8-2009

Scalable Link-Based Relay Selection for Anonymous Routing

Micah Sherr
University of Pennsylvania, msherr@cis.upenn.edu

Matthew Blaze
University of Pennsylvania, blaze@cis.upenn.edu

Boon Thau Loo
University of Pennsylvania, boonloo@central.cis.upenn.edu

Reprinted from:
Micah Sherr, Matt Blaze, and Boon Thau Loo. Scalable Link-Based Relay Selection for Anonymous Routing. In 9th Privacy Enhancing Technologies Symposium (PETS 2009), August 2009, Lecture Notes on Computer Science 5672, pp. 73–93
DOI: 10.1007/978-3-642-03168-7_5
The original publication is available at springerlink.com.
URL: http://www.springerlink.com/content/y25gv3x584387786/fulltext.pdf

Scalable Link-Based Relay Selection for Anonymous Routing

Abstract
The performance of an anonymous path can be described using many network metrics – e.g., bandwidth, latency, jitter, loss, etc. However, existing relay selection algorithms have focused exclusively on producing paths with high bandwidth. In contrast to traditional node-based path techniques in which relay selection is biased by relays' node-characteristics (i.e., bandwidth), this paper presents the case for link-based path generation in which relay selection is weighted in favor of the highest performing links. Link-based relay selection supports more flexible routing, enabling anonymous paths with low latency, jitter, and loss, in addition to high bandwidth. Link-based approaches are also more secure than node-based techniques, eliminating "hotspots" in the network that attract a disproportionate amount of traffic. For example, misbehaving relays cannot advertise themselves as "low-latency" nodes to attract traffic, since latency has meaning only when measured between two endpoints. We argue that link-based path selection is practical for certain anonymity networks, and describe mechanisms for efficiently storing and disseminating link information.

Comments
Reprinted from:
Micah Sherr, Matt Blaze, and Boon Thau Loo. Scalable Link-Based Relay Selection for Anonymous Routing. In 9th Privacy Enhancing Technologies Symposium (PETS 2009), August 2009, Lecture Notes on Computer Science 5672, pp. 73–93
DOI: 10.1007/978-3-642-03168-7_5

The original publication is available at springerlink.com.
URL: http://www.springerlink.com/content/y25gv3x584387786/fulltext.pdf

Scalable Link-Based Relay Selection for Anonymous Routing

Micah Sherr, Matt Blaze, and Boon Thau Loo

University of Pennsylvania
{msherr,blaze,boonloo}@cis.upenn.edu

Abstract. The performance of an anonymous path can be described using many network metrics – e.g., bandwidth, latency, jitter, loss, etc. However, existing relay selection algorithms have focused exclusively on producing paths with high bandwidth. In contrast to traditional *node-based* path techniques in which relay selection is biased by relays' node-characteristics (i.e., bandwidth), this paper presents the case for *link-based* path generation in which relay selection is weighted in favor of the highest performing links. Link-based relay selection supports more flexible routing, enabling anonymous paths with low latency, jitter, and loss, in addition to high bandwidth. Link-based approaches are also more secure than node-based techniques, eliminating "hotspots" in the network that attract a disproportionate amount of traffic. For example, misbehaving relays cannot advertise themselves as "low-latency" nodes to attract traffic, since latency has meaning only when measured between two endpoints. We argue that link-based path selection is practical for certain anonymity networks, and describe mechanisms for efficiently storing and disseminating link information.

1 Introduction

Anonymous communication networks have been gaining in popularity in recent years. As they scale to support large user bases and diverse applications, there is an increasing need not only for these networks to ensure that high performance routes are selected, but also to provide flexibility to tradeoff between performance and anonymity in order to meet the requirements of different applications.

In response to these challenges, there have been a variety of proposals [6,22,35] that are aimed at improving the performance of anonymous routes. These proposals have primarily used *node characteristics* such as self-advertised bandwidth [6,3] as the main criteria for selecting intermediate relay nodes.

In this paper, we argue that an alternative – one that offers strong security guarantees and flexibility – is to utilize *link-based* path selection strategies. In link-based selection, the sender (also called the *initiator*) selects high performing links to construct her anonymous paths. The initiator ranks randomly generated (but not instantiated) paths according to their predicted end-to-end (e2e) performance, estimated by aggregating the costs of their constituent links. From its set of candidate paths, the initiator selects (and subsequently constructs) a path using a probability distribution weighted by the e2e cost estimates. As

I. Goldberg and M. Atallah (Eds.): PETS 2009, LNCS 5672, pp. 73–93, 2009.

with recently proposed node-based strategies [35], our link-based algorithm allows the sender to bias her paths towards either anonymity or performance. Link-based routing is appropriate for anonymity networks in which the performance of anonymous paths is determined by the network topology rather than local effects at end nodes (e.g., congestion, queuing delay, etc.).

Link-based path selection offers several advantages over node-based techniques. First, link-based path selection supports various metrics such as latency, bandwidth, jitter, and loss. The flexibility provided by link-based solutions enables anonymity networks to support a wide variety of network applications that have previously been considered incompatible with these networks. For example, real-time applications (in particular, VoIP clients) require connections with specific latency, jitter, and loss properties. Existing node-based path selection algorithms cannot accurately predict the link properties of their generated routes and are therefore unfit for particular classes of network communication.

Second, link-based strategies are less susceptible to manipulation. In a node-based scheme, a malicious node can easily advertise favorable node characteristics in order to increase the likelihood of being selected as a relay node [3]. Given that *link* metrics are defined only with respect to a pair of relays, the same attack strategy is harder to succeed without the infiltration of a large number of attackers. For instance, a host cannot truthfully promote itself as a "low-latency node", as such a claim may be accurate only for its nearby peers.

The contributions of this paper are as follows:

The case for link-based strategies: Using realistic network traces [14,38,40], we demonstrate that link-based selection not only achieves a high degree of flexibility by supporting a variety of metrics, it is also significantly more resilient to manipulation as compared to node-based strategies. To quantify the anonymity properties of relay selection, we introduce *node prevalence*, a metric that measures the probability that a relay participates in an anonymous path. For example, using a snapshot of available bandwidths from the Tor [7] network's directory servers, we note that the highest bandwidth Tor relay is expected to participate in nearly 40% of anonymous paths when Tor's default relay selection algorithm is used. In comparison, the most popular node using our link-based selection strategy on a comparable dataset (in which bandwidth is described as a link characteristic) is present in just 2.5% of paths. We show that our techniques leak little information about the communicating parties, protecting their anonymity even against powerful and colluding adversaries.

Practical link-based selection implementation: A potential disadvantage of link-based path selection is the need to maintain pairwise link information. We demonstrate that network *coordinate embedding systems* [4,5,23] provide a lightweight and scalable mechanism for maintaining link-based metrics while requiring only minimal communication overhead at each node. In coordinate systems, each node is mapped to n-dimensional coordinates such that the Cartesian distance between two nodes' coordinates corresponds to the network distance (e.g., latency, bandwidth, jitter, or loss) between them. Participants of coordinate embedding systems update their coordinate by periodically conducting

measurements between themselves and randomly selected peers. Each node maintains a single coordinate for each link metric and updates a directory service whenever its coordinates change. Coordinate embedding systems effectively *linearize* the amount of information required to represent pairwise link characteristics, since the coordinates of N nodes are sufficient to estimate pairwise distances.

