CS182 2006, Reynolds and Astrachan

I) Introduction (1 page)
II) Motivation (3 pages)
    1) Types of anonymity
        A) Onion routing examines routing anonymity
    2) Who needs it?
        A) Well, criminals, naturally!
        B) But also: the government and the military
        C) Social activists, political dissenters, etc.
        D) All who enjoy freedom of speech: anonymity eliminates chilling effects
        E) Those vulnerable to traffic analysis
III) Technical Description (8 pages)
    1) Original description
        A) Basic concepts
            a) Proxy redirection
            b) Public-key cryptography
        B) Onion layers
        C) Reply onion
    2) Second-gen implementation (TOR)
        A) Stack integration
IV) Current Applications and Deployment (18 pages)
    1) Technical criticism (5 pages)
        A) Overhead
        B) Network attacks
            a) Replay
            b) Timing
            c) Intersection
    2) Legal (5 pages)
        A) Legal precedent
            a) Details in PGP decision
            b) Similarity/contrast to prior technologies
    3) Usability analysis (8 pages)
        A) Server
        B) Client
        C) Conclusions
V) Future Applications (10 pages)
    1) Usability/integration improvements (2 pages)
        A) Practicality of wide deployment
        B) Why doesn't everybody use encryption?
    2) Message-level encryption (1 page)
        A) Total communication anonymity
        B) Endpoint-only authentication
    3) Attack-specific solutions (1 page)
        A) Timing attacks: artificial network saturation
    4) Garlic routing (3 pages)

5) BitTorrent-style distributed data retrieval (3 pages)

VI) Bibliography

1) D. Chaum. *The Dining Cryptographers Problem*, Journal of Cryptology, vol 1, no 1, 1988.
2) R. Dingledine, M. Freedman, D. Molnar. *The Free Haven Project*. Workshop on Design Issues in Anonymity and Unobservability, 2000.
3) R. Dingledine, N. Mathewson, P. Syverson. *Tor: The Second-Generation Onion Router*. Proceedings of the 13th USENIX Security Symposium, 2004.
4) D. Goldschlag, M. Reed, and P. Syverson. *Hiding Routing Information*, Workshop on Information Hiding, 1996.
5) ___. *Anonymous Connections and Onion Routing*, IEEE Journal on Selected Areas in Communications, vol 16, no 4, May 1998.
6) P. Syverson. *Making Anonymous Communication*. Presentation at the National Science Foundation, 2004.
7) P. Syverson, G. Tsudik, M. Reed, and C. Landwher. *Towards an Analysis of Onion Routing Security*. Workshop on Design Issues in Anonymity and Unobservability, 2000.

Onion Routing Executive Summary
CS182 2006, Reynolds and Astrachan

Onion routing is one solution to the general problem of anonymous routing online. This scheme protects the identity of the sender from all other parties in the transmission, including attackers and the recipient, and protects the recipient's identity from everyone except the sender.

A fundamental question in examining such anonymizing schemes is who would use them and why. Many observers assume that the only beneficiaries of anonymous communication are criminals. The government, especially the military and intelligence communities, have an obvious need for the security and privacy that anonymizing schemes, especially supplemented by message-level encryption, provide. In addition, political dissidents in societies where freedom of speech is at a premium can use anonymizing schemes to protect themselves from unfair treatment. While their societies may judge them as criminals, some designers of anonymizing schemes believe that free expression is a basic human right, and therefore provide tools to spread this freedom to

restrictive societies.

Onion routing accomplishes anonymity by redirecting a message through a random selection of several proxies without revealing any information about the endpoints of the communication. Each proxy knows only to whom it should forward the message; the knowledge revealed at each “bounce”, then, is the bare minimum necessary for the communication to take place. Onion routing is founded upon the concept of a layered “routing onion”. At the core of the onion is the message. Accompanying the message is a label that contains a single routing instruction: it tells the sender of the onion core who the final recipient is. This pair of message and routing instruction are wrapped in an onion layer: they are encrypted together using the final sender's public key. This wrapped-up data parcel may now be considered as a new message; the next layer of the onion outward is formed by pairing routing information with this new message and wrapping it in an encryption layer determined by the second-to-last sender's public key. The layering process continues in this recursive fashion backward, all the way to the initiator of the communication. As the message travels through the network, each proxy unwraps the outermost layer with its private key, revealing the address of the next recipient, and a message that can only be read by that recipient. In this way, all information about the message's path more than one bounce backward or forward from any given proxy is kept secret.

The original design of onion routing was developed and patented by researchers in the US military and was an explicit user-level data packaging scheme. TOR is an open-source second-generation implementation of the concept of onion routing that dodges patent issues and integrates the scheme into the TCP stack. It remains to be seen

how transparent that actual implementation is; as part of this project we will attempt to set up cross-platform TOR clients and servers. Right now, there are approximately 200 TOR proxy servers available online.

There are several apparent improvements that could be made to onion routing as it now exists. As onion routing uses encryption only to hide the message and routing information from intermediate parties, the last bounce transmits the message in plaintext. Integrating public key (and open source) cryptography such as GnuPG with the TCP-level onion routing concept would allow for transparent anonymous and secure communication. Onion routing may be compromised by what is known as an "intersection attack". Such attacks become much less feasible as the pool of potential proxies grows, but another proposed solution to this problem is garlic routing. Each layer of a garlic bulb wraps up an arbitrary number of message/path cloves, each one intended for a different recipient. The proxy may choose any one of these cloves to continue the communication; this results in a nondeterministic path through the network, reducing the ability for an attacker to analyze the communication. Parallelizing the choice of garlic cloves—that is, sending out every clove rather than just one—results in a BitTorrent-like distribution scheme that provides redundancy, anonymity, increased total bandwidth, and extreme network load. While neither participant has experience coding low-level network services, we intend to examine these and other future applications and derivatives of onion routing with an eye toward practical and technical implementation issues, as well as their social and legal implications.