The Mis-entropists: New Approaches to Measures in Tor

Angèle M. Hamel
Dept. of Physics and Computer Science, Wilfrid Laurier University Waterloo, ON Canada
ahamel@wlu.ca

Jean-Charles Grégoire
INRS Montréal, QC Canada
gregoire@emt.inrs.ca

Ian Goldberg
Cheriton School of Computer Science, University of Waterloo Waterloo, ON Canada
iang@cs.uwaterloo.ca

ABSTRACT

Measuring path compromise in the anonymity system Tor is a problem of important and immediate interest. We discuss the drawbacks to the traditional anonymity network measure—entropy—and explore the advantages of an attack-based measure. Our perspective is that of an adversary with a bandwidth budget (a fixed amount of bandwidth he or she can compromise), and we derive theoretical and numerical results that illustrate path compromise under these conditions.

1. INTRODUCTION

> We need a better understanding of how much anonymity the Tor network provides against a partial network adversary who observes and/or operates some of the network. ... The simplistic metric also gets misused in academic research papers to produce misleading results like 'I can sign up 2% of the Tor relays and capture 50% of the paths,' when what they mean is 'I can provide half the capacity in the Tor network and capture half the paths.' ... The goal here isn't to come up with the one true metric for summarizing Tor's safety. Rather, we need to recognize that there are many threat models to consider at once, so we need many different views into what types of safety the network can offer. [6]

This Tor Blog entry highlights both the need for effective measures of the risks of using Tor in certain circumstances, and the difficulty of providing such measures. It also calls for different views on safety. Here we provide one such different view based on the capability of the adversary and the resulting probability of path compromise.

Messages sent through Tor are typically routed through a path consisting of three nodes: entry, middle, and exit. Of particular interest to the user is, how vulnerable is my path to compromise? That is, how likely is it that a nonglobal passive adversary can tell which services I am accessing? Path compromise is generally acknowledged to be a problem chiefly when the adversary controls the entry and exit nodes. Thus the question, "how likely is it that my path is compromised?" can be interpreted as, how likely is it that the adversary controls the entry and exit nodes on my path? This approach has already been advocated by various authors [11, 12].

Here we use the concept of "adversarial bandwidth fraction" or ABWF. In this model, the adversary has a bandwidth budget; that is, the adversary can afford to compromise a fixed percentage of bandwidth, say $n\%$. Of importance, too, is the choice of nodes made by the adversary within that budget. This bandwidth budget idea is related to the approach in Murdoch and Watson [11] in which they consider path construction where an adversary can control not just a given number of nodes, but also nodes of a given total bandwidth capacity. Assuming the adversary optimizes their bandwidth budget, i.e. he or she chooses the set of nodes that allows them the highest probability of compromising paths, what is the level of danger inherent in the system?

There is already a standard measure for anonymity: entropy. Why then pose this question at all? As we will discuss below, entropy has a number of shortcomings. In particular, it is based on holistic properties of the system, rather than the actions of the attacker. We have coined the term mis-entropist, as a play on *misanthropist*, to be one who distrusts entropy. It is time for a fresh look, to strip down to the bare wood, to question assumptions, and to see what properties we can obtain. As the Tor Blog reminds us, the perfect measure still eludes us.

Our contributions: We discuss entropy, summarizing a number of issues that have been advanced by other authors, and we propose that an attack-based measure would be superior. To this end, we calculate the probability of path compromise given various bandwidth budgets available to the adversary, and propose ways to use this information.

The paper is structured as follows: Section 2 looks at the classical measure of entropy and discusses some of its drawbacks. Section 3 explores the path selection algorithm for Tor and the primary ingredients to it: nodes and bandwidth. Section 4 discusses several different ways of calculating path probability depending on how bandwidth is taken into account and to what extent. It also explores the probability of picking a compromised pair of nodes, given an "adversarial bandwidth fraction." Section 5 delivers our results. Sections 6 and 7 round out the paper and address related work and conclusions.

2. ENTROPY

The classical measure for anonymity is (Shannon) entropy. It is defined as $-\sum_{i=1}^{K} P_i \log_2 P_i$, where there are K sep-

arate events and P_i is the probability of the ith event occurring. Entropy was proposed as a measure for anonymity networks by Serjantov and Danezis [14] and Diaz *et al* . [5], both at PET 2002. Although it was introduced to measure the anonymity in a mix network, it has been co-opted for use in measuring uncertainty in other aspects of anonymity networks—in particular, the randomness in path selection in Tor [1, 12]. Entropy has a lot of good statistical properties and is commonly used in a number of disciplines. Furthermore, it is relatively easy to compute.

While this measure has been widely used, there are some drawbacks to it. Exact measures of probability can be challenging. Further, estimates of probability are just that: estimates, and the entropy formula is only as good as the probabilities it uses.

As pointed out by Toth *et al.* [20], the entropy measure is a measure of the entire system (global anonymity) rather than a measure of the anonymity for a particular user (local anonymity). Shmatikov and Wang [16] advocate min-entropy instead of Shannon entropy, pointing out

> [Shannon] Entropy does not always capture the right anonymity property. Consider a distribution of 100 potential destinations, in which all but one are equally likely with probability 0.009, and a single destination has probability 0.109. [Shannon] Entropy of this distribution is 6.40, close to the theoretical maximum of 6.64.

Clauß and Schiffner [4] also make the same point that the influence of outliers is strong, i.e. if one element is very likely the entropy can be very high, even though the elements are not evenly distributed. Syverson [18] takes issue with the entire concept of indistinguishability as the basis for a measure. He introduces the idea of trust in anonymity network and advocates for security measures that reflect the difficulty an adversary has in overcoming them, and measures that do not depend on the value of variables that are difficult to predict or determine.

