Network Security

Professor Gaurav S. Kasbekar
Department of Electrical Engineering
Indian Institute of Technology, Bombay

Week - 01
Lecture - 02
Motivation and Overview

Hello, we will first provide some motivation for this course, and then overview some of the topics to be discussed in this course. The motivation for this course is that communication networks are very useful, but malicious users attack them in several ways. To illustrate this, we'll provide several examples of recent security attacks. One widely discussed attack was the interference in the 2016 U.S. elections. So, recall that in the 2016 U.S. elections, the Republican candidate was Donald Trump, and the Democratic candidate was Hillary Clinton.

So, there was some interference in these elections. Servers of the Democratic Party were hacked, and several emails and other documents were supplied to Wikileaks, which made them public. And it was alleged by U.S. intelligence agencies that these attacks were ordered by Russia's leadership to harm the chances of the Democratic presidential candidate and boost those of the Republican candidate. For more details, you can refer to the Wikipedia page, which is mentioned here. So, this is an example of a high profile attack.

Another high profile attack is Stuxnet. So, we know that Iran has a nuclear program, and some countries are opposed to this program. Stuxnet is a sophisticated computer worm believed to have been used to cause damage to Iran's nuclear program. The worm was installed to modify some industrial control systems that are used to control the centrifuges used for uranium enrichment in Iran's nuclear program. So, this worm installed itself on a computer that was connected to these industrial control systems.

And it sent commands, because of which the speed of the centrifuges increased suddenly, and then it reduced suddenly. So, this caused mechanical stresses in the centrifuges, and that damaged them. So, this is about more than 1,000 of Iran's centrifuges were believed to be damaged because of this worm. And this worm was introduced using a USB type. It was believed to have been jointly developed by American and Israeli programmers.

For more information, you can refer to the Wikipedia page on Stuxnet. In the Indian context, a widely discussed attack was the recent breach of Indian bank's data. So, several banks had outsourced their ATM transaction processing to Hitachi payment systems, and hackers managed to penetrate the network of Hitachi payment systems. Because of this, several customers reported unauthorized use of their cards and even lost money. As many as 3.2 million debit cards were compromised, and many of them had to be replaced or their security codes changed.

So, this led to a lot of financial losses. So, this is an article which provides more information about this attack. Then, another widely discussed attack was the Facebook account hacking attack, which happened in September 2018. In this attack, the Facebook accounts of 50 million users were compromised. To launch this attack, hackers exploited the "View As" feature in Facebook, which allows a user to view a page as if it were being viewed by another user.

So, hackers exploited this feature to get access to some security tokens or keys, and these security tokens were then used to gain access to Facebook accounts. So, another attack in the context of Internet of Things devices, or IoT devices, was the Mirai malware. This was used for attacks on several websites after August 2016. So, Mirai malware infects a large number of IoT devices. So, Internet of Things extends internet connectivity to resource constrained devices such as sensors, actuators, appliances, and so on.

So, this Mirai malware was designed to infect a lot of IoT devices, and after infecting them, the malware used them for carrying out what are known as distributed denial of service attacks on websites. So, in such distributed denial of service attacks, a large number of IoT devices sent some malicious traffic to websites, causing them not to be able to provide their service properly. Denial of service attacks are where some malicious computer sends some traffic to a victim computer so that it is not able to provide its service properly, and distributed DoS attacks are ones where the attack is spread over a lot of computers. In this case, a large number of infected IoT devices send some malicious traffic to the victim websites. The advantage of using a distributed DoS attack in contrast to an ordinary DoS attack is that the distributed DoS attacks are more difficult to detect because the malicious traffic comes in from a large number of computers at different locations, and they are also difficult to block.

So, the Mirai malware caused such infected devices to send a lot of malicious traffic and led to this distributed denial of service attack. So, the way this attack was implemented was

that devices infected by Mirai, they continuously scanned the internet to find IP addresses of IoT devices. They launched what is known as a dictionary attack. They had a table of 60 common factory default usernames and passwords. Many users did not change their login names and passwords from these factory defaults.

So, they were able to discover several devices which were using these usernames and passwords and infected them. Then, they use these infected devices to launch the distributed denial of service attack. There was a command and control server which controlled a large number of infected devices, and asked them to send malicious traffic to attack the victim website. This was a Mirai malware attack, and it infected a lot of IoT devices and launched a distributed denial of service attack. So, this had a large amount of economic impact.

For more information, you can see the Wikipedia page of Mirai malware. So, this is only a small sample of recent security attacks. For an extensive list of cyber attacks, you can see the Wikipedia page on cyber attacks, which is linked here. We now provide an overview of some of the topics to be discussed in this course. So, one of the topics we'll discuss is confidentiality.

