

CS 458 / 658

Computer Security and Privacy

Module 5
Internet Application Security and Privacy

Spring 2016

Module outline

- ① Basics of cryptography
- ② Secret-key encryption
- ③ Public-key encryption
- ④ Integrity
- ⑤ Authentication

5-2

Module outline

- ⑥ Security controls using cryptography
- ⑦ Link-layer security
- ⑧ Network-layer security
- ⑨ Transport-layer security and privacy
- ⑩ Application-layer security and privacy

5-3

Module outline

- ① Basics of cryptography
- ② Secret-key encryption
- ③ Public-key encryption
- ④ Integrity
- ⑤ Authentication

5-4

Cryptography

- What is cryptography?
- Related fields:
 - Cryptography (“secret writing”): Making secret messages
 - Turning **plaintext** (an ordinary readable message) into **ciphertext** (secret messages that are “hard” to read)
 - Cryptanalysis: Breaking secret messages
 - Recovering the plaintext from the ciphertext
- Cryptology is the science that studies these both
- The point of cryptography is to send secure messages over an insecure medium (like the Internet)

5-5

The scope of these lectures

- The goal of the cryptography unit in this course is to show you what cryptographic tools exist, and information about using these tools in a secure manner
- We won’t be showing you details of how the tools work
 - For that, see CO 487, or chapter 12 of the text
- The overview of public-key encryption and digital signatures in the textbook is poor.

5-6

Dramatis Personae

- When talking about cryptography, we often use a standard cast of characters
- Alice, Bob, Carol, Dave
 - People (usually honest) who wish to communicate
- Eve
 - A passive eavesdropper, who can listen to any transmitted messages
- Mallory
 - An active Man-In-The-Middle, who can listen to, **and modify, insert, or delete**, transmitted messages
- Trent
 - A Trusted Third Party

5-7

Building blocks

- Cryptography contains three major types of components
 - Confidentiality components
 - Preventing Eve from **reading** Alice's messages
 - Integrity components
 - Preventing Mallory from **modifying** Alice's messages without being detected
 - Authenticity components
 - Preventing Mallory from **impersonating** Alice

5-8

Kerckhoffs' Principle (19th c.)

- **The security of a cryptosystem should not rely on a secret that's hard (or expensive) to change**
- So don't have secret encryption methods
 - Then what do we do?
 - Have a large class of encryption methods, instead
 - Hopefully, they're all equally strong
 - Make the class **public** information
 - Use a secret **key** to specify which one you're using
 - It's easy to change the key; it's usually just a smallish number

5-9

Kerckhoffs' Principle (19th c.)

- This has a number of implications:
 - The system is at **most** as secure as the number of keys
 - Eve can just try them all, until she finds the right one
 - A **strong cryptosystem** is one where that's the best Eve can do
 - With weaker systems, there are shortcuts to finding the key
 - Example: newspaper cryptogram has 403,291,461,126,605,635,584,000,000 possible keys
 - But you don't try them all; it's way easier than that!

5-10

Strong cryptosystems

- What information do we assume the attacker (Eve) has when she's trying to break our system?
- She may:
 - Know the **algorithm** (the public class of encryption methods)
 - Know some **part of the plaintext**
 - Know a number (maybe a large number) of corresponding **plaintext/ciphertext pairs**
 - Have access to an encryption and/or decryption **oracle**
- And we still want to prevent Eve from learning the key!

5-11


Module outline

- ① Basics of cryptography
- ② Secret-key encryption
- ③ Public-key encryption
- ④ Integrity
- ⑤ Authentication

5-12

Secret-key encryption


- Secret-key encryption is the simplest form of cryptography
- Also called symmetric encryption
- Used for thousands of years
- The key Alice uses to encrypt the message is the same as the key Bob uses to decrypt it



5-13

Secret-key encryption

- Eve, not knowing the key, should not be able to recover the plaintext



5-14

Perfect secret-key encryption

- Is it possible to make a completely unbreakable cryptosystem?
- Yes: the **One-Time Pad**
- It's also very simple:
 - The key is a truly random bitstring of the same length as the message
 - The "Encrypt" and "Decrypt" functions are each just XOR

5-15

One-time pad

- Q: Why does “try every key” not work here?
- It's very hard to use correctly
 - The key must be **truly random**, not pseudorandom
 - The key must **never be used more than once!**
 - A “two-time pad” is **insecure!**
- Q: How do you share that much secret key?
- Used in the Washington / Moscow hotline for many years

5-16

Computational security

- In contrast to OTP's “perfect” or “information-theoretic” security, most cryptosystems have “computational” security
 - This means that it's certain they can be broken, given enough work by Eve
- How much is “enough”?
- At **worst**, Eve tries every key
 - How long that takes depends on how long the keys are
 - But it only takes this long if there are no “shortcuts”!

5-17

Some data points

- One computer can try about 17 million keys per second
- A medium-sized corporate or research lab may have 100 computers
- The BOINC project has 13 million computers



- Remember that most computers are idle most of the time (they're waiting for you to type something); getting them to crack keys in their spare time doesn't actually cost anything extra

5-18

40-bit crypto

- This was the US legal export limit for a long time
- $2^{40} = 1,099,511,627,776$ possible keys
- One computer: 18 hours
- One lab: 11 minutes
- BOINC: 5 ms


5-19

56-bit crypto

- This was the US government standard (DES) for a long time
- $2^{56} = 72,057,594,037,927,936$ possible keys
- One computer: 134 years
- One lab: 16 months
- BOINC: 5 minutes

5-20

Cracking DES



"DES cracker" machine of Electronic Frontier Foundation

5-21

128-bit crypto

- This is the modern standard
- $2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456$ possible keys
- One computer: 635 thousand million million million years
- One lab: 6 thousand million million million years
- BOINC: 49 thousand million million years

5-22

Well, we cheated a bit

- This isn't really true, since computers get faster over time
 - A better strategy for breaking 128-bit crypto is just to wait until computers get 2^{88} times faster, then break it on one computer in 18 hours.
 - How long do we wait? Moore's law says 132 years.
 - If we believe Moore's law will keep on working, we'll be able to break 128-bit crypto in 132 years (and 18 hours) :-)
 - Q: Do we believe this?