Furthermore, entropy is good for comparative measures—"Mine is bigger than yours!"—but it is less valuable in absolute terms. What to make of entropy values? Consider three cases, each involving about 10000 alternatives, but spread in three different ways: uniform, where all cases are equally likely; chosen few, where a handful are 100 times more likely than others, and exponential, where frequency of occurrence is in inverse proportion with the likelihood. We can calculate the entropy in each of these sitations as:

uniform: 13.2877
chosen few: 12.8194
exponential: 7.82193

Those numbers are quite distinct, yet essentially meaningless. The more heterogeneous the model, the lower the number, yet the impact of a compromised node depends strongly on the probability of its being chosen in a path, which entropy—in its multiple variants—does not reflect.

The authors mentioned above have taken issue with entropy as a measure of anonymity. But it also fails to measure the effect of an attack. It makes the assumption that more diversity/inequality/disorder/randomness in paths leads to better safety. But, as Murdoch and Watson [11] suggest, any measure based on uniformity of path selection misses key information:

> 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.

Thus we take the view that the attacker and his or her capabilities are key—particularly with reference to bandwidth—and derive results that indicate the direct effects of attack.

To summarize, entropy measures diversity and answers questions such as "how disordered is the system?" However, when the focus is on the attacker, and on compromised paths, the question to be answered is more properly "what is the danger a path is compromised?" and entropy does not speak to such questions.

3. PATH SELECTION

3.1 Nodes

The nodes in a Tor path, and their capabilities and limitations, are of fundamental interest. Exit nodes are special in that they connect directly to a server (e.g. web server) and thus may have exit policies restricting the types of traffic exiting from them. If traffic requires a certain port, the user will have to pick an exit node which supports the service they require.

The entry node position is occupied by a type of node called a guard node. Guard nodes were first proposed by Wright *et al.* [22] (where they are called "helper nodes"). As stated above, the problem occurs when an adversary controls both the entry and exit nodes. Since in Tor the path is changed every 10 minutes, a user connected for n minutes is exposed to $\lceil n/10 \rceil$ different paths and thus $\lceil n/10 \rceil$ possibilities for compromise. In general the issue is not so much the amount of traffic observed by the adversary, but the fact the adversary has observed any traffic at all. It is generally considered better to pick a few—the usual recommendation is three—possibilities for entry nodes and stick with those. These possibilities are called *guard nodes.*

Our analysis considers the situation before a user selects guard nodes. If all the guard nodes turn out to be honest, then there is no possibility of path compromise. The user does not know this *a priori*, of course.

According to the Tor specification [19], we thus have four kinds of nodes to select from in forming paths: guard nodes, exit nodes, guard+exit nodes (suitable for either the guard or the exit position), and non-flagged (i.e. middle) nodes. In our analysis we will treat these nodes differently.

Note, however, that all four kinds of nodes can occupy the middle position. Allowing such an occurrence can seem a bit wasteful of scarce resources, but does permit more randomness in path construction. However, Tor does impose a penalty in the form of weights. Depending on the position and scarcity of types of nodes they will be weighted by an additional factor during path selection. This weighting algorithm will be explained in more detail below.

In addition to being designated as guard, exit, or guard+exit nodes, nodes can also be *fast* or *stable* or both; see Figure 1. The Tor Directory Protocol [19] defines a node as fast "if it is active, and its bandwidth is either in the top 7/8ths for known active routers or at least 20 KB/s." It defines a node as stable "if it is active, and either its Weighted MTBF [Mean Time Between Failure] is at least the median for known active routers or its Weighted MTBF corresponds to at least 7 days. To calculate weighted MTBF, compute the weighted mean of the lengths of all intervals when the router was observed to be up, weighting intervals by α^n, where n is the amount of time that has passed since the interval ended, and α is chosen so that measurements over approximately one month old no longer influence the weighted MTBF much."

A fast path contains only fast nodes; a stable path contains only stable nodes.

3.2 Bandwidth

The bandwidth available is the lifeblood of the Tor network. More bandwidth means better efficiency and a greater number of users, and may contribute to a reduction in latency. However, there is a complex, and not completely understood, relationship between bandwidth and anonymity. Does more bandwidth in the system mean greater anonymity? Can bandwidth in the wrong places actually reduce anonymity (e.g. by routing too much traffic through a handful of nodes)? What influence does the bandwidth at individual nodes have on anonymity?

Bandwidth is one factor in the path selection algorithm in Tor. Initially, this was self-reported bandwidth; however, recognizing that malicious nodes can lie, this was changed to a consensus bandwidth [13]. The bandwidth used in the path selection calculation is capped at 10 MB/s to prevent too much influence from nodes that may be lying. See Figure 2 for an indication of how bandwidth has grown in Tor.

3.3 Path Selection Algorithm

The path selection algorithm first selects an exit node, then the guard node (from a list of guard nodes that the user has already selected), then the middle node. There is an additional list of constraints, such as no two nodes can belong to the same *family*, available in the Tor Path Specification [7].

Murdoch and Watson [11] identify three different random selection algorithms for node selection in Tor, which include catering to weights, as described above:

- 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 bandwidths.
- Adjusted bandwidth-weighted random selection (ABWRS), where weights specified in the Tor Directory Specification [19] are applied to the bandwidths, and nodes are selected with probability proportional to those weighted bandwidths. These weights are discussed below.

The additional weighting factors are multiplied by the bandwidths of the exit nodes and guard nodes. The purpose of the weighting factor is to discourage the use of scarce resources in times of shortage, i.e. if the system is short of exit nodes, it makes it more likely an exit node will be used as an exit than as a middle node. The weights are of the form W_{xy} where x is the position (guard (g), middle (m), exit (e)) of the node, and y is the type of node (guard nodes are denoted by g, exit nodes by e, and guard+exit nodes by d). Let G be the total bandwidth for Guard-flagged nodes, M be the total bandwidth for non-flagged nodes, E be the total bandwidth for Exit-flagged nodes, and D be the total bandwidth for Guard+Exit-flagged nodes. Full details are in the Tor Directory Protocol [19].

