

Metrics for Security and Performance in Low-Latency Anonymity Systems

Steven J. Murdoch and Robert N. M. Watson

Computer Laboratory, University of Cambridge, UK
<http://www.cl.cam.ac.uk/users/{sjm217,rnw24}>

Abstract. In this paper we explore the tradeoffs between security and performance in anonymity networks such as Tor. Using probability of path compromise as a measure of security, we explore the behaviour of various path selection algorithms with a Tor path simulator. We demonstrate that assumptions about the relative expense of IP addresses and cheapness of bandwidth break down if attackers are allowed to purchase access to botnets, giving plentiful IP addresses, but each with relatively poor symmetric bandwidth. We further propose that the expected latency of data sent through a network is a useful performance metric, show how it may be calculated, and demonstrate the counter-intuitive result that Tor's current path selection scheme, designed for performance, both performs well and is good for anonymity in the presence of a botnet-based adversary.

1 Introduction

The Tor network [1] is the most widely deployed anonymous communication system, whose estimated 250 000 users include companies, human rights workers and law enforcement. The network's security is therefore critical for the safety and commercial concerns of its users. In common with other deployed low-latency anonymity networks, Tor is vulnerable to an attacker who is able to monitor a user's communication both as it enters and leaves the system. Through traffic analysis, the attacker can use timing characteristics to confirm which incoming connection corresponds to an outgoing one, and so discover the user behaviour Tor seeks to hide. It is thus important to understand how the routing of connections through the Tor network affects the risk of their compromise.

A frequently stated problem with Tor is that it significantly slows down web browsing speed. This is, in part, a consequence of the volunteer-operated nature of servers – many are on slow connections or shared with other activity. It is therefore important to make the best use of the limited capacity available, when selecting the path over which a user's traffic will be routed through the Tor network. In order to prevent an attacker manipulating path selection, the servers on each path are selected by the initiator. For best performance, the path selection algorithm must therefore fairly distribute connections based on server capacity, using only information known to the initiator, while being difficult for an adversary to game.

In this paper we present an analysis of path selection algorithms for Tor, including the currently used ones and proposed improvements. We consider their anonymity and performance consequences, based on simulations and models driven by data collected from the deployed Tor network.

In contrast to previous work, we examine a more realistic threat model, in which attackers are limited not only by number of nodes they control, but also their bandwidth capacity. We find that by introducing this generalisation the relative security of different path selection algorithms is substantially changed.

Previous work has shown that alternative algorithms have improved performance, when only the client tested is modified. We show that, by modelling a network where the new algorithm is fully deployed, performance substantially deviates from previous predictions.

The remainder of the paper is structured as follows: Section 2 introduces the basic operation of Tor, and the threat model we will consider; Section 3 discusses related work on alternative path selection algorithms, metrics for evaluating their anonymity, and attacks which exploit their weaknesses; Section 4 describes, in detail, how the existing Tor path selection algorithm operates; Section 5 introduces our metric for path selection security and presents simulation results; finally, Section 6 evaluates the performance of path selection algorithms, when all clients are assumed to be using the new scheme.

2 Design of Tor

The Tor anonymity network, the latest generation of the Onion Routing project, aims to anonymise TCP traffic while maintaining a low enough latency to be usable for interactive protocols such as web browsing. When the Tor server software on a node (an *onion router*) is first configured, it generates a public/private key pair and sends the public half and other routing details to the *directory authorities*, in the form of a router descriptor. The directory authorities communicate between themselves and establish the subset of onion routers which they are all aware of and sign the resulting *consensus directory*.

Users install the Tor client software on their computer and configure their applications to use Tor as a SOCKS proxy (an *onion proxy*). On receiving a connection, the client establishes the desired destination and selects a path consisting of three Tor nodes listed in the consensus directory (which was downloaded from a directory authority or a mirror). The client then connects to each node on the path (a hop) in turn and builds an encrypted tunnel secured by a session key established through an authenticated Diffie-Hellman exchange. Each connection is made through the previous tunnel, so an external observer can only see the connection to the first hop (the entry node).

The client requests that the last hop (the exit node) connects to the desired destination, then splits data to be sent into 512 byte *cells*, encrypts them under all session keys for the path, and sends them to the entry node. Each node on the path removes one layer of encryption, establishes which is the next hop, and sends the cell on. Once the cell reaches the exit node, the final layer of encryption

is removed and the data sent to the destination server. Replies from the server follow the reverse process.

The full details of the procedure are covered in the Tor specification [2], but all that is required to understand for the remainder of the paper is that an encrypted tunnel can be built through the Tor network. Our focus will be on how the onion proxy selects the onion router nodes for each path. Section 4 will describe the path selection algorithms in detail, but we will first discuss what we assume about our attacker’s goals and abilities.

2.1 Threat Model

Tor imposes no restrictions on who can join the network, which has led to the rapid growth of network capacity, but increases the risk that some nodes in the network are malicious. We assume that an attacker’s goal is to link senders with receivers – that is, to de-anonymise the endpoints of traffic. A *global-passive* adversary, conventionally assumed in the study of high-latency anonymity networks, can break the anonymity properties of currently deployed low-latency systems by correlating traffic patterns; however in many situations such a powerful attacker is unrealistic. Instead, we will consider a weaker attacker who is only capable of monitoring traffic on nodes he has injected into the network.

Tor does not intentionally delay messages, or introduce dummy traffic other than limited message padding, so traffic patterns remain almost unchanged as they pass through the network. For this reason, it was assumed and subsequently demonstrated [3, 4] that by observing both ends of a connection, timing patterns are enough to confirm a suspected link between sender and receiver. This attack has also been shown to work even if only a small proportion (e.g. 1 in 2000) of packets can be observed [5]. We will thus assume that a connection is de-anonymised if both the first and last hop on the path are malicious.

