Distributed Denial of Service Attack

R.Roja[1]
[1]Student
[1]Department of Computer Science and Engineering
[1]Saveetha School of Engineering

Abstract— Distributed denial of service (DDoS) attack is a coordinated attack, generally performed on a massive scale on the availability of services of a target system or network resources. Owing to the continuous evolution of new attacks and ever-increasing number of vulnerable hosts on the Internet, many DDoS attack detection or prevention mechanisms have been proposed. In this paper, we present a comprehensive survey of DDoS attacks, revealing methods and tools used in wired networks. The paper also highlights open issues, research challenges and possible solutions in this area.

Keywords: Distributed denial of service, vulnerable hosts, and wired networks.

I. INTRODUCTION

A distributed denial-of-service (DDoS) attack is one in which a numerous of compromised systems attack a single target, thereby causing denial of service for users of the targeted system. The flood of incoming messages to the target system degree forces it to shut down, thereby denying service to the system to agreeing users. In a typical DDoS attack, the assailant begins by exploiting vulnerability in one computer system and making it the DDoS master. The attack master, also known as the botmaster, identifies and infects other wound systems with malware. Eventually, the attacking instructs the controlled machines to throw an attack against a specified target.

There are two types of DDoS attacks: a network-centric attack which overloads a service by using up bandwidth and an application-layer attack which overloads a service or database with application calls. The overflow of packets to the target causes a denial of service. While the media tends to focus on the target of a DDoS attack as the victim, in reality there are many destroy in a DDoS attack -- the final target and as well the systems controlled by the intruder. Although the owners of co-opted computers are typically not noticing that their computers have been compromised, they are nevertheless likely to suffer a degradation of service and not work well. A computer under the control of an intruder is known as a zombie or bot. A group of co-opted computers is known as a botnet or a zombie army.

II. DISTRIBUTED DENIAL OF SERVICE ATTACKS

Sometimes a cracker uses a network of zombie computers to sabotage a specific Web site or server. The idea is pretty simple -- a cracker tells all the computers on his botnet to contact a specific server or Web site repeatedly. The sudden increase in traffic can cause the site to load very slowly for legitimate users. Sometimes the traffic is enough to shut the site down completely. We call this kind of an attack a Distributed Denial of Service (DDoS) attack.

Some particularly tricky botnet use uncorrupted computers as part of the attack. Here's how it works: the cracker sends the command to introduce the attack to his zombie army. Each computer within the army sends an electronic connection request to an innocent computer called a reflector. When the reflector receives the request, it looks like it originates not from the zombies, but from the ultimate victim of the attack. The reflectors send information to the victim system, and finally the system's performance suffers or it shuts down completely as it is inundated with multiple unsolicited responses from several computers at once.

From the perspective of the victim, it looks like the reflectors attacked the system. From the perspective of the reflectors, it seems like the victimized system requested the packets. The zombie computers remain hidden, and even more out of sight is the cracker himself.

The list of DDoS attack victims includes some pretty major names. Microsoft suffered an attack from a DDoS called My Doom. Crackers have targeted other major Internet players like Amazon, CNN, Yahoo and eBay. The DDoS names range from mildly amusing to disturbing:

A. Ping of Death - Bot create huge electronic packets and sends them on to victims

B. Mail bomb - bots send a massive amount of e-mail, crashing e-mail servers

C. Smurf Attack - bots send Internet Control Message Protocol (ICMP) messages to reflects, see above illustration

D. Teardrop - bots send pieces of an not according packet; the victim system tries to recombine the fragment into a packet and crashes as a result

Once an army begins a DDoS attack against a creature system, there are few things the system administrator can do to guide catastrophe. He could choose to limit the amount of traffic allowed on his server, but this restricts accordant Internet connections and zombies alike. If the administrator can determine the origin of the attacks, he can filter the traffic. Unfortunately, since many zombie computers change (or spoof) their addresses, this isn't always easy to do.

III. TYPES OF ATTACKS

DDoS attacks can be broadly divided into three types

A. Volume Based Attacks

Includes UDP floods, ICMP floods, and other spoofed-packet floods. The attack's goal is to cause the bandwidth of the attacked site, and magnitude is measured in bits per second (Bps).