The weights are understood to satisfy balancing constraints (e.g. bandwidth of the guard nodes must equal the bandwidth of the middle nodes and the bandwidth of the exit nodes). For example, if the guard and exit nodes are not scarce, we have

$$\begin{aligned} W_{gd} &= 1/3 \\ W_{ed} &= 1/3 \\ W_{md} &= 1/3 \\ W_{ee} &= (E+G+M)/(3*E) \\ W_{me} &= 1 - W_{ee} \\ W_{mg} &= (2*G-E-M)/(3*G) \\ W_{gg} &= 1 - W_{mg} \end{aligned}$$

4. PROBABILITY OF PATH COMPROMISE

Next, given a set of compromised nodes, we calculate the probability of path compromise (i.e. the probability that both the guard and exit node are compromised) under each of the three path selection algorithms described above.

Calculating the probability of picking a compromised pair in a uniform situation, as in SRS, where each node is equally likely to be chosen, is straightforward. Suppose that c pairs of nodes are compromised. Then we can ask the question, what is the probability that a particular pair is one of these c? If the pairs are chosen with uniform probability for the paths then each one has an equal chance.

For example, given two possible entry nodes A_1 and A_2 and three possible exit nodes B_1, B_2, and B_3, then there are six possible pairs. If we know two pairs are compromised, then there is a 1/3 chance of having a compromised pair.

However, if the pairs are not chosen with uniform probability, i.e. if they are chosen with a bandwidth dependent

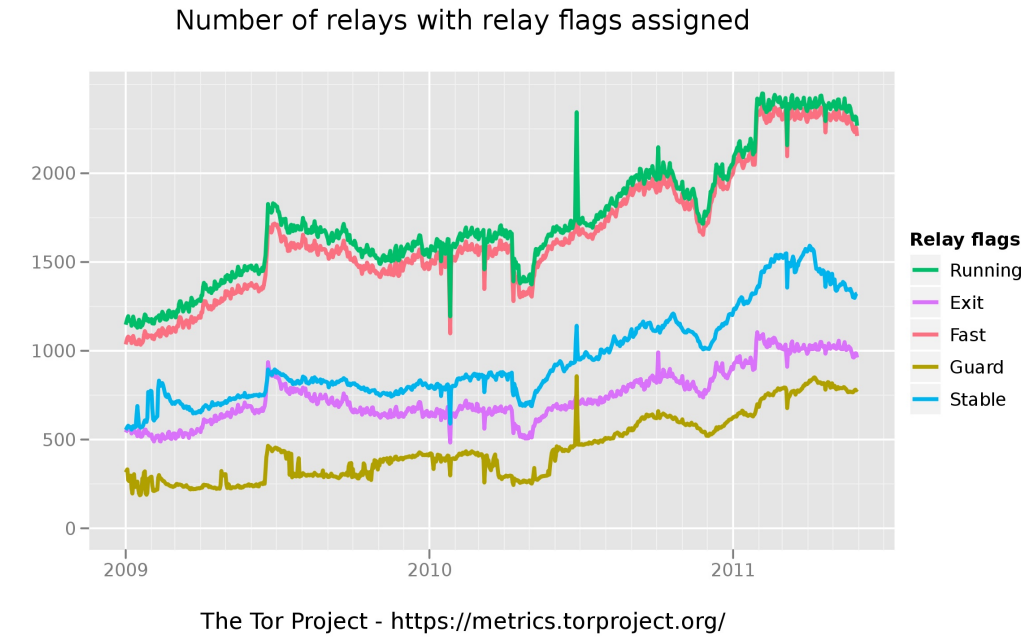

Figure 1: Graph showing growth in number of nodes in Tor from January 1, 2009 to May 31, 2011. As can be seen in the graph, almost all nodes are fast; a solid, but significantly smaller proportion, are stable.

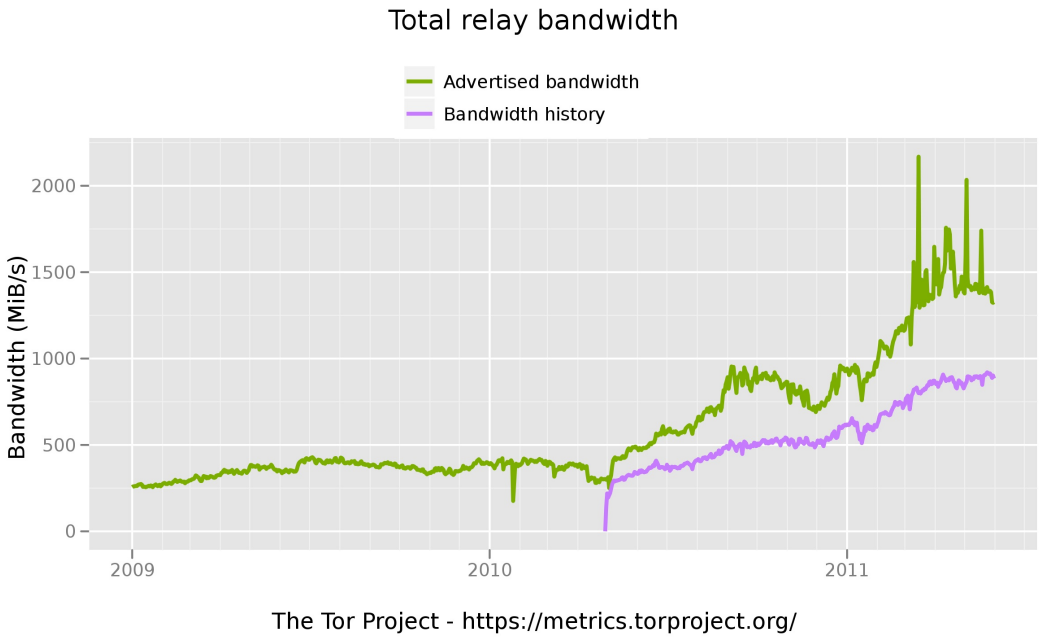

Figure 2: Graph showing the growth of total bandwidth available in Tor from January 1, 2009 to May 31, 2011.

probability as they are in the BWRS and ABWRS cases, then given that c pairs are compromised, it is challenging to calculate the probability of picking a compromised pair.