To motivate this topic, consider a user communicating using Wi-Fi with an access point. This shows an access point. It is a base station kind of device with which a user communicates. So, the user communicates with the access point using wireless transmissions, and these transmissions can be received by anyone who is nearby and tunes to the channel. Anyone can use a device known as a sniffer and intercept the wireless transmissions that are being sent between the user and the access point. Other broadcast media, such as cable internet and ethernet, also suffer from the same problem.

That is, any transmissions that are sent on this medium can be intercepted by other users who are connected to the same medium. Such media are known as broadcast media; that is, whatever information is sent on the medium can be received by all the users who are connected to the medium. In this context, since an intruder can intercept all the communication that is being exchanged between the two devices, user and access point, how do we ensure that intruders are not able to obtain confidential information, such as credit card numbers and passwords, which are sent between the user and access point? So, to ensure this, the transmitter encrypts information before sending it, and the receiver decrypts the information. Encryption means the information is disguised so that even if an interceptor collects that information, they are not able to understand it.

And the receiver decrypts it; that is, it does a conversion from the encrypted information back to the original print text. We'll review cryptography, which are techniques which allow us to do this encryption and decryption, and we'll discuss several algorithms that are used to encrypt and decrypt information in networks. Now, to further discuss confidentiality, suppose all messages that are sent by a transmitter, Alice, to a receiver, say Bob, or the channel between Alice and Bob, can be intercepted by an intruder, say Trudy. For Alice and Bob to be able to achieve confidentiality, clearly Alice and/or Bob must know some secret information which Trudy does not know. So, such secret information, which only Alice or Bob or both know, is known as a key.

The question is how do Alice and Bob agree upon the key prior to communication? So, Alice and Bob need to know the key in order to encrypt and decrypt information. So, in some contexts, it is very easy for Alice and Bob to agree upon the key. For example, suppose a bank and customer want to securely exchange messages over a network; then the bank can first send a key to the customer by post, and then the customer can use the key to communicate securely with the bank. But this is possible because there is an out-of-band channel available, namely communication by post.

Now, consider a more challenging scenario where the only way that Alice and Bob can possibly communicate is over a network. But this network is insecure. Any message that is sent over the network can be sniffed by intruders. In this case, how do Alice and Bob agree on a secret key? So, it is not very clear whether it is possible for Alice and Bob to communicate securely or not. In the 1970s it was shown by researchers Diffie and Hellman that the answer to this question is yes, Alice and Bob can communicate securely.

And the way they can communicate is through public key cryptography, which allows them to agree upon a secret key. To use public key cryptography, we require some components. For example, organizations called certification authorities-these components are required to enable public key cryptography. We'll discuss public key infrastructure, which is a set of components to enable public key cryptography. After Alice and Bob agree on a secret key, they can use that secret key to encrypt information and decrypt information transmitted on the channel between them.

So, Alice and Bob can use symmetric key cryptography after they have agreed upon a secret key. Symmetric key cryptography is less computationally expensive than public key cryptography, but it requires Alice and Bob to know some shared secret key. We'll discuss symmetric key cryptography as well, which can be used after Alice and Bob have agreed

upon a secret key. Another topic which we'll discuss in this course is authentication. How does a user prove its identity to an email server?

And conversely, how does the email server prove its identity to a user? More generally, if you have two users, Alice and Bob, communicating over a network, how do they authenticate each other? A possible solution is: Alice sends a password to Bob and Bob sends a password to Alice. But clearly we can see that this solution does not work, and Intruder Trudy can tune to the channel between Alice and Bob to obtain the password, and later on Trudy can authenticate. So, we might think that Alice can encrypt the password and send it to Bob.

So, does this solution work? But the downside is that Trudy can obtain the encrypted password, and then use it later on to authenticate to Bob. So, this solution doesn't work as well. We'll study more sophisticated techniques for achieving authentication in this course. Another fundamental building block for network security is message integrity.

So, to motivate message integrity, recall that for computation of routs in a network, a router sends what are known as link-state messages to other routers. In this picture here, these devices are end systems, that is they generate and download information. And in contrast, these devices shown as blue icons, these are routers. They transport information generated by end devices to other end devices. These routers send link-state messages to other routers.

For example, this router might send a link-state message to this router. And this link-state message has the list of attached links and their costs to all other routers. What is the benefit of sending such link-state messages? The benefit is that routers can compute routes to other routers and end systems using these link-state messages. Now in this context, an intruder may distribute bogus link-state messages or alter the contents of a link-state message.

So, when a router B receives a link state message with source address that of A, B needs to check that it was actually A who created the message and no one else created the message. B also needs to check that even though the message was created by A, was it modified during transit from A to B? So, more generally, suppose Bob receives a message with source address that of Alice, the message might be encrypted or it may be in plain text form. Bob needs to check whether Alice actually created the message Bob also needs to verify whether the message was modified during transit from Alice to Bob.