B. Protocol Attacks

Includes SYN floods, fragmented packet attacks, Ping of Death, Smurf DDoS and more. This type of attack consumes actual server resources, or those of intermediate communication, voyage such as firewalls and load balancers, and is measured in Packets per second.

C. Application Layer Attacks

Includes Slowloris, Zero-day DDoS attacks, DDoS attacks that target Apache, Windows or Open BSD vulnerabilities and more. Comprised of seemingly legitimate and innocent requests, the goal of these attacks is to crash the web server, and the magnitude is measured in Requests per second.

D. Specific DDoS Attacks Types

Some specific and particularly popular and dangerous types of DDoS attacks include:

1) UDP Flood

This DDoS attack leverages the User Datagram Protocol (UDP), a session less networking protocol. This type of attack floods random ports on a remote host with numerous UDP packets, causing the host to repeatedly check for the application listening at that port, and (when no application is found) reply with an ICMP Destination Unreachable packet. This process saps host resources, and can ultimately lead to inaccessibility.

2) ICMP (Ping) Flood

Similar in principle to the UDP flood attack, an ICMP flood overwhelms the target resource with ICMP Echo Request (ping) packets, generally sending packets as fast as possible without waiting for replies. This type of attack can consume both outgoing and incoming bandwidth, since the victim's servers will often attempt to respond with ICMP Echo Reply packets, resulting a significant overall system slowdown.

3) SYN Flood

A SYN flood DDoS attack exploits a known weakness in the TCP connection sequence (the "three-way handshake"), wherein a SYN request to initiate a TCP connection with a host must be answered by a SYN-ACK response from that host, and then confirmed by an ACK response from the requester. In a SYN flood scenario, the requester sends multiple SYN requests, but either does not respond to the host's SYN-ACK response, or sends the SYN requests from a spoofed IP address. Either way, the host system continues to wait for acknowledgement for each of the requests, binding resources until no new connections can be made, and ultimately resulting in denial of service.

4) Ping of Death

A ping of death ("POD") attack involves the attacker sending multiple malformed or malicious pings to a computer. The maximum packet length of an IP packet (including header) is 65,535 bytes. However, the Data Link Layer usually poses limits to the maximum frame size - for example 1500 bytes over an Ethernet network. In this case, a large IP packet is split across multiple IP packets (known as fragments), and the recipient host reassembles the IP fragments into the complete packet. In a Ping of Death scenario, following malicious manipulation of fragment content, the recipient ends up with an IP packet which is larger than 65,535 bytes when reassembled. This can overflow memory buffers allocated for the packet, causing denial of service for legitimate packets.

5) Slowloris

Slowloris is a highly-targeted attack, enabling one web server to take down another server, without affecting other services or ports on the target network. Slowloris does this by holding as many connections to the target web server open for as long as possible. It accomplishes this by creating connections to the target server, but sending only a partial request. Slowloris constantly sends more HTTP headers, but never completes a request. The targeted server keeps each of these false connections open. This eventually overflows the maximum concurrent connection pool, and leads to denial of additional connections from legitimate clients.

6) Zero-day DDoS

"Zero-day" is simply unknown or new attacks, exploiting vulnerabilities for which no patch has yet been released. The term is well-known amongst the members of the hacker community, where the practice of trading Zero-day vulnerabilities has become a popular activity.

Sources of DDoS Attacks

DDoS attacks are quickly becoming the most prevalent types of attacks, growing rapidly in the past year in both number and volume, according to recent market research. The trend is towards shorter attack duration, but bigger packet-per-second attack volume, and the overall number of attacks reported has grown markedly, as well.

During the Q4-2011, one survey found 45% more DDoS attacks compared to the parallel period of 2010, and over double the number of attacks observed during Q3-2011. The average attack bandwidth observed during this period was 5.2G bps, which is 148% higher than the previous quarter.

Another survey of DDoS attacks found that more than 40% of respondents experienced attacks that exceeded 1Gbps in bandwidth in 2013, and 13% were targeted by at least one attack that exceeded 10G bps.

From a motivational perspective, recent research found that ideologically motivated DDoS attacks are on the rise. The research also mentioned financial reasons (e.g., competitive feuds) as another common reason for such attacks.

