulation of the signal occurs. The cable modem termination system adds framing using MPEG-2 transmission convergence protocol enabling multiple services to be sent on the same channel. The MAC layer is where access to the upstream path is managed.

2.3.2 DOCSIS Protocol Stack

Figure 2.1 shows a protocol stack as related to DOCSIS at each layer of the OSI model. The first four layers are DOCSIS specific. The higher-level protocols (TCP, IP, UDP, etc.) are carried by DOCSIS layers across the cable network and are used for communications with the Internet [Fellows, 2001].

2.3.3 DOCSIS MAC Layer

The DOCSIS media access control (MAC) layer is the focus of this study. The MAC layer controls access to the upstream channel of the transmission medium for all cable modems. Using the standard client-server model, the upstream channel is the network path that carries traffic generated from cable modems (clients) to the cable modem termination system (server). In order for one cable modem to communicate with another, access has to be granted by the cable modem termination system for that cable modem to place data on the wire. Even if the destination cable modem is located on the same local network as the sending cable modem, the cable modem termination system has to grant access. Using the client-server model again, the sending cable modem, after access has been granted, transmits the data to the cable modem termination system (server) which will then transmit the data back down the downstream channel to the destination cable modem (client).

The request/grant mechanism is implemented via a bandwidth allocation map (MAP). Figure 2.2 shows the basic format of a DOCSIS MAP frame. The contention slots are used by the cable modems to request access to the upstream channel from the cable modem termination system. The data slots are where cable modems insert data after access has been granted by the cable modem termination system to the requesting cable modem. The maintenance slots are used for initialization and synchronization with the channel when a cable modem powers on, and periodically to maintain timing.




Figure 2.2. DOCSIS bandwidth allocation map (MAP). The contention slots are used by the cable modems to request access to the upstream channel from the cable modem termination system. The data slots are used after access has been granted to a cable modem to transmit data. The maintenance slots are used for initialization and synchronization when a cable modem powers on.

When a cable modem has data to send, it must wait for the cable modem termination system to send a MAP message. It must wait for the MAP because the cable modem termination system has not granted that cable modem access to the channel. The requesting cable modem will then utilize one of the contention slots of the MAP to request a grant from the cable modem termination system to transmit its data. It should be noted that the contention slot portion of the MAP can be accessed by any cable modem on the local network at any time. Therefore, collisions can occur when two or more cable modems request access via the same contention slot. If this occurs, the cable modems that experienced the collision will back off a random interval before they attempt to send another request. Once the cable modem termination system receives requests, it will notify the requesting cable modems in a subsequent MAP which mini-slots they have been granted access to for data transmission. This guarantees an interval in which cable modems can transmit collision-free. In order to maximize the available bandwidth in DOCSIS-based networks, DOCSIS also allows a bandwidth request to be piggybacked on previously granted data slot. This alleviates the requesting cable modem from having to

wait for the next MAP to request bandwidth and improves the efficiency of the upstream network channel.

DOCSIS also utilizes a mechanism referred to as concatenation in order to maximize available bandwidth. Concatenation allows for a cable modem to combine several smaller packets and transmit those combined packets as if they were one. The greatest performance improvement from this mechanism is observed in TCP throughput. When concatenation is utilized, multiple TCP ACK packets can be combined (i.e. concatenated). This is possible due to the smaller size of TCP ACK packets compared to most other packets. Rather than separate transmission of individual TCP ACK packets, multiple ACKs can be sent in the same DOCSIS frame, maximizing downstream throughput.

Two other messages that are transmitted by the cable modem termination system on the downstream channel that cable modems look for are the upstream channel descriptor (UCD) message and the time synchronization message (SYNC). The UCD provides the necessary information to the cable modem to determine if its capabilities (i.e. frequency range, modulation types, symbol rates, etc.) match that of the upstream channel it is attempting to access. The SYNC provides common timing for all modems to reference [Fellows, 2001].

2.4 Voice over Internet Protocol

Voice over Internet Protocol (VoIP) is the routing of voice communication traffic over the Internet or any Internet Protocol (IP) based network. The two major competing standards for VoIP are the Internet Engineering Task Force (IETF) standard Session

Internet Protocol (SIP) and the Internet Telecommunication Union (ITU) standard H.323. Initially H.323 was the most popular protocol, however it has since been surpassed by SIP. This was primarily due to the latter's better traversal of network address translation and firewalls, although recent changes introduced for H.323 have removed this advantage [“Voice over IP,” 2007].

2.5 Denial-of-Service Attacks

Denial-of-service (DoS) attacks are attempts by a malicious user or group of users to render a computer network, system, service, or resource unavailable to its intended users. The motive for launching such attacks varies, but the ultimate end-state is the inability of legitimate users to conduct normal business due to the unavailable resource. There are three basic types of DoS attacks [“Denial-of-Service attack,” 2007]:

1. Consumption of computation resources, such as bandwidth, disk space, or CPU time.
2. Disruption of configuration information, such as routing information.
3. Disruption of physical network components.

In a wired network environment, a common DoS attack is a ping flood attack. A ping flood DoS attack overwhelms the targeted system or network with Internet Control Management Protocol (ICMP) Echo Request packets (ping). In order for this attack to be effective, the attacker must have a network connection with greater capacity than the target network or system. For example, an attack launched from a Fast Ethernet network against a network or system utilizing a DSL connection would be effective. The Fast Ethernet network provides a maximum capacity of 100 Mbps. The standard DSL connection has a downstream capacity of 30 Mbps, but only a 5.12 Mbps upstream

capacity. A ping flood attack from the Fast Ethernet segment would be capable of consuming both the downstream and upstream capacity of the DSL connection.

Network and system administrators can defend against ping flood DoS attacks. Deployment of a firewall can limit or completely deny ICMP echo requests from accessing a network or individual system. This addresses the threat of ping flood DoS attacks, but simultaneously inhibits the monitoring of latency by legitimate users (latency of a network can be observed utilizing ICMP echo requests).

Denial-of-service attacks can be directed at wireless networks just as easily as they can at wired. At the application and transport layer, the attacks are carried out in the same manner. The differences of DoS attacks focused at wireless mediums versus wired can be found at the network, MAC, and physical layers. DoS attacks at the 802.11 MAC layer can be categorized into two vulnerability categories: identity and media-access control. Identity vulnerabilities consist of attacks that manipulate the deauthentication, disassociation, and transmit power control services. Media access control vulnerabilities consist of attacks that don't directly manipulate network services provided by the 802.11 standard, but directly attack the 802.11 protocol.

Stations in an 802.11 network implicitly trust the source address provided by any station it receives a message from. This implicit trust is the framework for the deauthentication and disassociation DoS attacks in 802.11 networks. A malicious station can spoof a valid station's address and manipulate the deauthentication and disassociation services. When a station joins an 802.11 network, it must associate itself to an AP within the network. Prior to association, the AP must authenticate that the station is indeed an

authorized user of the network. The station must send an authentication request to the access point. The AP will respond with an authentication response, validating the request and permitting the requesting station to continue with the association process. The station then sends an association request to the access point. The AP will respond with the association response message, completing the association process. The station is now authorized and capable of sending traffic on the network. A deauthentication attack is possible as soon as the association response is sent by the AP to the requesting station. A malicious station, “listening” to the authentication and association messages, spoofs the valid stations MAC address. It creates a deauthentication message using the spoofed address, sending the message to the access point. Once the AP receives the message, it will respond with a verification of deauthentication. At this point, the valid station is no longer authenticated, and subsequently, not associated to the network. It will not be able to send data on the network until it reauthenticates and reassociates with the access point.




