EECS 122, Lecture 29

Today's Topics:

Brief Intro to Security

Comprehensive Review

Kevin Fall, kfall@cs.berkeley.edu

Network Security

- You would somehow like to have your data (or that of others) be secure. This often means you want to:
  - know who really sent it
  - know nobody else read it
- More specifically, protect from:
  - eavesdropping, masquerading, replay, traffic analysis, exploit-based attacks, denial-of-service

Protecting Yourself

- These attacks are often classified as
  - Active:
    - somebody actually generates or modifies network traffic
    - easier to detect, harder to prevent
  - Passive:
    - somebody just collects and analyses network traffic
    - harder to detect, easier to prevent

Common Approaches to Security

- One type of approach is based on physical security
  - usually expensive and of limited scope
  - may provide the best assurances
- Other (more common) approaches
  - hide your data somehow
    - just don't tell ("security through obscurity")
    - scramble it using some math (cryptography)!

Cryptography

- Cryptographers develop mathematical codes to hide or sign data
- Cryptanalysts attempt to compromise the codes developed by the cryptographers
  - Cryptographers job is usually to find a thought-to-be hard mathematical problem and develop a coding scheme based on it
  - Cryptanalysts job is usually to find a flaw (not always with the math, but often with the initial assumptions!)

Basic Cryptographic Concepts

plaintext —E(M,K)→ ciphertext —D(E(M),K')→ plaintext

- E: encryption function, D: decryption function, M: cleartext
- K and K' are called keys (bit strings), where K may be equal to K' (called shared keys)

The Main Issues

- For the cryptographer, the main issues:
  - choice of the transformation (D and E)
    - is the underlying mathematical basis efficient for decoding and encoding with keys and hard without them?
    - do you publish the algorithm or not?
  - generation and distribution of keys
    - might like to use random numbers, but computers aren t exactly random devices
    - how do you get a secret from one person to another if you don t already have keys!?

The Main Issues

- For the cryptanalyst, the main issues:
  - what is already known?
    - algorithm, plaintext-ciphertext pairs, any information about generation of the keys
  - types of attacks
    - ciphertext only (freq analysis, brute force)
    - known plaintext
    - chosen plaintext

Secret Key Functions

- generally a single key [symmetric] shared among parties (K=K); (still must be distributed somehow)
- rekeying is common, distributed in-band
- example uses:
  - challenge/response authentication
  - secure storage in insecure media
  - cryptograph checksums

Public Key Functions

- Fascinating class of functions using more than 1 key (usually 2) [asymmetric]
  - public keys usually published
  - private keys are kept secret
  - they are mathematically related
- Provides auth and/or privacy:
  - E(M, priv(A)) ---> signed by A
  - E(M, pub(A)) ---> only A can read it
  - E(E(M, priv(A)), pub(B)) -> from A to only B

Hashes and Message Digests

- Essentially one-way "digests" of messages. With msg M, function H:
  - hard: find different M, M' where H(M)=H(M)
  - hard: given H(M), find M
  - easy: given M, produce H(M)
- Common example is MD5
  - hashes arbitrarily-long messages to 128-bit signature (RFC1321)
  - used for file verification, IPSEC, etc

The Basics of DES

- 56-bit keys, 64-bit block cypher, key usually expressed as 56+8 parity bits
  - not really long enough; why chosen?
  - Estimated 3DES is 2^56 times better
- symmetric, single-key
  - requires exchange of key
  - (maybe use public key for exchange)
- reasonably fast (de)crypt functions

The Basics of DES (cont'd)

- How it works (in a nutshell)
  - permute (shuffle) the 64 bits
  - using the 56-bit key, take 16 different 48-bit subsets of the 56-bit key as "per-round" keys
  - the input to each round is the output of the previous one, plus the 48-bit round key
  - swap the two 32-bit halves, then permute them (inverse of initial permutation)
- Details start on p. 372 in text

The Basics of RSA

- can be used for encryption and signing
- asymmetric "public-key" cryptography, using 512 bits (or more) for key
- how to start it off
  - pick big primes p, q; form n=pq; create an encryption key e such that e and (p-1)(q-1) are relatively prime [only comm. factor is 1]
  - decryption key d=1/e mod ((p-1)(q-1)); the inverse of e in mod((p-1)(q-1))

The Basics of RSA (cont'd)

- Public key is {e,n}, private key is {d,n}; p,q no longer needed but must not be disclosed
- With a message m and ciphertext c:
  - Encryption: c = m^e mod n
  - Decryption: m = c^d mod n
  - works because m^(ed) = m & m<n (reqd)
- If you can factor n (to get p,q), you can break RSA, but we think this is hard

Random Numbers

- Crytographic exchanges often involve the use of random numbers (session keys)
- computers typically produce pseudorandom numbers, initialized by a random number *seed*
  - completely predictable given initial conditions, numbers merely look random
  - sometimes use human input (typing) as a seed to a RNG, or some other device

Random Number Generation

- The RNG has been a key failure in some otherwise decent systems:
  - too small (16 random bits->64K keys)
  - using the clock, which doesn't increment fast enough (not that large a space to search)
  - divulging the seed
- See RFC1750

In the first half of the semester...