5-23

An even better strategy

- Don't break the crypto at all!
- There are always weaker parts of the system to attack
 - Remember the Principle of Easiest Penetration
- The point of cryptography is to make sure the information transfer is not the weakest link

5-24

Rubber hose cryptanalysis



5-25


Types of secret-key cryptosystems

- Secret-key cryptosystems come in two major classes
 - Stream ciphers
 - Block ciphers

5-26

Stream ciphers

- A stream cipher is what you get if you take the One-Time Pad, but use a pseudorandom keystream instead of a truly random one



- **RC4** is the most commonly used stream cipher on the Internet today

5-27

Stream ciphers

- Stream ciphers can be very fast
 - This is useful if you need to send a **lot** of data securely
- But they can be tricky to use correctly!
 - What happens if you use the same key to encrypt two different messages?
 - How would you solve this problem without requiring a new shared secret key for each message? Where have we seen this technique before?
- WEP, PPTP are great examples of how **not** to use stream ciphers

5-28


Block ciphers

- Note that stream ciphers operate on the message one bit at a time
- What happens in a stream cipher if you change just one bit of the plaintext?
- We can also use block ciphers
 - Block ciphers operate on the message one block at a time
 - Blocks are usually 64 or 128 bits long
- **AES** is the block cipher everyone should use today
 - Unless you have a really, really good reason

5-29

Modes of operation

- Block ciphers work like this:




- But what happens when the plaintext is larger than one block?
 - The choice of what to do with multiple blocks is called the **mode of operation** of the block cipher

5-30

Modes of operation

- The simplest thing to do is just to encrypt each successive block separately.
 - This is called Electronic Code Book (**ECB**) mode

- But if there are repeated blocks in the plaintext, you'll see the same repeating patterns in the ciphertext:




5-31

Modes of operation

- There are much better modes of operation to choose from
 - Common ones include Cipher Block Chaining (**CBC**), Counter (**CTR**), and Galois Counter (**GCM**) modes

- Patterns in the plaintext are no longer exposed
- But you need an **IV** (Initial Value), which acts much like a salt



5-32

Key exchange

- How do Alice and Bob share the secret key?
 - Meet in person; diplomatic courier
 - In general this is very hard
- Or, we invent new technology...

Module outline

- ① Basics of cryptography
- ② Secret-key encryption
- ③ Public-key encryption
- ④ Integrity
- ⑤ Authentication

5-34

Public-key cryptography

- Invented (in public) in the 1970's
- Also called asymmetric cryptography
 - Allows Alice to send a secret message to Bob **without** any prearranged shared secret!
 - In secret-key cryptography, the same (or a very similar) key encrypts the message and also decrypts it
 - In public-key cryptography, there's one key for encryption, and a **different** key for decryption!
- Some common examples:
 - RSA, ElGamal, ECC, NTRU, McEliece


5-35

Public-key cryptography

- How does it work?
 - Bob gives everyone a copy of his public encryption key. Alice uses it to encrypt a message, and sends the encrypted message to Bob
 - Bob uses his private decryption key to decrypt the message.
 - Eve can't decrypt it; she only has the encryption key.
 - Neither can Alice!
- So with this, Alice just needs to know Bob's public key in order to send him secret messages
 - These public keys can be published in a directory somewhere

5-36

Public-key cryptography



5-37

Public key sizes

- Recall that if there are no shortcuts, Eve would have to try 2^{128} things in order to read a message encrypted with a 128-bit key
- Unfortunately, all of the public-key methods we know **do** have shortcuts
 - Eve could read a message encrypted with a 128-bit RSA key with just 2^{33} work, which is **easy**!
 - If we want Eve to have to do 2^{128} work, we need to use a much longer public key

5-38

Public key sizes

Comparison of key sizes for roughly equal strength

<u>AES</u>	<u>RSA</u>	<u>ECC</u>
80	1024	160
116	2048	232
128	2600	256
160	4500	320
256	14000	512

5-39

Hybrid cryptography

- In addition to having longer keys, public-key cryptography takes a long time to calculate (as compared to secret-key cryptography)
 - Using public-key to encrypt large messages would be too slow, so we take a hybrid approach:
 - Pick a random 128-bit key K for a secret-key cryptosystem
 - Encrypt the large message with the key K (e.g., using AES)
 - Encrypt the key K using a public-key cryptosystem
 - Send the encrypted message and the encrypted key to Bob
 - This hybrid approach is used for almost every cryptography application on the Internet today

5-40

Is that all there is?

- It seems we've got this "sending secret messages" thing down pat. What else is there to do?
 - Even if we're safe from Eve reading our messages, there's still the matter of Mallory
 - It turns out that even if our messages are encrypted, Mallory can sometimes modify them in transit!
 - Mallory won't necessarily know what the message says, but can still change it in an undetectable way
 - e.g. bit-flipping attack on stream ciphers
 - This is counterintuitive, and often forgotten
 - The textbook even gets this wrong!
- How do we make sure that Bob gets the same message Alice sent?

5-41

Module outline

- ① Basics of cryptography
- ② Secret-key encryption
- ③ Public-key encryption
- ④ Integrity
- ⑤ Authentication

5-42

Integrity components

- How do we tell if a message has changed in transit?
- Simplest answer: use a checksum
 - For example, add up all the bytes of a message
 - The last digits of serial numbers (credit card, ISBN, etc.) are usually checksums
 - Alice computes the checksum of the message, and sticks it at the end before encrypting it to Bob. When Bob receives the message and checksum, he verifies that the checksum is correct

5-43

This doesn't work!