Figure 2.3. Message traffic in a deauthentication attack. A client station sends an authentication request. An access point sends an authentication response. The client station then sends an association request. The access point sends the association response. A malicious station at some point after association sends a deauthentication response, spoofing the valid client's address. The access point sends the deauthentication response. During the invalid deauthentication series, the client station attempts to send data resulting in an unsuccessful transmission due to the deauthentication initiated by the malicious station.

The disassociation attack takes advantage of a similar vulnerability as the deauthentication attack. As stated previously, 802.11 inherently trusts the source address of all traffic it receives. Just as the malicious station in the deauthentication attack spoofs the address of a valid station, a malicious station can do the same and initiate the disassociation attack. The distinctive difference between the two attacks is the number of messages that are required for a wrongly disassociated station to reassociate. In the

deauthentication attack, it takes four messages for the targeted station to reauthenticate to the network. When a station is successfully targeted by a disassociation attack, it may only take two messages to reassociate to the AP (subsequently, the network).

The 802.11 standard supports clients to enter a power saving mode in order to save energy. The client enters a sleep state where messages can neither be sent to the client nor received. Prior to entering the sleep state, the client announces to the AP that it intends to enter the sleep state so that the AP can begin buffering all traffic destined for the client. While the client is sleeping, it will occasionally “wake up” and poll the AP for any buffered traffic. Once the AP has delivered the buffered traffic to the sleeping client, it discards the data in the buffer. The AP provides a synchronization message that keeps the sleeping clients synchronized by sending a broadcast message identifying which clients have buffered traffic.

One form of a DoS attack via vulnerabilities created by the power save option is a malicious station spoofing the polling message of the sleeping client. A malicious station could contact the access point, masquerading as the sleeping client, and poll the AP for buffered traffic. The AP would trust that the poll message is truly from the sleeping client, supposedly deliver the traffic, and subsequently discard the traffic under the assumption that the traffic was correctly delivered to the sleeping client. When the sleeping client awakens and polls the access point, the AP will no longer have the buffered traffic it discarded and the client will not receive the traffic originally intended for it. The client could then return to the sleep state, allowing for this attack to continue

as long as the malicious station continues to execute it and the client remained in the power save state.

A second form of a power save option enabled DoS attack is a malicious station spoofing the broadcast message that is sent by the AP which identifies which stations have buffered traffic. The broadcast message is known as the traffic indication map (TIM). By spoofing the TIM message, a malicious station may convince a client that there is no buffered traffic for that station when in truth there is. The client, thinking there is no buffered traffic, returns to the sleep state without receiving the buffered traffic. Although the buffered traffic is not lost, there is potential for the access point's buffer to reach capacity producing unwanted results (i.e. dropped message traffic).

The third form of power save option enabled DoS attack is again spoofing the TIM message. This time, the malicious station can modify the synchronization information provided by the TIM so that the clients that receive this message will fall out of synchronization with the actual access point, subsequently not waking up at the appropriate time. Just as the previously mentioned TIM attack, a potential negative result of this is the capacity of the AP's buffer maxing out.

As network traffic increases, the performance of that network tends to decrease. One of the reasons for the decrease in network performance is the collisions that occur and the protocols implemented to deal with those collisions. 802.11 networks are no different than any other network standard. Great efforts are made to avoid collisions. Unfortunately, a problem that is encountered frequently in 802.11 networks is the hidden node problem. In order to appropriately address the hidden node problem, a combination

of physical carrier-sense and virtual carrier-sense mechanisms are employed together to manage access to the communications channel. Both of these mechanisms may be exploited by an attacker [Bellardo, 2003].

The physical carrier-sense mechanism employed by 802.11 networks breaks the separates the communications channel into four time windows. For the purpose of this paper, we will only discuss the first two time windows which are the Short Interframe Space (SIFS) and the Distributed Coordination Function Interframe Space (DIFS). Prior to any frame being sent onto the channel, the sending station must observe the channel and ensure that no traffic is being transmitted during one of the time windows. The SIFS window is for frames sent as part of a preexisting frame exchange [Bellardo, 2003]. The DIFS window is for stations who wish to initiate a new frame exchange. The period following the DIFS is subdivided into slots to in order to avoid multiple stations from transmitting as soon as the DIFS window expires. The transmitting stations randomly select which slot they will transmit in with equal probability of selecting any slot. If a collision occurs, the sending station utilizes an exponential backoff algorithm before retransmitting [Bellardo, 2003].

A malicious station has the potential to monopolize a communications channel if that station sent a short signal at the end of every SIFS window. This creates a denial-of-service to all stations on the channel. The 802.11 contention algorithm is a dual persistence algorithm. A station wishing to transmit data must wait the equivalent of two DIFS windows before it can transmit. If during that window a transmission is sensed from another station, the station wishing to send traffic must back off the amount of time

determined by the algorithm. So, a malicious station sending a short signal towards the end of every SIFS period would ultimately monopolize the channel, forcing all other stations to back off until the attack was over. Although this attack accomplishes what the malicious station wants, it does so with a price. Since the SIFS window is only 20 μ s long, the malicious station would have to transmit its signal approximately 37,000 times per second to occupy the channel which is not very efficient.




Figure 2.4. Graphical depiction of the virtual carrier-sense attack in action. The gradient portion of the attacker's frame indicates time reserved by the duration field although no data is actually sent. Continually sending the attack frames back to back prevents other nodes from sending legitimate frames [Bellardo, 2003].

The network allocation vector (NAV) is a value maintained by each station on an 802.11 network that identifies a time period that a station will not attempt to access the communications channel despite the channel being assessed as available (empty of traffic). Each 802.11 frame contains a duration field that identifies the number of microseconds that the channel is reserved [Bellardo, 2003]. A station will not attempt to

transmit until its NAV reaches zero. The request-to-send (RTS)/clear-to-send (CTS) communication exchange that takes place to synchronize two stations attempting to communicate on the 802.11 channel explicitly utilizes the NAV to address any hidden stations that may be interfering with transmissions on the channel.

The RTS/CTS handshake (which is the virtual carrier-sense mechanism) implemented by 802.11 to address hidden stations creates a vulnerability that allows a malicious station to modify the duration field in an 802.11 frame, making the value in the duration field extremely large. By doing so, the malicious station prevents clients who adhere to the virtual carrier-sense mechanism for channel control from accessing the channel. A malicious station has the option of using any frame for its attack, but it is in its best interest to utilize an RTS frame since most nodes will always respond to an RTS frame with a CTS. By influencing a good station to respond with a CTS frame, the malicious station has reduced the amount of resources it has to utilize to execute the attack since the station responding with the CTS will propagate the attack for the malicious station. In comparison to the SIFS monopolization attack, a malicious station only has to transmit 30 times a second due to the NAV's maximum value (32,767 which is approximately 32 milliseconds).

2.6 Security

In this thesis, when the topic of security is discussed we are referring to security issues as related to protocol implementation rather than physical security of a network. Physical access of modern networks is relatively simple given the necessary resources. Network and security administrators can easily implement extremely strict or lax security

procedures. What is not easily addressed is the underlying security issues found within the protocols that manage how stations that already have physical access to a network access the transmission medium.

2.6.1 802.3 (Switched Ethernet)

Securing Switched Ethernet entails applying limited access to data packets. Since Ethernet broadcasts data packets to all stations on its network segment, all stations are physically capable seeing those packets. With a properly implemented security mechanism, all stations can still physically see the data packets but are not capable of reading or understanding them. Such a security mechanism is referred to as encryption. A drawback of applying encryption to any network is the additional overhead in both packet size and processing.

