

Privacy at the communication layer

Claudia Diaz

K.U.Leuven ESAT/COSIC

June 16, 2009

BCRYPT



Belgian Fundamental Research on Cryptology and Information Security

KATHOLIEKE UNIVERSITEIT
LEUVEN

... before we start

- Website slides:
 - <http://www.cs.kuleuven.be/~berendt/teaching/Privacy09/>
- Location next talks (22/06/2009): **different building!**
 - Landbouwinstituut (Kasteelpark Arenberg 20)
 - Room 00.215

Overview

- Motivation
- Trusted and semi-trusted relays
- Mix systems
- Low-latency anonymous communications
- Concluding remarks

Motivation

Traffic analysis

- Even if communication is encrypted, traffic data can reveal a lot of information: source, destination, timing, volume, etc.
- Examples from WW II (signals intelligence):
 - Traffic analysis was used by the British at Bletchley Park to assess the size of Germany's air-force
 - Japanese traffic analysis countermeasures contributed to the surprise of their 1941 attack on Pearl Harbour
 - Increased volume: possible imminent action (example: D-day)
 - Identifying people by their typing
- Amateur plane-spotters revealed the CIA's 'extraordinary rendition' programme

Some examples where sensitive data can be inferred from traffic data

- Scenario 1: Home monitoring system that sends information to oncology department
 - Source-destination enough to determine that the person living in the house has cancer
- Scenario 2: Fridge that orders food from the Kosher shop
 - Source-destination enough to determine religious beliefs
- Scenario 3: Device to hear streaming radio from a highly conservative news station
 - Source-destination enough to determine political beliefs
- Companies searching for patents/info on a subject
- Undercover police who infiltrate criminal groups and want to report back
- Two people who have looong conversations every evening (probably having an affair)
- Sensor going off

Traffic analysis

- Traffic data has less volume than content:
 - Coarser, but highly valuable information
 - Can be used to select targets for more intensive surveillance
 - Formats that are easy to process for machines
 - Harder to conceal
- Diffie and Landau, in their book on wiretapping: “traffic analysis, not cryptanalysis, is the backbone of communications intelligence“
- Anonymity at the communication layer is assumed as building block for other privacy protection technologies

Censorship resistance

- SafeWeb (anonymous communication system) was funded by In-Q-Tel (mission: *identify and invest in companies developing cutting-edge technologies that serve the USA national security interests*)
- Goal of funding SafeWeb was to “*help Chinese and other foreign web users get information banned in their own company (sic)*”



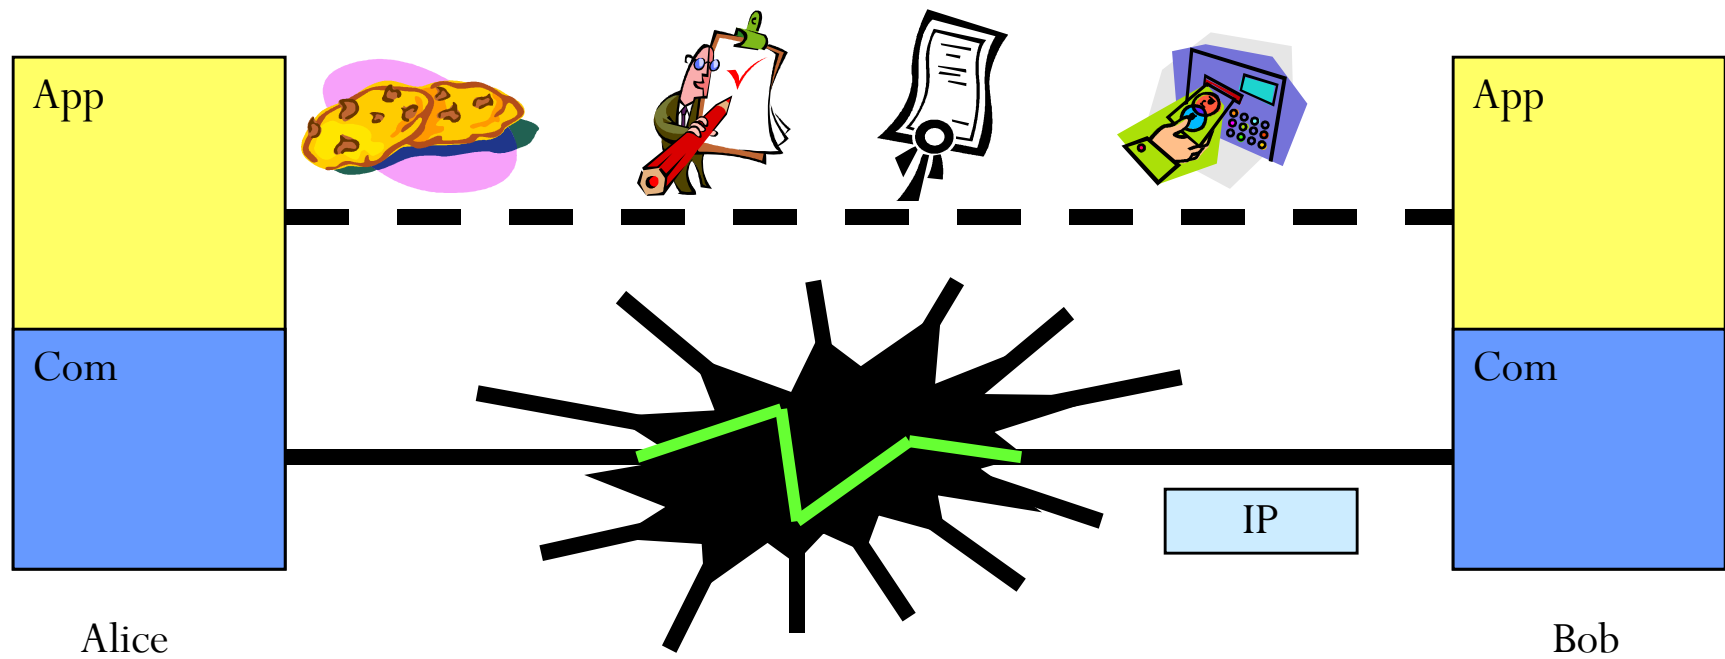
Searching for “Tiananmen Square” on Google.com and Google.cn