In contrast to Tor, JAP [6] requires that operators promise not to engage in malicious behaviour before being admitted. JAP is also a *cascade* system, in which traffic from all users flow over the same path. If an attacker were able to compromise or monitor the single entry and exit nodes for a given cascade then all of its users could be de-anonymised. Our threat model is only appropriate to a network where data may follow arbitrary paths (a *free route system*), and where nodes may freely join and leave. We will not further discuss JAP-style networks in this paper.

3 Related Work

3.1 Path Selection Algorithms for Tor

At a high level, Tor’s path selection algorithm works in two stages. For each hop on the path, Tor first builds a list of nodes which meet requirements (such as reachability and stability) and second, picks a random node according to a weighting scheme. When the design document [1] was written, Tor uniformly

weighted random node selection. This was initially adequate, but as heterogeneity of node bandwidth capabilities increased, path selection was changed to take into account bandwidth capacity of nodes. The basic algorithm is described by Bauer *et al.* [3], and our updated description is in Section 4.

Several proposals from the literature, on improving the Tor selection algorithm, have already been applied to the mainline Tor distribution. These include guard nodes [4] and bandwidth/uptime caps [3]. One further notable paper is by Snader and Borisov [7], in which two proposals are made. Firstly, they suggest that bandwidth estimates used for making routing decisions be measured by opportunistically sampling actual throughput, rather than nodes reporting their own capacity. Secondly, they propose a tunable algorithm for selecting nodes, which weights faster nodes more heavily, depending on user preferences for anonymity versus performance. In this paper we will primarily discuss the latter proposal, but will return to the former in Section 5.

While the current Tor path selection algorithm picks nodes with a probability proportional to their contribution to the total network bandwidth, the Snader Borisov (S-B) tunable variant only uses advertised node capacity to produce a rank ordering of nodes. The probability that a particular node will be selected depends solely on its position within this ordering. More precisely, let the family of functions f_s be defined as:

$$f_s(x) = \frac{1 - 2^{sx}}{1 - 2^s} \quad (\text{for } s \neq 0) \quad (1)$$

$$f_0(x) = x \quad (2)$$

To select each node, the n candidates are sorted in descending order of bandwidth and a number x selected uniformly at random from the interval $[0, 1)$. The selected node is at index $\lfloor n \times f_s(x) \rfloor$. This is equivalent to selecting a random number according to the cumulative distribution function (CDF) defined by the inverse of $f_s(x)$.

The value of s is selected by the client according to their preference for faster nodes. For $s = 0$, nodes are selected uniformly (intended for users with high anonymity requirements) but the higher s is, the more will faster nodes be preferred (for users willing to compromise anonymity for better performance). A practical upper limit is suggested to be 10, at which with $n = 1000$, the highest ranked node will be selected with probability 6%. By modifying a single client to adopt the new strategy, the authors experimentally confirm that higher values of s leads to better performing connections.

A potential problem with parametrised node selection is that if it is possible to identify a user's selection preference, their actions can be linked. Snader and Borisov used a naïve Bayesian classifier in order to establish how accurately a user's selection preference could be fingerprinted, based on path selection. They found that with a training set of 100 000 paths, the probability of correctly identifying s does not exceed 21%.

3.2 Metrics for Path Selection

The general consensus from the literature is that the further a path selection algorithm deviates from uniform weighting of nodes, the lower the anonymity it provides. Following this intuition, Bauer *et al.* adopt normalised Shannon entropy as their definition of anonymity. i.e. if the probability of selecting node x_i is q_i , for $1 \leq i \leq n$, the entropy is:

$$H = - \sum_{i=1}^n q_i \log_2 q_i \quad (3)$$

H is maximal at $\log_2(n)$ when the selection probability is uniform over all nodes, hence entropy can be normalised, giving a quantity $0 \leq S \leq 1$ representing how skewed the probability distribution is:

$$S = \frac{H}{\log_2(n)} \quad (4)$$

Entropy and normalised entropy have been used extensively in the study of high-latency remailers, although over users rather than the paths. For example, a traffic analysis attack which attempts to establish the sender of a message will result in a probability distribution over all system users. The users for which the probability is non-zero make up the anonymity set. One simple metric for anonymity system security is the size of the anonymity set, but this does not capture the non-uniformity. For this reason, Danezis and Serjantov [8] proposed entropy of the anonymity set as the effective size, and Diaz *et al.* [9] proposed normalised entropy as the system's degree of anonymity.

Snader and Borisov have adopted a different metric, the Gini coefficient, but it also measures the deviation from uniform path selection. The Gini coefficient G equals the normalised area between the CDF of the probability distribution being measured, and the uniform CDF. Thus $G = 0$ represents uniform distribution over all candidate nodes and $G = 1$ indicates that the same single node will always be chosen.

3.3 Attacks on Path Selection

A number of previous publications have taken advantage of Tor's path selection algorithm in developing attacks. Øverlier and Syverson [4] showed that, by causing a client to repeatedly generate fresh paths, an attacker can quickly gain a high probability of controlling both the first and last hop. This was performed in the context of hidden services, a feature of Tor which permits the pseudonymous operation of a server, but the same concepts apply to normal connections.

In order to de-anonymise connections through Tor, an attacker can simply add enough nodes to maintain a high probability that a connection will start and end with a malicious node – effectively a Sybil attack [10]. However, a much more economical variant, proposed by Øverlier and Syverson [4], is to exploit the fact that load-balancing calculations are based on nodes' self-reported bandwidth.