For example, now suppose A_1 has 9/10 of the bandwidth and A_2 has 1/10. Further suppose that B_1 has 1/2 the bandwidth and B_2 and B_3 each have 1/4. Then the pairs have probabilities 18/40, 2/40, 9/40, 9/40, 1/40, 1/40, and the chance that you pick a compromised pair is no longer 1/3. For example, the pair A_2B_3 is rarely picked, so even if it is compromised, it hardly matters; the issue becomes to calculate this probability.

In the following subsections we derive equations for each of these situations.

4.1 SRS

If we ignore bandwidth then the probability of a specific pair of entry and exit nodes appearing can be calculated using the number of nodes of each type. The user re-uses guard nodes, but this calculation is performed before the user selects them. Define the following:

- g =number of guard nodes
- e =number of exit nodes
- m =number of non-flagged nodes (i.e. only good for middle)
- d = number of guard+exit nodes

Thus the formula for number of paths can easily be calculated as:

THEOREM 4.1. *Let g be the number of guard nodes, e be the number of exit nodes, m be the number of non-flagged nodes, and d be the number of guard+exit nodes. Then the number of possible paths is*

$$(ge + gd + de + d(d-1)) * (g + d + e + m - 2) \quad (1)$$

Proof: There are two choices for the entry node: either it is from the set of guard nodes or from the set of guard+exit nodes. If it is from the set of guard nodes there are g choices. Then we have free choice on the final node to be from the set of exit nodes or from the set of guard+exit nodes. So there are $g(e+d)$ choices for the two nodes. Now the middle could be any node except the two already picked, and there are $g + d + e + m - 2$ choices for the middle node.

If instead the entry node is from the guard+exit nodes there are d choices. Then the final node could be from the exit nodes or from the guard+exit nodes (except the one already chosen). So the number of choices for the two nodes is $d(e + d - 1)$. Again, the middle nodes can be chosen from $g + d + e + m - 2$ choices. Putting it all together we have $(ge + gd + de + d(d-1)) * (g + d + e + m - 2)$ paths. $\diamond$

The following corollary is immediate:

COROLLARY 4.2. *1. The probability that a path will start from a given node from the set of guard nodes is*

$$\frac{(e+d)}{(ge + gd + de + d(d-1))}. \quad (2)$$

2. The probability that a path will start from a given node from the set of guard+exit nodes is $\frac{(e+d-1)}{(ge+gd+de+d(d-1))}$.

3. The probability that a path will end with a given node from the set of exit nodes is $\frac{(g+d)}{(ge+gd+de+d(d-1))}$.

4. The probability that a path will end with a given node from the set of guard+exit nodes is $\frac{(g+d-1)}{(ge+gd+de+d(d-1))}$.

Table 1: General Statistics for Consensus Sample set.

Total number of routers	2328
Total number of "Authority" routers	10
Total number of "Bad directory" routers	0
Total number of "Bad exit" routers	7
Total number of "Exit" routers	997
Total number of "Fast" routers	2273
Total number of "Guard" routers	810
Total number of "Named" routers	1686
Total number of "Stable" routers	1410
Total number of "Running" routers	2328
Total number of "Valid" routers	2328
Total number of "V2Dir" routers	1126
Total number of "Hidden Service" routers	851

Example. We downloaded a consensus of Tor on April 22, 2011 (see summary in Table 1, while the full distribution of bandwidths is shown in Figure 3). From this consensus we determined there were 2328 valid nodes altogether, of which $e = 721$ were pure exit nodes, $g = 541$ were pure guard nodes, $d = 269$ were guard+exit nodes, and $m = 797$ were non-flagged nodes. Then we could calculate the probabilities in the corollary as:

- The probability that a path will start from a given node from the set of guard nodes is $\frac{(990)}{(801631)} = 0.001235$.
- The probability that a path will start from a given node from the set of guard+exit nodes is $\frac{989}{801631} = 0.001234$.
- The probability that a path will end with a given node from the set of exit nodes is $\frac{810}{801631} = 0.001$.
- The probability that a path will end with a given node from the set of guard+exit nodes is $\frac{809}{801631} = 0.001$.

We can also examine this from a path compromise perspective. Recall that a path is compromised if the guard and exit nodes are compromised. Thus we can use the principle of inclusion–exclusion to conclude that the number of possible pairs of entry–exit nodes is

$$(g + d)(e + d) - d \quad (3)$$

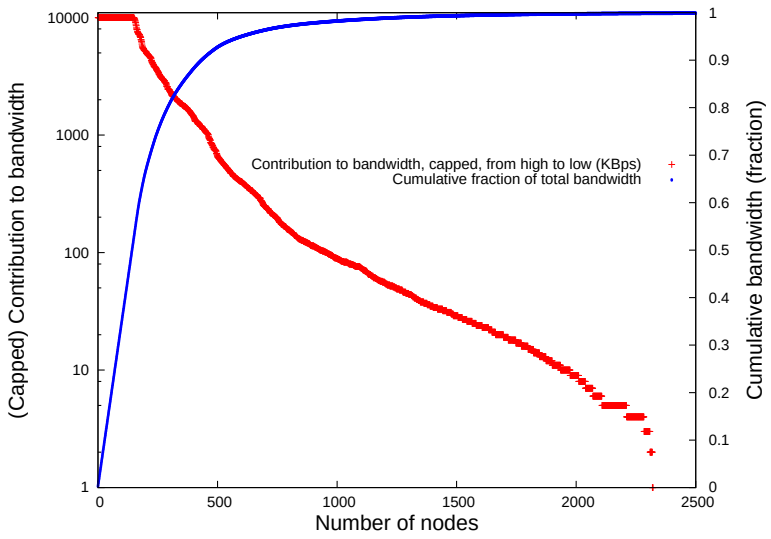

Figure 3: Capped Bandwidth distribution from April 22, 2011 Tor consensus data.

