

Throwing Your Weight Around: Fixing Tor’s Positional Weighting

Aaron Johnson
U.S. Naval Research Laboratory
aaron.m.johnson@nrl.navy.mil

Aaron D. Jaggard
U.S. Naval Research Laboratory
aaron.jaggard@nrl.navy.mil

Paul Syverson
U.S. Naval Research Laboratory
paul.syverson@nrl.navy.mil

ABSTRACT

We analyze deficiencies in Tor’s positional weighting system, identifying cases in which the system either fails to produce valid weights or fails to properly load balance across positions. We describe how an attacker can take advantage of these failures to reduce Tor’s performance, thereby also easing censorship and surveillance through a denial-of-service attack. Our attacks exploit incorrectly determined positional-weight equations by adding new capacity to the network or, for even more covertness, by just minor changes in the status of existing malicious relays. Our analysis of past Tor consensus shows that these attacks could have reduced the throughput of the network by as much as 45% due only to their triggering of Tor’s flawed position weights. Rather than a mere patch to Tor’s currently ad hoc scheme, we then propose a new, systematic method for deriving positional weights and propose two goal sets generated using that method. We derive new sets of weights, prove that they satisfy these goal sets, and give examples of how they would change the weights from the current system. Tor could use our results to quickly fix the main deficiencies of its positional weights as well as adopt a better approach long-term.

KEYWORDS

Tor, onion routing, denial-of-service attack

1 INTRODUCTION

Tor is a circuit-based anonymous-communication network comprised of thousands of volunteer-run relays and used by millions of individuals every day. Tor uses onion routing to protect communications metadata from network observers and even the relay operators. It is a popular tool to avoid network surveillance and censorship, and consequently some adversaries attempt to block all access to it. Tor has multiple mechanisms to resist blocking, such as bridge relays, which are not in the public list of Tor relays, and obfuscated transport protocols, which make it hard to recognize Tor traffic on the wire.

When attempting to surveil or censor users employing communications protections, however, adversaries may sometimes find it more useful to degrade performance of protected communication rather than to simply block it outright. Besides simply limiting undesirable traffic, this can serve as a disincentive for users. Casual users are more likely to switch to less secure communication methods if they experience a large performance hit from using protections such as Tor. Such induced self-censorship also has the effect that remaining traffic is more likely to be of interest to a

surveilling adversary, thus reducing the amount of noise in any surveillance data collected for analysis. For example, in 2013, rather than blocking encrypted protocols, Iran throttled them to about 20% of network speed otherwise available [1]. Such attacks may also be harder to distinguish than outright blocking from naturally occurring network performance issues, which can make the attacks less likely to be countered as well as plausibly deniable.

As a privacy resource for millions, proper provisioning of Tor resources is important even absent any adversary manipulation. Performance affects usability, and usability is a security property [7]. Tor must account for the widely varying capacity of its relays to effectively allocate resources when creating its circuits. This allocation issue is complicated by network entrance and exit policies in Tor that take into account differing capabilities and goals among relays and their operators [28]. This is reflected in Tor’s design through the assignment of the Guard and Exit flags to relays. Guard relays are the points at which Tor circuits enter the Tor network (except for the minority of circuits using bridges), and exit relays are those from which Tor circuits are permitted to connect to servers on the wider Internet.

Because of these different classes of relays, Tor needs to account, not just for the capacity of individual relays, but for the capacity of relay classes. For example, if only a small amount of the overall network capacity has the Exit flag, that class creates a performance bottleneck. Managing allocation of relay-class capacity is the job of Tor’s positional-weighting scheme, which is described in Tor’s directory protocol specification [30].

In this paper, we will look at Tor’s current positional-weighting scheme and its deficiencies, which can both naturally cause problems and be actively exploited by adversaries. We also describe alternative schemes. Rather than adopt a penetrate-and-patch approach to fixing the current scheme, our alternatives are based on a more principled consideration of general goals of security and performance.