Bauer *et al.* [3] demonstrated that by artificially inflating bandwidth claims, an attacker could compromise 46% of connections while controlling only 6 out of the 66 nodes on a private Tor network. In this paper we will expand on the results of Bauer *et al.*, simulating this attack while varying path selection algorithm. We also employ node parameters from the real Tor network, rather than a private one, and use a version of the Tor path selection algorithm adapted to respond to Bauer’s attack by capping per-node advertised bandwidth.

4 Path Selection in Tor

In this section, we discuss the Tor path selection algorithm and proposed variants. Because the Tor path selection algorithms have changed over time, we have chosen a version from late 2007 based on information from a combination of the Tor path selection specification [11] and the Tor source code.

Tor clients base their path decisions on two types of information: a database of Tor node properties provided by Tor directory authorities in the consensus directory, and the specific requirements of the requested connection. We detail the path selection algorithm in Tor in terms of node properties, path requirements, and node selection weighting.

4.1 Node Properties

Table 1. Node properties used in selecting paths

Network address	The IP address of the node
Node family	Administrator-configured equivalence class
Node bandwidth	Average, burst, and observed capacity; for most purposes, the average bandwidth is used, capped to 10 MB/s in order to limit bandwidth-based attacks
Node uptime	Time since node came online
Node status	Whether the node is <i>running</i> or is currently <i>hibernating</i> and so unwilling to accept connections
Exit policy	What, if any, types of exit node use is permitted
Publication timestamp	When this information was last received by a directory authority
Tor version	Version of Tor on the node

Table 1 lists the node properties which are reported by nodes to the Tor directory authorities, and to some extent monitored for accuracy. The path selection algorithm takes, as input, the following properties derived from the consensus directory:

- A node is *valid* if the version of Tor running on the node is not one of a number of known-bad versions, which render the node unsuitable for use as an entry or exit node.
- A node is *active* if the publication timestamp in the consensus directory is no more than 20 hours old, the node is valid and is marked as running.
- A node is *stable* if its self-reported uptime is among the top half of Tor network nodes, or whose uptime at least 30 days¹.
- A node is *fast* if its capped advertised bandwidth is among the top 7/8 of Tor network nodes, or whose capped advertised bandwidth is at least 100 kB/s.

4.2 Path Requirements

For any path, no two nodes may be in the same equivalence class, i.e. if their IP addresses have the same most significant 16 bits (in the same /16, following CIDR notation [12]), or if both nodes declares the other to be in the same family. This helps to avoid two nodes on a path being in the same administrative domain, as well as limiting the effectiveness of injecting many nodes using easily attainable /24 address ranges.

A *stable path* is one made solely of stable nodes. Stable paths will be used for connections that are expected to be long-lived, currently determined using a table of port numbers. For example, SSH and IRC connections will use stable paths, with the intent of providing more reliable service.

A *fast path* is one made up solely of fast nodes. Fast paths will be used for almost all connections, except for setting up and connecting to hidden services. The intent here is to provide the latter greater anonymity by allowing more nodes from the network to be used for these security-critical functions.

4.3 Node Selection Weighting

Once the list of acceptable candidate nodes is built, three random selection algorithms may be used to select from them:

- *Simple random selection (SRS)*, in which the node is selected from a set of candidate nodes with uniform probability.
- *Bandwidth-weighted random selection (BWRS)*, in which the node is selected from a set of candidate nodes with probability proportional to their individual capped bandwidth.
- *Adjusted bandwidth-weighted random selection (ABWRS)*, where bandwidth-weighted selection is used, but exit nodes are entirely eliminated from selection if less than one third of overall network bandwidth is advertised by exit nodes, or weighted in order to reduce the chances of selecting an exit node for other points on the resulting path.

¹ More recent versions of Tor directory authority implement basic validity checking on uptime; as we do not specifically consider time as a resource, this does not change our analysis. However, high-uptime malicious nodes are not difficult for an appropriately resourced attacker to create.

Nodes are selected randomly from the set of suitable nodes using one of two path selection algorithms: SRS for paths requiring additional security, and BWRS and ABWRS for fast paths. Entry and middle node selection are adjusted to avoid overload of potential exit nodes. Nodes are selected beginning with the exit node, typically the most constrained in choice due to the need to use a node with a suitable exit-policy, followed by the entry node, and then middle nodes.

5 Measuring the Probability of Path Compromise

Our metric of security is probability, with respect to cost and selection algorithm, that an attacker compromises a connection by controlling the first and last hop. While many factors could be considered in a definition of investment, we address the cost of attack in terms of two elements that act as inputs to the path selection algorithms: the number of nodes available to the attacker, and the bandwidth available for each node.

After one node in a family is selected, other nodes in the same family will be ignored during further path selection so we assume that attackers will invest in nodes only in different families in order to avoid wasting resources. While not traditionally considered, this approach is consistent with an attacker controlling a botnet, because bots are geographically and network topologically diverse. IP addresses and bandwidth could be further mapped into real-world costs, such as the costs of co-location and hardware, if desired.

Our analysis assumes optimal performance by Tor in determining the correct uptime and throughput of nodes; if nodes are able to lead Tor directory authorities to incorrectly publish information on their performance, this may lead to a greater success rate for the attacker, but is beyond the scope of this work. While Tor is currently unable to fully assess whether the advertised bandwidth of a node is accurate, proposals have been made to do so [7]; likewise, Tor increasingly tracks the reputation of nodes in the network in order to evaluate their stability claims.

