

Effective Restricting of Misbehaving Users Among Anonymous Users in a Network

¹M. Mastan, ²S. Jagadeswari

^{1,2}Dept. of CSE, NIMRA Institute of Science & Technology, Ibrahimpattanam, Vijayawada, Krishna, AP, India

Abstract

In this work, we propose a new platform to enable service providers, such as web site operators, on the Internet to block past abusive users of anonymizing networks (for example, Tor) from further misbehaviour, without compromising their privacy, and while preserving the privacy of all of the non-abusive users. Our system provides a privacy-preserving analog of IP address banning, and is modelled after the well-known Nymble system.

Nymble is a system that provides a blocking mechanism to a server to protect it from misbehaving users connecting through anonymizing networks such as Tor. Anonymous networks allow anyone to visit the public areas of the network. Here users access the Internet services through a series of routers. This hides the user's identities and IP address from the server. This may be an advantage for the misbehaving users to destroy popular websites. To avoid such activities, servers may try to block the misbehaving user, but it is not possible in case of anonymous networks. In such cases, if the abuser routes through an anonymizing network, administrators block all known exit nodes of anonymizing networks, denying anonymous access to misbehaving and behaving users. To overcome this problem, a nymble system is designed in which servers can blacklist the misbehaving users without compromising their anonymity. This paper explains the idea that the different service providers have different blacklisting policies. For example, Wikipedia might want to block a user one day for the first misbehaviour, one week for the second one, etc. In order to do this, we have to develop a dynamic link ability window whose length can be increased exponentially. Thus, at the start of each linkability window, all service providers must reset their blacklists and forgive all prior misbehavior.

Keywords

Anonymous Blacklisting, Privacy, Revocation, Pseudonymous Systems, Anonymous Credential Systems.

1. Introduction

Anonymity networks provide users with a means to communicate privately over the Internet. The Tor network is the largest deployed anonymity network; it aims to defend users against traffic analysis attacks by encrypting users' communications and routing them through a worldwide distributed network of volunteer-run relays. As of October 2009, there were 1,532 running Tor relays, operating in 57 different countries, with an estimated 90,000 to 130,000 users (depending mostly on the time of day), connecting from 126 countries, at any given time.

The ability to communicate without fear of network surveillance makes it possible for many users to express ideas or share knowledge that they might otherwise not be willing to reveal for fear of persecution, punishment or simply embarrassment (for a prime example, see the submissions instructions page at Wiki leaks). On the other hand, some users use the veil of anonymity as a license to perform mischievous deeds such as trolling forums or cyber-vandalism. For this reason, some popular websites (for example, Wikipedia and Slashdot [8, 9]) proactively ban any user

connecting from a known anonymous communications network from contributing content, thus limiting freedom of expression.

The privacy offered by Tor is directly related to the size of its anonymity set; i.e., the number of users on the network. The fewer Tor users there are, the easier it is to figure out which one of them initiated a particular connection. As a result, if users are discouraged from using the system, the privacy of those who do continue to use it suffers in consequence. Similarly, the anonymity afforded to each user is related to the number of volunteers running Tor nodes. One side effect of Tor exit nodes being banned from popular web services is that the operators of these relays get banned from these services as well, because their connections come from the same IP address as their Tor relay. This state of affairs provides a fairly strong incentive for many would-be operators not to volunteer to run relays.

Therefore, a real need exists for systems that allow anonymous users to contribute content online, while preserving the ability of service providers to selectively (and subjectively) ban individual users without compromising their anonymity. Not only would such a system benefit the estimated hundreds of thousands of existing Tor users, but it might also be a boon to wider acceptance of Tor [7]. Indeed, the need for an anonymous blacklisting mechanism has been acknowledged by several key people involved with The Tor Project [7]. Thus, it is reasonable to expect that the operators of Tor might be willing to provide the infrastructure necessary to realize such a system, a situation that would greatly reduce the burden on service providers and lead to greater adoption.

Several schemes (e.g., [10]) have been proposed with the goal of allowing anonymous blacklisting of Tor users. The original systems (e.g., [4,6]) attempt to recreate the common practice of IP address banning, without actually revealing a user's IP address; however, these systems suffer from some troubling security issues stemming from the use of trusted third parties (TTPs) who can easily collude to violate a user's anonymity. The most well-known of these is Nymble [4], which is the system after which we model our own.