- With most checksum methods, Mallory can easily change the message in such a way that the checksum stays the same
- We need a “cryptographic” checksum
- It should be hard for Mallory to find a second message with the same checksum as any given one

5-44

Cryptographic hash functions

- A **hash function** h takes an arbitrary length string x and computes a fixed length string $y = h(x)$ called a **message digest**
 - Common examples: MD5, SHA-1, SHA-2, SHA-3 (AKA Keccak, from 2012 on)
- Hash functions should have three properties:
 - Preimage-resistance:
 - Given y , it's hard to find x such that $h(x) = y$ (i.e., a “preimage” of x)
 - Second preimage-resistance:
 - Given x , it's hard to find $x' \neq x$ such that $h(x) = h(x')$ (i.e., a “second preimage” of $h(x)$)
 - Collision-resistance:
 - It's hard to find any two distinct values x, x' such that $h(x) = h(x')$ (a “collision”)

5-45

What is “hard”?

- For SHA-1, for example, it takes 2^{160} work to find a preimage or second preimage, and 2^{80} work to find a collision using a brute-force search
 - However, there are faster ways than brute force to find collisions in SHA-1 or MD5
- Collisions are always easier to find than preimages or second preimages due to the well-known birthday paradox

5-46

Cryptographic hash functions

- You can't just send an unencrypted message and its hash to get integrity assurance
 - Even if you don't care about confidentiality!
- Mallory can change the message and just compute the new message digest herself

5-47

Cryptographic hash functions

- Hash functions provide integrity guarantees only when there is a secure way of sending and/or storing the message digest
 - For example, Bob can publish a hash of his public key (i.e., a message digest) on his business card
 - Putting the whole key on there would be too big
 - But Alice can download Bob's key from the Internet, hash it herself, and verify that the result matches the message digest on Bob's card
- What if there's no external channel to be had?
 - For example, you're using the Internet to communicate

5-48

Module outline

- 1 Basics of cryptography
- 2 Secret-key encryption
- 3 Public-key encryption
- 4 Integrity
- 5 Authentication


5-49

Message authentication codes

- We do the same trick as for encryption: have a large class of hash functions, and use a shared secret key to pick the “correct” one
- Only those who know the secret key can generate, or even check, the computed hash value (sometimes called a **tag**)
- These “keyed hash functions” are usually called **Message Authentication Codes**, or **MACs**
- Common examples:
 - SHA-1-HMAC, SHA-256-HMAC, CBC-MAC

5-50

Message authentication codes



5-51

Combining ciphers and MACs

- In practice we often need both confidentiality and message integrity
- There are multiple strategies to combine a cipher and a MAC when processing a message
 - Encrypt-then-MAC, MAC-then-Encrypt, Encrypt-and-MAC
- Encrypt-then-MAC is the recommended strategy
- Ideally your crypto library already provides an **authenticated encryption mode** that securely combines the two operations so you don't have to worry about getting it right
 - E.g., GCM, CCM (used in WPA2, see later), or OCB mode

5-52

Repudiation

- Suppose Alice and Bob share a MAC key K , and Bob receives a message M along with a valid tag T that was computed using the key K
 - Then Bob can be assured that Alice is the one who sent the message M , and that it hasn't been modified since she sent it!
 - This is like a "signature" on the message
 - But it's not quite the same!
 - Bob can't show M and the tag T to Carol to prove Alice sent the message M

5-53

Repudiation

- Alice can just claim that Bob made up the message M , and calculated the tag T himself
- This is called **repudiation**; and we sometimes want to avoid it
- Some interactions should be repudiable
 - Private conversations
- Some interactions should be non-repudiable
 - Electronic commerce

5-54

Digital signatures

- For non-repudiation, what we want is a true **digital signature**, with the following properties:
- If Bob receives a message with Alice's digital signature on it, then:
 - Alice, and not an impersonator, sent the message (like a MAC)
 - the message has not been altered since it was sent (like a MAC), and
 - Bob can prove these facts to a third party (additional property not satisfied by a MAC).
- How do we arrange this?
 - Use similar techniques to public-key cryptography

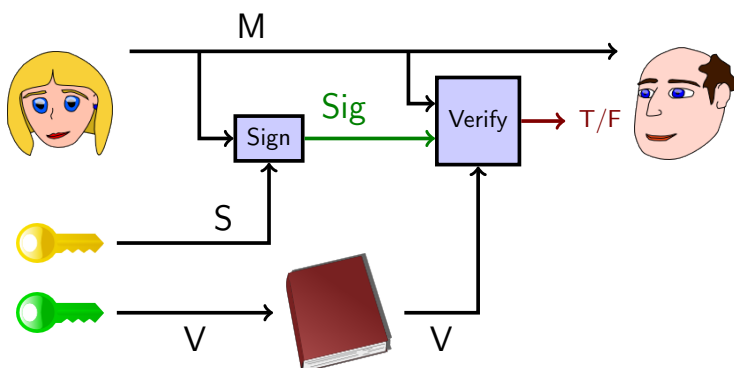
5-55

Making digital signatures

- Remember public-key crypto:
 - Separate keys for encryption and decryption
 - Give everyone a copy of the encryption key
 - The decryption key is private
- To make a digital signature:
 - Alice signs the message with her private **signature key**
- To verify Alice's signature:
 - Bob verifies the message with his copy of Alice's public **verification key**
 - If it verifies correctly, the signature is valid

5-56

Making digital signatures



5-57

Hybrid signatures

- Just like encryption in public-key crypto, signing large messages is slow
- We can also hybridize signatures to make them faster:
 - Alice sends the (unsigned) message, and also a signature on a **hash** of the message
 - The hash is much smaller than the message, and so it is faster to sign and verify
- Remember that authenticity and confidentiality are separate; if you want both, you need to do both

5-58

Combining public-key encryption and digital signatures

- Alice has two **different** key pairs: an (encryption, decryption) key pair and a (signature, verification) key pair
 - So does Bob
- Alice uses Bob's encryption key to encrypt a message destined for Bob
- She uses her signature key to sign the ciphertext
- Bob uses Alice's verification key to check the signature
- He uses his decryption key to decrypt the ciphertext
- Similarly for reverse direction