2.6.2 802.11 MAC Layer

Several versions of DoS attacks that can be experienced at the MAC layer in 802.11 networks were discussed in section 2.5. The number one countermeasure that could be implemented is the explicit authentication of management/control frames. The lone attack that this countermeasure would not be successful against is the SIFS monopolization which does not rely on the modification or spoofing of management/control frames. The deauthentication, disassociation, and NAV attacks have additional countermeasures that can be implemented beyond explicit authentication. The lone countermeasure for the power save option is explicit authentication. Although countermeasures have been identified for these attacks, the sheer numbers of devices that would require modifications has hindered any attempt to implement the countermeasures.

A proposed countermeasure for both the deauthentication and disassociation attacks is to delay the AP's response to such requests. By delaying the AP's response and subsequently having the AP monitor inbound traffic from the alleged station requesting deauthentication or disassociation, an AP can determine whether or not a station truly wishes to deauthenticate or disassociate. If the AP receives inbound traffic after a deauthenticate or disassociate request is received, then the AP knows that the request is from a malicious station (since the order of receipt is not correct). Subsequently, the AP would ignore the request.

There are two drawbacks to this countermeasure. An additional vulnerability is created and is observed when a station moves from one AP's BSA to another. Due to the imposed delay, packets may not be properly routed to the appropriate AP since the old AP may still consider the station associated with it. The second drawback is the malicious station could take advantage of the delay when a targeted station truly is moving from one BSA to another. The malicious station could continue to spoof the mobile station keeping the association with the spoofed AP valid.

A proposed countermeasure for the NAV DOS attack is to place a maximum allowable value for the duration field of the 802.11 frame. This would keep the valid stations from being wrongfully denied access to the medium from an invalid duration value. Although this countermeasure addresses the attack, it does not completely alleviate it. All a malicious station would need to do is increase its transmission from over 30 packets per second to 90 packets per second. By doing so, denial of service will

be achieved on the network. Again, the true countermeasure for this attack is explicit authentication that would guard against modification of the duration field.

Why does explicit authentication not effectively counter the SIFS monopolization DOS attack? The SIFS attack does not rely on the spoofing of addresses in order to modify management/control frames to deny service to a network and its users. A malicious station simply has to transmit at the end of the SIFS window, subsequently forcing all other stations wishing to transmit to exponentially back off. The attack is a result of the prioritization and ordering standardized by the virtual carrier-sense mechanism implemented by the 802.11 standard. In order to counter this attack, the behavior of stations waiting to access the communications channel would have to be changed from the current behavior. Hence, the virtual carrier-sense mechanism would have to be modified.

THE ATTACK DEFINED

The research conducted for this thesis entailed simulating a DoS attack launched against a DOCSIS network segment from a single malicious node located on the Internet side of a cable modem termination system providing connectivity to the DOCSIS network segment. The single node's bandwidth is equal to or greater than the downstream service rate of the DOCSIS network segment. The node does not need explicit authorization to the DOCSIS network segment but does have implicit authorization via available network monitoring capabilities (i.e. the ICMP echo service).

The goal of the DOCSIS attack is to consume the upstream contention slots (refer to figure 2.2) with illegitimate bandwidth requests reducing the availability of contention request slots for legitimate network traffic. As the number of upstream contention slots available for legitimate requests decreases, the average performance experienced by best effort traffic degrades. Specifically, the number of collisions and the average channel access time will increase, the number of packets piggybacked will decrease, and the number of contention requests and concatenated packets will increase.

The DOCSIS attack we define and evaluate in this thesis is a ping flood. Figure 3.1 graphically depicts the architecture of the attack. The parameters for the attack include number of TCP connections, attack rate interval, and number of nodes attacked. As figure 3.1 depicts, a node with Internet access launches a ping flood, DoS attack on a DOCSIS network. The rate at which the node attacks the network is labeled R^A . R^A is defined as the rate at which ICMP echo requests are sent by the attacking node to the targeted nodes under attack. The number of nodes attacked is labeled N^A . N^A is defined

as the number of nodes in which the attacking node has targeted for the DoS attack. These nodes will receive ICMP echo requests at a rate of R^A from the attacking node via the downstream channel. Subsequent ICMP echo replies will be sent from the nodes receiving ICMP echo requests to the attacking node via the upstream channel.

To establish a level of background network traffic, TCP connections are established between cable modems on the DOCSIS network segment and a node on the Internet side of the cable modem termination system. Since most DOCSIS networks assign unused data slots to be used for contention requests, sufficient background traffic must exist to consume the majority of data slots. In other words, the attack is most effective when the network is moderately busy. In a practical implementation of the attack, attack packets will use a spoofed source address




Figure 3.1. Graphical depiction of simulated network general layout. The attack node is located on any network outside of the DOCSIS network. The nodes labeled CM-1 through CM-N do not support any management or monitoring services and are “vulnerable” to the DoS attack. Nodes CM-N+1 to CM-Total support management and monitoring services (TCP traffic replicating varying network nodes, loss monitor, VoIP monitor).

In a typical wired network based on Ethernet technology, a ping flood DoS attack must generate enough network traffic so that all of the available bandwidth is consumed. In a DOCSIS network, we anticipate successful DoS without consuming all of the available bandwidth (both downstream and upstream). The rules employed at the MAC layer in DOCSIS networks create a phenomenon during a ping flood DoS attack where the contention slots are consumed by cable modems that have an ICMP reply packet ready for upstream transmission. This will increase the average channel access time. Non-attack network traffic that relies on contention request for upstream bandwidth will suffer. The required available bandwidth between the attacking node and the victim

nodes is a fraction of the upstream channel capacity. The attack exploits the inefficient upstream data transmission mechanism. As the intensity of the attack increases, the nodes under attack will make use of piggybacking and concatenation which reduces the reliance on contention-based requests and therefore offsets the impact of the attack. The attack has an optimal point that reduces the attack rate while maximizing damage to the network.

METHODOLOGY

4.1 Overview

The research described in this thesis uses an open-source, discrete event simulator called the network simulator or NS to simulate a DOCSIS network under various network loads and attack intensities [NSNAM, 2007]. The simulator configuration is detailed in section 4.2. The simulated model is discussed in section 4.3.

4.2 Simulator & Network Configuration

NS is a discrete event simulator targeted for network research. NS provides substantial support for simulating TCP, routing, and multicast protocols over wired and wireless networks [NSNAM, 2007]. Additionally, a module was added to the base NS program to provide the capability of simulating a network based on the DOCSIS protocol. Validation of this module is detailed by Martin and Westall in “Validating an 'ns' Simulation Model of the DOCSIS Protocol” [Martin, 2006].

The network which this thesis is based on is depicted in figure 4.1. It is actually comprised of two distinct networks. On the left side of figure 4.1 is the wired, wide area network (WAN). It consists of five nodes labeled N2, N3, N4, N5, and L0. On the right side of figure 4.1 is the wired, DOCSIS-based large area network (LAN). It consists of two nodes labeled N1 and L1. There are also 400 cable modems which are connected to node N1. Both networks are connected via node N0 which is a Cable Modem Termination System (CMTS). All of the links between N0 and the nodes on the WAN-side of the network represent the Internet. The link between N0 and N1 represent a private network segment provided by a cable service provider.