Censorship resistance

- Link between anonymous communications and censorship resistance
 - If the censor does not know the content and destination of the communication, blocking becomes harder
 - A rigorous study of the relationship between anonymous communication and censorship resistance has not yet been made, although the link is commonly made by researchers
- Anonymous communication services such as Tor, JAP and Anonymizer provide mechanisms for censorship resistance:
 - Tunneling, fresh addresses (arms race)

System and adversary models

Anonymity – Data and Communication Layers





Classical Security Model

- Confidentiality
- Integrity
- Authentication
- Non repudiation
- Availability



Alice




Bob




Eve

Passive / Active

Anonymity – Concept and Model



Anonymity Adversary



Recipient?

Third Parties?

- Passive / Active
- Partial / Global
- Internal / External

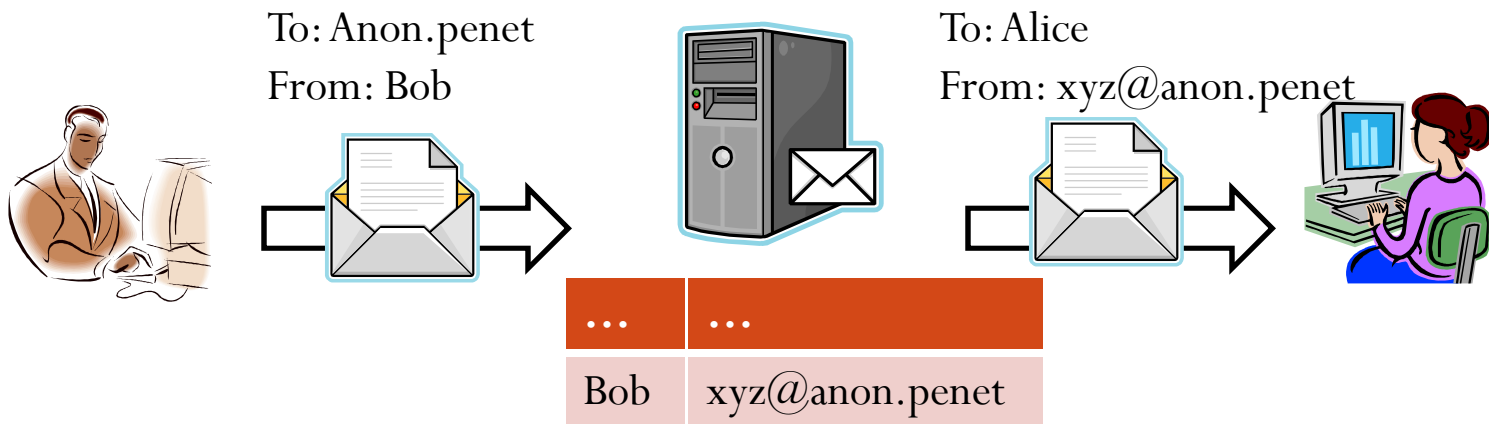
Anonymity Adversary

- The adversary will:
 - Try to find who is communicating with whom
 - Observe
 - All links (*Global Passive Adversary*)
 - *Some links*
 - Modify, delay, delete or inject messages.
 - Control some nodes in the network.
- The adversary's limitations
 - Cannot break cryptographic primitives.
 - Cannot see inside nodes he does not control.