5.1 Experimental Design

We have created a Tor path simulator, which accepts as input the existing consensus directory tracked by the Tor directory authorities, and a set of artificially introduced malicious nodes with given parameters. The simulator generates paths with specified constraints, such as the requirement for being fast or stable, and evaluates the probability that each path selected is compromised.

We captured a snapshot of the Tor consensus directory on 7 June 2007, which includes 1 484 router descriptors, of which 1 044 are considered active and 521 are considered both active and stable. While the set of nodes available for exit traffic varies by protocol and destination, 325 nodes in the data set are appropriate for general HTTP exit traffic.

We then ran the simulator with various path requirements and attacker node investments in order to evaluate the protection that the path selection algorithms

offered. We consider only the overall probability that any given connection will be compromised, not whether any given end-user will have their connection compromised. This requires us to observe the guard requirements for entry nodes, but not to consider the conditional probability of connection compromise given a client’s previous path choices.

We compared four path selection algorithms:

B/W: Tor’s default bandwidth-weighted algorithm based on a combination of the BWRS and ABWRS algorithms.

Uniform: Tor’s uniform selection path algorithm, which uses SRS for more security-sensitive paths, such as with hidden services.

S-B(1): The S-B path algorithm with parameter $s = 1$, reflecting a desire for increased anonymity.

S-B(15): The S-B path algorithm with parameter $s = 15$, reflecting a desire for increased performance.

The simulator selects 1000 random paths through the network and determines whether the first and last hop are malicious, generating an approximate probability of compromise given the attacker model and path selection algorithm. We implement two attacker investment strategies that reflect possible cost models: one in which each additional node adds a fixed amount of bandwidth, and a second in which the attacker is able to invest a fixed total amount of bandwidth over a variable number of nodes.

The former reflects a world in which many sites contribute malicious nodes with bandwidth, such as a botnet², and the latter in which a single site is able to use diverse network addresses over a single link of fixed bandwidth, such as in a co-location centre³. These represent extremes in strategy, but allow us to evaluate the effectiveness of different path selection algorithms inside, and outside, of their assumptions.

5.2 Results

In our first experiment, shown in Figure 1, we compare the rate of path compromise when injecting 20 kB/s and 256 kB/s malicious nodes over the four path selection algorithms. The uniform and S-B(1) algorithms are unaffected and relatively unaffected, respectively, by the differing bandwidth of injected nodes. With low bandwidth injected nodes, the compromise rate for the bandwidth-weighted path selection algorithm is similar to S-B(15), but with higher bandwidth injected nodes it is comparable to uniform selection.

² Dagon [13] has determined that botnets see both high geographic and network topology distribution. One small botnet consists of 1965 members distributed over 305 unique ASes, 70 unique /8s, 1084 unique /16s, and 1872 unique /24s.

³ Because collocation centres frequently carry network traffic for many different customers and participate fully in BGP, we consider them to also have considerable address space access, able to span multiple /16 networks.

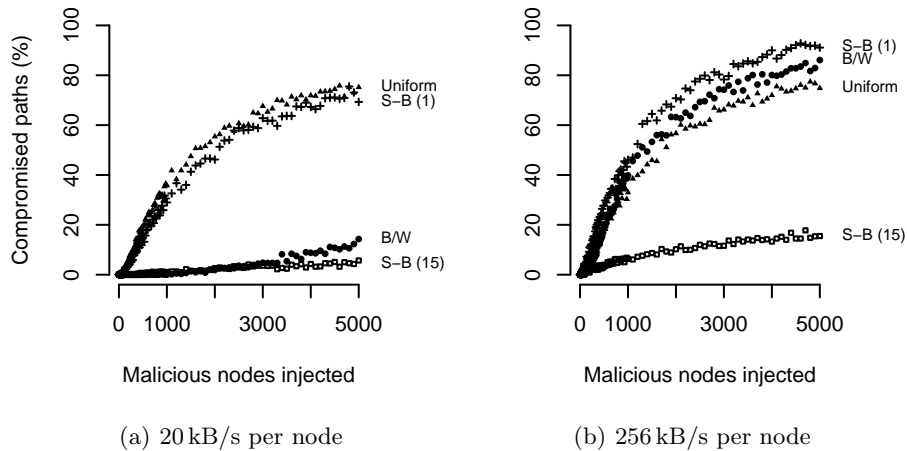


Fig. 1. Percentage of 1 000 generated paths compromised by injecting nodes at 20 kB/s per node or 256 kB/s per node, using the Tor bandwidth-weighted, uniform, S-B(1) and S-B(15) path selection algorithms

In our second experiment, shown in Figure 2, we compare compromise rate, while injecting 20 kB/s nodes for HTTP and SSH connections. Tor requires the use of a stable path for SSH, due to the expectation of longer connection life spans, unlike HTTP. As a result of the smaller pool of stable nodes, the compromise curve for uniform selection has a significantly steeper slope with SSH than it does for HTTP. It requires one third fewer malicious nodes (and hence two thirds the bandwidth) for an attacker to compromise half of the connections using 20 kB/s malicious nodes.

In our third experiment, shown in Figure 3, we examine the probability of compromise when a fixed investment of bandwidth (100 MB/s) is shared over a varying number of nodes. The Tor uniform path selection algorithm exhibits the expected behaviour that rate of compromise corresponds simply to the number of nodes injected, and not their bandwidth. S-B(1) tracks the Tor uniform path selection algorithm, placing little weight on bandwidth. Because the Tor bandwidth cap is 10 MB/s per node, the full investment of bandwidth is only realised once at least ten injected nodes are present. The Tor bandwidth-weighted path selection algorithm therefore offers essentially a constant compromise rate, once the bandwidth per node drops below the bandwidth cap.