- Networking concepts
  - remote access to resources
  - controlled sharing
    - multiplexing: TDM, Stat Mux
  - protocols and layering
    - ISO reference model, encapsulation
    - service model, error detection
    - end-to-end argument
    - soft state

In the first half of the semester…

- Development of the Internet
  - interconnection of heterogeneous networks
  - simple best-effort service model
  - fully-connected graph of hosts (routing)
- Internet scaling issues
  - use of hierarchies in routing, addresses, DNS
  - use of caching in DNS

In the first half of the semester…

- Direct-link networks
  - signals, modulation, error detection
  - best-effort delivery between attached stations
  - possible error correction using codes
  - MAC protocols, Ethernet

In the first half of the semester…

- The Internet Protocol
  - IP service model
    - best-effort datagram model
    - error detection in header only
    - consistent, abstract packet, addressing
    - routing
    - signaling (ICMP)
    - multicasting, IGMP, multicast routing
    - IP futures with IPv6

And in the second half…

- Routing Protocols
  - interior and exterior versions
  - distance vector
  - link state
- Internet Scaling Issues
  - CIDR
  - IPv6 / very large addresses

And in the second half…

- The Transport Layer
  - access to processes/endpoints, ports
- User Datagram Protocol (UDP)
  - best-effort datagram service
  - error detection, no correction, pseudoheader
- Transport Control Protocol (TCP)
  - reliable stream service
  - error detection/correction
  - flow and congestion control

And in the second half…

- How to achieve reliability
  - ARQ, ACKs, retransmission
  - Stop & Wait, performance
  - the bandwidth-delay product
  - go-back-n, window-based protocols
  - retransmission timers
  - types of ACKs/NACKs, dup ACKs
  - window-based flow control

And in the second half…

- Introduction to Congestion
  - congestion and congestion collapse
  - implicit/explicit congestion notification
  - bw-product as ideal window size
- Queuing and Scheduling
  - FIFO/FCFS, burst loss
  - Fair Queuing and Round-Robin
  - Random Early Detection (RED)

And in the second half…

- Congestion control
  - closed versus open-loop
  - host vs network enforcement
  - effectiveness, fairness, fairness index
- Congestion Control in TCP
  - slow start & congestion avoidance
  - fast retransmit
  - RTT estimation and timeout, Karn algorithm
  - silly window syndrome, Nagle algorithm

And in the second half…

- Connection Management in TCP
  - bi-directional connections, 4-duples
  - SYN 3-way handshake, FIN exchange
  - initial sequence numbers
- Some Implementation Issues
  - user vs kernel space
  - buffers
  - lookup maps/tables
  - event management, timers

And in the second half…

- Introduction to the Telephone Network
  - circuit switching, calls, set-up/tear-down
  - the regulatory environment
  - basic structure
  - TDM, pleisiochronous operation, justification
  - SONET
  - data/control separation, signaling
  - CCIS and SS7
  - routing: DNHR, metastability, RTNR

And in the second half…

- Introduction to ATM
  - asynchronous operation, VC approach
  - cells and reasons for them
  - framing and adaptation layers, AAL5
  - issues in IP over ATM

And in the second half…

- Quality of Service
  - application types (elastic and inelastic)
  - traffic descriptors
  - leaky/token buckets and regulation
  - max-min fairness
  - performance bounds, delay distribution
  - choices in scheduler design
  - scheduling and packet drop strategies
    - GPS, WRR, DRR, PGPS/WFQ, VC, EDD

And in the second half…

- Real-world QoS examples
  - ATM QoS
    - CBR, VBR, ABR, UBR services
    - PCR, SCR, MBS, MCR traffic parameters
    - CDV (T), maxCTD, CLR QoS parameters
    - admission control, equivalent capacity (not really ATM specific per-se)
  - Internet QoS
    - IntServ: TSPECS, controlled load, guaranteed
    - DiffServ: general model, AF and EF PHBs
    - RSVP