2 Assumptions and Limitations

The link-based path selection strategies presented in this paper estimate the e2e performance of potential anonymous routes by aggregating the costs of their constituent hops. For example, the e2e latency of a possible anonymous path is estimated by summing the latencies between adjacent nodes in the path. To be effective, link-based routing requires that path performance (whether it be measured by bandwidth, latency, jitter, etc.) be due to network effects.

If, however, local effects at end nodes (e.g., congestion or queuing delay) dominate performance, then link-based path selection is less effective (since the savings gained from optimizing link costs is overshadowed by node effects). At the extreme, link-based selection becomes equivalent to node-based selection when the communication cost of routing between two nodes is determined solely by properties of the receiving host.

The performance and anonymity results in the remainder of this paper assume path performance is dictated by the network rather than end-host effects. Although we leave the determination of the dominant factors that influence performance in various anonymity networks as a future research direction, we briefly note that link-based relay selection is likely better suited for P2P anonymity networks rather than networks in which the client to relay ratio is very high (e.g., in the case of Tor), causing congestion to determine path performance.

3 Related Work

Previously proposed relay selection techniques have focused on improving the bandwidth of generated paths [6,35]. To produce high bandwidth routes, the Tor [7] path selection algorithm sorts relays in increasing order of bandwidth and computes the sum $B = \sum_{i=0}^{|N|-1} b_i$, where b_i is the bandwidth of node i. The initiator chooses r uniformly at random from $[0, B)$ and selects the node with index k as a relay, where k is the largest integer such that $\sum_{i=0}^{k-1} b_i \leq r$. The initiator repeats this procedure to select each relay in the anonymous circuit [6].[1]

Øverlier and Syverson first identified that Tor's path selection algorithm is susceptible to manipulation [24]. By falsely advertising high bandwidths, nodes under an adversary's control can exploit the weighted probability distribution and increase their chances of being selected. If multiple nodes under the attacker's control are selected as relays, the adversary can apply a circuit-linking

[1] In practice, Tor may apply different weights for entry and exit nodes. For simplicity, we assume that all nodes may function as entry or exit relays.

algorithm [3] or perform timing analysis [21] to discern whether two of its relays reside on the same path. (Tor is designed to restrict each relay to knowing only the previous and next hop [7].) If the attacker controls the first and last relays in an anonymous path, he defeats anonymity since the first and last relays respectively know the identities of the initiator and responder. Bauer *et al.* demonstrate that when an adversary controlled just six of 66 nodes in a Tor deployment on PlanetLab [25], the attacker compromised more than 46% of all anonymous paths [3].

Snader and Borisov [35] propose two modifications to Tor to defend against Øverlier *et al.*'s attack. First, to prevent nodes from reporting false bandwidths, relays report the observed bandwidths of peer relays to the directory server. When queried for a node's bandwidth, the directory server reports the median of the node's observed measurements. Second, Snader and Borisov introduce a more tunable weighting system in which the initiator can tradeoff between anonymity and performance. They define the family of functions

$$f_s(x) = \begin{cases} \frac{1-2^{sx}}{1-2^s} & \text{if } s \neq 0 \\ x & \text{if } s = 0 \end{cases} \tag{1}$$

where s is a parameter chosen by the initiator that allows it to tradeoff between anonymity and performance. After having ranked the relays by bandwidth, the initiator chooses the relay with index $\lfloor n \cdot f_s(x) \rfloor$, where n is chosen uniformly at random from $[0, 1)$. By applying higher values of s, the initiator is able to more heavily bias her selections towards bandwidth. If $s = 0$, a relay is chosen uniformly at random [35]. Each relay is selected independently and without replacement according to the distribution imposed by Eq. 1.

Snader and Borisov's defense relies on *opportunistic measurements* – relays report the observed bandwidths of their peers [35]. There are unfortunately disadvantages of such an approach. First, a relay can report opportunistic measurements only when it participates in an anonymous circuit with a peer. Transmitting the observation to a directory server effectively informs the server of the existence of the circuit as well as the identities of the two relays that constitute one of its hops. Given that directory servers may be malicious, revealing segments of the path is undesirable. Second, the directory cannot discern whether reported measurements are truthful. Colluding malicious relays may (falsely) report that members of their coalition have high bandwidth. If there are a sufficient number of attackers to influence the median of a relay's measurements, then Øverlier *et al.*'s attack becomes feasible. Finally, as noted in Murdoch and Watson's recent work [22], attackers may have access to large botnets and may therefore join the anonymity network with relays that have sufficient bandwidth to attract peers. The use of opportunistic measurements attempts to protect against false self-reported measurements, but does not prevent an attacker from acquiring high performing nodes to attract traffic. As we show below, link-based measurements inherently reduce the attacker's ability to influence path selection, as each node is restricted to advertising a single coordinate, which, in turn, is perceived as favorable only to its nearby peers.

Table 1. Link concatenation operators. The distance between successive links is denoted as $d_1, d_2, ..., d_n$.

Metric	Path cost
Latency / RTT	$\sum_{i=1}^{h} d_i$
Bandwidth	$\min(d_1, d_2, ..., d_h)$
Loss rate	$1 - \prod_{i=1}^{h}(1 - d_i)$
Jitter (variance) (assumes jitter of two successive links is independent)	$\sum_{i=1}^{h} d_i$
Autonomous System (AS) Traversals	$\sum_{i=1}^{h} d_i$

The use of coordinate systems to estimate e2e path performance was first proposed in our earlier position paper [33]. This paper presents novel path selection algorithms, and is the first work of which we are aware that analyzes the performance and anonymity properties of link-based relay selection.

4 Link-Based Path Selection

Existing approaches [6,7,35] to producing high performance anonymous paths have focused exclusively on *node characteristics* – performance metrics (i.e., bandwidth) that may be attributed to individual relays. Node-based relay selection strategies randomly select relays according to a nonuniform probability distribution biased by the relays' node characteristics.

In link-based path selection, the e2e performance of a path is computed by aggregating the cost of all links that comprise the path, where cost is defined in terms of *link characteristics* such as latency, loss, and jitter. (While bandwidth is a node-based characteristic, it can also be represented as a link characteristic by considering the measured available bandwidth on a link connecting two nodes.) The use of link rather than node characteristics enables more flexible routing, as initiators can construct anonymous routes that meet more specific communication requirements.

WEIGHTED *Path Selection* Our link-based path selection algorithm, WEIGHTED, operates in two phases. In the first phase, the initiator rapidly generates (but does not instantiate) candidate paths consisting of three relays chosen uniformly at random without replacement. The initiator computes the e2e cost of each generated candidate path using a *link concatenation operator* (see Table 1).[2] For example, the e2e bandwidth of a path is the minimum of the bandwidths of its links, whereas the latency of the route may be estimated by summing the latencies of its hops.

In the second phase, the initiator sorts the candidate paths by their cost estimates. Using the family of functions introduced by Snader and Borisov [35] (see Eq. 1), the initiator instantiates the candidate path with index $\lfloor n \cdot f_s(x) \rfloor$, where

[2] Our approach may be extended to define the performance of a path in terms of *multiple metrics* by assigning weights to each metric in a manner that reflects its importance as determined by the initiator. The e2e path cost estimate is then calculated as the weighted average over the cost estimates for each individual metric.

n is chosen uniformly at random from $[0, 1)$. As with Snader's and Borisov's algorithm, a larger value of s more heavily weighs path selection in favor of performance. When $s = 0$, each randomly generated path is equally likely to be chosen. For clarity, we will refer to the case in which $s = 0$ as using the UNIFORM selection strategy.