7) FAMOUS ATTACKS

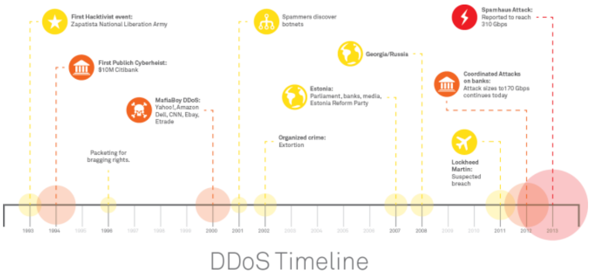

Fig. 1: FAMOUS ATTACKS

DDoS attacks are nothing new, but they've seen a revival over the past several years with attacks reaching historic sizes. Witness the 2012 attacks on U.S. banks and the 2013 SpamHaus attack that was recorded at 310Gbps.

This timeline shows notable DDoS developments over the past 20 years. While not an exhaustive list of every major attack, it does show the growth in DDoS attacks as a tool for registering dissent, perpetrating cybercrime and even hacker bragging rights.

The Week of Famous DDos Attacks

- February 7-11-2000
- CNN, Yahoo, Bay, Datek taken down for several hours at a time due to traffic flooding
- Under administrated computers at California college used as the slave attack computers

– Trinoo, Tribal Flood Network, TFN2K, and Stacheldraht suspected tools used in attacks

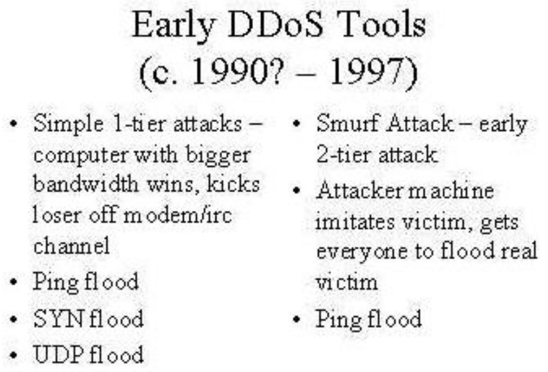

Early DDoS Tools
(c. 1990? – 1997)

- Simple 1-tier attacks – computer with bigger bandwidth wins, kicks loser off modem/irc channel
- Ping flood
- SYN flood
- UDP flood
- Smurf Attack – early 2-tier attack
- Attacker machine imitates victim, gets everyone to flood real victim
- Ping flood

Fig. 2: Early DDos Tools

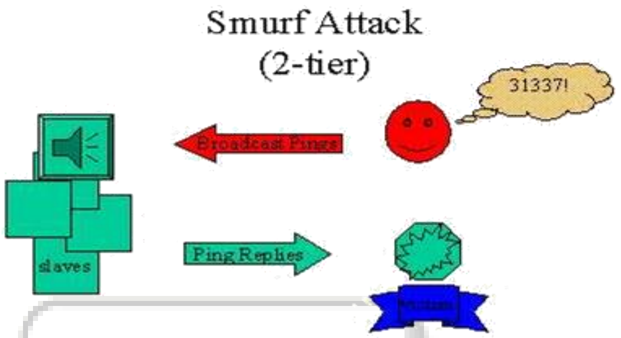

Smurf Attack
(2-tier)

Fig. 3: Smurf Attacks

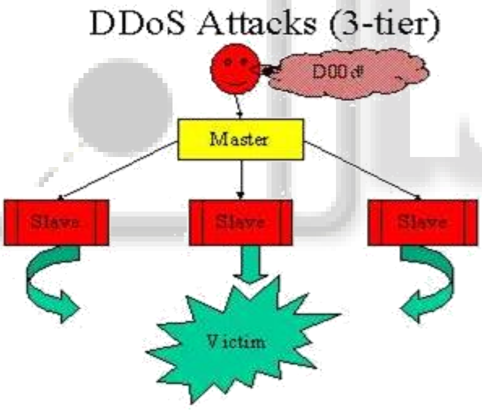

DDoS Attacks (3-tier)

Fig. 3: DDoS Attacks (3-tier)

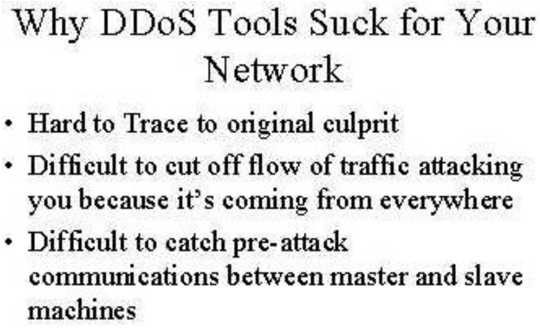

Why DDoS Tools Suck for Your Network

- Hard to Trace to original culprit
- Difficult to cut off flow of traffic attacking you because it's coming from everywhere
- Difficult to catch pre-attack communications between master and slave machines