The S-B algorithm with parameter 15 exhibits quite interesting behaviour: as long as the malicious node bandwidth is greater than the bandwidth of almost all Tor nodes, performance grows rapidly to a 50 percent compromise rate at 24 nodes. After that point, the probability of compromise drops rapidly because the bandwidth of malicious nodes drops and the probability of selecting them is greatly reduced.

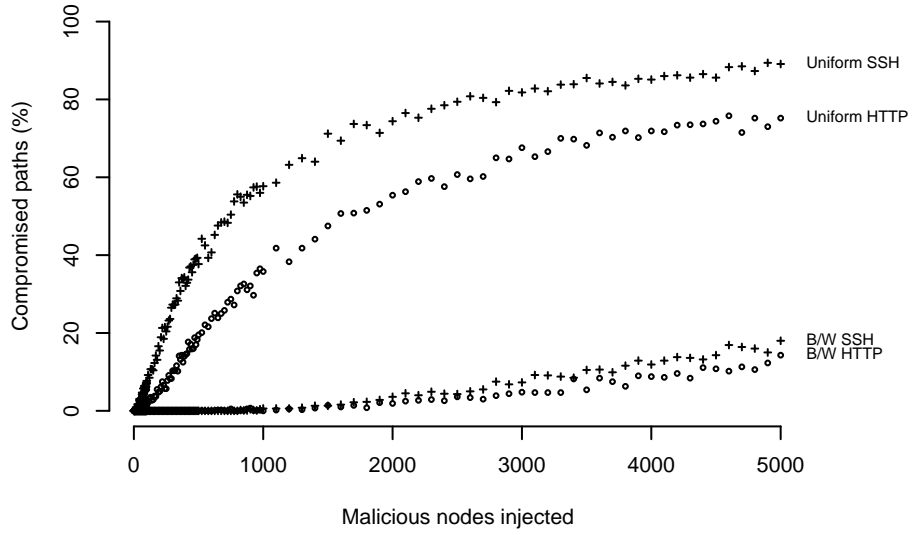


Fig. 2. Comparison of HTTP and SSH compromise probability when injecting 20 kB/s malicious nodes

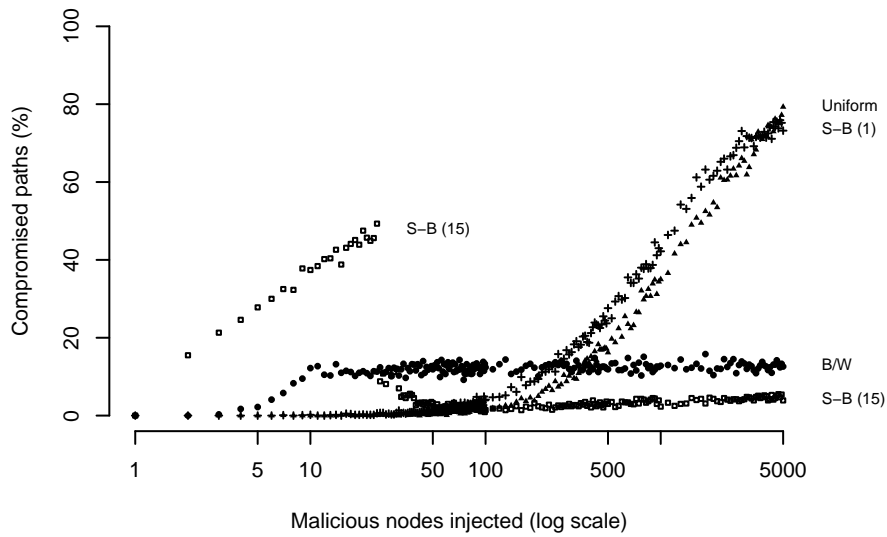


Fig. 3. Compromise rate as a bandwidth budget of 100 MB/s is distributed over a varying number of malicious injected nodes

These results lead to a surprising conclusion: the Tor bandwidth-weighted path selection algorithm and S-B(15), tuned for better performance, have a significantly lower compromise rate for low-bandwidth malicious nodes injected in large numbers, than the Tor uniform and S-B(1) path selection algorithms, which are tuned for better anonymity. Both the Tor uniform and the S-B(1) path selection algorithms resist attack from a small number of malicious high bandwidth nodes. However, botnets allow attackers to obtain large quantities of low-bandwidth nodes. This is in contrast to the higher bandwidth nodes apparently envisioned by the Tor designers when implementing the 10 MB/s bandwidth cap.

Metrics of anonymity based solely on uniformity of path selection, such as entropy or Gini coefficient do not capture these factors. The link between these metrics and practical security depends on the assumption that cost of injecting nodes is independent of the path selection algorithm and node parameters. In the case of Tor, where path selection probability depends strongly on bandwidth, this is analogous to assuming that an attacker has unlimited bandwidth, but is constrained by IP addresses. Where a more realistic view of the threat model is adopted, in which both bandwidth and IP addresses have a cost, path selection algorithms which are supposedly less secure under the entropy model may actually resist attack, provided bandwidth capacity can be accurately tracked.

5.3 Generalising the Attacker

In this section, we examined attackers with a variable number of nodes, and different constraints on bandwidth. We can generalise this approach by modelling an attacker as having a budget c , and where the cost of buying n nodes each using b bandwidth is $C(n, b)$. Such an attacker can choose to occupy any point (n, b) such that $C(n, b) \leq c$, and rationally will pick the point at which the probability of compromising paths is the highest.