5 The Case for Link-Based Selection

In this section, we present the case for link-based path selection. We demonstrate that link-based anonymous routing is flexible, enabling high performance paths, whether performance be quantified in terms of bandwidth, latency, or jitter. Additionally, we show that our selection strategy is more resilient to manipulation than previously established techniques, providing greater anonymity to the communication endpoints.

We first consider an *oracular* model in which all measurements (node or link) in the network are known to the initiator. This enables us to compare node- and link-based path selection strategies irrespectively of their measurement techniques. We revisit actual implementation strategies in Section 6.

5.1 Performance Analysis

Our performance analysis highlights two main benefits of link-based path selection over existing node-based techniques. First, link-based techniques support a variety of performance metrics, hence offering greater flexibility. In particular, the WEIGHTED selection strategy produces paths with low latency and jitter, few autonomous system (AS) traversals, and high bandwidth. Second, as with recently proposed node-based approaches [35], our link-based relay strategy enables the initiator to carefully tradeoff between anonymity and performance.

Our performance analysis is carried out using a trace-driven path simulator that takes as input an $N \times N$ matrix describing the pairwise network distances (i.e., latency, bandwidth, etc.) between relays. The pairwise link distances used as input to the simulator are obtained from actual network traces [14,40] as well as our own measurements carried out on the PlanetLab testbed [25]. Since the performance and security of link-based path selection is influenced by the underlying topology, we analyze the results of generating 150 anonymous paths between each of the $N(N-1)$ pairs of relays. That is, for each pair of relays, we generate anonymous paths between the pair using the remaining $N-2$ nodes in the dataset as potential relays. The simulator models a single pair of communicants at any given time; i.e., we assume node congestion does not impact path performance. To produce each path, WEIGHTED generates (but does not instantiate) 150 candidate paths before randomly selecting the chosen path according to the weighted (e.g., by bandwidth) probability distribution.

Table 2 describes the trace-driven datasets used as input to our simulator. The `King` [14] and `S`^{3}`-BW` [40] datasets are based on measurements obtained from prior publications and are commonly used in the networking research community; `PL-ASes` and `PL-Jitter` represent newer metrics that are novel to this work. Due

Table 2. Network datasets used to evaluate link-based relay selection

Dataset	Metric	Nodes	Description
`King` [14]	Latency	500	Pairwise latencies captured using the King method [12]
`S`^{3}`-BW` [40]	Available Bandwidth	365	Pairwise bandwidths from PlanetLab measured using `PathChirp` [29]
`PL-ASes`	AS Traversals	156	Pairwise number of AS crossings on PlanetLab measured using `traceroute`
`PL-Jitter`	Jitter (variance)	153	Pairwise jitter (variance of interarrival times of 30 pings) on PlanetLab
`Tor-BW`	Available Bandwidth	500	Available (also called "observed") bandwidth of 500 Tor nodes, obtained from Tor directory servers

to the lack of existing published traces on these metrics, we conducted our own measurements using geographically distributed PlanetLab nodes.

Since simulation time grows geometrically with network size, only the pairwise measurements for the first 500 relays from the `King` and `Tor-BW` datasets are used as input to the simulator. The remaining datasets contained fewer than 500 nodes, and are used in their entirety.

Bandwidth metric: Fig. 1 shows the bandwidth improvement resulting from using WEIGHTED on the `S`^{3}`-BW` dataset. When $s = 9$, WEIGHTED more than doubles the median available bandwidth over all pairwise paths to 42.3 Mbps, compared to 20.1 Mbps when relays are selected uniformly at random (UNIFORM). (Recall that relay selection is weighted more heavily towards performance when s is increased.) The ability to provide high performance bandwidth paths using link-based relay selection is particularly interesting, given that bandwidth is often perceived as a node characteristic [2,15]. Bandwidth may, of course, be represented as a link characteristic (as is the case in the `S`^{3}`-BW` dataset). This latter characterization enables more flexible routing, as bandwidth bottlenecks may result from Internet routing policies rather than node capacities.

Non-bandwidth metrics: Fig. 2-4 demonstrates WEIGHTED's ability to produce high performance paths for non-bandwidth metrics. The median e2e latency of the anonymous paths formed using `Uniform` is 277.2ms (Fig. 2). The median latency decreases by 20.6% to 220.1ms when $s = 3$ and by 52.7% to 131.2ms when $s = 15$. Additionally, WEIGHTED decreases the percentage of high latency paths: 93.0% of paths produced via UNIFORM have latencies of 250ms or greater compared to just 22.5% of routes generated using WEIGHTED with $s = 3$.

Jitter, defined as the variance in interarrival times (measured in ms) of 30 ping messages, significantly decreased using WEIGHTED. As shown in Fig. 3 (log scale), the median jitter decreased by 72% when $s = 3$ and by 97% when $s = 9$.

It may also be advantageous to minimize the number of AS crossings in an anonymous path, both to decrease the probability that a given AS can observe multiple hops in the path [8] and also to potentially achieve greater path performance (since routing within an AS is typically low-latency and high-bandwidth). Although analyzing the relationships between AS traversals, anonymity, and performance is beyond the scope of this paper, we include the metric here to emphasize the flexibility of link-based routing. Fig. 4 shows the cumulative distribution of AS traversals for anonymous paths. Using `Uniform`, 66% of anonymous paths

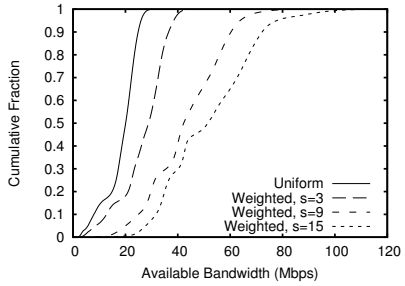

Fig. 1. E2e available bandwidth using the S^3-BW dataset

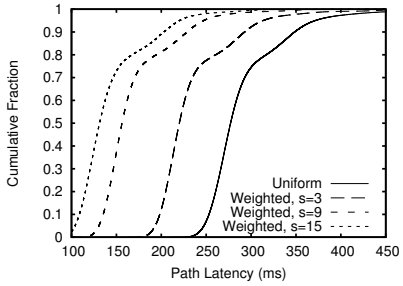

Fig. 2. E2e path latencies using the King dataset

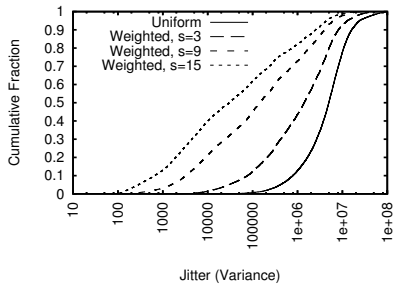

Fig. 3. E2e jitter using the PL-Jitter dataset

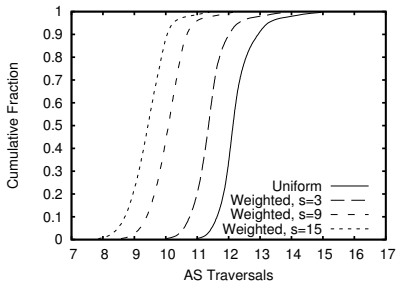

Fig. 4. E2e AS traversals using the PL-ASes dataset

traversed 12 or more ASes. When $s = 3$ and $s = 9$, only 10% and 0.3% of their respective paths crossed 12 or more ASes.