Trusted and semi-trusted relays

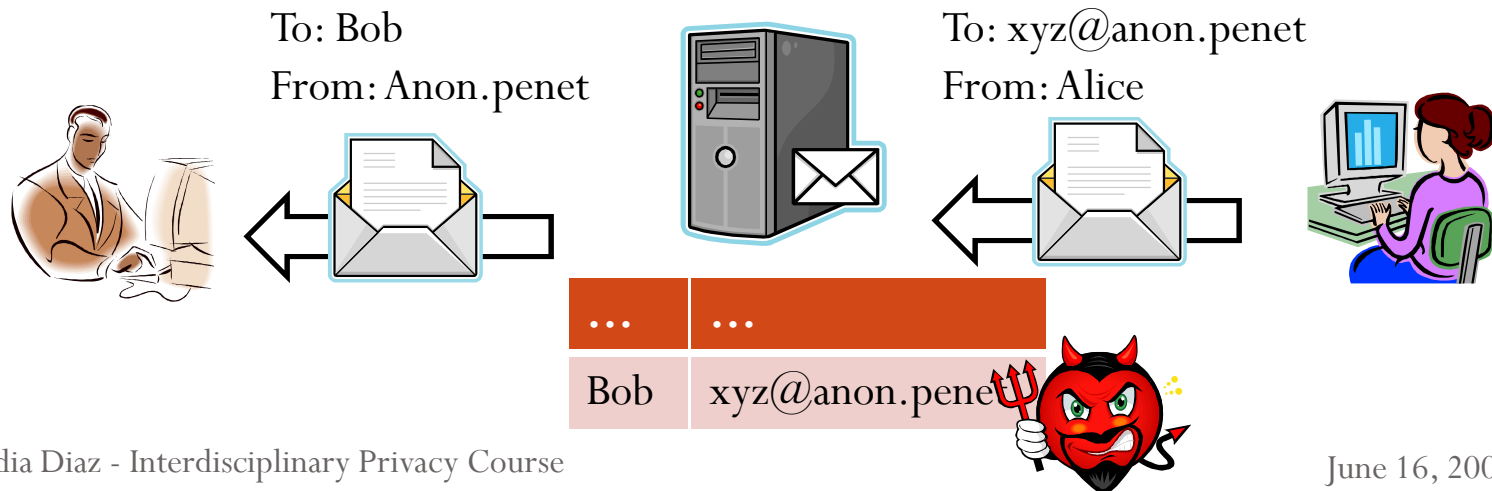
Anon.penet.fi (Helsingius 1993)

- Simple proxy, substituted email headers
- Kept table of correspondences nym-email





Anon.penet.fi (Helsingius 1993)

- Supported anonymous replies
- Threat model: recipient
 - Trivial to find correspondence by observing the server
- Brought down by legal attack in 1996
 - Lesson learned: do not keep tables of correspondences!
 - Protection of users, but also protection of services themselves




Type I Cypherpunk remailers (Hughes, Finney 1996)

- No tables (routing info in messages themselves)
 - Remailers can still be forced to decrypt a message
- PGP encryption (no attacks based on content) — attacks based on size are possible
- Chains of remailers (distribution of trust)
- Reusable reply blocks
 - Source of insecurity: replay attacks



Anonymizer and SafeWeb (mid-90s)

- Web proxies: strip identifying information and forward
- Less vulnerable to legal compulsion attacks: keeping long-term logs is not needed (communication always initiated by the user)
- Filtering of active content (attacks on these features have been found)
- Connection between user and server is encrypted (SSL)
- No padding, no mixing
 - Vulnerable to attacks that correlate traffic to and from the server
 - Vulnerable to fingerprinting attacks (matching against a database of traffic signatures)
 - These attacks are more effective if we consider a sequence of linked web pages



Mix systems


Chaumian Mix (Chaum 1982)

- Messages of fixed size
 - Large messages are divided into blocks
- Original designed used randomized RSA encryption
 - Encrypted with the public key of the mix
 - Later found vulnerable to some attacks (e.g., tagging attacks)
- Several mixes could be chained to distribute trust:
 - Sender \rightarrow Mix₁ : $\{\text{Mix}_2, \{\text{Rec}, \text{msg}\}_{K_{\text{Mix}_2}}\}_{K_{\text{Mix}_1}}$
- Goal: an adversary observing the input and output of the mix is not able to relate input messages to output messages
 - Bitwise unlinkability
 - The mix performs a decryption on input messages
 - Input/output of the mix cannot be correlated based on content or size
 - Prevent traffic analysis based on message I/O order and timing
 - Achieved by batching messages

