Mixes

Yanyan Zhuang
Department of Computer Science
http://www.cs.uccs.edu/~yzhuang/

Chaum Mixes / Mix Networks

- Originally designed for anonymous email
  - David Chaum, 1981
  - Concept has since been generalized for TCP traffic
- Hugely influential ideas
  - Onion routing
  - Traffic mixing
  - Dummy traffic (a.k.a. cover traffic)

Basic Notations

- Public key K
    - Private key K^{-1}
- Encryption of X: K(X)
    - Private key does the reverse $K^{-1}(K(X))=X$
- Signature
    - Large constant C, the owner of K^{-1} does $K^{-1}(C, X) = Y$
    - Everybody else can verify: $K(K^{-1}(C, X))=C, X$, i.e., Y has been signed by the holder of K^{-1}
- Use a random number R before encrypting
    - K(R, X)
    - Prevent guessing of X=Y by checking K(X) = K(Y)

How It Works (1)

- Participant X wants to send a message M to Y
    a. X prepares a message for delivery to Y by appending a random value R0 to the message → (R0, M)
    b. X seals it with the Y's public key Ky, appends Y's address Ay → Ky(R0, M), Ay
    c. X seals the result with the mix's public key K1, appending another R1 → K1(R1, **Ky(R0, M), Ay**)
- Mix opens it with his private key $K1^{-1}$
    a. Gets Ky(R0, M), Ay
    b. Mix now knows Y's address Ay, and he sends Ky(R0, M) to Y

How It Works (2)

- K1(**R1**, Ky(**R0**, M), Ay)
  - R1 is needed to prevent replay attack from X to mix
    - Mix opens K1(R1, Ky(R0, M), Ay) with its private key
    - Will only accept different R1's each time
  - R0 is needed to prevent an attacker from guessing messages
    - Assume attacker can observe all incoming and outgoing messages
    - If R0 is not used (i.e. only Ky(M) is sent to Y), the attacker can test whether Ky(M')=Ky(M) is true

How It Works: Cascade

- A cascade of mixes Mn, Mn-1, Mn-2...
  - X sends
    - Kn(Rn, Kn-1(Rn-1, ..., K2(R2, **K1(R1, Ky(R0, M), Ay)**)...))
  - The first mix Mn encrypts and gets
    - Kn-1(Rn-1, ..., K2(R2, **K1(R1, Ky(R0, M), Ay)**)...)

The Other Way Round?

- Return traffic: how can the destination respond to the sender?
  - Mix only sends Ky(R0, M) to Y
  - For Y to respond to X while still keeping the identity of X secret from Y
- Solution: During path establishment, the sender places keys at each mix along the path

Return Traffic (1)

- X forms an **untraceable return address** **K1(S1, Ax), Kx**
  - **S1** is a key that will also act as a random string for purposes of sealing
  - Ax is X's own real address
  - **Kx** is a public one-time key chosen for the current occasion only
  - X sends this return address to Y as part of the message sent
- Originally, X → Y: K1(R1, Ky(R0, M), Ay)
  - Mix sends Ky(R0, M) to Y
- With the untraceable return address, X → Y: K1(R1, Ky(R0, M, **K1(S1, Ax), Kx**), Ay)
  - Mix opens with its private key, gets Ky(R0, M, **K1(S1, Ax), Kx**), Ay
  - Mix sends Ky(R0, M, **K1(S1, Ax), Kx**) to Y

Return Traffic (2)

- Continue with previous...
  - Mix sends Ky(R0, M, **K1(S1, Ax), Kx**) to Y
- Y opens and gets M, K1(S1, Ax), Kx
  - Remember that Kx is a public one-time key chosen for the current occasion only

Return Traffic (3)

- Continue with previous...
    - Mix sends Ky(R0, M, K1(S1, Ax), Kx) to Y
- Y opens and gets M, K1(S1, Ax), Kx
    - Remember that Kx is a public one-time key chosen for the current occasion only
- Y sends K1(S1, Ax), Kx(S0, M’) to mix
    - M’ is the response message, S0 is a random string
    - Mix transforms it to Ax, S1(Kx(S0,M’))
        - Mix uses S1 that it finds after decrypting K1(S1, Ax) as a key to re-encrypt the message part Kx(S0, M’)
    - Mix sends S1(Kx(S0,M’)) to X
        - Only X can decrypt the resulting output S1(Kx(S0,M’)): **X created both S1 and Kx**
    - Kx assures that the mix cannot see the content of the reply-message

Election

- If registered voters are accepted for a roster
- For a single mix, each voter submits a ballot of the form $K1(R1, K, K^{-1}(C, V))$
  - K is the voter's pseudonym and V is the actual vote
  - K1 is mix’s public key
- Items in the final output batch are of the form $K, K^{-1}(C, V)$
- Each ballot is counted
  - Checking that the pseudonym K which forms its prefix, is also contained in the roster
  - The pseudonym properly decrypts the signed vote V