5.2 Anonymity Analysis

Our anonymity analysis aims to compare the anonymity properties of link-based and node-based relay selection under various attacker strategies. As with the existing literature, we consider an anonymous route to be compromised if the attacker controls its first and last relay [35]. (Resiliency to the *predecessor attack* [28,39] is discussed in Appendix B.)

We model an attacker that controls or monitors $f \cdot N$ of a N-node network, where $0 \leq f < 1$. We further assume that the adversary has complete network information and may select *a priori* which of the $f \cdot N$ nodes it controls (e.g., those with highest bandwidth). While this is a particularly strong threat model, it enables us to explore the limitations of our techniques by allowing the attacker to select the most "attractive" relays in a realistic network topology.[3] Due to the

[3] Prior work utilizes attacker models in which the adversary may supplement the network with additional malicious relays [22]. Link-based path selection is difficult to accurately assess using such models, as the performance and anonymity of anonymous paths depend upon the precise locations of all relays.

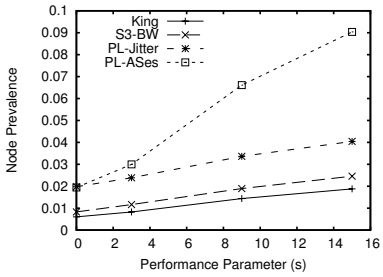

Fig. 5. The maximum of relays' node prevalences in the `King`, `S`^{3}`-BW`, `PL-Jitter`, and `PL-ASes` datasets using WEIGHTED

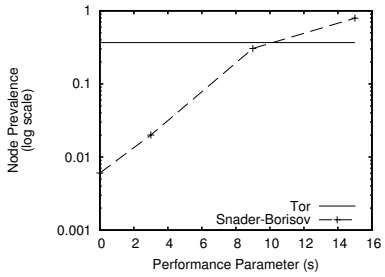

Fig. 6. Maximum node prevalences of relays in the `Tor-BW` dataset using the default Tor routing algorithm and the Snader-Borisov refinement

ease at which an adversary may acquire high performance nodes using a botnet, we view our threat model as conservative, but realistic.

Node Prevalence. To quantitatively compare link- and node-based relay selection, we introduce a new measure of anonymity, *node prevalence*, defined as the probability that a relay is selected as a participant in an anonymous path. Since link-based routing selects relays based on e2e performance estimations (including links containing the initiator or responder), we compute the node prevalence of a relay as the average probability of selection over all combinations of initiators and responders.[4] Intuitively, relays with high node prevalences are valuable to attackers since, by definition, they have a greater chance of being selected in anonymous paths.

In node-based techniques, high-bandwidth nodes are consistently perceived as attractive to all initiators, leading to relays with high node prevalences. In contrast, the likelihood that a node will be attractive for all paths using link-based approaches is fairly small, since a node's attractiveness is a function of the locations of the initiator, responder, and already chosen relays in the path. The ability of link-based relay selection to prevent "hotspots" leads to increased anonymity since a small coalition of malicious relays cannot easily attract a disproportionate amount of traffic.

Fig. 5 plots the *maximum* of all relays' node prevalences – the frequency at which the most popularly chosen node is present in anonymous paths. Even when WEIGHTED is tuned for high performance ($s = 15$), the most popular relay is present in less than 5% of paths in the `King`, `S`^{3}`-BW`, and `PL-Jitter` datasets,

[4] Snader *et al.* [35] propose the use of the Gini Coefficient [11] as a summary statistic of the inequality of relay selection. In contrast, node prevalence measures the popularity of a particular node. By calculating the node prevalence of each relay, we can study the worst-case anonymity of a particular path selection technique, which happens when the adversary has under its control the relays with highest node prevalences (i.e., those used most often in anonymous paths).

and less than 10% of routes using the PL-ASes trace. The corresponding performance of the paths is shown in Fig. 1 through 4.

In comparison, node-based relay selection yields substantially higher node prevalences. Fig. 6 shows the maximum node prevalence for the default Tor path selection strategy [6] and Snader and Borisov's proposed refinement [35] using the Tor-BW dataset. (Tor's routing algorithm takes no performance parameter and is shown as a straight line.) For both strategies, high bandwidth relays are attractive to all initiators. In particular, the highest bandwidth node is present in 36.9% of all paths produced using the default Tor algorithm. The tunable Snader-Borisov strategy has a modest maximum node prevalence of 2.0% when $s = 3$, but results in much poorer anonymity for greater values of s. When $s = 15$, 79.2% of paths contain the node with the greatest bandwidth. Although Fig. 5 and Fig. 6 cannot be directly compared since they use different underlying topologies and metrics, it is apparent from the figures that while there are no statically-attractive relays as perceived by WEIGHTED, node-based techniques result in hotspots that are present in a large fraction of paths.

Attack Strategies. We next consider various strategies available to the attacker. As described above, we utilize a conservative attacker model in which the adversary can choose *a priori* which relays he will compromise (up to some fraction f of the network). We further assume that the attacker has complete network knowledge (i.e., pairwise distances) to which to base his decision.

BestLinks: *Compromising Attractive Links.* In the BestLinks strategy, the attacker compromises the endpoints of the most attractive links. Mirroring the behavior of the initiator, the attacker ranks smaller distances more favorably if the metric is latency, jitter, loss, or AS traversals, and views larger distances as more advantageous for bandwidth. Given an ordering of links, the two endpoints of each link are assigned to the attacker until he controls $f \cdot N$ relays.

The effectiveness of the BestLinks strategy is depicted in Fig. 7. The x-axis denotes the fraction of nodes controlled by the attacker (f), while the y-axis plots the resultant percentage of paths that are compromised. As can be observed from the Fig., WEIGHTED successfully protects most anonymous paths even when the attacker controls 50% of the network. When paths are weighted heavily in favor of performance ($s = 15$) and 30% of the network is controlled by the attacker, only 12.4% of the anonymous paths in the King dataset become compromised (Fig. 7*(a)*). Similarly, for bandwidth (Fig. 7*(b)*), 16.4% of paths are compromised when 30% of the network is malicious. Results for the PL-ASes and PL-Jitter datasets are comparable, and are omitted for brevity.