Chaumian Mix (Chaum 1982)

- Phase 1: collect inputs
- Parameter T (threshold): $T=4$ in example


Mix



Chaumian Mix (Chaum 1982)

- Phase 2: mix and flush


Mix



MixMaster (Cottrell, mid-90s)

- Type II Cypherpunk remailer: most widely deployed remailer
- Evolving since 1995
- Improved crypto format of messages to prevent tagging attacks and protect the integrity of messages
- MixMaster did not support replies
- Pool mixing: increased anonymity wrt Chaumian mixes

Threshold = 4, Pool = 2




MixMinion (Danezis et al., 2003)

- Type III Cypherpunk remailer: state-of-the-art in remailers
- Anonymous replies through SURBs (Single Use Reply Blocks)
 - Prevent replay attacks
 - Forward and backward messages are indistinguishable
- Improved cryptographic packet format
 - Two headers further divided into subheaders
 - Protection from tagging
- Trail of keys that are updated with one-way functions to provide forward security


Stop-and-Go mixes (Kesdogan 1998)

- Reordering strategy based on independently delaying each message
 - Anonymity level depends on volume of traffic
 - In threshold and pool mixes, it is the delay that depends on the volume of traffic
- Delays generated by the user from an Exponential distribution (proven optimal by Danezis)
- Timestamping to prevent active attacks
 - Trusted Time Service




Blending (n-1) attacks

1. Empty the mix from legitimate messages
2. Let the target message into the mix
3. Fill the mix with attacker-generated messages, while preventing other legitimate messages from entering the mix




Mix



Blending (n-1) attacks

4. At the time of flushing the adversary recognizes his own messages. The unknown message is the target
 - Variants of this attack break the anonymity the other types of mixes
 - The effects of the attack can be mitigated with randomization and dummy traffic

Mix



Dummy traffic




- Fake messages/traffic introduced to confuse the attacker
- Undistinguishable from real traffic
- These messages may be generated
 - By users
 - By mixes
- Dummies improve the anonymity by making more difficult the traffic analysis
- Necessary for unobservability
- Can also be used to detect n-1 attacks: Heartbeat Traffic [Dan03]
- Dummy traffic is expensive (bandwidth)
 - Unclear how to use it in an optimal way

Long-term intersection attacks

- Family of attacks with many variants:
 - Disclosure attack (Agrawal, Kesdogan)
 - Hitting set attack (Kesdogan)
 - Statistical disclosure attack (Danezis, Serjantov)
 - Extensions to SDA (Dingledine and Mathewson)
 - Two-Sided SDA (Danezis, Diaz, Troncoso)
 - Perfect-Matching disclosure attack (Troncoso et al.)
- Assumptions:
 - Alice has persistent communication relationships (she communicates repeatedly with her friends)
 - Large population of senders, and a different subset mixes their messages with hers in each round

Long-term intersection attacks

- Method:
 - Combine many observations (looking at who receives when Alice sends)
- Intuition:
 - If we observe rounds in which Alice sends, her likely recipients will appear frequently
- Result:
 - We can create a vector that expresses Alice's sending profile
 - Hard to conceal persistent communications (also in low-latency systems!)



Verifiable mixing

- Mixes can be used for implementing e-voting schemes
- In e-voting applications, it is important to make sure that
 1. Votes are anonymous
 2. All votes are counted
- N-1 and intersection attacks hard to deploy in e-voting scenarios
- Mixes must prove that the outputs are a permutation of the (cryptographically transformed) inputs
- Whole body of research to attempt to create mix systems that are:
 - Robust against malicious servers that fail to deliver some votes
 - No entity learns anything except for the vote tally
 - Provide universal verifiability (correctness of the tally)
 - Provide receipt-freeness to prevent coercion/selling of votes


Low-latency anonymous communications

ISDN Mixes (Pfitzmann 1991)

- Anonymization of ISDN phone conversations
- Practical design from an engineering point of view
 - Signaling channel used to establish keys
 - Dummy traffic on the subscriber lines (no additional bandwidth needed)
- Protection against very powerful global adversaries, who control everything in the system but one mix
 - Synchronous establishment and teardown of connections
- The design was later extended to Web Mixes for IP networks


Crowds (Reiter, Rubin 1998)