Reasoning: There are $g + d$ choices for the entry node and $e + d$ choices for the exit node. The total choices are $(g+d)(e+d) = ge+gd+ed+d^2$; however, some of the node pairs counted by d^2 are those that use the same guard+exit node for entry and exit, which is not allowed in Tor, thus we need to subtract those duplicate instances, of which there are d.

If we know that some of these are compromised—say, g_C guard nodes, e_C exit nodes, and d_C guard+exit nodes—then the probability of picking a pair of compromised nodes is

$$\frac{(g_C + d_C)(e_C + d_C) - d_C}{(g + d)(e + d) - d} \tag{4}$$

by reasoning similar to the above.

4.2 BWRS

Recall there are four types of nodes: guard, exit, guard+exit, and non-flagged (only good for the middle position). Let B_i be the bandwidth of node x_i. Let Sg be the set of guard nodes, Se the set of exit nodes, Sd the set of guard+exit nodes, and Sm the set of non-flagged nodes.

In all cases the contribution of the middle node is essentially irrelevant. From a calculation point of view, once the guard and exit nodes are chosen, the middle node is chosen from what remains, so we must subtract the bandwidth used up by the other nodes. The probability of a given middle node m being in a path is

$$\frac{B_m}{(\sum_{i\in Se} B_i + \sum_{i\in Sg} B_i + \sum_{i\in Sd} B_i + \sum_{i\in Sm} B_i) - B_g - B_e}. \tag{5}$$

For conciseness, we let $\mathcal{B} = (\sum_{i\in Se} B_i + \sum_{i\in Sg} B_i + \sum_{i\in Sd} B_i + \sum_{i\in Sm} B_i) - B_g - B_e$ and write $B_m/\mathcal{B}$. However, if we want the probability of compromise, we sum the probability for all given nodes, which is

$$\frac{(\sum_{i\in Se} B_i + \sum_{i\in Sg} B_i + \sum_{i\in Sd} B_i + \sum_{i\in Sm} B_i) - B_g - B_e}{(\sum_{i\in Se} B_i + \sum_{i\in Sg} B_i + \sum_{i\in Sd} B_i + \sum_{i\in Sm} B_i) - B_g - B_e} \tag{6}$$

which is $\mathcal{B}/\mathcal{B} = 1$. Of course, the fact that the contribution of the middle nodes is just 1 can be seen intuitively as the probability of path compromise is simply the probability of guard node compromise multiplied by the probability of exit node compromise; whether the middle node is compromised is actually not relevant.

Now we consider the choice of guard and exit nodes. The probability of x_g being chosen as the first node is

$$\frac{B_g}{\sum_{i\in Sg} B_i + \sum_{i\in Sd} B_i} \tag{7}$$

because the guard node could come from either Sg or Sd. This follows from a direct counting argument, and Bauer *et al.* [2] has the same declaration—the available bandwidth for that node is the numerator and the total available bandwidth of all nodes is the denominator, so this indicates what fraction of the total bandwidth the node controls. The probability of x_e being chosen as the exit node is

$$\frac{B_e}{\sum_{i\in Se} B_i + \sum_{i\in Sd} B_i} \tag{8}$$

again because it could come from either Se or Sd.

It is tempting to say that the probability that a path contains x_g as the entry node, x_m as the middle node, and x_e as the exit node is:

$$\left(\frac{B_g}{\sum_{i\in Sg} B_i + \sum_{i\in Sd} B_i}\right)\left(\frac{B_e}{\sum_{i\in Se} B_i + \sum_{i\in Sd} B_i}\right)\left(\frac{B_m}{\mathcal{B}}\right).$$

However, this is not quite right—for example, if g is from Sd then e could be chosen from $(Se \cup Sd)\backslash g$ and not from $Se \cup Sd$.

To take care of this we could partition the consideration into cases; however, since we need the sum over all possibilities, this can be accomplished using a counting argument.

More formally, suppose the set of compromised nodes $C = G_C \cup E_C \cup D_C$ is partitioned into a set of G_C guard nodes, a set E_C of exit nodes, and a set D_C of guard+exit nodes. An elementary counting argument can be used to establish:

THEOREM 4.3. *The probability that a path constructed according to BWRS will start with a compromised node and end with a compromised node is*

$$\frac{\left(\sum_{i\in G_C} B_i + \sum_{i\in D_C} B_i\right)\left(\sum_{i\in E_C} B_i + \sum_{i\in D_C} B_i\right)}{\mathcal{S}} - \frac{\left(\sum_{i\in D_C} B_i^2\right)}{\mathcal{S}} \quad (9)$$

where $\mathcal{S} = \left(\sum_{i\in Sg} B_i + \sum_{i\in Sd} B_i\right)\left(\sum_{i\in Se} B_i + \sum_{i\in Sd} B_i\right) - \left(\sum_{i\in Sd} B_i^2\right)$.