For comparison, Fig. 8 shows the percentage of compromised paths for node-based selection strategies when the attacker uses the BestNodes attacker strategy on the Tor-BW dataset. Analogous to BestLinks, BestNodes ranks nodes according to their advertised bandwidths, with the attacker controlling the $f \cdot N$ nodes with greatest bandwidth. BestNodes is particularly successful against the default Tor algorithm. When the attacker controls the top 10% of relays, he is able to compromise 54.7% of anonymous paths. The Snader-Borisov ("SB") algorithm fares better for low values of s. However, the strategy becomes vulnerable

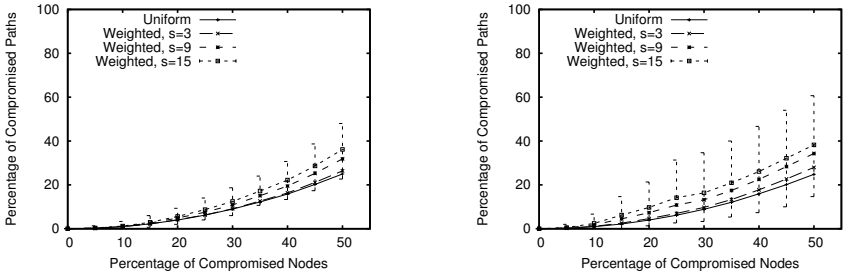

(a) `King` *(b)* `S3-BW`

Fig. 7. The percentage of compromised paths as a function of the fraction of compromised nodes when the attacker uses the `BestLinks` strategy and the initiator uses WEIGHTED with the *(a)* `King` and *(b)* `S3-BW` datasets. Points represent the mean value with error bars (for $s = 15$) indicating the 5th and 95th percentiles. Error bars are omitted for $s \neq 15$ for readability. In all cases, the 5th-95th percentile ranges for $s \neq 15$ were less than that for $s = 15$.

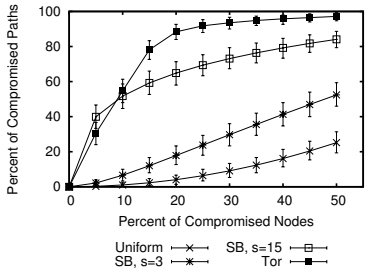

Fig. 8. The percentage of compromised paths when the attacker uses `BestNodes` and the initiator uses node-based strategies with the `Tor-BW` dataset

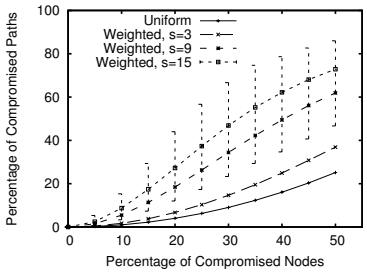

Fig. 9. The percentage of compromised paths when the attacker uses the `Confirmation` strategy and the initiator uses WEIGHTED with the `King` dataset

when performance is more highly valued. An adversary who operates the top 30% of high bandwidth nodes controls 73.1% of paths when $s = 15$.

`MedianDist`: *Compromising Nodes with Shortest Median Distances.* Alternatively, the attacker may choose the $f \cdot N$ nodes that have the smallest median distance between itself and all other nodes. Intuitively, `MedianDist` locates relays that are likely to be chosen due to their proximity to other relays. Fig. 10 plots the effectiveness of such a strategy when used with the `King` dataset. When weighted most heavily in favor of performance ($s = 15$), only 13.1% of paths are compromised when the attacker controls 30% of the network. Results for the remaining link-based topologies are consistent with `King` and are omitted for brevity. Although `MedianDist` is more effective than `BestLinks`, link-based relay selection significantly limits the ability to compromise paths, even against our powerful attacker.

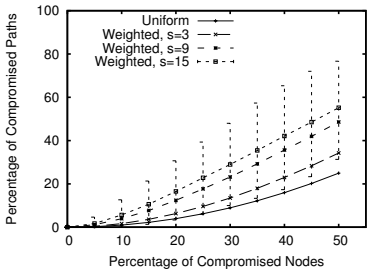

Fig. 10. The percentage of compromised paths as a function of the fraction of compromised nodes when the attacker uses the `MedianDist` strategy and the initiator uses WEIGHTED with the `King` dataset

`Confirmation`: *Determining whether Alice is Communicating with Bob.* The previous attacks attempt to compromise arbitrary paths in the anonymous network. In contrast, an attacker may apply the `Confirmation` attack to test whether a fixed pair of nodes (Alice and Bob) is anonymously communicating. Here, the attacker compromises the nearest node (e.g., having smallest RTT) to Alice that has not yet been compromised, and does the same with respect to Bob, and continues compromising nodes in this manner until he has controls $f \cdot N$ nodes. That is, the attacker compromises the nodes that are nearest to Alice and Bob to maximize the probability that he controls the first and last relays in their anonymous path (assuming such a path exists).

The results of using the `Confirmation` strategy against the `King` dataset are shown in Fig. 9. The figure plots the results of experiments between all pairwise initiators and responders. In each experiment, the attacker compromises the $f \cdot N$ nodes in the manner described above to target the particular initiator and responder pair. When routes are weighted heavily in favor of performance ($s = 9$), an attacker who controls 30% of the network and who can target particular initiator and responder pairs, can discern 34.4% of anonymous paths. As discussed in Appendix A, a slightly modified WEIGHTED strategy better protects against the `Confirmation` attack at the cost of a small degree of performance.

`Relay-in-the-Middle`: *Deducing Communication Endpoints.* If the adversary controls the middle relay (R2) in a three-relay anonymous path, she trivially knows the first (R1) and last (R3) relays as well. Since WEIGHTED ranks candidate paths based on e2e path estimates (i.e., the cost of Alice→R1→R2→R3→Bob), the attacker can estimate the cost of α →R1→R2→R3→ β for all possible initiator and responder pairs $\alpha, \beta \in N \setminus \{R1, R2, R3\}$, $\alpha \neq \beta$. By applying Eq. 1, she can compute the probability that a given candidate initiator/responder pair selected the subsequence R1→R2→R3 in its anonymous path. Although the size of the anonymity network and the performance parameter s may reduce the practicality and usefulness of this attack, we describe a countermeasure in Appendix A.

`Cluster`: *Joining the Network with a Cluster of Nodes.* An attacker may attempt to attract anonymous paths by joining the network with a large cluster of nodes that share a local network, offering low latency and high bandwidth connections between malicious peers. The efficacy of `Cluster` is described in Appendix D.

6 Practical Link-Based Path Selection

In this section, we explore the practical considerations of scalably deploying link-based anonymous path selection over the Internet.

6.1 Link Cost Estimation

Our analysis in Section 5 assumes that the initiator has knowledge of pairwise distances between potential relays. In practice, maintaining pairwise distances will require $O(N^2)$ in communication and network state, hence imposing a significant overhead on the anonymity network.

One practical solution to the above challenge is via the use of *network coordinate systems* that enable the pairwise distances between all participating nodes to be estimated to high accuracy with low overhead. Network coordinate systems, such as Vivaldi [5], PIC [4], NPS [23], and Big Bang Simulation [31] map each relay to n-dimensional coordinates such that the Euclidean distance between two relays' coordinates corresponds to the actual network distance between the pair. Although their individual implementations differ, coordinate systems use distributed algorithms in which each participant periodically measures the distance between itself and a randomly selected peer. By comparing the empirical measurement with the Euclidean distance between the two nodes' coordinates, the relay can adjust its coordinate either towards (in the case of over-estimation) or away from (for under-estimation) the neighbor's coordinate.

Network coordinate systems are well-suited for link-based relay selection, effectively linearizing the quantity of information that must be stored and communicated. By downloading the coordinates of N relays, an initiator can estimate the pairwise distances between them. These systems are lightweight, requiring little bandwidth overhead, and adapt quickly to changes in the network [5]. Additionally, these systems have proved to operate efficiently at Internet scale. For example, the Vuze BitTorrent client [36] currently operates a coordinate system consisting of more than one million nodes [16]. Finally, as we describe below, there exist well-established techniques for securing these systems to ensure the accuracy of advertised coordinates, preventing misbehaving relays from falsifying their coordinates to attract traffic.