Traffic Mixing

- Hinders timing attacks
  - Messages may be artificially delayed
  - Temporal correlation is warped
- Problems
  - Requires lots of traffic
  - Adds latency to network flows

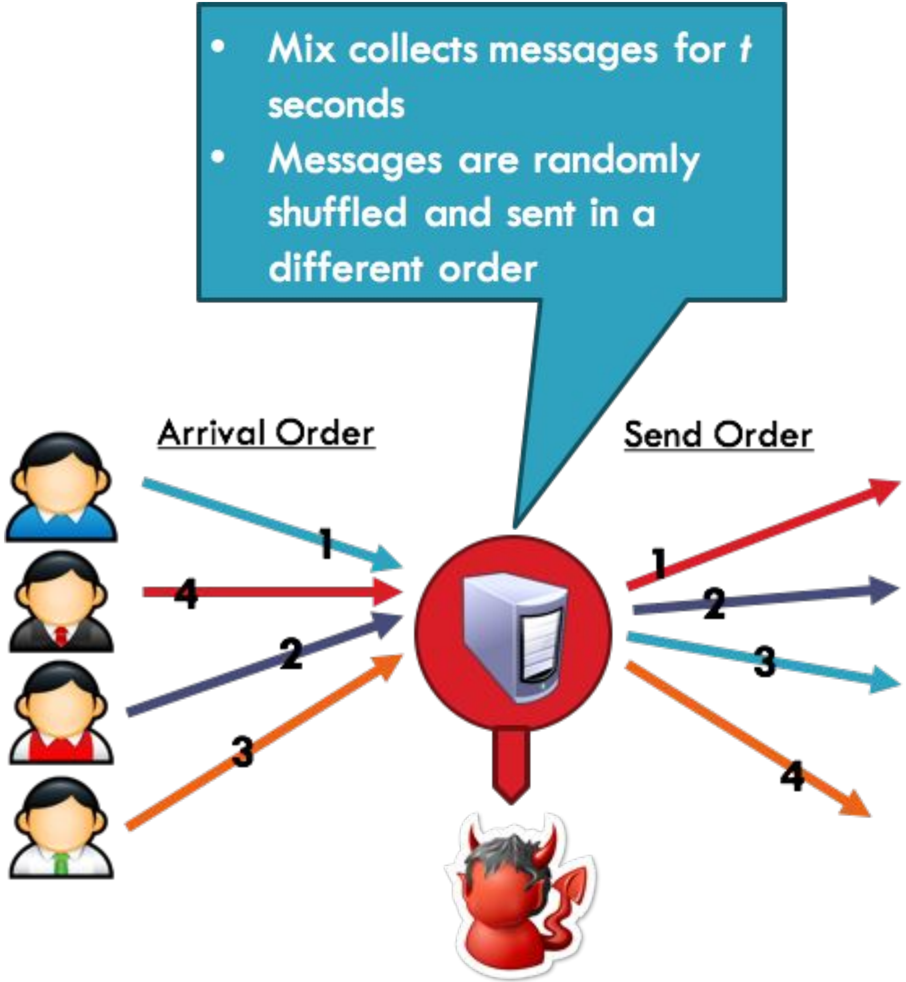

Putting it together (1)

- Routing protocols that create hard-to-trace communications
  - Using a chain of proxy servers known as mixes
  - A mix takes in messages from multiple senders, shuffle them, and send them back out in random order to the next destination (possibly another mix)
    - Breaks the link between source/destination, making it harder for eavesdroppers to trace end-to-end communications
  - A mix only knows the node that it immediately received the message from, and the immediate destination to send the shuffled messages to
    - Resistant to malicious mix nodes

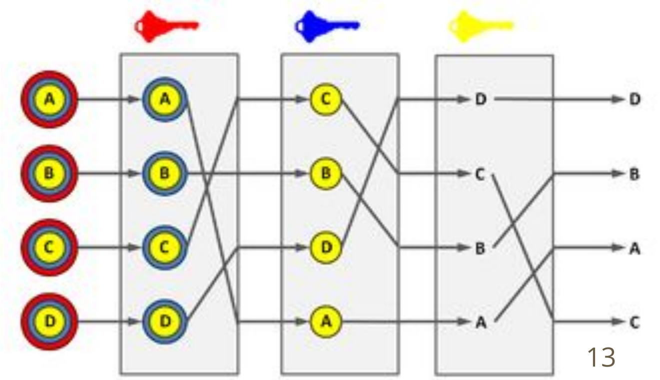

Putting it together (2)

- Each message is encrypted to each proxy using public key cryptography
  - Encryption is layered like a Russian doll (except that each "doll" is of the same size) with the message as the innermost layer
  - Each mix strips off its own layer of encryption to reveal where to send the message next