Figure 4 shows the path compromise probability for a series of points in the (n, b) space. Lines overlaid show the bandwidth and node tradeoffs available to the three attackers discussed in this section – previous graphs represent slices in Figure 4. The attacker shown in Figure 1(a) and Figure 2, with a variable number of 20 kB/s nodes can be modelled as $b \leq 20$; the attacker shown in Figure 1(b) is similar, with $b \leq 256$. An attacker with a constant bandwidth, shown in Figure 3, can be modelled as $n \times b \leq 100\,000$.

No one path selection algorithm gives the minimum compromise rate for all points, so no algorithm is clearly the most secure. Instead, the best-performing scheme depends on the threat model, which in turn defines allowable points in the (n, b) space. Once this is done, each strategy can be examined to establish the attacker’s maximum compromise rate, and the selection algorithm with the minimum selected.

6 Modelling the Performance of Tor

Our metric for performance is the expected processing time for a cell. One other proposed metric for performance, as measured by Snader and Borisov, is the

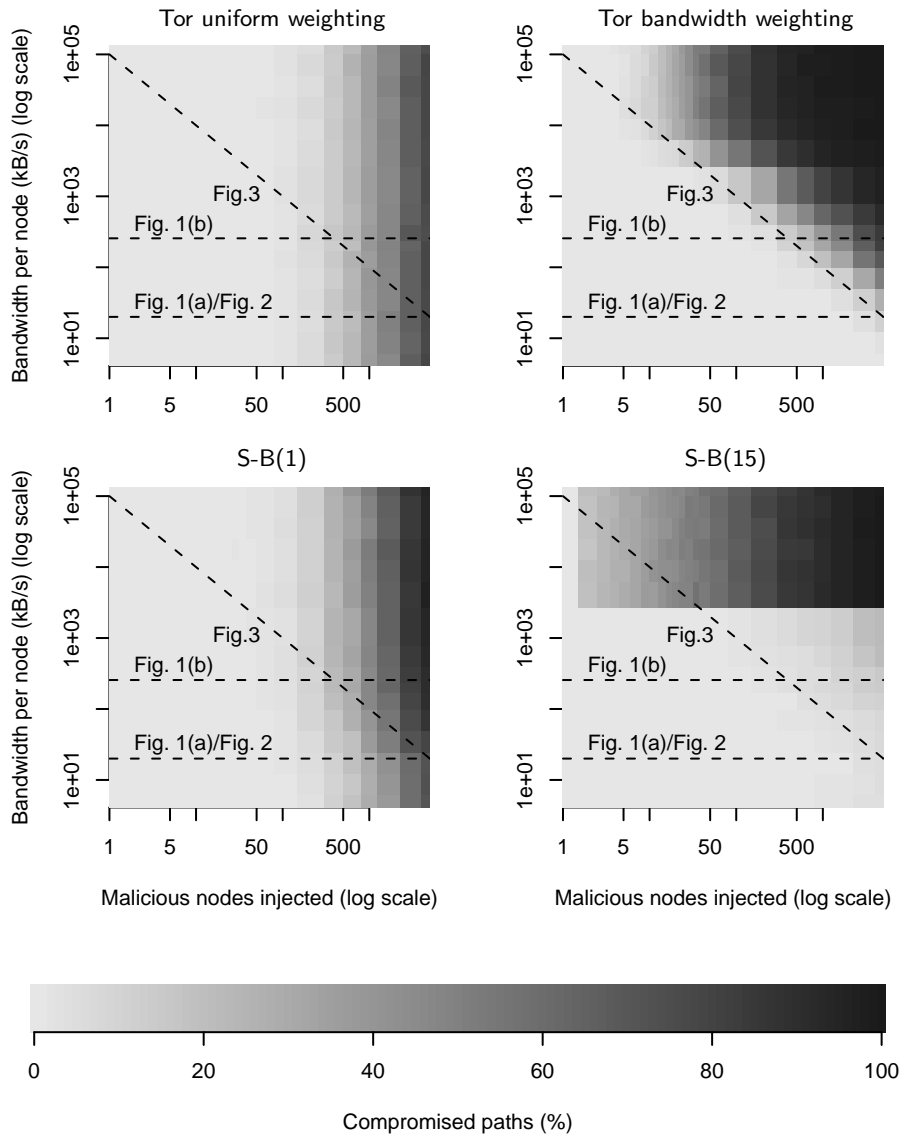


Fig. 4. Compromise rate versus nodes injected and bandwidth per node, for each path selection strategy. Dashed lines indicate the slice of points used for previous figures

throughput of a node which adopts the modified selection scheme. This is a good measure for user desires, but it does not take into account system-level effects that would result if all nodes adopted the same strategy. In contrast, this section models the effect of modifying the full network.

In the previous section, we evaluated the performance of four Tor path selection algorithms with respect to their susceptibility to compromise, concluding that they vary in response to the injection of malicious nodes. However, that analysis is not useful in isolation, as the Tor design is intended to balance competing demands for anonymity and performance. Here, we evaluate the performance of the path selection algorithms with respect to network performance.

6.1 Queueing Theory Background

We will consider an infinite length queue Q , with mean arrival rate λ cells per second, and whose requests are processed by a single server, each taking on average \bar{x} seconds. Initially we will make no assumptions about the distribution of processing duration, but the input process is assumed to be Poisson i.e. we have a M/G/1 queue. The Poisson assumption is known not to perfectly match actual usage, but with a sufficiently large number of users it should be close enough for our results to be useful.