Performance Impact of Coordinate Systems. To quantify the accuracy of coordinate systems, a well established metric is the *median error ratio* of each node – the median of the percentage differences between the estimated and actual distances between itself and all other relays in the network. Note that these errors are due to the presence of network triangle inequality violations (TIVs) that cannot be accurately modeled using Euclidean geometry.

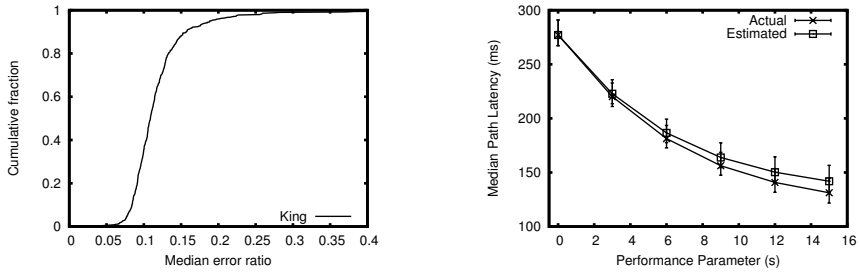

Fig. 11. *Left:* CDF of median error ratios for the `King` dataset. *Right:* Path performance for the `King` dataset using actual network distances ("Actual") and coordinate-based distance estimations ("Estimated"). Points denote median values with errorbars representing the standard deviation.

Fig. 11*(left)* plots the CDF of the median error ratios of all relays after stabilization for the `King` dataset. Coordinates were calculated using Vivaldi [5] with a 5 dimensional coordinate system. The median of the relays' median error ratios (the median error in link estimation) is just 10.9% (6.1ms). The use of coordinate systems for non-latency metrics is considered later in this section.

Using the `King` dataset, Fig. 11*(right)* shows the resulting impact these estimation errors have on the actual e2e performance of anonymous paths. The figure compares the e2e performance obtained using actual distances ("Actual") against the performance that results from using coordinate-based estimations ("Estimated"). The use of the coordinate system to estimate distances produces paths with low-latency. For example, when $s = 15$, the median e2e path latency is 131.2ms using actual network distances; the use of virtual coordinates incurs a modest 8% increase in latency, resulting in paths with a median e2e latency of 141.9ms (still far below the 277.1ms median obtained by UNIFORM).

In addition to performance, the use of coordinate systems has implications to anonymity. We investigate the relationship between coordinate-based link estimation and anonymity in Appendix C. To summarize the results, the use of coordinate systems does not decrease anonymity relative to using actual distances. For example, an attacker who controls 30% of the network is able to compromise 29.0% of anonymous paths when he uses the `MedianDist` attack and the initiator uses actual distances with $s = 15$. Using coordinate-based distance estimations and keeping all other parameters fixed, the attack compromises 28.5% of paths.

6.2 Other Practical Considerations

We briefly outline other practical considerations of link-based relay selection.

Securing coordinate systems. The distributed nature of coordinate systems make them vulnerable to manipulation if not properly defended. Malicious relays may advertise false coordinates or delay measurement probes, either to make themselves appear more favorable or to cause disorder in the system. Fortunately, practical techniques exist that mitigate such attacks. For example, the Veracity

system protects the accuracy of coordinate systems when up to 40% of the network is malicious [34,32]. Given the large number of available coordinate protection schemes [4,30,13,41,34,32], we consider the challenge of securing coordinate systems to be orthogonal and out-of-scope of this paper.

Pairwise bandwidth estimation. Coordinate systems have been known to estimate with high accuracy (i.e. low error ratios) link metrics that tend to be additive in nature when used to compute metrics across multiple links. Examples of metrics that work well include latency and AS traversal. However, these systems have been shown to be inaccurate at estimating pairwise bandwidth between any two nodes, due to the high incidence of TIVs in bandwidth measurements.

However, we note that there have been a number of recent promising proposals that enable one to estimate pairwise bandwidth accurately and scalably. For example, there have been several attempts to identify links that cause severe network TIVs [37,18,17], enabling initiators to avoid them when forming paths. Separate work [26,27] has directly addressed the problem of bandwidth embeddings, introducing techniques for embedding bandwidth distances in tree structures. Their results show that pairwise PlanetLab bandwidths can be embedded with a median error ratio of approximately 0.25 [26]. Finally, as a third alternative, rather than rely on coordinate embedding systems, initiators can anonymously query network measurement services such as IDMaps [9] or iPlane [19] to estimate the bandwidth of network links.

Locating the Responder. To estimate e2e path performance, the initiator must predict the distance between the exit relay and the responder. The initiator cannot estimate the cost of this final hop if the responder does not participate in the coordinate system. Instead, the initiator can locate the closest relay to the responder using publicly available network information services. For example, OASIS [10], ClosestNode [1], and iPlane [19] all provide interfaces for resolving the closest server to any given IP address. The initiator can anonymously query such services to locate the relay that is nearest to the responder. The closest relay can then proxy requests between the exit relay and the responder.

Alternatively, initiators can disregard the link between the exit relay and the responder when selecting anonymous paths. As discussed in Appendix A, such an approach incurs only a modest decrease in performance.

7 Conclusion

This paper makes the case for link-based relay selection for flexibly tuning the performance and anonymity properties of anonymous paths. In comparison to node-based techniques in which performance may be quantified only in terms of node properties (i.e., bandwidth), link-based selection enables the generation of high performance paths across multiple metrics: latency, jitter, loss, and bandwidth. Using realistic network traces, we validate that our link-based WEIGHTED strategy reduced by 71% the number of paths with end-to-end latencies greater than 250ms (in comparison to selecting relays uniformly at random), and doubled the median available bandwidths of anonymous paths.

We also show that link-based relay selection is also significantly more resilient to manipulation than traditional node-based techniques. For example, when applying the default Tor path selection algorithm to a subset of bandwidth data obtained from the Tor network, an adversary who controls the top 30% of highest bandwidth relays is able to compromise 93.5% of anonymous paths. In comparison, using WEIGHTED on a network trace in which bandwidth is measured as a link characteristic, an attacker who controls the same percentage of anonymizing relays compromises less than a third of anonymous paths.

Acknowledgments

The authors are grateful to Roger Dingledine for his many helpful suggestions. We would like to thank the anonymous reviewers for their insightful feedback, and Steven Murdoch who shepherded the paper. This work is partially supported by NSF Grants CNS-0831376, CNS-0524047, CNS-0627579, and CNS-0721845 and ONR MURI N00014-07-1-0907.