The WAN-side link settings are as follows:

Link Type: Duplex
Maximum Capacity: 100 Mbps
Queuing Algorithm: Drop Tail
Propagation Delay: 24 ms

The DOCSIS network settings are as follows:

Downstream Channel Rate: 30 Mbps
Upstream Channel Rate: 5.12 Mbps
Fragmentation: No
Concatenations: Yes
Piggybacking: Yes
Queue Size: 50
Contention Slots: 12
Management Slots: 3
Map Time: .002 seconds
Map Frequency: .002 seconds
Number of Cable Modems: 400




Figure 4.1. Network map and configuration settings used as the test network for this thesis. On the left side of the figure is the WAN-side nodes (Internet-residing nodes). On the right side are cable modems belonging to the private, DOCSIS-based network.

Sets of simulated runs were devised distinguished by combinations of R^A , N^A , and the number of active TCP connections. Five variations based on N^A were observed: 0, 10, 50, 100, and 200 nodes under attack. Each variation of N^A was observed with six variations of R^A : .05, .25, .5, 1, 2, and 4 second intervals between attack packets (longer interval = less intense attack). Within each set, six runs are performed distinguished by the number of active TCP connections: 0, 2, 5, 10, 15, and 20 active TCP connections. The varying of TCP connections enabled us to observe the network at what would be

perceived as average to above average (15-20) and below average (10 or less) loads of non-malicious network traffic.

To monitor the loss present on the network, a loss monitor (L0) is attached to a cable modem on the DOCSIS segment (CM1) with a partner loss monitor (L1) attached to a node on the WAN-side network (N2). These loss monitors exchange packets of size 210 bytes every .02 seconds, maintaining the loss rate of the network according to their packet exchange. The loss monitor also maintains the necessary statistics for calculating the MOS value.

As each set was simulated, nine statistics (focus statistics) were isolated to observe the effects of varying the variation parameters: aggregate downstream (DS) attack packet bandwidth, aggregate upstream (US) attack packet reply bandwidth, mean opinion score value, downstream utilization percentage, upstream utilization percentage, percentage of upstream packets sent via piggybacking, percentage of upstream packets sent via concatenation, percentage of upstream packets sent via contention requests, and collision rate.

The aggregate downstream attack packet bandwidth is measured in bits per second. It is calculated by multiplying the number of attack packets traced inbound to the DOCSIS network segment from the downstream channel by eight (each packet is eight bits in size) and dividing that product by the time stamp of the last attack packet received. This value represents the portion of the downstream channel's maximum capacity that is consumed by the DoS attack.

The aggregate upstream attack packet reply bandwidth is also measured in bits per second. It is calculated by multiplying the number of observed replies from attack packets received that are seen outbound from the DOCSIS network segment on the upstream channel by eight and dividing that product by the time stamp of the last reply transmitted. This value represents the portion of the upstream channel's maximum capacity that is consumed by the DoS attack.

The mean opinion score (MOS) value is a numerical representation (1-5) of how a media transmission's quality is perceived. The following list describes each value:

- 5—Excellent quality, imperceptible impairment
- 4—Good quality, perceptible impairment but not annoying
- 3—Fair quality, slightly annoying impairment
- 2—Poor quality, annoying impairment
- 1—Bad quality, very annoying impairment

Calculation of the MOS is handled by a loss monitor procedure that was added to the NS DOCSIS module. The procedure utilizes formulas for calculating MOS value which are based on computations detailed by Cole in “Voice over IP performance monitoring” [Cole, 2001]. When observing VoIP quality and using MOS values to determine whether a call is of “toll quality” or not, a minimum MOS value of '4' is the telephone industry standard [Miller, 2005]. Depending on which CODEC is used for the observed communications channel, MOS values of 3.6 can be considered toll quality [Keneipp, 2000]. For the purposes of this thesis, any MOS value less than 3.0 is considered less than toll quality.

The downstream utilization percentage is calculated in the NS DOCSIS module code. The code tracks the observed downstream bandwidth consumed, divides that by the available downstream bandwidth, and multiplies by 100.

The upstream utilization percentage is calculated in the NS DOCSIS module code. The code tracks the observed upstream bandwidth consumed, divides that by the available upstream bandwidth, and multiplies by 100.

The percentage of packets sent via piggybacking and concatenation and the ratio of contention requests to total packets sent are all tracked by the NS DOCSIS module. All three percentages are calculated by dividing the number of each type observed by the total number of packets transmitted, then multiplying by 100.


Data to determine the collision rate was captured for three groups of cable modems: all cable modems, only the cable modems under attack, and only the cable modem supporting the VoIP monitor. To calculate the collision rates, the number of collisions observed for the specified group were divided by the total number of frames sent by that group.

The configuration of the baseline set for this simulation is zero nodes under attack with an attack packet interval of .5 seconds. This set was run six times with the number of active TCP connections increasing as detailed in section three above. The observed focus statistics for this set will be used to compare focus statistics gathered during subsequent simulated sets after the variation parameters are changed.

ANALYSIS

5.1 Baseline Results

Figures 5.1, 5.2, and 5.3 depict the focus statistics observed during the simulation of the baseline configuration set. Figure 5.1 depicts the MOS value. Figure 5.2 depicts the downstream and upstream utilization rate. Figure 5.3 depicts the percentage of packets sent upstream via piggybacking, contention requests, and concatenation. Once two or more upstream TCP connections are active, more user data is transferred than periodic management messages. Figure 3 illustrates that this moves the system to use primarily piggybacked requests. Aggregate downstream and upstream attack packet bandwidth is not depicted in a graph due to all occurrence in the baseline configuration resulting in zero for both statistics. Table 5.1 provides actual values of each statistic observed during the simulation of the baseline configuration.



1

Figure 5.1. MOS value observed during simulation of baseline configuration. 400 cable modems, range of active TCP connections a follows: 0, 2, 5, 10, 15, and 20.




Figure 5.2. Downstream and upstream utilization rate observed during simulation of baseline configuration. 400 cable modems, range of active TCP connections a follows: 0, 2, 5, 10, 15, and 20.




Figure 5.3. Percentage of packets sent via piggybacking, contention request, and concatenation observed during simulation of baseline configuration. 400 cable modems, range of active TCP connections a follows: 0, 2, 5, 10, 15, and 20.

BASELINE CONFIGURATION								
Number of TCP Connections	Aggregate US Attack Bandwidth	Aggregate DS Attack Bandwidth	Piggy Back	Contention	Concatenation	MOS	US UTIL	DS UTIL
0	0	0	0.48	99.08	0.44	4.41	5.24	3.23
2	0	0	79.43	17.96	2.61	4.4	49.37	1.76
5	0	0	76.66	15.84	7.5	4.34	50.46	1.82
10	0	0	75.55	12.56	11.89	4.2	50.87	1.96
15	0	0	72.37	10.84	16.79	3.6	51.4	2.11
20	0	0	70.03	9.65	20.33	3.63	51.84	2.25

Table 5.1. Observed focus statistics for simulation of baseline configuration.

5.1.1 No Background Traffic (Zero Active TCP Connections)

Both the aggregate downstream and upstream attack bandwidth produced by the baseline configuration was zero. As shown in Figure 5.1 (with zero active TCP connections), the observed MOS value was 4.41. This value is well above industry standards of good quality (MOS value of 4.0). The downstream utilization was 3.23 percent. The upstream utilization was 5.24 percent. The lower utilization rates are due to network traffic being limited to only management and monitoring traffic. The percentage of packets sent upstream via piggybacking was .48. In order for piggybacking to occur, sufficient levels of network traffic must be generated. If a cable modem is not sending data via granted slots, there are no data slots to piggyback. Therefore, piggybacking percentage will be lower and contention request percentage higher. The percentage of packets sent via contention request was 99.08. Just as piggybacking requires a certain level of network traffic, so does concatenation. The percentage of packets sent via concatenation was .44.