5-59

Relationship between key pairs

- Alice's (signature, verification) key pair is long-lived, whereas her (encryption, decryption) key pair is short-lived
 - Gives perfect forward secrecy (see later)
- When creating a new (encryption, decryption) key pair, Alice uses her signing key to sign her new encryption key and Bob uses Alice's verification key to verify the signature on this new key
- If Alice's communication with Bob is interactive, she can use secret-key encryption and does not need an (encryption, decryption) key pair at all (see TLS or SSH)

5-60

The Key Management Problem

- One of the hardest problems of public-key cryptography is that of **key management**
- How can Bob find Alice's verification key?
 - He can know it personally (**manual keying**)
 - SSH does this
 - He can trust a friend to tell him (**web of trust**)
 - PGP does this
 - He can trust some third party to tell him (**CA's**)
 - TLS / SSL do this

5-61

Certificate authorities

- A CA is a trusted third party who keeps a directory of people's (and organizations') verification keys
- Alice generates a (signature, verification) key pair, and sends the verification key, as well as a bunch of personal information, both signed with Alice's signature key, to the CA
- The CA ensures that the personal information and Alice's signature are correct
- The CA generates a **certificate** consisting of Alice's personal information, as well as her verification key. The entire certificate is signed with the CA's signature key

5-62

Certificate authorities

- Everyone is assumed to have a copy of the CA's verification key, so they can verify the signature on the certificate
- There can be multiple levels of certificate authorities; level n CA issues certificates for level n+1 CAs
 - Public-key infrastructure (PKI)
- Need to have only verification key of root CA to verify certificate chain

5-63

Putting it all together

- We have all these blocks; now what?
- Put them together into **protocols**
- This is HARD. Just because your pieces all work, doesn't mean what you build out of them will; you have to *use* the pieces correctly
- Common mistakes include:
 - Using the same stream cipher key for two messages
 - Assuming encryption also provides integrity
 - Falling for replay attacks or reaction attacks
 - *LOTS* more!

5-64

Module outline

⑥ Security controls using cryptography

⑦ Link-layer security

⑧ Network-layer security

⑨ Transport-layer security and privacy

⑩ Application-layer security and privacy

5-65

Security controls using cryptography

- Q: In what situations might it be appropriate to use cryptography as a security control?
- A: Remember that there needs to be some separation, since any secrets (like the key) need to be available to the legitimate users but not the adversaries
- In some situations, this may make secret-key crypto problematic
- If your web browser can decrypt its file containing your saved passwords, then an adversary who can read your web browser probably can, too
- Q: How is this solved in practice?

5-66

Program and OS security

- Using secret-key crypto can be problematic for the above reason
 - But public-key is OK, if the local machine only needs access to the public part of the key
 - So only encryption and signature verification; no decryption or signing
 - E.g., apps can be installed only if digitally signed by the vendor (BlackBerry) or upgraded only if signed by the original developer (Android)
 - OS may allow execution of programs only if signed (iOS)

5-67

Encrypted code

- There is research into processors that will only execute encrypted code
- The processor will decrypt instructions before executing them
- The decryption key is processor-dependent
- Malware won't be able to spread without knowing a processor's key
- Downsides?

5-68

Encrypted data

- Harddrive encryption protects data when laptop gets lost/stolen
- It often does not protect data against other users who legitimately use laptop
- Or somebody installing malware on laptop
- Or somebody (maybe physically) extracting the decryption key from the laptop's memory

5-69

OS authentication

- Authentication mechanisms often use cryptography
 - E.g., salted hashes (see Module 3)
- Unfortunately, people are bad at doing cryptography in their heads, so some mechanisms require hardware token



Photo from <http://itc.ua/>

5-70

Network security and privacy

- The primary use for cryptography
 - “Separating the security of the medium from the security of the message”
- Entities you can only communicate with over a network are inherently less trustworthy
 - They may not be who they claim to be

5-71

Network security and privacy

- Network cryptography is used at every layer of the network stack for both security and privacy applications:
 - Link
 - WEP, WPA, WPA2
 - Network
 - VPN, IPsec
 - Transport
 - TLS / SSL, Tor
 - Application
 - ssh, Mixminion, PGP, OTR

5-72

Module outline

- 6 Security controls using cryptography
- 7 Link-layer security
- 8 Network-layer security
- 9 Transport-layer security and privacy
- 10 Application-layer security and privacy

5-73

Link-layer security controls

- Intended to protect **local area networks**
- Widespread example: WEP (Wired Equivalent Privacy)
- WEP was intended to enforce three security goals:
 - Confidentiality
 - Prevent an adversary from learning the contents of your wireless traffic
 - Access Control
 - Prevent an adversary from using your wireless infrastructure
 - Data Integrity
- Unfortunately, **none** of these is actually enforced!

5-74

WEP description

- Brief description:
- The sender and receiver share a secret k
 - The secret k is either 40 or 104 bits long
- In order to transmit a message M :
 - Compute a checksum $c(M)$
 - this does not depend on k
 - Pick an IV (a random number) v and generate a keystream $RC4(v, k)$
 - XOR $\langle M, c(M) \rangle$ with the keystream to get the ciphertext
 - Transmit v and the ciphertext over the wireless link

5-75

Problem number 0: Widely shared “secrets”



http://www.theregister.co.uk/2014/06/25/brace_yourselves_brazil_dill_in_world_cup_wifi_spill/

5-76

WEP description

- Upon receipt of v and the ciphertext:
 - Use the received v and the shared k to generate the keystream $RC4(v, k)$
 - XOR the ciphertext with $RC4(v, k)$ to get $\langle M', c' \rangle$
 - Check to see if $c' = c(M')$
 - If it is, accept M' as the message transmitted
- Problem number 1: v is 24 bits long
 - Why is this a problem?