References

1. ClosestNode.com, `http://www.closestnode.com/`
2. Akella, A., Seshan, S., Shaikh, A.: An Empirical Evaluation of Wide-area Internet Bottlenecks. In: Conference on Internet Measurement (IMC) (2003)
3. Bauer, K., McCoy, D., Grunwald, D., Kohno, T., Sicker, D.: Low-Resource Routing Attacks against Tor. In: Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society, pp. 11–20 (2007)
4. Costa, M., Castro, M., Rowstron, R., Key, P.: PIC: Practical Internet Coordinates for Distance Estimation. In: International Conference on Distributed Computing Systems (2004)
5. Dabek, F., Cox, R., Kaashoek, F., Morris, R.: Vivaldi: a Decentralized Network Coordinate System. SIGCOMM Comput. Commun. Rev. 34(4), 15–26 (2004)
6. Dingledine, R., Mathewson, N.: Tor Path Specification (January 2008), `http://www.torproject.org/svn/trunk/doc/spec/path-spec.txt`
7. Dingledine, R., Mathewson, N., Syverson, P.: Tor: The Second-Generation Onion Router. In: Proc. of the 13th USENIX Security Symposium, pp. 303–320 (2004)
8. Feamster, N., Dingledine, R.: Location Diversity in Anonymity Networks. In: WPES 20: Proceedings of the 2004 ACM workshop on Privacy in the electronic society, pp. 66–76 (2004)
9. Francis, P., Jamin, S., Jin, C., Jin, Y., Raz, D., Shavitt, Y., Zhang, L.: IDMaps: A Global Internet Host Distance Estimation Service. IEEE/ACM Trans. Netw. 9(5), 525–540 (2001)
10. Freedman, M.J., Lakshminarayanan, K., Mazières, D.: OASIS: Anycast for Any Service. In: Proc. 3rd USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI 2006) (2006)
11. Gini, C.: Measurement of Inequality of Incomes. The Economic Journal 31(121), 124–126 (1921)
12. Gummadi, K.P., Saroiu, S., Gribble, S.D.: King: Estimating Latency Between Arbitrary Internet End Hosts. In: ACM SIGCOMM Workshop on Internet Measurment (IMW) (2002)

13. Kaafar, M.A., Mathy, L., Barakat, C., Salamatian, K., Turletti, T., Dabbous, W.: Securing Internet Coordinate Embedding Systems. In: ACM SIGCOMM (August 2007)
14. "king" data set, http://pdos.csail.mit.edu/p2psim/kingdata/
15. Lakshminarayanan, K., Padmanabhan, V.N.: Some Findings on the Network Performance of Broadband Hosts. In: IMC 2003: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement, pp. 45–50 (2003)
16. Ledlie, J.T.: A Locality-Aware Approach to Distributed Systems. PhD thesis, Harvard University (September 2007)
17. Lee, S., Zhang, Z.-L., Sahu, S., Saha, D.: On Suitability of Euclidean Embedding of Internet Hosts. In: SIGMETRICS 2006/Performance 2006: Proceedings of the joint international conference on Measurement and modeling of computer systems, pp. 157–168 (2006)
18. Lumezanu, C., Levin, D., Spring, N.: PeerWise Discovery and Negotiation of Shorter Paths. In: Workshop on Hot Topics in Networks (HotNets) (2007)
19. Madhyastha, H.V., Isdal, T., Piatek, M., Dixon, C., Anderson, T., Krishnamurthy, A., Venkataramani, A.: IPlane: An Information Plane for Distributed Services. In: Symposium on Operating Systems Design and Implementation (OSDI 2006) (2006)
20. Muller, M.E.: A Note on a Method for Generating Points Uniformly on N-Dimensional Spheres. Communications of the ACM 2(4), 19–20 (1959)
21. Murdoch, S.J.: Hot or Not: Revealing Hidden Services by Their Clock Skew. In: CCS 2006: Proceedings of the 13th ACM Conference on Computer and Communications Security, pp. 27–36 (2006)
22. Murdoch, S.J., Watson, R.N.M.: Metrics for Security and Performance in Low-Latency Anonymity Systems. In: Borisov, N., Goldberg, I. (eds.) PETS 2008. LNCS, vol. 5134, pp. 115–132. Springer, Heidelberg (2008)
23. Ng, T.S.E., Zhang, H.: A Network Positioning System for the Internet. In: Proceedings of the 2004 USENIX Annual Technical Conference (June 2004)
24. Øverlier, L., Syverson, P.: Locating Hidden Servers. In: IEEE Symposium on Security and Privacy (2006)
25. PlanetLab, http://www.planet-lab.org
26. Ramasubramanian, V., Malkhi, D., Kuhn, F., Abraham, I., Balakrishnan, M., Gupta, A., Akella, A.: A Unified Network Coordinate System for Bandwidth and Latency. Technical Report MSR-TR-2008-124, Microsoft Research (Sept. 2008)
27. Ramasubramanian, V., Malkhi, D., Kuhn, F., Balakrishnan, M., Gupta, A., Akella, A.: On the Treeness of Internet Latency and Bandwidth. In: SIGMETRICS/Performance (June 2009)
28. Reiter, M.K., Rubin, A.D.: Crowds: Anonymity for Web Transactions. In: ACM Transactions on Information and System Security (1998)
29. Ribeiro, V., Riedi, R., Baraniuk, R., Navratil, J., Cottrell, L.: pathChirp: Efficient Available Bandwidth Estimation for Network Paths. In: Passive and Active Measurement Workshop (2003)
30. Saucez, D., Donnet, B., Bonaventure, O.: A Reputation-Based Approach for Securing Vivaldi Embedding System. In: Dependable and Adaptable Networks and Services (2007)
31. Shavitt, Y., Tankel, T.: Big-bang Simulation for Embedding Network Distances in Euclidean Space. In: IEEE Infocom (April 2003)
32. Sherr, M., Blaze, M., Loo, B.T.: Veracity: Practical Secure Network Coordinates via Vote-based Agreements. In: USENIX Annual Technical Conference (USENIX 2009) (June 2009)

33. Sherr, M., Loo, B.T., Blaze, M.: Towards Application-Aware Anonymous Routing. In: USENIX Workshop on Hot Topics in Security (HotSec) (August 2007)
34. Sherr, M., Loo, B.T., Blaze, M.: Veracity: A Fully Decentralized Service for Securing Network Coordinate Systems. In: IPTPS (February 2008)
35. Snader, R., Borisov, N.: A Tune-up for Tor: Improving Security and Performance in the Tor Network. In: 15th Annual Network and Distributed System Security Symposium (NDSS) (February 2008)
36. Vuze bittorrent client, `http://azureus.sourceforge.net/`
37. Wang, G., Zhang, B., Ng, T.S.E.: Towards Network Triangle Inequality Violation Aware Distributed Systems. In: ACM SIGCOMM Conference on Internet Measurement (IMC 2007), pp. 175–188 (2007)
38. Wong, B., Slivkins, A., Sirer, E.G.: Meridian: a Lightweight Network Location Service without Virtual Coordinates. In: SIGCOMM (2005)
39. Wright, M., Adler, M., Levine, B., Shields, C.: The Predecessor Attack: An Analysis of a Threat to Anonymous Communications Systems. ACM Transactions on Information and System Security (TISSEC) 4(7), 489–522 (2004)
40. Yalagandula, P., Sharma, P., Banerjee, S., Basu, S., Lee, S.: S^3: A scalable Sensing Service for Monitoring Large Networked Systems. In: SIGCOMM Internet Network Management Workshop (2006)
41. Zage, D.J., Nita-Rotaru, C.: On the Accuracy of Decentralized Virtual Coordinate Systems in Adversarial Networks. In: CCS (2007)

Appendix A: A Revised WEIGHTED Algorithm

The WEIGHTED algorithm introduced in Section 4 ranks candidate paths by the expected e2e path cost. Unlike node-based relay selection strategies, the e2e path cost includes the links from the initiator to the first relay and from the last relay to the responder, potentially leaking information about the communication participants.