- Anonymity for web browsing
- Group of users form a “crowd”
- Initiator chooses a random member of the crowd and forwards the web request to her
- The recipient of the request flips a biased coin and forwards the request to another member with probability p and to the end server with probability $1-p$
- A tunnel is established between the initiator and the exit crowd member (static paths)




Crowds (Reiter, Rubin 1998)

- Communication between members is encrypted with symmetric keys
 - BUT: all members can see the request in clear
- Adversary model:
 - Assumed adversary cannot control all links
 - Instead, the adversary controls a subset of the crowd and/or the end server
- Probability that predecessor is the initiator or just a forwarder
 - We can measure initiator anonymity as a function of the fraction of corrupted nodes and the probability of forwarding



Crowds (Reiter, Rubin 1998)

- Predecessor attacks
 - If initiator repeatedly accesses the same resource over different sessions, it will appear as predecessor of the first adversarial member more often than other crowd members
 - Anonymity degrades with
 - Amount of linkable requests made in different sessions
 - Size of the crowd



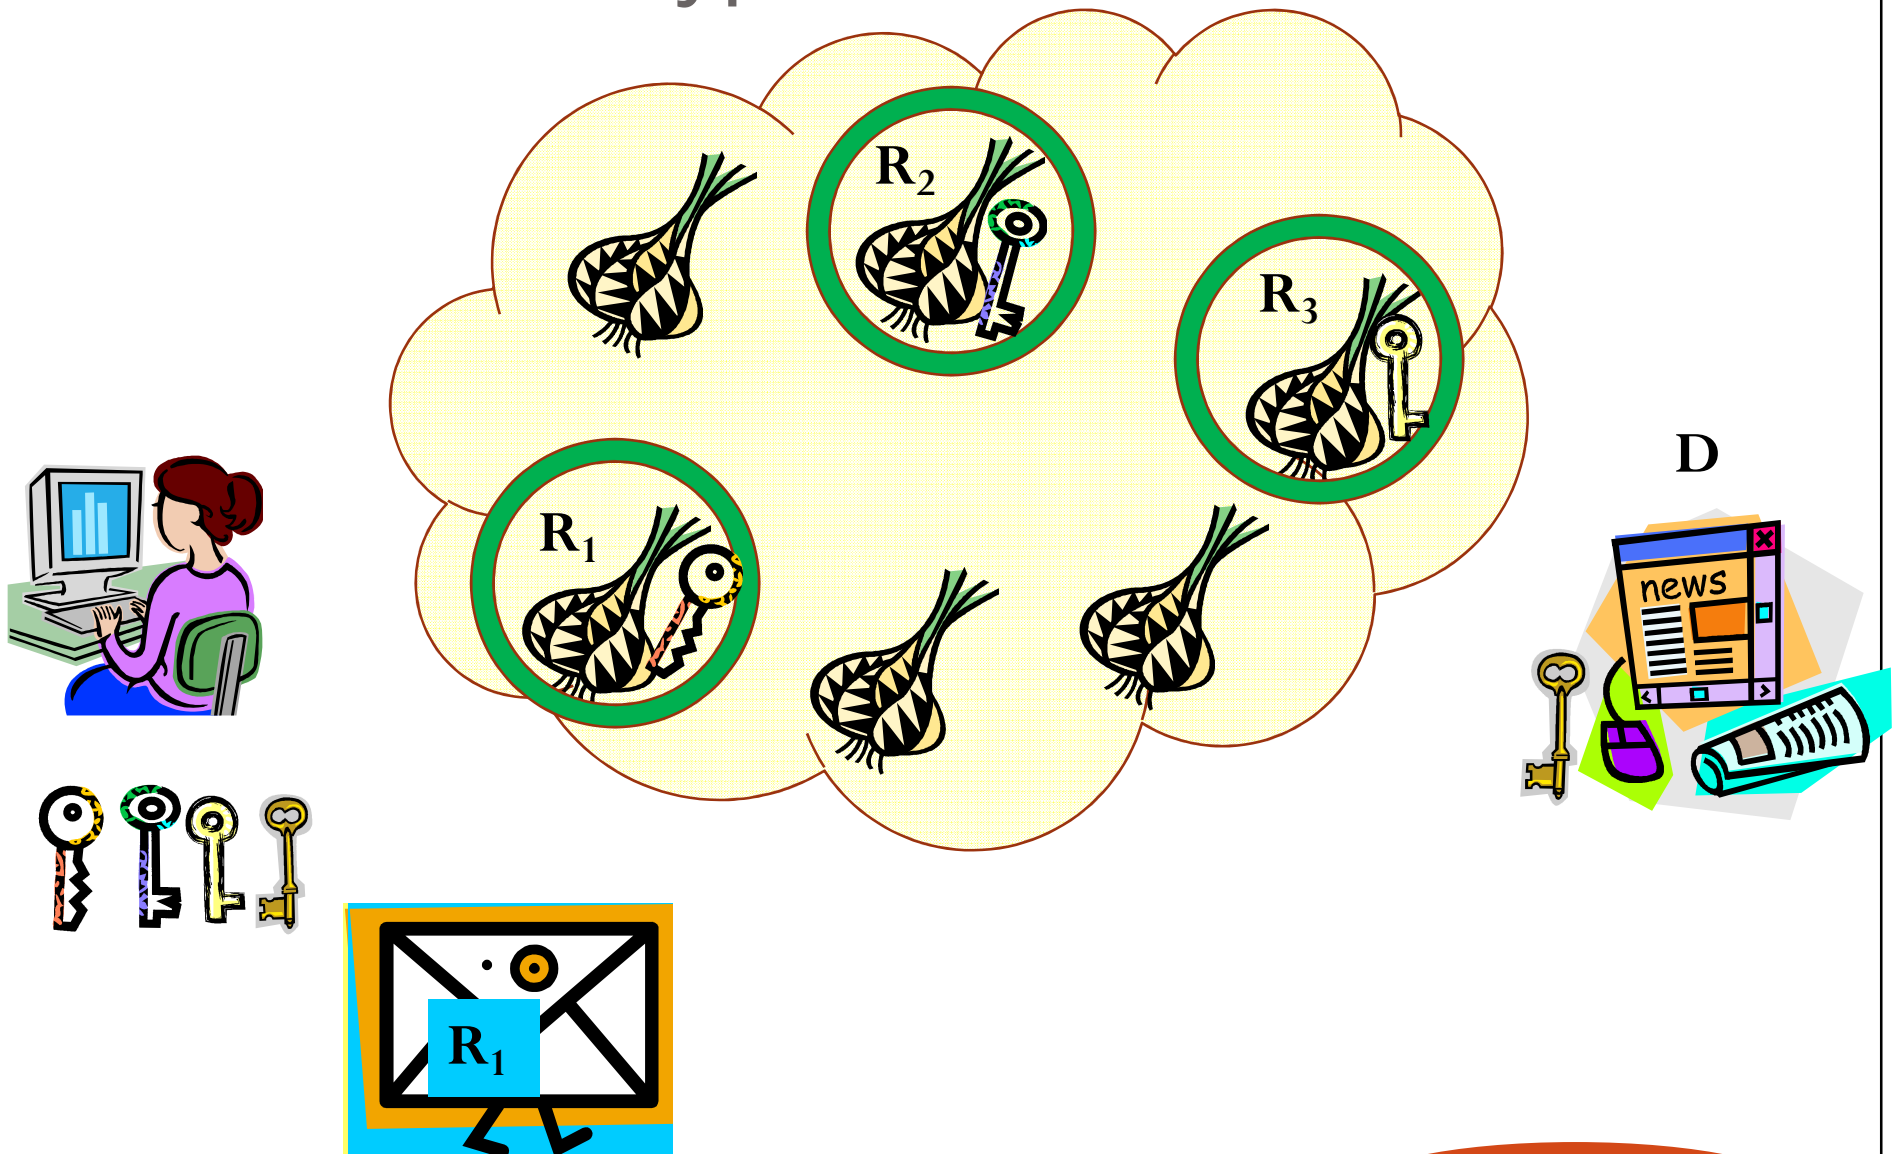
Anonymizing web traffic not trivial

- Difficult to conceal traffic pattern
- Difficult to pad
 - Lots of padding: scalability / cost problem
 - Little padding: not enough to conceal pattern
- Vulnerable to strong adversaries (entry+exit)
- Fingerprinting attacks
 - Adversary observes only user side
- Internet exchanges: global adversary