5-77

WEP data integrity

- Problem 2: the checksum used in WEP is CRC-32
 - Quite a poor choice; there's already a CRC in the protocol to detect random errors, and a CRC can't help you protect against malicious errors.
- The CRC has two important properties:
 - It is independent of k and v
 - It is **linear**: $c(M \text{ XOR } D) = c(M) \text{ XOR } c(D)$
- Why is linearity a pessimal property for your integrity mechanism to have when used in conjunction with a stream cipher?

5-78

WEP access control

- What if the adversary wants to inject a new message F onto a WEP-protected network?
- All he needs is a single plaintext/ciphertext pair
- This gives him a value of v and the corresponding keystream $RC4(v, k)$
- Then $C' = \langle F, c(F) \rangle \text{ XOR } RC4(v, k)$, and he transmits v, C'
- C' is in fact a correct encryption of F , so the message must be accepted

5-79

WEP authentication protocol

- How did the adversary get that single plaintext/ciphertext pair required for the attack on the previous slide?
 - Problem 3: It turns out the authentication protocol gives it to the adversary **for free!**
- This is a major disaster in the design!
- The authentication protocol (described on the next slide) is supposed to prove that a certain client knows the shared secret k
- But if I watch you prove it, I can turn around and execute the protocol myself!
 - “What’s the password?”

5-80

WEP authentication protocol

- Here’s the authentication protocol:
 - The access point sends a challenge string to the client
 - The client sends back the challenge, WEP-encrypted with the shared secret k
 - The wireless access point checks if the challenge is correctly encrypted, and if so, accepts the client
- So the adversary has just seen both the plaintext and the ciphertext of the challenge
- Problem number 4: this is enough not only to inject packets (as in the previous attack), but also **to execute the authentication protocol himself!**

5-81

WEP decryption

- Somewhat surprisingly, the ability to modify and inject packets also leads to ways the adversary can **decrypt** packets!
 - The access point knows k ; it turns out the adversary can trick it into decrypting the packet for him and telling him the result.
- Note that none of the attacks so far:
 - Used the fact that the stream cipher was RC4 specifically
 - Recovered k

5-82

Recovering a WEP key

- Since 2002, there have been a series of analyses of RC4 in particular
 - Problem number 5: it turns out that when RC4 is used with similar keys, the output keystream has a subtle weakness
 - And this is (often) how WEP uses RC4!
- These observations have led to programs that can recover either a 104-bit or 40-bit WEP key in **under 60 seconds**, most of the time
 - See the optional reading for more information on this

5-83

Replacing WEP

- Wi-fi Protected Access (WPA) was rolled out as a short-term patch to WEP while formal standards for a replacement protocol (IEEE 802.11i, later called WPA2) were being developed
- WPA:
 - Replaces CRC-32 with a real MAC (here called a MIC to avoid confusion with a Media Access Control address)
 - IV is 48 bits
 - Key is changed frequently (TKIP)
 - Ability to use 802.1x authentication server
 - But maintains less-secure PSK (Pre-Shared Key) mode for home users
 - Able to run on most older WEP hardware

5-84

Replacing WEP

- The 802.11i standard was finalized in 2004, and the result (called WPA2) has been required for products calling themselves “Wi-fi” since 2006
- WPA2:
 - Replaces the RC4 and MIC algorithms in WPA with the CCM authenticated encryption mode (using AES)
 - Considered strong, except in PSK mode
 - Dictionary attacks still possible

5-85

Module outline

- 6 Security controls using cryptography
- 7 Link-layer security
- 8 Network-layer security
- 9 Transport-layer security and privacy
- 10 Application-layer security and privacy

5-86

Network-layer security

- Suppose every link in our network had strong link-layer security
- Why would this not be enough?
- We need security **across** networks
 - Ideally, **end-to-end**
- At the network layer, this is usually accomplished with a Virtual Private Network (VPN)

5-87

Virtual Private Networks

- Connect two (or more) networks that are physically isolated, and make them appear to be a single network
 - Alternately: connect a single remote host (often a laptop) to one network
- Goal: adversary between the networks should not be able to read or modify the traffic flowing across the VPN
 - But DoS and some traffic analysis still usually possible

5-88

Setting up a VPN

- One host on each side is the **VPN gateway**
 - Could be the firewall itself, or could be in DMZ
 - In the laptop scenario, it will of course be the laptop itself on its side
- Traffic destined for the “other side” is sent to the local VPN gateway
- The local VPN gateway uses cryptography (encryption and integrity techniques) to send the traffic to the remote VPN gateway
 - Often by **tunnelling** (see next slide)
- The remote VPN gateway decrypts the messages and sends them on to their appropriate destinations

5-89

Tunnelling

- Tunnelling is the sending of messages of one protocol inside (that is, as the payload of) messages of another protocol, out of their usual protocol nesting sequence
- So TCP-over-IP **is not** tunnelling, since you're supposed to send TCP (a transport protocol) over IP (a network protocol; one layer down in the stack)
- But IP-over-TCP **is** tunnelling (going up the stack instead of down), as are IP-over-IP (same place in the stack), and PPP (a link layer protocol; bottom of the stack) over DNS (an application layer protocol; top of the stack)

5-90

IPsec

- One standard way to set up a VPN is by using **IPsec**
- Many corporate VPNs use this (open) protocol
- Two modes:
 - **Transport** mode
 - Useful for connecting a single laptop to a home network
 - Only the contents of the original IP packet are encrypted and authenticated
 - **Tunnel** mode
 - Useful for connecting two networks
 - The contents **and the header** of the original IP packet are encrypted and authenticated; result is placed inside a new IP packet destined for the remote VPN gateway

5-91

Other styles of VPNs

- In addition to IPsec, there are a number of other standard ways to set up a VPN
- Microsoft's PPTP was an older protocol
 - It had about as many design flaws as WEP
 - Most users now migrating to IPsec
- VPNs based on ssh
 - Tunnel PPP over ssh
 - That is, IP-over-PPP-over-ssh-over-TCP-over-IP
 - Some efficiency concern, but extremely easy to set up on a standard Unix/Linux box
 - OpenSSH v4 supports IP-over-SSH tunnelling directly