5.1.2 With Background Traffic (Greater than Zero Active TCP Connections)

The addition of background traffic (i.e. active TCP connections) produced a high MOS value of 4.41 at zero active TCP connections and a low MOS value of 3.6 at 15

active TCP connections. Downstream utilization remained relatively constant ranging from 1.76 to 3.23, varying as active TCP connections were increased. Upstream utilization also remained relatively constant ranging from 49.37 to 51.84, increasing in conjunction with the increase of active TCP connections (at zero TCP connections, there was only an upstream utilization of 5.24). The percentage of packets sent via piggybacking ranged from 79.41 to 69.9, decreasing in conjunction with the increase of active TCP connections. The percentage of packets sent via contention requests ranged from 99.08 to 9.65, decreasing in conjunction with the increase of active TCP connections. There was a significant drop from zero to two active TCP connections. This behavior can be attributed to piggybacking and concatenation requiring other traffic to exist in order to function. The percentage of packets sent via concatenation ranged from .44 to 20.33, increasing in conjunction with the increase in active TCP connections.

5.2 Impact of Increasing N^A

N^A was increased as follows: 0, 10, 50, 100, and 200. The maximum of 200 nodes attacked is equivalent to 50 percent of the cable modems on the DOCSIS segment. As N^A was increased, the aggregate attack packet bandwidth for both the downstream and upstream channel increased. The increase was relatively constant across the entire range observed with both doubling as N^A was doubled. This behavior was expected and is graphically depicted in figures 5.4 and 5.5. The attack packet size is 64 bytes, however DOCSIS is required to encapsulate the ICMP message in a 188 byte MPEG frame. Therefore, the anticipated downstream attack bandwidth is:

$$\text{Bandwidth} = (N^A * 188 * 8) / R^A$$

The upstream bandwidth that is consumed is less since the frame size is now 64 bytes (plus framing overhead).




Figure 5.4. Aggregate downstream attack packet bandwidth, 400 cable modems, .5 second attack interval, across the number of nodes attacked (x-axis) in bits per second (y-axis). Also graphed is the anticipated bandwidth given the rate of attack.




Figure 5.5. Aggregate upstream attack packet bandwidth, 400 cable modems, .5 second attack interval, across the number of nodes attacked (x-axis) in bits per second (y-axis).

As N^A was increased, the first MOS value below the industry minimum for toll quality was observed when ten nodes were attacked with 15 TCP connections present (value: 3.58). The first MOS value below the benchmark established for this thesis was seen when 200 nodes were attacked with 20 TCP connections present (value: 1.96). The difference between the MOS value at 200 nodes attacked with 15 TCP connections (value: 3.34) and 200 nodes attacked with 20 TCP connections is significant compared to all other decreases. Prior to the observed decrease from 3.34 to 1.96 (delta of 1.38), the largest delta was .63 observed between ten and 15 TCP connections while 50 nodes were attacked. Overall, MOS value experienced the most change with 200 nodes attacked, as expected. The observed behavior is graphically depicted in figure 5.6.1. This is a key result which captures the susceptibility to exploitation of the upstream channel in DOCSIS systems. The MOS value dropped by 50% (from 3.58 to 1.96) with only a two percent increase in downstream utilization (from 2% to 4%).




Figure 5.6.1. MOS value observed with 400 cable modems, .5 second attack interval, number of nodes attacked range from zero to 200. First MOS value observed below the thesis benchmark of 3.0 is at 200 nodes attacked with 20 TCP connections active (light blue).

Figures 5.6.2 and 5.6.3 depict the behavior of the average collision rate for the cable modems under attack and the collision rate for the cable modem supporting VoIP monitoring with 15 and 20 TCP connections. As N^A was increased, the average collision rate for cable modems under attack gradually increased up to 100 nodes attacked. Doubling the nodes attacked from 50 to 100 nodes, a collision rate change of 33 percent was observed. Doubling the nodes attacked from 100 to 200 nodes, a collision rate change of over 80 percent was observed. The same behavior was observed with the VoIP cable modem. From 50 to 100 nodes attacked, the rate change observed was just over 40 percent. From 100 to 200 nodes attacked, the rate change was again over 80 percent.




Figure 5.6.2. Collision rate (y-axis) for cable modems under attack, 400 cable modems, .5 second attack rate interval, range of nodes attacked zero to 200 (x-axis). Over 80 percent change in rate from 100 to 200 nodes attacked, 15 TCP connections.




Figure 5.6.3. Collision rate (y-axis) for VoIP cable modem, 400 cable modems, .5 second attack rate interval, range of nodes attacked zero to 200 (x-axis). Over 80 percent change in rate from 100 to 200 nodes attacked, 15 TCP connections.

As N^A was increased, no significant difference in downstream and upstream utilization was observed across the range of zero to 200 for N^A for 15 and 20 TCP connections. The range of downstream utilization was 2.11 to 4.16, increase observed at each increase of N^A . The range of upstream utilization was 51.4 to 58.56, increase observed at each increase of N^A . This behavior was expected and is very significant. The maximum change in downstream utilization rate was observed between 100 and 200 nodes attacked, but was a modest 30 percent. All other changes were just under 20 percent. By doubling the number of nodes attacked, we only increased the change in downstream utilization rate by 10 percent at the most intense level observed. Furthermore, the 30 percent increase only increased the downstream utilization rate to an extremely low value of 4.16 percent. That leaves over 95 percent of the downstream bandwidth available for other network traffic and still achieving the desired MOS value of under 3.0. This behavior is observed in figure 5.7.

Another significant point observed is the upstream utilization rate. Just as the change in downstream utilization rate was relatively insignificant, so was the change in upstream utilization rate. The range of upstream utilization was from 51.4 to 58.56 percent. The change in rate increased with the increase in the number of nodes attacked. The significance in this change is the total change in rate was only 14 percent over the entire increase of N^A . Just as the benchmark of 3.0 for MOS value was broken without overwhelming the downstream channel, the upstream channel retained a relatively large portion for non-attack traffic while still achieving DoS of the VoIP service benchmark. This behavior is also observed in figure 5.7.




Figure 5.7. Downstream and upstream utilization for 400 cable modems, .5 second attack interval, number of nodes attacked range from zero to 200. Significance observed is the extremely low percentage of the downstream channel required to achieve the DoS attack objective, as well as, the impact on the upstream channel.

As N^A was increased, the observed piggybacking, contention request, and concatenation behavior was what was anticipated. The more nodes targeted by the DoS attack produced less packets sent via piggybacking. This behavior is graphically depicted in figure 5.9. The ratio of contention requests to packets sent increased as N^A increased. This behavior is graphically depicted in figure 5.10. The percentage of packets sent via concatenation increased as N^A increased. This behavior is graphically depicted in figure 5.11. Figure 5.8 graphically depicts the observed behavior for all three metrics for the two scenarios (with an attack rate of .5 and the number of upstream TCP connections set at 15 and then 20). The results suggests that less than 15% of the attack packets sent

access the transmission channel either via piggybacking or concatenation. The ratio of contention requests to total packets sent for both scenarios is over 50%. Figure 5.8.1 depicts the behavior observed strictly of VoIP packets. Compared to attack packets, VoIP packets utilize less piggybacking and more concatenation for improving efficiency.