Fig. 4: Why DDoS Tools Suck for Your Network

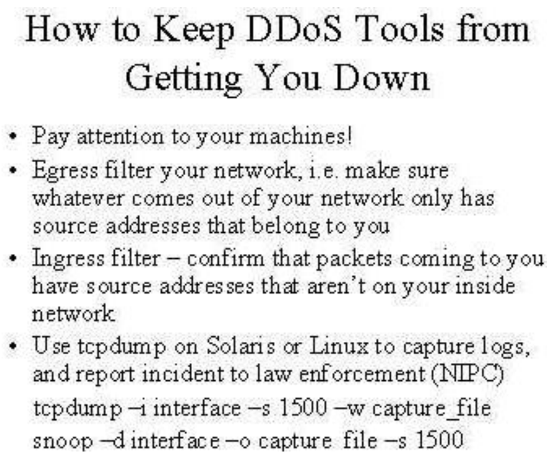

How to Keep DDoS Tools from Getting You Down

- Pay attention to your machines!
- Egress filter your network, i.e. make sure whatever comes out of your network only has source addresses that belong to you
- Ingress filter – confirm that packets coming to you have source addresses that aren't on your inside network
- Use tcpdump on Solaris or Linux to capture logs, and report incident to law enforcement (NIPC)

```
tcpdump –i interface –s 1500 –w capture_file
snoop –d interface –o capture_file –s 1500
```

Fig. 4: How to keep DDoS Tools from Getting You Down

You can detect the onset of a Denial of Service Attack in the following ways:

Check your website periodically to see if it is running more sluggishly than usual. Also, watch to see if you are sporadically receiving the "Service Unavailable" message shown above. As the hacker starts ramping up his attack, you may notice these symptoms.

Check your web stats. If you don't currently have web stats software setup, consider using Smarter Stats (my personal favorite) or the free Analog. These packages provide many graphs and reports, and it is a good idea to familiarize yourself with them, and look at them often so that you know how the graphs look when your site is functioning normally. In this way, you can easily spot anomalies which indicate a DoS attack (not to mention a whole bunch of other problems, but that's another story). Here is what a normal graph looks like (for my site, anyway). As you can see, I get a fairly consistent number of page views every weekday (approximately 30K) Now, here is a graph showing a DoS attack. Instead of seeing approximately 30K page views per day, I'm getting over 1 million page views per day. Obviously, something is wrong. When I first saw this anomaly, I paid it little attention, thinking that my ISP might have been having some problem with my web log files. I made a big mistake here by not quickly addressing this issue. Also, note that the hacker gradually ramped up his attack over time. Later on, I will discuss how a hacker can set this up. Here is a graph showing me which IP addresses used my site the most. This shows relatively normal conditions, although it looks like a few IP addresses may be using an inordinate amount of bandwidth. These are likely some kind of Search Engine Spiders, and I'll show you later how to find out more about these IP addresses. Here is the same report when my site is under attack. Note that there are a couple of IP addresses here that are using an inordinate amount of bandwidth. As you find these IP addresses, you should keep a list, as I will be showing you how to contact the ISPs and website owners later on. Remember that these organizations are also victims of your attacker. If you have some kind of Remote Desktop Access (e.g. Remote Desktop Connection, VNC, PC Anywhere, etc.) you can log in to your web server, open a command prompt, and type "net stat -an" (without the quotes). (If you find command line applications like net stat unwieldy, you might try using TCP View by legendary programmer Mark Russinovich.) Here is