5-92

Module outline

⑥ Security controls using cryptography

⑦ Link-layer security

⑧ Network-layer security

⑨ Transport-layer security and privacy

⑩ Application-layer security and privacy

5-93

Transport-layer security and privacy

- Network-layer security mechanisms arrange to send individual IP packets securely from one network to another
- Transport-layer security mechanisms transform arbitrary TCP connections to add security
 - And similarly for “privacy” instead of “security”
- The main transport-layer security mechanism:
 - TLS (formerly known as SSL)
- The main transport-layer privacy mechanism:
 - Tor

5-94

TLS / SSL

- In the mid-1990s, Netscape invented a protocol called Secure Sockets Layer (SSL) meant for protecting HTTP (web) connections
 - The protocol, however, was general, and could be used to protect **any** TCP-based connection
 - HTTP + SSL = **HTTPS**
- Historical note: there was a competing protocol called S-HTTP. But Netscape and Microsoft both chose HTTPS, so that’s the protocol everyone else followed
- SSL went through a few revisions, and was eventually standardized into the protocol known as **TLS** (Transport Layer Security, imaginatively enough)

5-95

TLS at a high level

- Client connects to server, indicates it wants to speak TLS, and which **ciphersuites** it knows
- Server sends its certificate to client, which contains:
 - Its host name
 - Its verification key
 - Some other administrative information
 - A signature from a Certificate Authority (CA)
- Server also chooses which ciphersuite to use

5-96

TLS at a high level (cont.)

- Client validates server's certificate
 - Is its signature from a CA whose public key is embedded in the client (i.e., browser)?
 - Does the host name in the certificate match the host name of the web site that client wants to download?
- Client and server run a key agreement protocol to establish keys for symmetric encryption and MAC algorithms from the chosen ciphersuite
 - Server signs its protocol messages with its signature key
- Communication now proceeds using chosen symmetric encryption and MAC algorithms

5-97

Certificate validation

This server could not prove that it is **convention.gop**; its security certificate is from **funyuns.com**. This may be caused by a misconfiguration or an attacker intercepting your connection.

[Proceed to convention.gop \(unsafe\)](#)

5-98

Security properties provided by TLS

- Server authentication
- Message integrity
- Message confidentiality
- Client authentication (optional)

5-99

The success of TLS

- Though designed as a security mechanism, TLS (including SSL) has become the most successful **Privacy Enhancing Technology (PET)** ever
- Why?
 - It comes with your browser
 - Which encouraged web server operators to bother paying \$\$ for their certificates
 - It just works, without you having to configure anything
 - Most of the time, it even protects the privacy of your communications
 - Increasingly important due to the success of WiFi

5-100

Privacy Enhancing Technologies

- So far, we've only used encryption to protect the **contents** of messages
- But there are other things we might want to protect as well!
- We may want to protect the **metadata**
 - Who is sending the message to whom?
 - If you're seen sending encrypted messages to Human Rights Watch, bad things may happen
- We may want to hide the **existence** of the message
 - If you're seen sending encrypted messages at all, bad things may happen

5-101


Tor

- **Tor** is another successful privacy enhancing technology that works at the transport layer
 - Hundreds of thousands of users
 - March 2015: \approx 2 million users
- Normally, a TCP connection you make on the Internet automatically reveals your IP address
 - Why?
- Tor allows you to make TCP connections **without** revealing your IP address
- It's most commonly used for HTTP (web) connections

5-102

How Tor works


- Scattered around the Internet are about 7,000 Tor nodes, also called **Onion Routers**
- Alice wants to connect to a web server without revealing her IP address



5-103

How Tor works


- Alice picks one of the Tor nodes (n_1) and uses public-key cryptography to establish an encrypted communication channel to it (much like TLS)



5-104

How Tor works


- Alice tells n_1 to contact a second node (n_2), and establishes a new encrypted communication channel to n_2 , tunnelled within the previous one to n_1



5-105

How Tor works


- Alice tells n_2 to contact a third node (n_3), and establishes a new encrypted communication channel to n_3 , tunnelled within the previous one to n_2



5-106

How Tor works

- And so on, for as many steps as she likes (usually 3)
- Alice tells the last node (within the layers of tunnels) to connect to the website



5-107

Sending messages with Tor

- Alice now shares three secret keys:
 - K_1 with n_1
 - K_2 with n_2
 - K_3 with n_3
- When Alice wants to send a message M , she actually sends $E_{K_1}(E_{K_2}(E_{K_3}(M)))$
- Node n_1 uses K_1 to decrypt the outer layer, and passes the result $E_{K_2}(E_{K_3}(M))$ to n_2
- Node n_2 uses K_2 to decrypt the next layer, and passes the result $E_{K_3}(M)$ to n_3
- Node n_3 uses K_3 to decrypt the final layer, and sends M to the website

5-108

Replies in Tor

- When the website replies with message R , it will send it to node n_3
 - Why?
- Node n_3 will **encrypt** R with K_3 and send $E_{K_3}(R)$ to n_2
- Node n_2 will encrypt that with K_2 and send $E_{K_2}(E_{K_3}(R))$ to n_1
- Node n_1 will encrypt that with K_1 and send $E_{K_1}(E_{K_2}(E_{K_3}(R)))$ to Alice
- Alice will use K_1 , K_2 , and K_3 to decrypt the layers of the reply and recover R

5-109

Who knows what?

