Poster: Blindspot: Indistinguishable Anonymous Communications

Joseph Gardiner, Shishir Nagaraja
Security Lancaster Research Centre
Lancaster University
United Kingdom, LA1 4WA
Email: {j.gardiner1,s.nagaraja}@lancaster.ac.uk

***Abstract*—Communication anonymity is a key requirement for individuals under targeted surveillance. Practical anonymous communications also require indistinguishability — an adversary should be unable to distinguish between anonymised and non-anonymised traffic for a given user. We propose Blindspot, a design for high-latency anonymous communications that offers indistinguishability and unobservability under a (qualified) global active adversary, which we believe is the first system to provide both properties. Blindspot creates anonymous routes between sender-receiver pairs by subliminally encoding messages within the pre-existing communication behaviour of users within a social network. Specifically, the organic image sharing behaviour of users. Thus channel bandwidth depends on the intensity of image sharing behaviour of users along a route. A major challenge we successfully overcome is that routing must be accomplished in the face of significant restrictions — *channel bandwidth is stochastic*. We show that conventional social network routing strategies do not work. To solve this problem, we propose a novel routing algorithm. We evaluate Blindspot using a real-world dataset. We find that it delivers reasonable results for applications requiring low-volume unobservable communication.**

I. Introduction

In anonymous communication networks, anonymity is provided under the assumption that users solely communicate via anonymous channels. This is unrealistic as widespread adoption of anonymous communications is hindered by usability problems (performance overheads) as well as lack of awareness. Therefore, users of anonymous communication networks also have to communicate over conventional channels. Thus the statistical characteristics of the anonymous channel of a user must be *equivalent* to (or *indistinguishable* from) non-anonymous communication channels. *Indistinguishability* is the property that anonymous and non-anonymous communications are not differentiable by an adversary. Indeed, *without indistinguishability, anonymous communication networks have a basic usability problem.*

A. Current Approaches

Current approaches to this challenge take two directions. One line of work is on censorship resistance mechanisms. These works (such as Collage [1]) all introduce extra communication endpoints (such as cloud file storage services) into traffic, meaning that there is some form of distinguishable behaviour to monitor. Approaches that make use of protocol-mimicking (attempting to hide one type of traffic by making it appear as another) fall victim to incomplete emulation of the protocols that they are attempting to mimic [2].

A second direction is the design of anonymous communication networks. The unlinkability property alone is inadequate for a number of reasons. First, it does not hide the volume of communications and hence leaks enough information to a global adversary who can compile an ordered list of targets [3]. Second, it does not defend against traffic confirmation attacks (for example [4]) where an adversary injects traffic load patterns to determine communicating end points. To combat these problems full unobservability — passive attacker cannot distinguish whether or not a user is communicating — is needed. Systems such as Drac [5] provide unobservable anonymous communications. However, the attacker can still distinguish between the following communicating states of a user: whether or not 'unobservable' anonymous channels are in use, so the protection accorded is still short of what users need.

B. Our Approach

Credible defences against targeted surveillance attacks require both indistinguishability and unobservability. To address these requirements, we propose *Blindspot*, an anonymous communication network that leverages conventional communication channels on social networks. Specifically, the image-exchange behaviour of users. In Blindspot, nodes communicate by broadcasting messages to their neighbours. Messages are steganographically embedded within an image the user uploads. When uploaded by the communicating (sender) node, the message carrying image is available to all its neighbours. Each participating node checks for incoming messages by monitoring images uploaded by its neighbours. It provides *indistinguishability* by ensuring that the image upload behaviour of participating users remains unchanged and provides probabilistic *unobservability* through the use of steganography. No extra communication traffic or endpoints that would have existed without the system in operation are introduced. To provide unlinkability, universal re-encryption [6] is used as it allows re-encryption without prior knowledge of the message paths.

To the best of our knowledge, Blindspot is the first system to provide both properties. This may make it a useful building block of some forms of low volume, delay tolerant anonymous communications, such as announcing a meeting of a social club (corporate gathering, undercover organisation, dissident organisation, or a protest group).

Blindspot routes through the pre-existing social network, thus exploiting trust relationships to secure routing. We leverage this in four ways. First, all routing is through pre-existing social fabric interconnecting nodes, therefore no new communication endpoints are introduced. Second, it leverages the diffusion properties of the social network to efficiently and anonymously route to the destination. Third, Blindspot does not alter statistical characteristics of conventional user traffic. This is achieved by piggy-backing on existing image sharing behaviour, no extra images are added nor are extra uploads scheduled. Channel traffic characteristics, as observable to the adversary, are not altered. So routing depends wholly on the innate image sharing behaviour of the users. This is a non-trivial challenge. Blindspot operates on a network topology where channel bandwidth is severely constrained. And, the latency is a function of stochastic user behaviour along the entire route. Fourth, social networks can be used to detect Sybil attacks by detecting a small cut that separates honest and sybil groups. Blindspot benefits from Sybil resistance properties of social network topologies. Specifically, this restricts the power of adversaries that seek to inject large numbers of misbehaving participants into the network.