what the net stat results should look like under normal circumstances: The first column of the net stat results shows the protocol. For the purposes of our current discussion, we only care about the TCP protocol. However, if you are confronted with a more sophisticated attacker, you will want to learn more about other protocols, particularly UDP and ICMP. The second column of the net stat results shows your computer's IP address. Note that your computer has several IP addresses. The IP address is followed by a colon and a port number. Different services use different port numbers. For example, your web server is likely using port 80, whereas Remote Desktop Connection uses port 3389. For the purposes of this article, we only care about connections to port 80 (web server) or port 443 (web server with SSL). The third column shows the IP address of a remote computer. For example, if a customer is looking at your website (on port 80), and the customer has IP address 24.200.167.248, and your web server has IP address 10.114.242.92, then you will see the following line:
TCP 10.114.242.92:80 24.200.167.248:60 734 ESTABLISHED. Note that the customer's machine is using port 60734. All you need to know for now is that well-defined services like your Web Server listen on well-defined ports like port 80. In contrast, clients who are connecting to your web server also need a port number to make this connection, but the port number is some random high number. This is why the "Foreign Address" column has such funky numbers after the colon, whereas the Local Address column has (mostly) well-known port numbers. In the "State" column, you see a bunch of lines that end with "LISTENING", "ESTABLISHED", "TIME_WAIT", etc. These tell you about the current state of that connection (socket). If a socket is "LISTENING", it a program is waiting for some remote computer to connect to it via the network. If a socket is "ESTABLISHED", it means a client is connected to your machine (e.g. a customer is visiting your website). If the socket says "TIME_WAIT", the socket may be setting up a connection, or it may be tearing down a connection. In any case, its waiting for something. For this article, I'm mainly focusing on simple Denial of Service (DoS) attacks. When dealing with a more high-tech attacker, I urge you to investigate which services are "LISTENING" on your server (did the hacker set something up), do you see many, many TIME_WAIT's (could be a more sophisticated DoS attack), etc. Now, here is what the net stat output looks like during a DoS attack. As you can see, 216.36.50.65 is making many, many connections to my web server (thousands of connections). You need to be a little careful here. When someone looks at a page on your website, their web browser will open several connections to your web server. So it is not uncommon to see a legitimate user with 8-20 connections to your web server. Also, many companies have a single public IP address for a large number of their employees use. So if many employees at that company are simultaneously visiting your website, you may see a whole bunch of legitimate connections from the same IP address. However, if you scroll through hundreds and hundreds of lines of your netstat results, and you are seeing the same IP address (in the third column) again and again, then this indicates a DoS attack. If you see such an IP address, write it down, as we'll later use it to track down and contact the owners of this other computer so that we can notify them that their machine has been compromised and is being used to attack your web server.

You can also check your log files for suspicious activity. Here is what a normal log file looks like: IIS (and other web servers) allow you to configure the log files in many ways, so your log files may look different, but will contain the information I discuss below. The important fields to look at are c-ip (the ip address of the person or program that is visiting your site), cs-uri-stem (which tells you which page they are visiting, and cs (User-Agent), which tells you which program is being used to visit your website (e.g. IE, Firefox, Google Chrome, etc.) Note that in the table above, we see several different IP addresses for our clients (89.248.174.2, 217.196.17.20, 66.249.67.153, etc.), different WebPages (/Financial/Default.aspx, /financial/StyleSheet.css, etc.) and different user agents (IE 8 and a Google "bot") Here is what the log file might look like during a Denial of Service attack: Note that we are seeing the same IP address (69.163.239.247) again and again (for hundreds of lines), that this bot or user is accessing the same web page again and again (/ is the root page of my website), and that the program accessing our website is called Apache Bench. Each time a user visits one of your pages, you should expect to see about 1-10 lines in the log files for that user's visit, and you should see the user's web browser downloading an html page (on one line), and any art or links embedded in that page downloaded (each also on its own line). What you don't expect to see is the **same** item downloaded again and again and again by the same IP addresses hundreds of times. This is what a Denial of Service (DoS) looks like. Also note that the suspicious lines also show the use of Apache Bench. From apache.org, "Apache Bench is a tool for benchmarking your Apache Hypertext Transfer Protocol (HTTP) server". That is, if you were trying to find out how many customers your website could handle in a short period of time, you would run this tool from the command line (it's called "ab"), the tool would simulate many, many visitors coming to your site at the same time, and then you could watch for the "Service Unavailable" message shown at the top of this page, and you would know how many visitors could simultaneously access your website before it crashed. Unfortunately, this helpful tool can also be used by a hacker to bring your website down. If the hacker runs this tool from several locations against your website (making it a **Distributed** Denial of Service Attack), he will be able to overwhelm your website with the phony traffic from Apache Bench. At this point, you will want to keep a list of these IP addresses, and I will show you later how we can use them to stop and/or mitigate the attack, but first we need to understand who the players are in this fiasco. We may need to contact some of these people, so we need to be sure we understand who the victims are, and who the culprits are.