The average time a request will wait in the system (firstly waiting in the queue, then being serviced) is therefore:

$$t = \bar{x} + \bar{w} \tag{5}$$

Where \bar{w} is the average time a request will remain in the queue.

Let the *utilisation factor* for Q be $\rho = \lambda\bar{x}$. A queue for which $0 \leq \rho < 1$ has a finite value of \bar{w} .

From the Pollaczek-Khinchin result [14, p16], we can calculate \bar{w} as follows:

$$\bar{w} = \frac{\lambda\bar{x}^2}{2(1-\rho)} \tag{6}$$

6.2 Calculating Waiting Time for a Family of M/D/1 Queues

We will model a single-hop anonymity network as a family of M/D/1 queues, that is an infinite length queue Q_i , for $1 \leq i \leq n$ with a Poisson input process with rate λ_i and constant processing time of x . The user-base of the whole network can be treated as a Poisson process of rate Λ . Each client will select paths according to a path selection algorithm, which results in node i being selected with probability q_i . The traffic at any individual node λ_i will therefore be $q_i\Lambda$. From Equation 5 and Equation 6, the average time t_i a cell will wait in queue Q_i is:

$$t_i = x_i + \frac{q_i\Lambda x_i^2}{2(1 - q_i\Lambda x_i)} \tag{7}$$

The expected waiting time for a cell is t_i weighted by the probability of Q_i being selected, i.e.

$$T = \sum_{i=1}^n q_i t_i = \sum_{i=1}^n \frac{q_i x_i (2 - q_i x_i \Lambda)}{2(1 - q_i x_i \Lambda)} \quad (8)$$

6.3 Waiting Time for Tor's Bandwidth-Weighted Algorithm

We will first consider the average waiting time for the Tor network, with the bandwidth-weighted algorithm. Here, the probability q_i of selecting a node is the ratio between that node's bandwidth $1/x_i$ and the network total M .

$$q_i = \frac{1}{x_i M} \quad (9)$$

From Equation 8 and Equation 9, the time a cell is expected to wait in the network is:

$$T_{\text{b/w weighted}} = \sum_{i=1}^n \frac{2 - \frac{\Lambda}{M}}{2(M - \Lambda)} = n \frac{2 - \frac{\Lambda}{M}}{2(M - \Lambda)} \quad (10)$$

6.4 Waiting Time for Tor's Uniform Path Selection Algorithm

In contrast, when bandwidth is not considered, the probability of selecting a node Q_i is n^{-1} . Hence from Equation 8 the expected waiting time becomes:

$$T_{\text{uniform}} = \sum_{i=1}^n \frac{n^{-1} x_i (2 - n^{-1} x_i \Lambda)}{2(1 - n^{-1} x_i \Lambda)} \quad (11)$$

6.5 Waiting Time for the S-B Selection Algorithm

The selection probability q_i for the S-B scheme depends on the position of the node Q_i in the bandwidth ranking. From this, $T_{\text{S-B}}$ can be found by substituting the expression for q_i into Equation 8. The CDF for the probability distribution of q_i is the inverse of Equation 1 and Equation 2. If x_i are sorted in ascending order the selection probability is therefore:

$$q_i = s^{-1} \left(\log_2 \left(1 - \frac{n-i+1}{n} (1-2^s) \right) - \log_2 \left(1 - \frac{n-i}{n} (1-2^s) \right) \right) \quad (12)$$

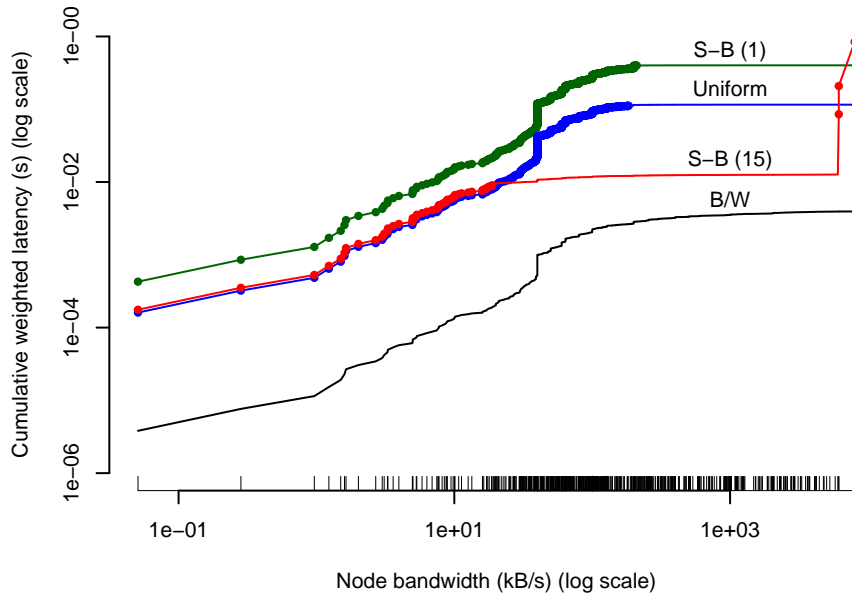


Fig. 5. Cumulative weighted expected waiting time at each node in the network. The rightmost point shows the expected latency for the whole network. Dots are drawn at nodes which will be overloaded

6.6 Comparing Path Selection Algorithms

We can calculate the expected latency for the whole network from Equation 8, by establishing the selection probabilities q_i for each of the path selection algorithms. Also, we require x_i which can be obtained from our Tor consensus directory snapshot. Finally, we need the input rate for the network. Each Tor node reports how much traffic it has processed over a number of 900 second periods, so by taking the median for each node, and summing these, we can approximate Λ . We find that the utilisation factor for the network is 49.6%.