This is known as the message integrity problem. That is, when Bob receives a message from Alice, he needs to verify whether Alice actually created the message and that the

message was not modified while being forwarded from Alice to Bob. We'll study different techniques for achieving message integrity. A fundamental building block that is used to achieve message integrity is cryptographic hash functions. These are one way functions.

It is very easy to compute the hash function corresponding to some input, but it is difficult to find the input corresponding to a particular output of the hash function. We'll discuss cryptographic hash functions in this course, which can be used to achieve message integrity. We'll also discuss several other applications of cryptographic hash functions. Quite related to message integrity is the problem of creating digital signatures. Recall that manual signatures are extensively used by all of us on checks, credit cards, receipts, legal documents, letters, and so on.

A person makes a manual signature to indicate that he or she created a document or to indicate that he or she agrees with or acknowledges the contents of the document. A digital signature is used to achieve the same objectives for documents in digital form. What properties must a digital signature have? Similar to a manual signature, a digital signature must be verifiable and non-forgeable. What do we mean by verifiable and non-forgeable?

Verifiability means that it must be possible to prove that a person's signature on a document is indeed that person's signature. So, it must be possible for some other person to verify that a person's signature on a document is indeed that person's signature. And what does non-forgeability mean? Non-forgeability means that no one should be able to create a person's digital signature except the person himself or herself. For example, only Alice should be able to create Alice's digital signature.

Bob should not be able to create Alice's digital signature. We'll discuss various schemes for implementing digital signatures. Recall that some desirable properties of secure communication are confidentiality, message integrity, and authentication. And as we said, we'll discuss several mechanisms that can be used to achieve the above properties. We'll then discuss several practical systems that use these mechanisms to provide security in the internet.

In particular, we'll discuss PGP - Pretty Good Privacy, and Secure Multipurpose Internet Mail Extension, or S/MIME, for securing email. We'll discuss mechanisms for achieving security at the transport layer of the protocol stack. Specifically, we'll discuss secure sockets layer (SSL), and transport layer security (TLS), for securing TCP connections. TCP is the transmission control protocol. It is a popular transport layer protocol used in the internet.

We'll also discuss IPSec and virtual private networks for network layer security. We all use Wi-Fi extensively. We'll discuss various mechanisms for securing wireless LANs or Wi-Fi. In particular, we'll discuss 802.11i and 802.11w, which are protocols for securing Wi-Fi. We'll also discuss protocols for securing wireless cellular networks, including 2G, 3G, 4G, and 5G cellular networks.

Next we'll discuss operational security. What do we mean by operational security? Most organizations, such as universities and companies, have networks connected to the public internet. So, in this context, attackers may attempt to launch various attacks on the networks of an organization. For example, they may attempt to infect machines with malware.

They may try to obtain corporate secrets. They can map the internal network configurations, and later use that information to launch other attacks. They can launch denial of service attacks on an organization's network. We'll discuss various mechanisms for preventing such attacks on organization networks. In particular, we'll discuss firewalls.

A firewall is illustrated in this picture. A firewall is a device, which sits at the boundary of an organization's network. The firewall resides between the organization's network, which is on the left side, and the public internet. It filters the packets which go through the firewall. So, for example, some of the packets that are sent from the public internet to the organization's network may be filtered out by the firewall.

And when some sensitive information is sent from computers inside the organization's network into the public internet, some of that sensitive information might be blocked by the firewall. So, related to firewalls, are other devices known as intrusion detection systems. So, firewalls and intrusion detection systems can be used to detect and/or prevent attacks on the networks of organizations. We'll discuss firewalls and intrusion detection systems in this course. Next we'll discuss security of the Internet of Things.

The Internet of Things, or IoT, extends internet connectivity from traditional devices such as desktops, laptops, and smartphones to devices and appliances that were traditionally not connected to the internet. For example, sensors, vehicles, and appliances. IoT enables a lot of useful applications. Some examples are agriculture. In this picture, these circles, white circles in this area, they are sensors.

And this cloud shown here, that is a farm. These sensors monitor various quantities of interest such as soil moisture, temperature, and humidity in a farm, at various points in a farm. And these measurements can be used to control irrigation and application of

fertilizers, and so on. For example, this sensor might detect that the soil moisture in its vicinity has become low. So, it sends an alert to the user, and the user can spray some water in this area to raise the soil moisture level back to the normal level, or this might be automatic.

So, the sensor might send an alert to a controller which might activate the irrigation, in this part. So, using this technique, we can have precision agriculture, which means that resources such as water and fertilizer, they are applied judiciously only in those parts in which they are required. Another application of IoT is healthcare. Sensors can be used for remote health monitoring. In this application, sensors are attached to a patient's body, and these sensors monitor various quantities, such as blood pressure, heart rate, and so on.