In pseudonymous credential systems [5] users log into Web sites using pseudonyms, which can be added to a blacklist if a user misbehaves. Unfortunately, this approach results in pseudonymity for all users, and weakens the anonymity provided by the anonymizing network. Anonymous credential systems employ group signatures. Basic group signatures allow servers to revoke a misbehaving user's anonymity by complaining to a group manager. Servers must query the group manager for every authentication, and thus, lacks scalability. Traceable signatures allow the group manager to release a trapdoor that allows all signatures generated by a particular user to be traced; such an approach does not provide the backward unlinkability that we desire, where a user's accesses before the complaint remain anonymous. Backward unlinkability allows for what we call subjective blacklisting, where servers can blacklist users for whatever reason since the privacy of the blacklisted user is not at risk. In contrast, approaches without backward unlinkability need to pay careful attention to when and why a user must have all their connections linked, and users must

worry about whether their behaviors will be judged fairly. Subjective blacklisting is also better suited to servers such as Wikipedia, where misbehaviors such as questionable edits to a Webpage, are hard to define in mathematical terms. In some systems, misbehavior can indeed be defined precisely. For instance, double spending of an “e-coin” is considered a misbehavior in -cash systems following which the offending user is deanonymized. Unfortunately, such systems work for only narrow definitions of misbehavior—it is difficult to map more complex notions of misbehavior onto “double spending” or related approaches.

II. Problem Statement

The Tor network is an overlay network; each Onion Router (OR) runs as a normal user-level process without any special privileges. Each onion router maintains a TLS connection to every other onion router. Each user runs local software called an Onion Proxy (OP) to fetch directories, establish circuits across the network, and handle connections from user applications. These onion proxies accept TCP streams and multiplex them across the circuits. The onion router on the other side of the circuit connects to the requested destinations and relays data. Each onion router maintains a long-term identity key and a short-term onion key. The identity key is used to sign TLS certificates, to sign the OR’s router descriptor (a summary of its keys, address, bandwidth, exit policy, and so on), and (by directory servers) to sign directories. The onion key is used to decrypt requests from users to set up a circuit and negotiate ephemeral keys. The TLS protocol also establishes a short term link key when communicating between ORs. Short-term keys are rotated periodically and independently, to limit the impact of key compromise.

In this way the Tor network forms between the user and a Web server. ANONYMIZING networks such as Tor route traffic through independent nodes in separate administrative domains to hide a client’s IP address. Unfortunately, some users have misused such networks—under the cover of anonymity, users have repeatedly defaced popular Web sites such as Wikipedia. Since Web site administrators cannot blacklist individual malicious users’ IP addresses, they blacklist the entire anonymizing network. Such measures eliminate malicious activity through anonymizing networks at the cost of denying anonymous access to behaving users. In other words, a few “bad apples” can spoil the fun for all. (This has happened repeatedly with Tor.)

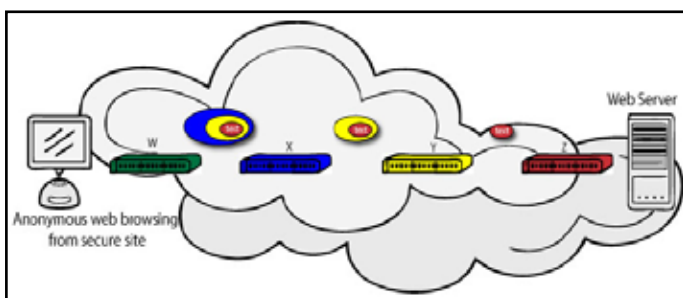


Fig. 1: Tor Diagram Between Web Server and a User

III. Proposed System Nymble Overview

In this paper we propose a secure Nymble system [3], where users acquire an ordered collection of nymbles, a special type of pseudonym [1], to connect to Websites. Without additional information, these nymbles are computationally hard to link and hence, using the stream of nymbles simulates anonymous access to services. Web sites, however, can blacklist users by obtaining a seed for a particular nymble, allowing them to link future nymbles

from the same user. Servers can therefore blacklist anonymous users without knowledge of their IP addresses while allowing behaving users to connect anonymously. Our system ensures that users are aware of their blacklist status before they present a nymble, and disconnect immediately if they are blacklisted.

Although our work applies to anonymizing networks in general, we consider Tor for purposes of exposition. In fact, any number of anonymizing networks can rely on the same Nymble system, blacklisting anonymous users regardless of their anonymizing network(s) of choice. The purpose of the Nymble project is to allow for responsible, anonymous access online. It provides a mechanism for server administrators to block misbehaving users while allowing for honest users to stay anonymous; in fact even the blocked users remain anonymous. The name “Nymble” comes from a play on the word “pseudonym” and “nimble”. Instead of giving users a simple pseudonym, the Nymble system assigns users “nymbles”; that is, a pseudonym with better anonymity properties.