II. Evaluation

A. *Current Results*

The system has been tested using a Java simulator. The simulator takes as input a graph and the individual upload behaviour of each node. We then choose pairs of nodes to communicate, such that a source sends to a destination. The simulator works in a scale of days, with monthly image upload counts for each node split evenly amongst 30 days, for a total of 64 months. Each pair of nodes exchange 1 message per month.

The primary dataset that we use is the Flickr dataset from Nagaraja et al [7]. This dataset contains the social graph and monthly image upload behaviour of 7200 users of the Flickr social network.

To test the performance of the algorithm on a differently structured network, we also use a network generated using the Barabasi-Albert (BA) model [8], with a parameter of 5. The Barabasi-Albert model produces a scale free network. We generate a network of 7200 nodes and apply the Flickr upload counts to this in a 1:1 mapping. As can be seen in Figure 1, even under a high level of congestion, the algorithm achieves a delivery rate of over 85% in both of the test networks. In both networks, the average delay for received messages is 1 day. This remains constant as the load increases.

The system has also been tested under the removal of nodes in both a targeted and random manner. Removing nodes at random from the network has little effect on performance. Targeting the high degree nodes reduces the delivery rate by around 30-40%, although even with 50% of the noes removed the delivery rate remains above 50%.

B. *Future Work*

1) Further Experiments: As a next step, we will test the system on a larger network, with an up-to-date dataset of image upload behaviour. We will also test with different levels of participation in order to measure the effect of sparsity of the overlay network on performance.

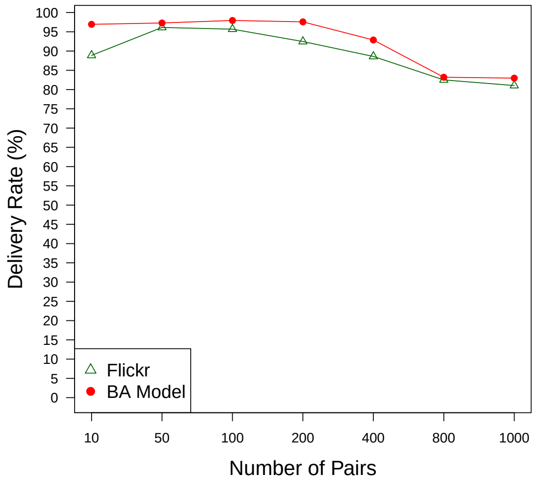

Fig. 1. Performance under increasing number of communicating pairs

2) Improved Key Usage: Currently, multiple keys are in use to enable encryption for providing anonymity. The routing algorithm also requires the usage of neighbourhood keys to improve delivery rates. We will look to reduce the key requirements, through means such as identity-based cryptography.

3) System expansion: Currently, the Blindspot is only designed to function over a single social network using images as a medium. One approach to improving the performance of the system will be to explore the system being built as an overlay to multiple social networks simultaneously, with multiple data carriers (for example text and video).

References

[1] S. Burnett, N. Feamster, and S. Vempala, "Chipping away at censorship firewalls with user-generated content," in *Proceedings of the 19th USENIX Conference on Security*, ser. USENIX Security'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 29–29. [Online]. Available: http://dl.acm.org/citation.cfm?id=1929820.1929859

[2] A. Houmansadr, C. Brubaker, and V. Shmatikov, "The parrot is dead: Observing unobservable network communications," in *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, 2013.

[3] G. Danezis and B. Wittneben, "The economics of mass surveillance and the questionable value of anonymous communications." in *WEIS*, 2006.

[4] M. G. Reed, P. F. Syverson, and D. M. Goldschlag, "Anonymous connections and onion routing," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 4, 1998. [Online]. Available: citeseer.ist.psu.edu/reed98anonymous.html

[5] G. Danezis, C. Diaz, C. Troncoso, and B. Laurie, "Drac: an architecture for anonymous low-volume communications," in *Proceedings of the 10th international conference on Privacy enhancing technologies*, ser. PETS'10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 202–219.

[6] P. Golle, M. Jakobsson, A. Juels, and P. Syverson, "Universal re-encryption for mixnets," in *Topics in Cryptology CT-RSA 2004*, ser. Lecture Notes in Computer Science, T. Okamoto, Ed. Springer Berlin Heidelberg, 2004, vol. 2964, pp. 163–178.

[7] S. Nagaraja, A. Houmansadr, P. Piyawongwisal, V. Singh, P. Agarwal, and N. Borisov, "Stegobot: A covert social network botnet," in *Information Hiding*. Springer, 2011, pp. 299–313.

[8] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, pp. 47–97, Jan. 2002.