IV. PREVENTION

A. How to prevent damage from DDoS attacks

Hackers want to disrupt your business, harm your company’s reputation, and keep your organization from performing its mission. If they succeed, they win, and you lose. By being prepared for DDoS attacks and working with Prolexic to eliminate vulnerabilities before an attack, you can prevent DDoS hackers from damaging your brand and interfering with your organization’s mission.

A range of industries is commonly targeted by DDoS attackers – financial services, e-Commerce, online gaming, travel and hospitality, healthcare, media and telecommunications, government agencies, and cloud-based software and services. Organizations in these industries and others with much to lose from a website outage need to have a plan in place for DDoS prevention and mitigation.

B. Can all DDoS attacks be prevented?

DDoS prevention comes in multiple forms. Often simply engaging Prolexic to provide your DDoS protection services and displaying the Protected by Prolexic logo can prevent a DDoS attack. Experienced DDoSers know it is a waste of their time to try to outsmart the ever-ready and highly trained DDoS mitigation experts in our Security Operations Center.

Many DDoSers have learned that their attacks against Prolexic clients will quickly be rerouted through our high-capacity global network to our scrubbing centers, from which only clean traffic will be sent back to our client’s network.

While Prolexic can’t prevent all DDoS attacks, but we can protect you from them, allowing your website to continue serving customers even while under attack.

C. DDoS prevention and mitigation planning

There are a number of ways to prevent the worst effects of a DDoS attack or mitigate them completely. Beyond having the right protective systems in place and ensuring sufficient overflow capacity is available, much depends on an active, well-informed incident response.

We help our clients plan and prepare so that their organizations are ready to respond calmly and effectively if they ever become the target a DDoS attack. Learn more about preventing chaos during a DDoS attack in the white paper Plan vs. Panic Making a DDoS Mitigation Playbook Part of Your Incident Response Plan.

D. DDoS prevention, detection and effective mitigation

The following Prolexic services can help prevent damage from DDoS attacks:

- DDoS detection and monitoring
- DDoS protection and mitigation
- Anti-DDoS intelligence
- Network protection

V. CONCLUSION AND FUTURE SCOPE

DDoS attack causes either disruption or degradation on victim’s shared resources, as a result preventing legitimate users from their access right on those resources. DDoS attack may target on a specific component of computer, entire computer system, certain networking infrastructure, or even entire Internet infrastructure. Attacks can be either by exploits the natural weakness of a system, which is known as logical attacks or overloading the victim with high volume of traffic, which is called flooding attacks. A distributed form of DDoS attack called DDoS attack, which is generated by many compromised machines to coordinately hit a victim. DDoS attacks are adversarial and constantly evolving. Once a particular kind of attack is successfully countered, a slight variation is designed that bypasses the defense and still performs an effective attack. In this paper, we covered an overview of the DDoS problem, available DDoS attack tools, defense challenges and principles, and a classification of available DDoS prevention mechanisms. This provides better understanding of the problem and enables a security administrator to effectively equip his arsenal with proper prevention mechanisms for fighting against DDoS threat. The current prevention mechanisms reviewed in this paper are clearly far from adequate to protect Internet from DDoS attack. The main problem is that there are still many insecure machines over the Internet that can be compromised to launch large-scale coordinated DDoS attack. One promising direction is to develop a comprehensive solution that encompasses several defense activities to trap variety of DDoS attack. If one level of defense fails, the others still have the possibility to defend against attack. A successful intrusion requires all defense level to fail.

REFERENCES

[1] Deborah Estrin, Ramesh Govindan, John Heidemann and Satish Kumar " Next Century Challenges: Scalable Coordination in Sensor Networks "In MobiCOM , August 1999.

[2] T.K.Srikanth and Sam Toueg. Optimal clock synchronization. J-ACM,34(3):626-645, July 1987.

[3] David L. Mills. Internet Time Synchronization: The Network Time Protocol. In Zhonghua Yang and T.AnthonyMarsland,editors, Global State and Time in Distributed Systems. IEEE Computer Society Press, 1994.