A. Nymble Properties

1. Anonymous blacklisting: A server can block the IP address of a misbehaving user without knowing the identity of the user or his/her IP address.
2. Privacy: Honest and misbehaving users both remain anonymous.
3. Backward anonymity: The blacklisted user’s previous activity remains anonymous/unlinkable, and is refused future connections.
4. Blacklist-status awareness: A user can check whether he/she has been blocked before accessing services at the server.
5. Subjective judging: Since misbehaving users are blocked without compromising their privacy, servers can provide their own definition of “misbehavior”.

Type	Initiator	Responder	Link
Basic	–	Authenticated	Confidential
Auth	Authenticated	Authenticated	Confidential
Anon	Anonymous	Authenticated	Confidential

Fig. 2: Different Types of Channels Utilized in Nymble

B. Anonymizing Networks - Tor

Tor is an anonymizing network that hides a client’s identity (actually, your computer’s IP address) from the servers that it accesses. Tor keeps a client’s IP-address anonymous by bouncing its data packets through a random path of relays. Each relay knows only of the relay that sent it data and the next relay in the random path. As long as the entry and exit nodes do not collude, the client’s connections remain anonymous. Tor provides anonymity, but some people abuse this anonymity. Since website administrators depend on blocking the IP addresses of misbehaving users, they are unable to block misbehaving users who connect through Tor their IP address is hidden after all. Frustrated by repeated offenses through the Tor network, the usual response for websites such as Slashdot and Wikipedia is to block the entire Tor network. This is hardly an optimal solution, as honest users are denied anonymous access to these websites through Tor (or any anonymizing network for that matter).

C. Nymble for Blacklisting Anonymous Users

By providing a mechanism for server administrators to block anonymous misbehaving users, we hope to make the use of anonymizing networks such as Tor more acceptable for server

administrators everywhere. All users remain anonymous—misbehaving users can be blocked without deanonymization, and their activity prior to being blocked remain unlinkable (anonymous).

We present a secure system called Nymble, which provides all the following properties: anonymous authentication, backward unlinkability, subjective blacklisting, fast authentication speeds, rate-limited anonymous connections, revocation auditability (where users can verify whether they have been blacklisted), and also addresses the Sybil attack to make its deployment practical. In Nymble, users acquire an ordered collection of nymbles, a special type of pseudonym, to connect to websites. Without additional information, these nymbles are computationally hard to link, and hence using the stream of nymbles simulates anonymous access to services.

Websites, however, can blacklist users by obtaining a seed for a particular nymble, allowing them to link future nymbles from the same user—those used before the complaint remains unlinkable. Servers can therefore blacklist anonymous users without knowledge of their IP addresses while allowing behaving users to connect anonymously. Our system ensures that users are aware of their blacklist status before they present a nymble, and disconnect immediately if they are blacklisted. Although our work applies to anonymizing networks in general, we consider Tor for purposes of exposition. In fact, any number of anonymizing networks can rely on the same Nymble system, blacklisting anonymous users regardless of their anonymizing network(s) of choice

- Blacklisting anonymous users. We provide a means by which servers can blacklist users of an anonymizing network while maintaining their privacy.
- Practical performance. Our protocol makes use of inexpensive symmetric cryptographic operations to significantly outperform the alternatives.

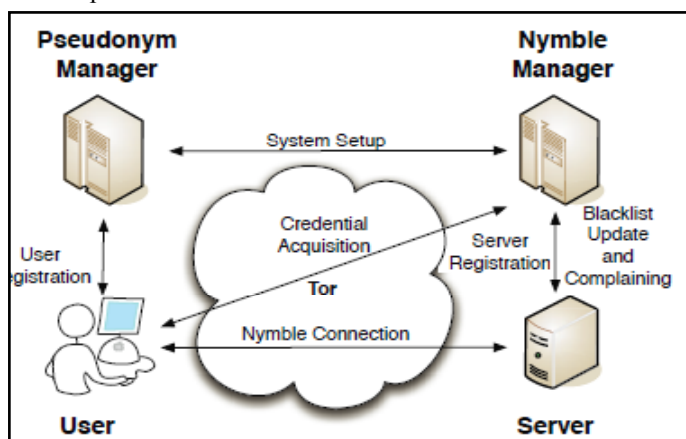


Fig. 2: The Nymble system architecture showing the various modes of interaction. Note that users interact with the NM and servers through the anonymizing network

- Open-source implementation. With the goal of contributing a workable system, we have built an open source implementation of Nymble, which is publicly available. We provide performance statistics to show that our system is indeed practical.