- Notice that node n_1 knows that Alice is using Tor, and that her next node is n_2 , but does not know which website Alice is visiting
- Node n_3 knows some Tor user (with previous node n_2) is using a particular website, but doesn't know who
- The website itself only knows that it got a connection from Tor node n_3
- **Note:** the connection between n_3 and the website is **not encrypted!** If you want encryption as well as the benefits of Tor, you should use end-to-end encryption **in addition**
 - Like HTTPS

5-110

Anonymity vs. pseudonymity

- Tor provides for **anonymity** in TCP connections over the Internet, both **unlinkably** (long-term) and **linkably** (short-term)
- What does this mean?
 - There's no long-term identifier for a Tor user
 - If a web server gets a connection from Tor today, and another one tomorrow, it won't be able to tell whether those are from the same person
 - But two connections in quick succession from the same Tor node are more likely to in fact be from the same person

5-111

The Nymity Slider

- We can place transactions (both online and offline) on a continuum according to the level of **nymity** they represent:
 - **Verinymity**
 - Government ID, SIN, credit card #, address
 - **Persistent pseudonymity**
 - Noms de plume, many blogs
 - **Linkable anonymity**
 - Prepaid phone cards, loyalty cards
 - **Unlinkable anonymity**
 - Cash payments, Tor

5-112

The Nymity Slider

- If you build a system at a certain level of nymity, it's **easy** to modify it to have a higher level of nymity, but **hard** to modify it to have a lower level
- For example:
 - It's easy to add a loyalty card to a cash payment, or a credit card to a loyalty card
 - It's hard to remove identity information if you're paying by credit card
- The lesson: design systems with a low level of nymity fundamentally; adding more is easy

5-113

Module outline

- 6 Security controls using cryptography
- 7 Link-layer security
- 8 Network-layer security
- 9 Transport-layer security and privacy
- 10 Application-layer security and privacy

5-114

Application-layer security and privacy

- TLS can provide for encryption at the TCP socket level
 - “End-to-end” in the sense of a network connection
 - Is this good enough? Consider SMTPS (SMTP/email over TLS)
- Many applications would like true end-to-end security
 - Human-to-human would be best, but those last 50 cm are really hard!
 - We usually content ourselves with desktop-to-desktop
- We’ll look at three particular applications:
 - Remote login, email, instant messaging

5-115

Secure remote login (ssh)

- You’re already familiar with this tool for securely logging in to a remote computer (the ugster machines)
- Usual usage (simplified):
 - Client connects to server
 - Server sends its verification key
 - The client **should** verify that this is the correct key
 - Client and server run a key agreement protocol to establish session keys, server signs its messages
 - All communication from here on in is encrypted and MACd with the session keys
 - Client authenticates to server
 - Server accepts authentication, login proceeds (under encryption and MAC)

5-116

Authentication with ssh

- There are two main ways to authenticate with ssh:
 - Send a password over the encrypted channel
 - The server needs to know (a hash of) your password
 - Sign a random challenge with your private signature key
 - The server needs to know your public verification key
- Which is better? Why?

5-117

Pretty Good Privacy

- The first popular implementation of public-key cryptography.
- Originally made by Phil Zimmermann in 1991
 - He got in a lot of trouble for it, since cryptography was highly controlled at the time.
 - But that's a whole 'nother story. :-)
- Today, there are many (more-or-less) compatible programs
 - GNU Privacy Guard (gpg), Hushmail, etc.

5-118

Pretty Good Privacy

- What does it do?
 - Its primary use is to protect the contents of email messages
- How does it work?
 - Uses public-key cryptography to provide:
 - Encryption of email messages (using hybrid encryption)
 - Digital signatures on email messages (hash-then-sign)

5-119

Recall

- In order to use public-key encryption and digital signatures, Alice and Bob must each have:
 - A public encryption key
 - A private decryption key
 - A private signature key
 - A public verification key

5-120

Sending a message

- To send a message to Bob, Alice will:
 - Write a message
 - Sign it with her own signature key
 - Encrypt both the message and the signature with Bob's public encryption key
- Bob receives this, and:
 - Decrypts it using his private decryption key to yield the message and the signature
 - Uses Alice's verification key to check the signature

5-121

Back to PGP

- PGP's main functions:
 - Create these four kinds of keys
 - encryption, decryption, signature, verification
 - Encrypt messages using someone else's encryption key
 - Decrypt messages using your own decryption key
 - Sign messages using your own signature key
 - Verify signatures using someone else's verification key
 - Sign other people's keys using your own signature key (see later)

5-122



to be extra safe, check that there's a big block of jumbled characters at tl

5-123

Obtaining keys

- Earlier, we said that Alice needs to get an authentic copy of Bob's public key in order to send him an encrypted message
- How does she do this?
 - In a secure way?
- Bob could put a copy of his public key on his webpage, but this isn't good enough to be really secure!
 - Why?

5-124

Verifying public keys

- If Alice knows Bob personally, she could:
 - Download the key from Bob's web page
 - Phone up Bob, and verify she's got the right key
 - Problem: keys are big and unwieldy!

```
mQGibDi5qEURBADitpDzvvzW+91j/zYgK78G3D76hvvvIT6gpTI1wg6WIJNLKJat
01yNpMIYNvpwi7EUd/1SN16t1/A022p7s7bDbE4T5NJda0IOAgWe0Z/plIJC4+o2
tD2RNuSkwDQcxzm8KUNZOJ1a4LvgRkm/oUubxyeY5omus7hcfNrB0wjC1wCg4Jnt
m7s3eNfMu72Cv+6FzBgFog8EANirkNdC1Q8oSMDihWj1ogiWbBz4s6HMxzAaqNf/
rCJ9qoK5SLFeoB/r5ksRwty9QKV4VdhhCIy1U2B9tST1EPYXJHQPZ3mwCxUnJpGD
8UgFM5uKXaEq2ppwArTm367k0tTpMQgXAN2HwiZv//ahQXH4ov30kBBVL5VfxMUL
UJ+yA/4r5HLTp2SbbqtPWdeW7uDwhe2dTqffAGuf0kuCpHwCTAhr83ivXzT/70M
```