Figure 5.8. Percentage of packets sent via piggybacking, contention requests without concatenation, and with concatenation for 400 cable modems, .5 second attack interval, number of nodes attacked range from zero to 200. The observed behavior was in-line with what was expected, as the number of nodes attacked increased, the ratio of contention requests to total packets sent and packets sent via concatenation increased while the number of packets sent via piggybacking decreased.




Figure 5.8.1. Percentage of VoIP packets sent via piggybacking, contention requests without concatenation, and with concatenation for 400 cable modems, .5 second attack interval, number of nodes attacked range from zero to 200. Compared to attack packet behavior, VoIP packets utilize less piggybacking and more concatenation.




Figure 5.9. Percentage of packets sent via piggybacking as the number of nodes attacked increased. 400 cable modems, .5 second attack interval, range of nodes attacked from zero to 200. The number of packets sent via piggybacking decreased as the number of nodes attacked was increased.




Figure 5.10. The ratio of contention requests to total packets sent as the number of nodes attacked increased. 400 cable modems, .5 second attack interval, range of nodes attacked from zero to 200. The ratio of contention requests to total packets sent increased, as expected, as the number of nodes attacked increased.




Figure 5.11. The percentage of packets sent via concatenation as the number of nodes attacked increases. 400 cable modems, .5 second attack interval, range of nodes attacked from zero to 200. The percentage of packets sent via concatenation increased, as expected, with the number of nodes attacked increased.

5.3 Impact of Increasing and Decreasing R^A

R^A was increased from the baseline configuration of .5 second intervals between attack packets as follows: 1, 2, and 4 seconds. R^A was decreased from the baseline configuration of .5 second intervals between attack packets as follows: .25, .05, and .01. Each increase in interval doubled the length of the previous interval. The first decrease in interval was half of the baseline configuration, the second decrease equaling a quarter of the previous attack interval, and the final decrease twenty percent of the previous interval.

As R^A was increased, MOS value remained above the prescribed threshold value of 3.0 for all runs with 15 TCP connections. As R^A was decreased, MOS value dropped to less than 1.3 for 15 TCP connections, an attack rate interval of .25 seconds, and 200 nodes attacked. A further decrease of R^A to .05 seconds resulted in an MOS value of 2.78 for 50 nodes attacked, 2.53 for 100 nodes attacked, and 1.73 for 200 nodes attacked. A final decrease to .01 seconds resulted in an MOS value of 2.6 for 50 nodes attacked and less than 1.0 for 100 or more nodes attacked. These expected behaviors are graphically depicted in figures 5.11.1 and 5.11.2.




Figure 5.11.1. Mean Opinion Score (y-axis) for 15 TCP connections, range of nodes attacked from zero to 200, and a attack rate interval range from 4 to .01 seconds (x-axis). An MOS value less than the threshold value of 3.0 is first observed at 200 nodes attacked and an attack rate interval of .25 seconds. A sub-3.0 MOS value is observed at 50, 100, and 200 nodes attacked with an attack rate interval of .05 and .01 seconds.




Figure 5.11.2. Mean Opinion Score (y-axis) for 20 TCP connections, range of nodes attacked from zero to 200, and a attack rate interval range from 4 to .01 seconds (x-axis). An MOS value less than the threshold value of 3.0 is observed at 200 nodes attacked and an attack rate interval of .5 seconds (baseline configuration). Additionally, sub-3.0 MOS values are observed at 100 and 200 nodes attacked with an attack rate interval of .25 seconds, as well as, 50, 100, and 200 nodes attacked with an attack rate interval of .05 and .01 seconds.

As R^A was decreased over the range of 4 to .01 seconds for runs configured with 15 and 20 TCP connections, the aggregate downstream attack packet bandwidth increased slightly over the 4 to .5 second range. From .25 to .01 seconds, the increase in aggregate downstream bandwidth observed was much higher. This dramatic change can be attributed to the high intensity behavior created by the extremely smaller interval between attack packets. This behavior is graphically depicted in figures 5.12 and 5.13.




Figure 5.12. The aggregate downstream attack packet bandwidth in bits per second (y-axis) as the attack rate interval is decreased from 4 seconds to .01 seconds (x-axis). As the attack rate interval increases (approaches 4 seconds), downstream bandwidth consumed decreases. As the attack rate interval decreases (approaches .01 seconds), downstream bandwidth consumed increases.




Figure 5.13. The aggregate downstream attack packet bandwidth in bits per second (y-axis) as the attack rate interval is decreased from 4 seconds to .01 seconds (x-axis). As the attack rate interval increases (approaches 4 seconds), downstream bandwidth consumed decreases. As the attack rate interval decreases (approaches .01 seconds), downstream bandwidth consumed increases.

As R^A was decreased over the range of 4 to .01 seconds for runs configured with 15 and 20 TCP connections, the aggregate upstream attack packet bandwidth increased steadily from 4 to .5 second intervals. At .25 seconds and faster, the aggregate upstream attack packet bandwidth began to decrease drastically for 50 or more nodes attacked. This sudden change can be attributed to the complete degradation of the upstream channel. Both attack and non-malicious packets suffered from the abundance of attack traffic. The lone exception is ten nodes attacked. The aggregate upstream attack packet bandwidth for ten nodes attacked actually increased. It should be noted that despite the increase, the MOS value associated with ten nodes attacked and .25 seconds and faster attack interval remained above 3.0. These behaviors are graphically depicted in figures 5.14 and 5.15.




Figure 5.14. The aggregate upstream attack packet bandwidth in bits per second (y-axis) as the attack rate interval is decreased from 4 seconds to .01 seconds (x-axis). As the attack rate interval increases (approaches 4 seconds), upstream bandwidth consumes decreases. As the attack rate interval decreases (approaches .01 seconds), Upstream bandwidth consumed increases.




Figure 5.15. The aggregate upstream attack packet bandwidth in bits per second (y-axis) as the attack rate interval is decreased from 4 seconds to .01 seconds (x-axis). As the attack rate interval increases (approaches 4 seconds), upstream bandwidth consumes decreases. As the attack rate interval decreases (approaches .01 seconds), Upstream bandwidth consumed increases.

As R^A was increased from .5 seconds to 4 seconds, both the downstream and upstream utilization decreased. Decreasing the rate from .5 seconds to .01 seconds resulted in an increase in both downstream and upstream utilization. This expected behavior is depicted in figures 5.16 through 5.19. The significance in downstream and upstream utilization is not only the direct correlation of increasing R^A with decreasing utilization, but also the portion of the available downstream and upstream channel required for the DoS attack to reach its goal. Only a minimal portion of the downstream channel is required (less than five percent when the first sub-3.0 MOS value is observed) and similarly in the upstream channel (just over half of the available upstream channel).




Figure 5.16. Downstream utilization for 15 TCP connections, range of nodes attacked from zero to 200, and an attack rate interval range of 4 seconds to .01 seconds. As the attack rate interval was increased (approached 4 seconds), downstream utilization decreased. As the attack rate interval was decreased (approaches .01 seconds), the downstream utilization increased.




Figure 5.17. Downstream utilization for 20 TCP connections, range of nodes attacked from zero to 200, and an attack rate interval range of 4 seconds to .01 seconds. As the attack rate interval was increased (approached 4 seconds), downstream utilization decreased. As the attack rate interval was decreased (approaches .01 seconds), the downstream utilization increased.