Nymble is based on two administratively-separate “manager” servers, the Pseudonym Manager (PM) and the Nymble Manager (NM). The PM is responsible for pairing a user’s IP address with a pseudonym deterministically generated based on the user’s IP address [1]. The NM pairs a user’s pseudonym with the target server. As long as the two managers are not colluding, the user’s

connections remain anonymous to the PM, pseudonymous to the NM (note that the user does not communicate directly with the NM, and connects to the NM through Tor), and anonymous to servers that the user connects to.

D. Pseudonym Manager

The user (in this case, Alice) must first demonstrate control over a resource that is the Alice’s IP-address. To do this Alice must first connect directly with the PM before receiving a pseudonym. The PM has knowledge of existing Tor routers, and thus can ensure that Alice is communicating with it directly. Note that the PM has no knowledge of the user’s destination [1], similar to the entry node in Tor. The PM’s sole responsibility is to map IP addresses to pseudonyms.

E. Nymble Manager

Alice then connects to the NM through Tor presenting her pseudonym and her target server. The NM does not know the IP address of the user, but the pseudonym provided by the PM guarantees that some unique IP address maps to the pseudonym. She receives a set of nymble tickets as her credential for the target server. These nymble tickets are unlinkable, and therefore Alice can present these nymble tickets (once each) to gain anonymous access at the target server. The nymble ticket provides cryptographic protection as well as a trap door that can be accessed using a linking token.

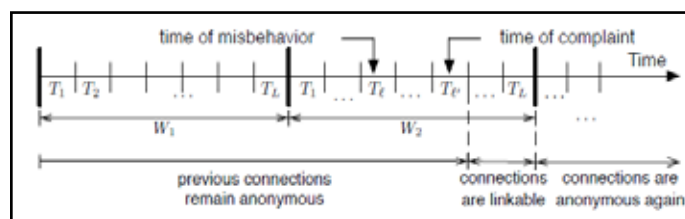


Fig. 3:

To provide the requisite cryptographic protection and security properties, the NM encapsulates nymbles within nymble tickets, and trapdoors within linking tokens. Therefore, we will speak of linking tokens being used to link future nymble tickets. The importance As illustrated in Figure 2, in our system, time is divided into linkability windows of duration W , each of which is split into smaller time periods of duration T , where the number of time periods in a linkability window $L = W/T$ is an integer. We will refer to time periods and linkability windows chronologically as T_1, T_2, \dots, T_L and W_1, W_2, \dots respectively. While a user’s access within a time period is tied to a single nymble ticket, the use of different nymble tickets across time periods grants the user anonymity between time periods—smaller time periods provide users with enough nymble tickets to simulate anonymous access. For example, T could be set to 5 minutes, and W to 1 day. The linkability window serves two purposes—it allows for dynamism since IP addresses can get reassigned to different well-behaved users, making it undesirable to blacklist an IP address indefinitely, and it ensures forgiveness of misbehavior after a certain period of time. of these constructs will become apparent as we proceed.

F. Blacklisting a User

Servers can present a user’s nymble ticket to the NM as part of a complaint. The NM extracts a “linking token” from the nymble ticket [15] that will allow the server to link future connections by the blacklisted user.

Blacklisting a User

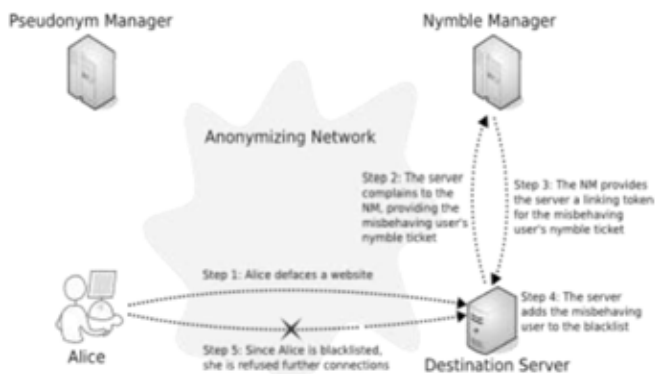


Fig. 4:

The NM also issues servers with blacklists, which users can examine before performing any actions at the server.

By checking servers' blacklists, blacklisted users are assured that their privacy is not compromised. We now explain the process of blacklisting in a little more detail. We first explain how nymble tickets are bound to certain "time periods" and "linkability windows."

This is done for two reasons:

A. Dynamism

IP-addresses can be reassigned to different, well-behaved users making it undesirable to permanently blacklist IP-addresses.