To prevent the `Relay-in-the-Middle` attack (see Section 5.2), an alternative strategy is to exclude the first (from the initiator) and last (to the responder) links when ranking paths. That is, the initiator ranks paths by the cost of the subsequence R1→R2→R3, where R1, R2, and R3 are the relays in a candidate path.

The revised WEIGHTED strategy has two advantages. First, it disassociates the communication endpoints from path selection. An adversary who knows the identities of R1, R2, and R3 cannot infer any information about the initiator and responder. Second, it does not require the responder to participate in the coordinate system (see Section 6.2) since the distance from the exit relay to the responder does not influence router selection.

The obvious cost of using the revised WEIGHTED strategy is that the first and last hops may be expensive, incurring poor performance even though the subsequence R1→R2→R3 may be efficient.

Our experimental evaluation indicates that the performance penalty due to the revised WEIGHTED strategy is minimal. Fig. 12 shows the performance of the vanilla and revised WEIGHTED strategies with $s = 9$. For comparison, the performance achieved using UNIFORM is also plotted. Although the unmodified

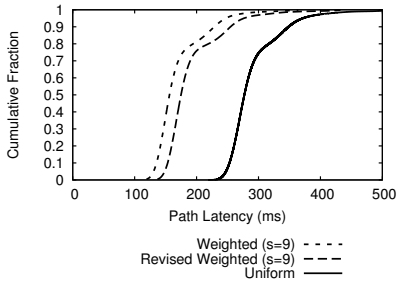

Fig. 12. The cumulative distribution of e2e path latencies when using variants of the WEIGHTED strategy on the `King` dataset. The performance achieved using `Uniform` is provided for comparison.

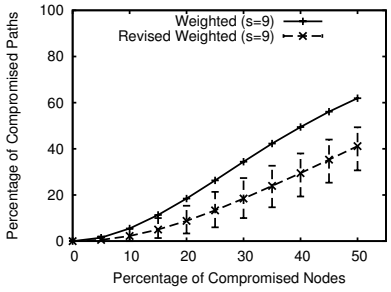

Fig. 13. The percentage of compromised paths as a function of the fraction of compromised nodes when the attacker users the `Confirmation` strategy and the initiator uses variants of WEIGHTED with the `King` dataset

WEIGHTED algorithm achieves the lowest median e2e path latency (156.3ms), the modified version also achieves significantly lower latencies (174.9ms) than UNIFORM (277.2ms).

Fig. 13 compares the resilience to the `Confirmation` attack using the vanilla and modified versions of WEIGHTED. Since the positions of the initiator and responder do not influence relay selection when revised WEIGHTED is used, the attacker's strategy is less effective. For example, when 30% of nodes are malicious, the attacker compromises 34.4% of paths when the initiator uses the unmodified WEIGHTED technique and only 18.5% against revised WEIGHTED.

Appendix B: Preventing the Predecessor Attack

An anonymized connection between an initiator and responder is often reset due to node churn, requiring it to be reconstructed using different relays [39]. The adversary can conduct a *predecessor attack* to discover the initiator by counting the number of times each relay precedes the attackers' relays in the anonymous path [28,39]. Since the initiator is always present in such circuits, it will have a higher count than the relays that are chosen randomly whenever the circuit is rebuilt.

Tor mitigates the predecessor attack by using a small number of fixed entry nodes called *guards* [6]. Link-based path selection is equally vulnerable to the predecessor attack, but may also be defended using guards. Guards must be chosen carefully since their locations affect the performance of a path. However, as described in Appendix A (see, in particular, Fig. 12), link-based routing produces high performance paths even if the first hop (connecting the initiator to the guard node) is not considered by the path selection algorithm. Link-based

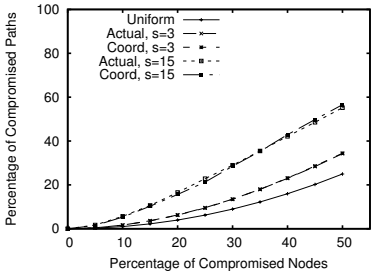

Fig. 14. The percentage of compromised paths on the `King` dataset when the attacker uses the `MedianDist` strategy. "Actual" denotes results obtained using actual network distances. "Coord" reflects performance when initiators estimate distances using the coordinate embedding system.

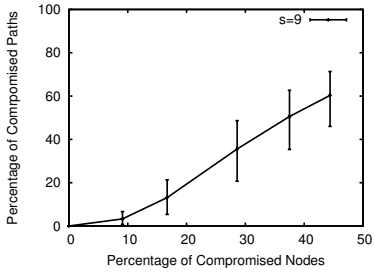

Fig. 15. The percentage of compromised paths as a function of the fraction of compromised nodes when the attacker uses the `Cluster` strategy and the initiator uses WEIGHTED with the `King` dataset using $s = 9$

routing may therefore adopt the same mitigation strategy as Tor [7]; namely, the initiator selects a relay (having a long uptime) to act as its entry guard for all anonymous paths.

Appendix C: The Impact of Coordinate Systems on Anonymity

Coordinate systems linearize pairwise distances by mapping each node to n-dimensional coordinates. Due to network triangle inequality violations that cannot be represented in Euclidean space, coordinate systems do not perfectly predict distances. As shown in Fig. 11*(right)*, the use of coordinate systems imposes a modest decrease in path performance. In this Appendix, we consider the effects of using coordinate systems on anonymity.

Fig. 14 shows the percentage of compromised paths using the `King` dataset when the attacker applies the `MedianDist` strategy. The figure compares performance results obtained using network distances ("Actual") to estimations based on the coordinate system ("Coord"). As can be observed from the Fig., the effectiveness of the attack does not substantially differ when actual and coordinate distances are used. For example, when $s = 15$ and actual distances are used, an attacker who controls 30% of the network can compromise 29.0% of paths. In comparison, the same attacker can compromise 28.5% of paths when virtual coordinates are used in place of actual distances.

Appendix D: The `Cluster` Attack

An attacker may attempt to compromise a large fraction of anonymous paths by joining the anonymity network using multiple nodes from the same LAN. Due to

the high bandwidths and low latencies within the LAN, paths composed entirely of malicious nodes from the LAN will have low e2e cost estimates and will be favored by the WEIGHTED algorithm.

Since our experimental datasets do not contain large clusters of similarly located nodes, it was necessary to adapt the attacker model to permit the attacker to insert nodes. To determine the location for the new nodes, we first use the Vivaldi [5] virtual embedding system to assign n-dimensional coordinates to each node in the existing topology such that the Cartesian distance between two nodes' coordinates corresponds to the network distance (e.g., latency) between them. To provide the attacker with a desirable location in the topology, we assign each malicious node a coordinate that is at most 5ms from the centroid of the network. Hence, any two malicious nodes are separated by at most 10ms. Locations from the centroid are randomly chosen according to Muller's uniform hypersphere point generation technique [20]. Network distances between a malicious node and another peer are estimated using the Cartesian distance between the nodes' coordinates.

Fig. 15 illustrates the efficacy of the `Cluster` attack when the initiator uses the WEIGHTED algorithm with $s = 9$ on the `King` dataset. When the attacker controls 28.6% of the network (i.e., he adds 200 nodes to the existing 500 node topology), he compromises just 35.6% of anonymous paths.

It is worth noting that the `Cluster` attack may be further mitigated by requiring that adjacent nodes in anonymous paths reside in separate autonomous systems or have a minimum latency between them.