When calculating the expected network latency, we face the problem that for the S-B path selection algorithm, some nodes have a utilisation factor greater than 1 (750 nodes for S-B(1) and 53 nodes for S-B(15)). In our model these nodes would have an infinite queue length and hence our expected latency for the network would also be infinite. The expected latency, using the Tor bandwidth-weighted algorithm, is 4 ms. In contrast, when only a single node was modified to use S-B(15), its latency approximately halved [7]. The fact that the network cannot sustain the same overall throughput with the S-B algorithm, strongly suggests that the current Tor path selection algorithm is superior.

To quantitatively compare path selection algorithms, we can instead drop the overall throughput and see at which point the expected queue length becomes non-infinite. Even at 1% throughput, the expected waiting time is still infinite, but now because of 4 and 1 low-bandwidth nodes for S-B(1) and S-B(15) respec-

tively. The throughput must be reduced to 0.25% for S-B(15) to generate a well defined expected latency of 13 ms; at the same point Tor's bandwidth-weighted algorithm gives 3 ms. At 0.025% throughput, S-B(1) has an expected latency of 48 ms whereas Tor's bandwidth-weighted algorithm still gives 3 ms.

Alternatively, for nodes which are overloaded, we can clip the expected latency to the maximum value of the rest of the network (analogous to a timeout). The result is shown in Figure 5. Here we can see that the expected network latency is still far higher for the S-B selection algorithms: 401 ms for S-B(1) and 275 ms for S-B(15), compared to 4 ms for Tor's bandwidth-weighted algorithm. Also, the timeout failures are substantial: 65% for S-B(1), because of the 750 nodes which are overloaded (72% of the network) and 44% for S-B(15), because the top three nodes are overloaded, which are used for 43.9% of selections.

7 Conclusion

In this paper, we have analysed the effectiveness of several Tor path selection algorithms with in terms of two metrics: probability of path compromise with respect to attacker investment, and expected latency. This analysis has demonstrated the surprising result that not only does Tor's default bandwidth-weighted path selection algorithm offer improved performance over the supposedly more secure Tor uniform path selection algorithm, but also offers improved anonymity in the presence of node-rich but bandwidth-poor attackers.

The vulnerability of supposedly secure path selection algorithms reflects a historical assumption that bandwidth is a low-cost commodity to acquire, but that large numbers of nodes in different equivalence classes are expensive. We believe that this assumption no longer holds due to the proliferation of botnets, which frequently have poor upstream bandwidth from each individual node, but high network and geographical diversity.

Acknowledgements

The authors gratefully acknowledge the Tor Project and Google, Inc. for supporting this research. We would like to thank Nikita Borisov, Richard Clayton, George Danezis, Roger Dingledine, Markus Kuhn, Nick Mathewson, Andrei Serjantov, and the anonymous reviewers for their feedback, and Claudia Diaz and Carmela Troncoso, who shepherded the paper. We would also like to thank David Dagon for providing the botnet data we used in our analysis.

References

1. Dingledine, R., Mathewson, N., Syverson, P.: Tor: The second-generation onion router. In: Proceedings of the 13th USENIX Security Symposium. (August 2004)
2. Dingledine, R., Mathewson, N.: Tor protocol specification. Technical report, The Tor Project (October 2007) <https://www.torproject.org/svn/trunk/doc/spec/tor-spec.txt>.

3. Bauer, K., McCoy, D., Grunwald, D., Kohno, T., Sicker, D.: Low-resource routing attacks against anonymous systems. Technical Report CU-CS-1025-07, University of Colorado at Boulder (2007)
4. Øverlier, L., Syverson, P.F.: Locating hidden servers. In: Proceedings of the 2006 IEEE Symposium on Security and Privacy, Oakland, CA, US, IEEE Computer Society (May 2006)
5. Murdoch, S.J., Zieliński, P.: Sampled traffic analysis by Internet-exchange-level adversaries. In: Privacy Enhancing Technologies (PET). Volume 4776 of LNCS., Ottawa, Canada, Springer (June 2007)
6. Berthold, O., Federrath, H., Köpsell, S.: Web MIXes: A system for anonymous and unobservable Internet access. In Federrath, H., ed.: Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability. Volume 2009 of LNCS., Springer (July 2000) 115–129
7. Snader, R., Borisov, N.: A tune-up for Tor: Improving security and performance in the Tor network. In: Network & Distributed System Security Symposium, Internet Society (February 2008)
8. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In Dingledine, R., Syverson, P., eds.: Privacy Enhancing Technologies (PET). Volume 2482 of LNCS., San Francisco, CA, US, Springer (April 2002) 259–263
9. Diaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In Dingledine, R., Syverson, P., eds.: Proceedings of Privacy Enhancing Technologies Workshop (PET). Volume 2482 of LNCS., San Francisco, CA, US, Springer (April 2002) 184–188
10. Douceur, J.: The Sybil Attack. In: Proceedings of the 1st International Peer To Peer Systems Workshop (IPTPS 2002). Volume 2429 of LNCS., Springer (March 2002)
11. Dingledine, R., Mathewson, N.: Tor path specification. Technical report, The Tor Project (October 2007) <https://www.torproject.org/svn/trunk/doc/spec/path-spec.txt>.
12. Fuller, V., Li, T.: Classless inter-domain routing (CIDR): The Internet address assignment and aggregation plan. RFC 4632, IETF (August 2006)
13. Dagon, D. Personal communication
14. Kleinrock, L.: Queueing Systems. Volume 2. John Wiley & Sons (1976)