B. Forgiveness

It ensures that bad behavior is forgiven after a certain amount of time. Nymble is a system that allows websites to selectively blacklist users of anonymizing networks such as Tor without knowing the user's IP-address. Users not on the blacklist enjoy anonymity while blacklisted users are not allowed future connections for duration of time while their previous connections remain unlinkable. Since Nymble allows websites to blacklist anonymous users of their choice, and since users are notified of their blacklist status, Nymble gives websites the power to define their own definition of "misbehavior". Our hope is that Nymble's properties well make the usage of anonymizing networks such as Tor more acceptable.

G. Ticket Revocation

Ticket revocation is initiated, when a client is compromised and all his secrets are disclosed to the adversary. In our system adversary takes the ticket associated secrets from the compromised client and start gaining the network services. When gateways have records in the revocation database, they immediately report the revocation to the home TA, which will update and distribute the revocation list for all gateways in the trust domain for reference.

IV. Conclusion

We present a system that allows websites to selectively block users of anonymizing networks such as Tor. Using our system, websites can blacklist users without knowing their IP addresses. Users not on the blacklist enjoy anonymity, while blacklisted users are blocked from making future accesses. Furthermore, blacklisted users' previous connections remain anonymous. Since websites are free to blacklist anonymous users of their choice, and since users are notified of their blacklisting status, our system avoids the complications associated with judging "misbehavior." We believe

that these properties will enhance the acceptability of anonymizing networks such as Tor by enabling websites to selectively block certain users instead of blocking the entire network, all while allowing the remaining (honest) users to stay anonymous.

References

- [1] B. Gedik, L. Liu, "Protecting location privacy with personalized k-anonymity: Architecture and algorithms", IEEE TMC, Vol. 7, No. 1, pp. 1.18, 2008.
- [2] C. Cornelius, A. Kapadia, P.P. Tsang, S.W. Smith, "Nymble: Blocking Misbehaving Users in Anonymizing Networks", Technical Report TR2008-637, Dartmouth College, Computer Science, Dec. 2008.
- [3] J. Feigenbaum, A. Johnson, P.F. Syverson, "A Model of Onion Routing with Provable Anonymity", Proc. Conf. Financial Cryptography, Springer, pp. 57-71, 2007.
- [4] Johnson, P.C., Kapadia, A., Tsang, P.P., Smith, S.W.: "Nymble: Anonymous IP-address blocking", In: Borisov, N., Golle, P. (eds.) Privacy Enhancing Technologies. Lecture Notes in Computer Science, Vol. 4776, pp. 113-133. Springer (June 2007)
- [5] D. Chaum, "Showing Credentials without Identification Transferring Signatures between Unconditionally Unlinkable Pseudonyms", Proc. Int'l Conf. Cryptology (AUSCRYPT), Springer, pp. 246-264, 1990.
- [6] Holt, J.E., Seamons, K.E., "Nym: Practical pseudonymity for anonymous networks", Internet Security Research Lab Technical Report 2006-4, Brigham Young University, Provo, UT, USA (June 2006)
- [7] Lewman, A. handrew@torproject.org: Re: Talking w/ local service CEOs [LJ, goog...], <http://marc.info/?l=tor-talk&m=126137307104914&w=2>, [Private e-mail message to hgrarpamp@gmail.com; 21-December-2009]
- [8] Dingleline, R. harma@freehaven.net: Re: Banned from Slashdot, <http://archives.seul.org/or/talk/Jun-2005/msg00002.html>, [Private e-mail message to Jamie McCarthy; 01-June-2005]
- [9] CmdrTaco: Slashdot FAQ - accounts, <http://slashdot.org/faq/accounts.shtml#ac900>, [Online; accessed 11-January-2010; modified 02-July-2002]
- [10] Brickell, E., Li, J.: Enhanced Privacy ID: A direct anonymous attestation scheme with enhanced revocation capabilities. In: Ning, P., Yu, T. (eds.) WPES. pp. 21-30. ACM (2007)



Mohammad Mastan received his M.Tech in CSE and presently pursuing his Ph.D. He is working as Assistant Professor in NIMRA Institute of Science & Technology Ibrahimpatnam, Vijayawada. He has Six years of experience in total. His research interests include Networks and Security.



Suddapalli Jagadeswari is currently pursuing her M.Tech (CSE) from NIMRA Institute of Science & Technology, Ibrahimpatnam, Vijayawada. She presently doing her project work on Networks. Her research interests include Networking, Adhoc Networks, Information Security and Cryptography.