Proof: Given a set of compromised nodes, the probability of a compromised node being chosen for a path is the sum of the probabilities of each given node being chosen, i.e. for the guard and guard+exit nodes being chosen for entry position, the probability is

$$\left(\frac{\sum_{i\in G_C} B_i + \sum_{i\in D_C} B_i}{\sum_{i\in Sg} B_i + \sum_{i\in Sd} B_i}\right),$$

and for exit and guard+exit nodes it is

$$\left(\frac{\sum_{i\in E_C} B_i + \sum_{i\in D_C} B_i}{\sum_{i\in Se} B_i + \sum_{i\in Sd} B_i}\right).$$

For the middle nodes, from the analysis before equations (5) and (6), we know the contribution from the middle nodes is just 1.

We cannot simply multiply these probabilities together, however, if we need to consider the probabilities of pairs of nodes. This is because such a product would imply the counting is *with replacement*, meaning it allows the same node to be used as guard node and as exit node, which is not allowed in Tor. Thus, using the same sort of analysis as in (4), we derive the ratio in the statement of the theorem. ⋄

4.3 ABWRS

An analysis parallel to that of the previous section provides the following weighted theorem:

THEOREM 4.4. *The probability that a path constructed according to ABWRS will start with a compromised node and end with a compromised node is*

$$\frac{\mathcal{G}_\mathcal{C} \times \mathcal{E}_\mathcal{C} - \mathcal{D}_\mathcal{C}}{\mathcal{G} \times \mathcal{E} - \mathcal{D}} \quad (10)$$

where

$$\begin{aligned}
\mathcal{G}_\mathcal{C} &= \left(W_{gg}\sum_{i\in G_C} B_i + W_{gd}\sum_{i\in D_C} B_i\right)\\
\mathcal{E}_\mathcal{C} &= \left(W_{ee}\sum_{i\in E_C} B_i + W_{ed}\sum_{i\in D_C} B_i\right)\\
\mathcal{D}_\mathcal{C} &= \left(W_{gd}W_{ed}\sum_{i\in D_C} B_i^2\right)\\
\mathcal{G} &= \left(W_{gg}\sum_{i\in Sg} B_i + W_{gd}\sum_{i\in Sd} B_i\right)\\
\mathcal{E} &= \left(W_{ee}\sum_{i\in Se} B_i + W_{ed}\sum_{i\in Sd} B_i\right)\\
\mathcal{D} &= \left(W_{gd}W_{ed}\sum_{i\in Sd} B_i^2\right)
\end{aligned}$$

5. RESULTS

In Section 2 we have discussed the drawbacks to entropy and how entropy is not suitable as an attack-based measure. Instead we propose the probability of path compromise given an adversarial bandwidth budget. In Section 4 we derived formulas for such probabilities under different bandwidth and weighting conditions. Here we calculate these probabilities using consensus data for Tor from April 22, 2011.

To explore an attack-based measure we want to explore from the perspective of the adversary: what is best for him or her? Note, again, that a measure like entropy does not take such perspectives into account, so does not consider what might be a realistic attack or probability of compromise. We consider the adversary bandwidth fraction (ABWF) situation. We assume the adversary has a fixed bandwidth budget, $n\%$. That is, he or she can control up to $n\%$ of the entire bandwidth. Of course, many sets of nodes could give that $n\%$. Of interest, then, is which of these sets gives the highest probability of compromise, and, further, what that probability is. Recall that a path is compromised if the adversary controls the entry node and exit node of the path.

Figure 4 represents this value, taken over all sets. The x-axis is labelled by the fraction of bandwidth controlled by the adversary, while the y-axis is labelled with the percentage of guard-exit pairs that are compromised. The point (b, p) on the graph means that p is the highest probability of choosing a compromised pair, chosen over all sets of compromised nodes with total bandwidth at most b. These numbers were derived from the same dataset used above and illustrated in Figure 3. As can be seen from the graph, the chance of compromise starts to increase rapidly when the fraction of compromised nodes is more than 25%.

To find the best combination of nodes to compromise the maximum number of paths with a given bandwidth, we explore all combinations of guard and exit nodes, starting with the smallest bandwidth offered. This obviously leads to combinatorial exploration of all possibilities, so we stop after a reasonable amount of time has elapsed since the last discovery. In all cases, the best case (in that sense) was discovered within one hour of computation, sometimes after several minutes.

This leads us to believe that the solution found, understandably made from a large number of nodes of each category, is optimal because using such nodes allows us to get the closest to the allocated bandwidth budget.

We can observe that our results are not far from a rough approximation to the probability of picking a compromised path in Tor by f^2, where f is the proportion of bandwidth controlled by the adversary, assuming those nodes could play either guard or exit roles. In practice we may not achieve exactly f^2 because nodes come in fixed, predetermined amounts (or denominations, if you will) and it maybe not be possible to find a set of nodes whose bandwidth sum exactly equals the bandwidth budget (cf. the knapsack problem).

In comparing these results to that of Murdoch and Watson [11], we find a lower expectation of path compromise. However, our points of view are different. One difference between our work and theirs is that their results flow from a simulation of a small number of paths and we are working from bandwidth compromised.

Our approach satisfies the criterion given by the quotation from Murdoch and Watson in Section 2—it reflects the limited bandwidth available to an attacker. Entropy misses this key information and, in a sense, measures the crowd and how easy it is to get lost in it, rather than how likely it is that members of that crowd are malicious.

5.1 Example

We ran a number of calculations on the consensus dataset from April 22, 2011 and produced the following results.