Figure 5.18. Upstream utilization for 15 TCP connections, range of nodes attacked from zero to 200, and an attack rate interval range of 4 seconds to .01 seconds. As the attack rate interval was increased (approached 4 seconds), upstream utilization decreased. As the attack rate interval was decreased (approaches .01 seconds), the upstream utilization increased.




Figure 5.19. Upstream utilization for 20 TCP connections, range of nodes attacked from zero to 200, and an attack rate interval range of 4 seconds to .01 seconds. As the attack rate interval was increased (approached 4 seconds), upstream utilization decreased. As the attack rate interval was decreased (approaches .01 seconds), the upstream utilization increased.

The collision rate observed behaved as expected, increasing as the attack rate interval decreased over the range of 4 seconds to .01 seconds. This behavior is graphically depicted in figures 5.20 and 5.21. A point of significance is the relatively low rate of collisions observed when the first sub-3.0 MOS value is observed (just over 25 percent rate of occurrence). Despite the low occurrence of collisions, the DoS attack was capable of degrading network performance such that the target MOS was achieved.




Figure 5.20. Collision rate (y-axis) for 15 TCP connections, number of nodes attacked range from zero to 200, and attack rate interval range of 4 to .01 seconds. Decrease in attack rate interval results in an increase in collision rate.




Figure 5.21. Collision rate (y-axis) for 20 TCP connections, number of nodes attacked range from zero to 200, and attack rate interval range of 4 to .01 seconds. Decrease in attack rate interval results in an increase in collision rate.

As R^A decreased over the range of 4 to .01 seconds, the observed ratio of contention requests to total packets sent behaved as anticipated increasing over the identified range. This behavior is graphically depicted in figures 5.22 and 5.23. The increase in contention request ratio is related to the increase in collision rate observed over the same range. Each time a cable modem is unable to obtain an upstream slot for transmission due to a collision, an additional contention request will be made. Over the range of R^A identified, the smaller interval between attack packets results in a more intense DoS attack. With more attack packets sent DS, responding cable modems contend for upstream contention request slots resulting in an increase in collisions and contention requests.




Figure 5.22. Contention request ratio for 15 TCP connections, number of nodes attacked range from zero to 200, attack rate interval range from 4 to .01 seconds. As the interval between attack packets decreases, the number of contention requests required to transmit a packet increases.




Figure 5.23. Contention request ratio for 20 TCP connections, number of nodes attacked range from zero to 200, attack rate interval range from 4 to .01 seconds. As the interval between attack packets decreases, the number of contention requests required to transmit a packet increases.

As R^A decreases from 4 to .05 seconds, the percentage of packets that are transmitted via piggybacking decreases, as well. The lone exception once again is observed at ten nodes attacked. Figures 5.24 and 5.25 graphically depict this behavior for 15 and 20 TCP connections. The cause of this behavior is related to the relationship between increases in contention request ratio and collision rate. A cable modem may attempt to piggyback a request, but due to the increased level of collisions, that request is likely to not be granted. The cable modem will then have to attempt to request future data slots via normal contention request slots. Therefore, a faster attack rate resulting in higher collisions not only produces more contention requests but fewer packets transmitted via piggybacking. The ten nodes attacked data point behavior helps explain the MOS value remaining above 3.0 despite the previously noted aggregate upstream attack packet bandwidth increasing, contrary to other data points.




Figure 5.24. Percentage of packets transmitted via piggybacking for 15 TCP connections, number of nodes attacked range from zero to 200, attack rate interval range from 4 to .01 seconds. As attack rate interval decreases from 4 to .01 seconds, percentage of packets sent via piggybacking decreases.




Figure 5.25. Percentage of packets transmitted via piggybacking for 20 TCP connections, number of nodes attacked range from zero to 200, attack rate interval range from 4 to .01 seconds. As attack rate interval decreases from 4 to .01 seconds, percentage of packets sent via piggybacking decreases.

As the percentage of packets sent via piggybacking decreases with an increasing R^A , the percentage of packets concatenated for transmission increases. Figures 5.26 and 5.27 graphically depict this behavior.




Figure 5.26. Percentage of packets transmitted via concatenation for 15 TCP connections, number of nodes attacked range from zero to 200, attack rate interval range from 4 to .01 seconds. As the attack rate interval approaches .01 seconds, the percentage of packets sent via concatenation increases.




Figure 5.27. Percentage of packets transmitted via concatenation for 20 TCP connections, number of nodes attacked range from zero to 200, attack rate interval range from 4 to .01 seconds. As the attack rate interval approaches .01 seconds, the percentage of packets sent via concatenation increases.

Tables 5.2 through 5.7 contain the data statistics relating to channel access with 15 and 20 TCP connections over an attack rate range of 4 to .01 seconds and number of nodes attacked range of zero to 200 nodes. These tables are graphically depicted and discussed in previous figures.

PIGGY BACKING							
	15 TCP, 4 Sec.	15 TCP, 2 Sec.	15 TCP, 1 Sec.	15 TCP, .5 Sec.	15 TCP, .25 Sec.	15 TCP, .05 Sec.	15 TCP, .01 Sec.
0	3.49	3.43	3.43	3.44	3.43	3.44	3.39
10	3.38	3.33	3.39	3.3	3.54	1.79	16.99
50	3.22	3.14	3.04	3.19	3.65	1.5	1.76
100	2.74	2.86	2.74	2.87	3.15	0.93	0.21
200	2.64	2.49	2.18	2.36	2.1	0.5	0.06

Table 5.2. Piggybacking data for 15 TCP connections, number of nodes attacked from zero to 200, attack rate interval range from 4 to .01 seconds.

PIGGY BACKING							
	20 TCP, 4 Sec.	20 TCP, 2 Sec.	20 TCP, 1 Sec.	20 TCP, .5 Sec.	20 TCP, .25 Sec.	20 TCP, .05 Sec.	20 TCP, .01 Sec.
0	2.58	2.58	2.63	2.62	2.59	2.61	2.57
10	2.61	2.52	2.57	2.62	2.92	1.48	14.57
50	2.45	2.34	2.39	2.71	3.33	1.28	1.25
100	2.3	2.21	2.27	2.43	3.1	0.9	0.21
200	2.14	1.98	2	2.04	2.09	0.45	0.06

Table 5.3. Piggybacking data for 20 TCP connections, number of nodes attacked from zero to 200, attack rate interval range from 4 to .01 seconds.

CONTENTION REQUEST RATIO							
	15 TCP, 4 Sec.	15 TCP, 2 Sec.	15 TCP, 1 Sec.	15 TCP, .5 Sec.	15 TCP, .25 Sec.	15 TCP, .05 Sec.	15 TCP, .01 Sec.
0	52.17	51.94	51.96	52.25	52.07	52.21	52.06
10	52.24	52.32	52.06	52.46	52.73	56.37	54.32
50	52.45	52.79	53.12	54.02	55.42	74.45	77.6
100	53.33	53.57	54.59	56.08	60.2	81.44	84.53
200	54.16	55.15	57.51	62.35	68.41	85.76	89.25

Table 5.4. Contention request ratio data for 15 TCP connections, number of nodes attacked range from zero to 200, attack rate interval range from 4 to .01 seconds.