And these measurements are sent to a system in a hospital. And then, if these measurements indicate that there is some abnormality, in that case, some action can be taken, such as some medical personnel can be deployed to the patient's home or a warning can be sent to the patient and so on. So, these sensors can be used to take actions based on the results of the measurements. Another set of applications is in smart buildings and homes. So, one example application is that lighting, heating, and air conditioning can be controlled automatically.

There are sensors which detect the presence of human beings, and lights, heating, air conditioning, etc., are switched on only in rooms where humans are present, and they're automatically switched off when there are no humans. Another example of an application is smartphones can be used to automatically switch on or switch off air conditioning or media systems. So, to summarize, all these applications are possible because sensors, actuators, and other devices are connected to the internet. So, these applications are enabled by the Internet of Things. We'll discuss security in the Internet of Things.

There are various challenges in IoT security. One challenge is that many of these connected devices are small, inexpensive, and battery operated. So, they have limited processing power, memory, and access to energy. As a result, traditional security solutions, which were designed for powerful devices, such as desktops, laptops, and smartphones, cannot be directly used to provide IoT security. We need to design custom solutions for IoT security.

So, in this course we'll discuss the security of the IoT. We'll design security solutions which minimize the communication overhead, memory, and processing requirements. Another topic that we'll discuss in this course is cryptocurrencies and blockchain. We know that

Bitcoin and other cryptocurrencies, for example, Litecoin and Ethereum, are being extensively used by users around the world. For example, it was estimated that in 2017, there were 2.9 to 5.8 million unique users using a cryptocurrency wallet, most of them using Bitcoin.

What is a cryptocurrency? A cryptocurrency uses cryptographic techniques to regulate the generation of units of the currency, and to implement their transfer. In contrast to the usual currency that we use, cryptocurrencies are not issued by a central authority, such as a bank. Some example uses of cryptocurrencies are as follows. They can be used as an investment. For example, we can invest some amount in Bitcoin, and if the value goes up, then we gain money.

Another example is international monetary transfers. Money can be transferred from one nation to another fast and with lower transaction fees than using traditional means. Another example used is that cryptocurrencies can be used as an alternative source of wealth that cannot be frozen by authorities such as governments. So, governments can freeze bank accounts, or they can freeze the properties of an individual, such as real estate, but they cannot freeze cryptocurrencies. Bitcoin is a decentralized system that is a distributed system based on the blockchain technology.

What is blockchain? Blockchain is a database of all the past transactions. In the context of Bitcoin, the blockchain records the creation of new Bitcoins and the transfers of Bitcoins. It is difficult for a malicious user to modify the transactions stored in the blockchain. It is extremely difficult to modify transactions, which results in the security of the blockchain.

And a copy of the blockchain is stored at multiple nodes connected to the internet. We'll study cryptocurrencies with a focus on Bitcoin, which is the most popular cryptocurrency currently. Another topic we'll study is anonymous connections and onion routing. To motivate this, suppose a user Alice wants to visit a controversial website, for example, a political activist site. Alice does not want to reveal her IP address to the website.

If the website is a political activist site, and if there is some legal action on the website, then Alice would get into trouble. So, Alice does not want to reveal her IP address to the website. Alice also does not want her local internet service provider to know that she's visiting the website. And also Alice does not want her local ISP to see the data she's exchanging with the website. So, how do we achieve these properties?

So, the third property is not difficult to achieve. We can use encryption to encrypt the data that Alice is exchanging with the website. This will prevent her local ISP from seeing the data she's exchanging. But despite encryption, the local ISP will know which website Alice is visiting, and also Alice's IP address will be known to the website. So, how do we achieve the first two properties?

We'll discuss various techniques for achieving such anonymous connections. In particular, we'll study TOR, which stands for "The Onion Router". It's a widely used scheme for enabling anonymous connections, such as the one in our example. Some example applications of TOR are as follows. Individuals use TOR to connect to news sites, political activist sites, sites like YouTube or Facebook, and so on, when these are blocked by their local internet providers or the countries they live in.

Journalists use TOR to communicate more safely with whistle-blowers and dissidents. In such contexts, clearly anonymity is very important. Militaries use TOR. Why do militaries need to use TOR? Suppose some military personnel are operating in a country, and insurgents may monitor internet traffic and discover all the hotels and other locations from which people are connecting to known military servers.

Then, they can find out the locations of military personnel, and then attack them. So, military field agents deployed away from home use Tor to mask the sites that they are visiting. They can build anonymous connections between themselves and the military servers, and this prevents insurgents from finding out their locations. So, these are some of the topics that we'll discuss in this course. This is not an exhaustive list.

For example, we'll discuss cloud security and post-quantum cryptography as well. Yeah, so this is not an exhaustive list, and there are several other topics that we will also discuss in this course. Thank you.