The associated weights, as discussed in Section 3.3 are:

$$W_{gd} = 0.1639, W_{ed} = 0.6722, W_{md} = 0.1639,$$
$$W_{ee} = 1, W_{me} = 0, W_{mg} = 0.3692, W_{gg} = 0.6308 \quad (11)$$

The nodes and their bandwidths were divided into categories as follows:

Node type	Number of Nodes	Total Bandwidth (KB/s)
Guard	540	1,246,537
Exit	723	183,093
Guard+Exit	270	836,524
Non-flagged	788	392,255

We then gave the adversary various budgets of (capped) bandwidth and computed "optimal" solutions, as described above. The types of nodes the adversary would be choosing from are either guard nodes, exit nodes or guard+exit nodes. The table below shows, for a given bandwidth budget, the total bandwidth of each type of node that is in the computed "optimal" solution:

Bw %	G bw (KB/s)	E bw (KB/s)	G+E bw (KB/s)
5%	40,031	20,935	62,339
10%	77,005	34,435	125,169
15%	113,795	50,427	185,699
20%	158,155	51,375	253,699
25%	186,845	61,375	326,229

Using equation (10) we can derive the probability of node compromise for each of these bandwidth budgets as:

Bandwidth %	Probability of compromise
5%	0.0032
10%	0.0119
15%	0.0259
20%	0.0455
25%	0.0697

6. RELATED WORK

As mentioned in Section 2, entropy was introduced as a measure for anonymity by Serjantov and Danezis [14] and Diaz *et al.* [5]. There have been several alternative measures proposed for mix networks. As already mentioned, Shmatikov and Wang [16] explore min entropy. Clauß and Schiffner [4] advocate Rényi entropy. Edman *et al.* propose a combinatorial measure [9], later improved [21] by Gierlichs *et al.* However, to the best of our knowledge, none of these have been applied to Tor paths. Snader and Borisov [17] introduce an inequality measure, the Gini coefficient, to measure the inequality in their Tor path selection algorithm, but they consider this measure only. Syverson [18] asserts that a metric should reflect the difficulty an adversary has in compromising the system. Finally, Murdoch and Watson [11] and Panchenko and Renner [12] also look at measures when the adversary controls the entry and exit nodes. Murdoch and Watson [11] further discuss the idea of a bandwidth budget available to the adversary and perform simulations using several path selection techniques. They calculate percentage of compromised paths in these simulations.

7. CONCLUSION AND FUTURE WORK

We exposed a number of critiques of entropy and discussed in particular how it fails to consider an attack-based approach. We calculated the probability of path compromise under three different Tor path selection algorithms, and combined this with the idea of adversarial bandwidth fraction to determine the likelihood of path compromise under different bandwidth budget conditions.

Additional constraints. Broadening the range of situations in which we measure compromise would also be an interesting avenue. Compromise is more likely in some situations than others. For example, what if your guard node is controlled by an adversary? Dingledine and Murdoch [8] said one of the problems was that guard nodes were hanging around too long. That is, once chosen, they were held onto by the user for months at a time, and were thus overloaded.

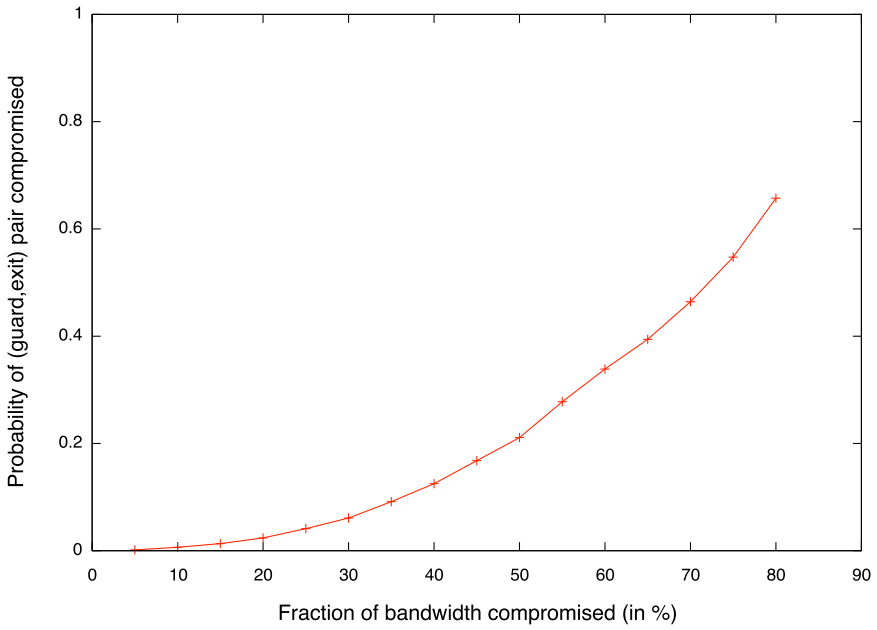

Figure 4: "Worst-case" probability of path compromise given an adversary's bandwidth budget

They suggest that they should be retained for a few weeks instead. But the frequency of change of guard nodes has implications for the likelihood of path compromise.

Another issue is that when the nodes operate under the same administration they are more vulnerable to compromise—if the administration is corrupt. There are further cases to consider, e.g. are they managed by the same operator, a different branch of the same operator, or operators running under the same jurisdiction (country). While it is clear that coordination becomes harder as we get further away from network operations, higher risks remain. Ideally a measure would allow for consideration of these types of constraints as well.

In our investigations we found that the top 10 countries in bandwidth offered (in MB/s) were the following:

Country	**Bandwidth**
Germany DE	240
Unites States US	212.5
Holland NL	180.7
Sweden SE	39.71
France FR	36
Austria AT	25.7
Great Britain GB	22
Romania RO	19.9
Switzerland CH	19.3
Luxembourg LU	8.8

Thus the amount of bandwidth for substantial compromise is clearly beyond what any single country offers at this point (for the April 22, 2011 data set). It would be interesting to investigate whether the country information could be used to decide on a "reasonable" level of compromise: a country could not reasonably be entirely under the control of an unique entity (speaking of the countries at the top of our list), nor could such an entity control but a small proportion of nodes in most (key) jurisdictions. For example, we see that approximately 5% of the bandwidth is in the country with the highest amount of bandwidth—Germany—so it may be reasonable to assume that no adversary would control more than that.