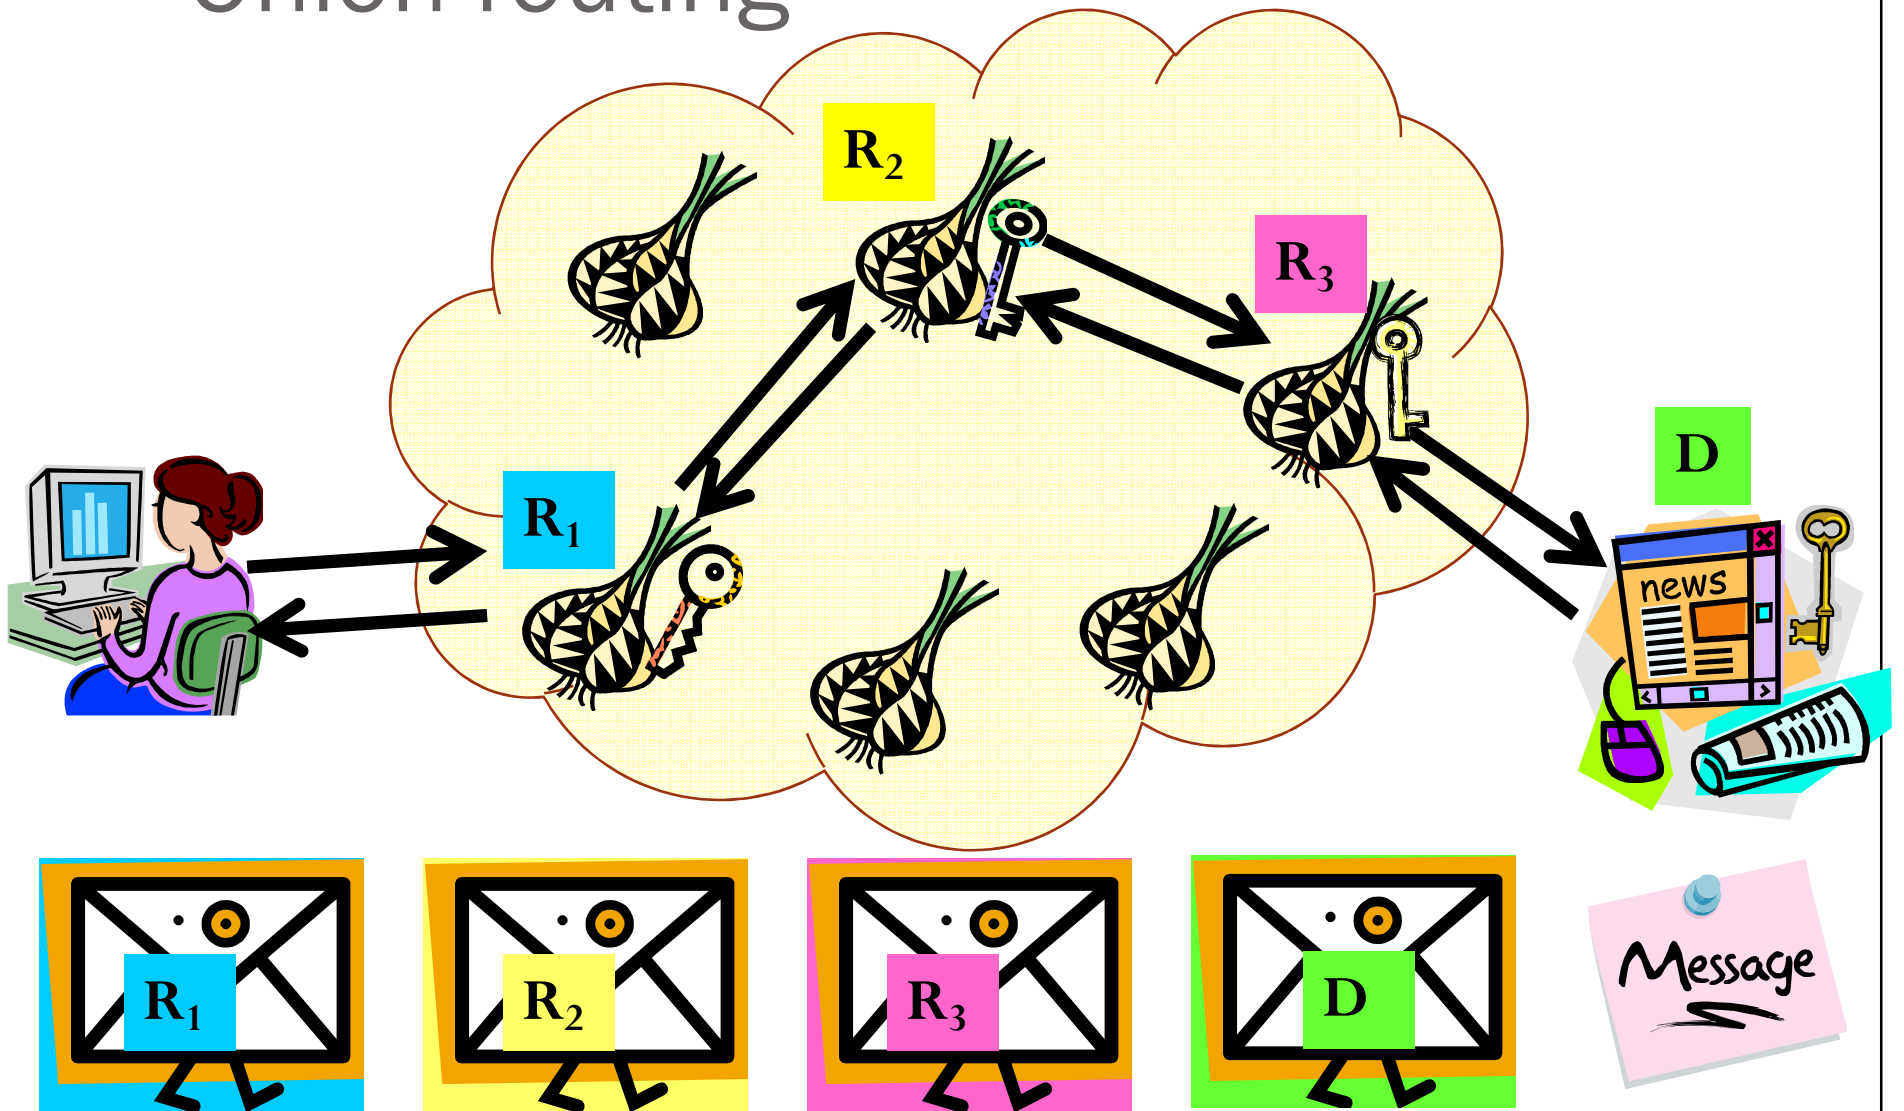
Onion Routing (Reed, Syverson, Goldschlag, 1998)