5-125

Fingerprints

- Luckily, there's a better way!
- A **fingerprint** is a cryptographic hash of a key
- This, of course, is *much shorter*:
 - B117 2656 DFF9 83C3 042B C699 EB5A 896A 2898 8BF5
- Remember: there's no (known) way to make two different keys that have the same fingerprint, provided that we use a collision-resistant hash function

5-126

Fingerprints

- So now we can try this:
 - Alice downloads Bob's key from his webpage
 - Alice's software calculates the fingerprint
 - Alice phones up Bob, and asks him to read his key's actual fingerprint to her
 - If they match, Alice knows she's got an authentic copy of Bob's key
- That's great for Alice, but what about Carol, who doesn't know Bob
 - At least not well enough to phone him

5-127

Signing keys

- Once Alice has verified Bob's key, she uses her signature key to sign Bob's key
- This is effectively the same as Alice signing a message that says "I have verified that the key with fingerprint B117 2656 DFF9 83C3 042B C699 EB5A 896A 2898 8BF5 really belongs to Bob"
- Bob can attach Alice's signature to the key on his webpage

5-128

Key signing parties



5-129

Web of Trust

- Now Alice can act as an introducer for Bob
- If Carol doesn't know Bob, but does know Alice (and has already verified Alice's key, and trusts her to introduce other people):
 - she downloads Bob's key from his website
 - she sees Alice's signature on it
 - she is able to use Bob's key without having to check with Bob personally
- This is called the Web of Trust, and the PGP software handles it mostly automatically

5-130

So, great!

- So if Alice and Bob want to have a private conversation by email:
 - They each create their sets of keys
 - They exchange public encryption keys and verification keys
 - They send signed and encrypted messages back and forth
- Pretty Good, no?

5-131

Plot twist

- Suppose (encrypted) communications between Alice and Bob are recorded by the "bad guys"
 - criminals, competitors, subpoenaed by the RCMP
- Later, Bob's computer is stolen by the same bad guys
- Or just broken into
 - Virus, trojan, spyware
- All of Bob's key material is discovered
 - Oh, no!

5-132

The bad guys can...

- Decrypt past messages
- Learn their content
- Learn that Alice sent them
- And have a mathematical **proof** they can show to anyone else!

- How private is that?

5-133

What went wrong?

- Bob's computer got stolen?

- How many of you have never...
 - Left your laptop unattended?
 - Not installed the latest patches?
 - Run software with a remotely exploitable bug?

- What about your friends?

5-134

What really went wrong

- PGP creates lots of incriminating records:
 - Key material that decrypts data sent over the public Internet
 - Signatures with proofs of who said what

- Alice had better watch what she says!
 - Her privacy depends on Bob's actions

5-135

Casual conversations

- Alice and Bob talk in a room
- No one else can hear
 - Unless being recorded
- No one else knows what they say
 - Unless Alice or Bob tells them
- No one can prove what was said
 - Not even Alice or Bob
- These conversations are “off-the-record”

5-136

We like off-the-record conversations

- Legal support for having them
 - Illegal to record other people's conversations without notification
- We can have them over the phone
 - Illegal to tap phone lines
- But what about over the Internet?

5-137

Crypto tools

- We have the tools to do this
 - We've just been using the wrong ones
 - (when we've been using crypto at all)
- We want perfect forward secrecy
- We want deniable authentication

5-138

Perfect forward secrecy

- Future key compromises should not reveal past communication
- Use secret-key encryption with a short-lived key (a **session key**)
- The session key is created by a modified Diffie-Hellman protocol
- Discard the session key after use
 - Securely erase it from memory
- Use long-term keys only to authenticate the Diffie-Hellman protocol messages
- Q: Why does this approach not have the very same forward secrecy problem as PGP?

5-139

Deniable authentication

- Do **not** want digital signatures
 - Non-repudiation is great for signing contracts, but undesirable for private conversations
- But we **do** want authentication
 - We can't maintain privacy if attackers can impersonate our friends
- Use Message Authentication Codes
 - We talked about these earlier

5-140

No third-party proofs

- Shared-key authentication
 - Alice and Bob have the same MK
 - MK is required to compute the MAC
 - How is Bob assured that Alice sent the message?
- Bob cannot prove that Alice generated the MAC
 - He could have done it, too
 - Anyone who can verify can also forge
- This gives Alice a measure of deniability

5-141

Using these techniques

- Using these techniques, we can make our online conversations more like face-to-face “off-the-record” conversations
- But there’s a wrinkle:
 - These techniques require the parties to communicate interactively
 - This makes them unsuitable for email
 - But they’re still great for instant messaging!

5-142

Off-the-Record Messaging

- Off-the-Record Messaging (OTR) is software that allows you to have private conversations over instant messaging, providing:
 - Confidentiality
 - Only Bob can read the messages Alice sends him
 - Authentication
 - Bob is assured the messages came from Alice

5-143

Off-the-Record Messaging

- Perfect Forward Secrecy
 - Shortly after Bob receives the message, it becomes unreadable to anyone, anywhere
- Deniability
 - Although Bob is assured that the message came from Alice, he can’t convince Charlie of that fact
 - Also, Charlie can create forged transcripts of conversations that are every bit as accurate as the real thing

5-144

Signal Protocol

- Signal is an app for iOS and Android
 - Original protocol based on OTR and used for encrypted SMS
- The Signal Protocol is now used by WhatsApp, Google Allo, and Facebook Messenger
- Provides forward secrecy, improved deniability

5-145

Signal Protocol

- Perfect forward secrecy
 - Similar to OTR, uses a “ratchet” technique to constantly rotate session keys
- Deniability
 - Uses “Triple Diffie-Hellman” deniable authenticated key exchange (DAKE)
 - Anyone can forge a conversation between Alice and Bob using only their public keys

5-146

Recap

- Basics of cryptography
- Secret-key encryption
- Public-key encryption
- Integrity
- Authentication
- Security controls using cryptography
- Link-layer security
- Network-layer security
- Transport-layer security and privacy
- Application-layer security and privacy

5-147