We also assumed a passive adversary. it would be interesting to examine this work under the assumption of an active attacker, as in Borisov *et al.* [3].

Policy. The probability of path compromise clearly depends on a number of quantitative factors such as total number of nodes, number of compromised nodes, bandwidth of nodes, and how often paths change. However, in addition to these factors are a number of qualitative factors we call *policy*. These are factors set by the user and include things like countries to avoid (e.g. most users may be happy to connect through Sweden, but Julian Assange may be less content!), how long a user connects for, and how often a user connects. As the user sets policies, the set of nodes from which it is choosing will shrink, and a graph representing the probability of detection vs. number of paths will change. Policy is individual and hard to quantify, but is there a way of incorporating it into a measure?

Alternate designs. We consider the default path selection algorithm in Tor. But we could have looked at others such as Snader and Borisov [17], Panchenko and Renner [12], Sherr *et al.* [15], Edman and Syverson [10], Zhang *et al.* [23]. This would also be interesting future work.

8. ACKNOWLEDGEMENTS

The authors thank Kevin Bauer, Tariq Elahi, Ryan Henry, Rob Smits, and Tao Wang for helpful comments on the manuscript. This work was supported by NSERC, MITACS and The Tor Project.

9. REFERENCES

[1] Kevin Bauer, Damon McCoy, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. Low-resource routing attacks against anonymous systems. Technical report, University of Colorado at Boulder, CU-CS-1025-07, 2007.

[2] Kevin Bauer, Damon McCoy, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. Low-resource routing attacks against Tor. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2007)*, pages 11–20, Washington, DC, USA, October 2007.

[3] Nikita Borisov, George Danezis, Prateek Mittal, and Parisa Tabriz. Denial of service or denial of security? How attacks on reliability can compromise anonymity. In *Proceedings of CCS 2007*, pages 92–102, October 2007.

[4] S. Clauß and S. Schiffner. Structuring anonymity metrics. In *DIM*, pages 55–62, November 2006.

[5] Claudia Diaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.

[6] R. Dingledine. The Tor Blog, Research Problem: measuring the safety of the Tor network, February 5, 2011. https://blog.torproject.org/blog/research-problem-measuring-safety-tor-network. Accessed April 2011.

[7] R. Dingledine and N. Mathewson. Tor path specification. https://gitweb.torproject.org/torspec.git?a=blob_plain;hb=HEAD;f=path-spec.txt. Accessed April 2011.

[8] R. Dingledine and S. Murdoch. Performance improvements on Tor or, why Tor is slow and what we're going to do about it. Tor Technical Report 2009. https://svn.torproject.org/svn/projects/roadmaps/2009- 03-11-performance.pdf. Accessed April 2011.

[9] M. Edman, F. Sivrikaya, and B. Yener. A combinatorial approach to measuring anonymity. *Intelligence and Security Informatics, 2007 IEEE*, pages 356–363, May 2007.

[10] Matthew Edman and Paul F. Syverson. AS-awareness in Tor path selection. In Ehab Al-Shaer, Somesh Jha, and Angelos D. Keromytis, editors, *Proceedings of the 2009 ACM Conference on Computer and Communications Security, CCS 2009, Chicago, Illinois, USA, November 9-13, 2009*, pages 380–389. ACM, 2009.

[11] Steven J. Murdoch and Robert N. M. Watson. Metrics for security and performance in low-latency anonymity networks. In Nikita Borisov and Ian Goldberg, editors, *Proceedings of the Eighth International Symposium on Privacy Enhancing Technologies (PETS 2008)*, pages 115–132, Leuven, Belgium, July 2008. Springer.

[12] A. Panchenko and J. Renner. Path selection metrics for performance-improved onion routing. In *SAINT*, pages 114–120, July 2009.

[13] M. Perry. TorFlow: Tor network analysis. In *HotPETS*, 2009.

[14] Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.

[15] M. Sherr, B.T. Loo, and M. Blaze. Towards application-aware anonymous routing. In *USENIX Workshop on Hot Topics in Security, Article 4*, 2007.

[16] Vitaly Shmatikov and Ming-Hsiu Wang. Measuring relationship anonymity in mix networks. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2006)*, October 2006.

[17] Robin Snader and Nikita Borisov. A tune-up for Tor: Improving security and performance in the Tor network. In *Proceedings of the Network and Distributed Security Symposium - NDSS '08*. Internet Society, February 2008.

[18] P. Syverson. Why I'm not an entropist. In *17th Security Protocols Workshop, Cambridge, UK*, 2009.

[19] Tor Directory Protocol. https://gitweb.torproject.org/torspec.git?a=blob_plain;hb=HEAD;f=dir-spec.txt. Accessed April 2011.

[20] Gergely Tóth, Zoltán Hornák, and Ferenc Vajda. Measuring anonymity revisited. In Sanna Liimatainen and Teemupekka Virtanen, editors, *Proceedings of the Ninth Nordic Workshop on Secure IT Systems*, pages 85–90, Espoo, Finland, November 2004.

[21] Carmela Troncoso, Benedikt Gierlichs, Bart Preneel, and Ingrid Verbauwhede. Perfect matching statistical disclosure attacks. In Nikita Borisov and Ian Goldberg, editors, *Proceedings of the Eighth International Symposium on Privacy Enhancing Technologies (PETS 2008)*, pages 2–23, Leuven, Belgium, July 2008. Springer.

[22] Matthew Wright, Micah Adler, Brian Neil Levine, and Clay Shields. Defending anonymous communication against passive logging attacks. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, pages 28–43, May 2003.

[23] N. Zhang, W. Yu, X. Fu, and S.K. Das. gPath: A game-theoretic path selection algorithm to protect Tor's anonymity. In *GameSec*, pages 58–71, November 2010.