- Developed at the Navy Research Lab
- Bi-directional, low latency communication
 - Onion routers do not perform “mixing”, instead they just forward packets
 - No dummy traffic
- Users select a set of routers that constitute the anonymous channel (source routed)
- A commercial implementation called ‘Freedom Network’ was deployed between 1999 and 2001 (Zero Knowledge Systems)
- Second-generation Onion Routing: Tor (from 2003)
 - Volunteer nodes: currently several thousands
 - Hundreds of thousands of users from all over the world
 - Usability:
 - Easy to install and use (Tor button)
 - Bad QoS because the network is overloaded


Onion encryption



Onion routing




TOR – adversary model




Network topology


- Mixes are combined in networks in order to
 - Distribute trust
 - Improve availability



Fully connected network



Restricted route network



Cascades vs Free Route topologies

- Flexibility of routing
 - Surface of attack
 - Advantage free routes
 - Availability
 - Advantage free routes
 - Intersection attacks
 - Advantage cascades (anonymity set smaller but no partitioning possible)
 - Trust
 - Advantage free routes (more choices available to user)
- Free routes: Tor, Mixmaster
- Cascades: Web Mixes (JAP)

Peer-to-peer vs client-server architectures

- Symmetric vs asymmetric systems
 - Surface of attack
 - Advantage peer-to-peer
 - Liability issues
 - Advantage client-server
 - Resources / incentives / quality of service
 - Advantage client-server
 - Availability
 - Advantage peer-to-peer
 - Sybil attacks
 - Advantage? Depending on admission controls (for peers/servers)

Concluding remarks


(Some) Attacks on anonymous communication systems

- Many adversary models are possible and realistic
- Passive attacks
 - Long-term intersection attacks
 - Traffic correlation / confirmation
 - Fingerprinting
 - Epistemic attacks (route selection)
 - Predecessor attacks
- Active attacks
 - N-1 attacks
 - Sybil
 - Traffic watermarking
 - Tagging
 - Replay
 - DoS

Conclusions

- Tradeoffs cost/anonymity
 - Cost: delay, overhead
 - Economics of privacy:
 - Crypto: little overhead → lots of security
 - Anonymity: lots of overhead → a little bit of security
- High-latency applications (email):
 - Well established primitive
 - Problems with persistent user behavior
- Low-latency applications
 - Insecure towards strong adversaries
 - Large scale systems: is P2P the solution?
- New scenarios with increased possibilities to obtain traffic data
 - Pervasive computing scenarios
 - Social networks
- Anonymous communications are fragile
 - If you want to propose a new system:
 - Check the literature
 - Check known attacks

Thank you!



Recommended bibliography on the subject:

<http://www.freehaven.net/anonbib/>

- Website slides:
 - <http://www.cs.kuleuven.be/~berendt/teaching/Privacy09/>
- Location next talk (22/06/2009): **different building!**
 - Landbouwinstituut (Kasteelpark Arenberg 20)
 - Room 00.215