CONTENTION REQUEST RATIO							
	20 TCP, 4 Sec.	20 TCP, 2 Sec.	20 TCP, 1 Sec.	20 TCP, .5 Sec.	20 TCP, .25 Sec.	20 TCP, .1 Sec.	20 TCP, .05 Sec.
0	52.37	52.44	52.39	52.47	52.67	52.75	52.63
10	52.44	52.7	52.63	52.9	53.08	55.79	56.58
50	52.84	53.31	53.39	53.97	55.22	73.72	78.66
100	53.46	53.89	54.77	56.46	59.02	81.46	84.74
200	54.5	55.41	57.11	60.74	66.68	86	89.53

Table 5.5. Contention request ratio data for 20 TCP connections, number of nodes attacked range from zero to 200, attack rate interval range from 4 to .01 seconds.

CONCATENATION							
	15 TCP, 4 Sec.	15 TCP, 2 Sec.	15 TCP, 1 Sec.	15 TCP, .5 Sec.	15 TCP, .25 Sec.	15 TCP, .1 Sec.	15 TCP, .05 Sec.
0	0.22	0.22	0.24	0.24	0.23	0.23	0.37
10	0.26	0.3	0.37	0.55	1.36	5	13.82
50	0.32	0.46	0.83	1.78	4.04	7.93	15.82
100	0.42	0.68	1.37	3.22	6.22	8.8	12.39
200	0.56	1.09	2.49	4.98	9.47	10.33	9.44

Table 5.6. Concatenation data for 15 TCP connections, number of nodes attacked range from zero to 200, attack rate interval from 4 to .01 seconds.

CONCATENATION							
	20 TCP, 4 Sec.	20 TCP, 2 Sec.	20 TCP, 1 Sec.	20 TCP, .5 Sec.	20 TCP, .25 Sec.	20 TCP, .1 Sec.	20 TCP, .05 Sec.
0	0.23	0.27	0.24	0.24	0.23	0.23	0.4
10	0.24	0.3	0.4	0.62	1.33	4.54	13.61
50	0.35	0.54	1.02	1.93	4.1	9.38	15.35
100	0.43	0.79	1.65	3.51	6.9	8.87	12.19
200	0.63	1.31	2.77	5.67	10.08	10.32	9.21

Table 5.7. Concatenation data for 15 TCP connections, number of nodes attacked range from zero to 200, attack rate interval from 4 to .01 seconds.

CONCLUSION

The data produced from the simulation supports all of our preliminary expectations. As the number of nodes attacked by a ping flood DoS attack on a DOCSIS network segment is increased, the upstream channel is "choked" by attack traffic contention requests. Additionally, access to the upstream channel is limited due to the increase in collisions which further increases the contention request ratio. Furthermore, as the interval between attack packet transmissions is decreased, the percentage of packets accessing the upstream channel via piggybacking decreases while packets accessing the same channel via concatenation increases.

The increase in contention requests and decrease in packets transmitted via piggybacking degrade network performance such that VoIP transmission quality is below an acceptable MOS value of 3.0 with several combinations of N^A and R^A . The relationship between the various combinations is the smaller the interval defined by R^A , the smaller the amount of nodes targeted for attack defined by N^A . Subsequently, the DoS attack defined and analyzed in our research supports the theory that in DOCSIS networks, the attack requires only a small portion of the available downstream bandwidth in order to severely impact upstream performance, especially when the focus of the attack is best effort VoIP sessions.

APPENDIX

All files associated with this thesis, as well as a "snapshot" of the NS simulation environment, are on the compact disc included with this manuscript. On the compact disc, if you navigate to the directory `./ns-cpsc854/project-docsis/`, you will be in the home directory for all simulations performed for this thesis. The directory `./GoRuns` contains all of the `goruns.dat` files used for configuring each simulation run. The directory `./thesisData` contains all of the data files in `tar.gz` form for all simulations performed during our research. To locate specific data results, the directories in `./thesisData` are organized with the following naming convention: `set4X_Y` where `X` is the number of nodes attacked and `Y` is the attack rate interval. The sets for .5 second attack rate have no `Y` value (i.e. `set400` would be zero nodes attacked with an attack rate interval of .5 seconds).

To recreate the results, a sample script would go as follows:

```
./cleanHouse  
./prepRun 0_4  
./goCPRruns.script 400_4 &
```

The above script would cleanup all unnecessary files, copy all `goruns.dat` files required for a run with zero nodes attacked with a four second attack interval into the appropriate directories, and execute the simulation for those runs saving the results in `tar.gz` form in the `./thesisData/set400_4` directory. The naming convention is simply 'number of nodes attack' and 'attack interval' separated with an underscore. Consolidated data is then placed in a file named `data_set400_4.out`.

REFERENCES

- ALOHAnet. (2007, June 6). In *Wikipedia, The Free Encyclopedia*. Retrieved 19:54, June 15, 2007, from <http://en.wikipedia.org/w/index.php?title=ALOHAnet&oldid=136304323>.
- Bellardo, J., Savage, S. (2003). "802.11 Denial-of-Service Attacks: Real Vulnerabilities and Practical Solutions." Retrieved 10:00, June 7, 2007, from http://sysnet.ucsd.edu/~bellardo/pubs/usenix-sec03-8021_dos-html/index.html.
- Cole, R., Rosenbluth, J (2001). Voice over IP performance monitoring. In *ACM SIGCOMM Computer Communication Review* (volume 31, issue 2, pp. 9-24). New York: ACM Press.
- Denial-of-Service attack. (2007, June 2). In *Wikipedia, The Free Encyclopedia*. Retrieved 21:53, June 6, 2007, from http://en.wikipedia.org/w/index.php?title=Denial-of-service_attack&oldid=135323193.
- DSL overtakes Cable in the U. S. (2006, June). In *Bandwidth Report*. Retrieved 9:25, June 6, 2007, from <http://www.websiteoptimization.com/bw/0606/>.
- Fellows, D., Jones, D. (2001). "DOCSIS Cable Modem Technology." *IEEE Communications Magazine*, vol. 39 (no. 3), pp. 202-209.
- Keneipp, R. (2000, May). "What's Toll Quality Voice." Retrieved 11:06, June 9, 2007, from <http://www.itworld.com/Net/2621/ITW849/>.
- Martin, J. , Westall, M. (2006, July). "Validating an 'ns' Simulation Model of the DOCSIS Protocol." Proceedings from 2006 SPECTS'06: International Symposium on Performance Evaluation of Computer and Telecommunication Systems. Calgary, CA: pp. 297-304.
- Meckler, Alan (2004, October 7). Jupiter Media. JupiterResearch Forecasts Voice Over IP Telephony Services To Reach 12.1 Million Households By 2009. Retrieved 9:26, June 6, 2007, from <http://www.jupitermedia.com/corporate/releases/04.10.07-newjupresearch.html>.
- Miller, M. (2005, July 19). "Do You Hear What I Hear?—Part IV: Measuring "Toll Quality"." Retrieved 10:45, June 9, 2007, from <http://www.voipplanet.com/backgrounders/article.php/3521171>.
- NSNAM. (2007, March 11). Main Page, NSNAM. Retrieved 17:15, June 7, 2007 from http://nnsam.isi.edu/nnsam/index.php/User_Information.

Peyravi, H (1999, March). "Medium access control protocols performance in satellite communications." *IEEE Communications Magazine*, vol. 37 (no. 3), pp. 62-71.

U. S. Broadband Penetration Breaks 80% Among Active Internet Users. (2007, February). In *Bandwidth Report*. Retrieved 9:24, June 6, 2007, from <http://www.websiteoptimization.com/bw/0703/>.

Voice over IP. (2007, June 6). In *Wikipedia, The Free Encyclopedia*. Retrieved 12:49, June 6, 2007, from http://en.wikipedia.org/w/index.php?title=Voice_over_IP&oldid=136206069.