In the next section, we present background necessary to understand Tor, its load balancing, and its positional-weighting scheme. Then, in Section 3, we describe some of the deficiencies of that scheme, in particular cases where it either fails to assign weights at all or where it suboptimally assigns weights. In Section 4, we describe attacks an adversary can perform by strategically adding to particular classes or causing relay capacity to shift from one class to another. For either approach, the attack causes the weighting-scheme case to change to a case that fails to assign weights to positions. Looking at past Tor consensus, we analyze the amount of added or shifted bandwidth necessary to cause the attack and the effect of the attack on network throughput. In Section 5, we consider alternatives to Tor’s ad hoc positional weighting based on two sets of specific prioritized goals. For each alternative, we define positional weights and prove that the resulting allocation satisfies the set of goals in prioritized order. In Section 6, we discuss ethical

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.
Proceedings on Privacy Enhancing Technologies 2023(4), 1–20
Distribution Statement A: Approved for public release, distribution is unlimited.
<https://doi.org/XXXXXXXX.XXXXXXX>



disclosure, recommendations for Tor, limitations, and relevance to onion-routing networks other than Tor. In the final two sections, we describe related work and present our conclusions.

2 BACKGROUND

The Tor network is coordinated by a set of Directory Authorities [30]. They collectively generate a consensus document each hour, which contains data about the network that clients need to use it, including a list of all the Tor relays and their individual bandwidths. The bandwidth listed for each relay is also called a consensus weight, and it is best understood as a unitless value because Tor’s measurement process only makes it accurate relative to that of other relays [22]. Tor clients regularly download the latest consensus and use it to select relays when creating new circuits.

Tor relays create different types of circuits for different purposes. Exit circuits are created to make connections to destinations on the Internet that are unaware of Tor. Onion circuits are created to connect to onion services, which are Tor-aware servers, and create end-to-end encrypted connections with their clients. Directory circuits are used to download data about the network, including the consensus and per-relay descriptors. Tor is largely optimized for exit circuits because they make up the majority of circuits in the network and carry most of the traffic [31]. Exit circuits are generally three-hop circuits consisting first of a *guard* relay, then a *middle* relay, and finally an *exit* relay.

The Directory Authorities indicate which relays may be chosen in the guard and exit positions by including in the consensus a set of flags for each relay. The Guard flag is required for relays to be chosen in the guard position, and relays are required to have a minimum bandwidth and uptime to receive this flag. Being chosen for the exit position requires an exit policy that allows connection to the desired port and IP address, but the Exit flag is assigned based on common ports to approximate which relays are likely to be chosen as exits.

That approximation for the Exit flag serves the main purpose of enabling load-balancing across the positions in exit circuits. Relays have different bandwidths and are able to serve in different circuit positions. Therefore, the total bandwidth available in each position varies, and also using a relay’s bandwidth in one position makes it unavailable in another position. For example, in the current Tor network, the aggregate bandwidth of relays with the Exit flag is typically significantly smaller than the aggregate bandwidth of relays with the Guard flag. Choosing relays with both flags for the guard position reduces their bandwidth available for the exit position and reduces the overall network throughput.

To solve this load-balancing problem, a Tor consensus defines a set of positional weights that apply to a relay based on its possession of the Guard and Exit flags. Let \mathcal{G} be the set of relays with the Guard but not the Exit flag, \mathcal{M} be the set with neither flag, \mathcal{E} be the set with the Exit but not the Guard flag, and \mathcal{D} be the set with both flags. Also, let g indicate the guard position, m indicate the middle position, and e indicate the exit position. A consensus defines weights $W_{xy} \in [0, 1]$ for each $x \in \{g, m, e\}$ and $y \in \{\mathcal{G}, \mathcal{M}, \mathcal{E}, \mathcal{D}\}$ (for numerical reasons, the W_{xy} values are actually given as integers in a fixed range and later normalized). Clients weight their random selection of a relay for a position by the positional weight that

applies to that relay in that position. Specifically, for a relay with bandwidth b and in class y , the client chooses that relay for position x with probability proportional to $b \cdot W_{xy}$ if the relay is eligible for position x and with probability 0 otherwise [8].

This design allows the positional weights to be set to balance the use of relays across the positions and optimize network throughput. The Directory Authorities use a set of equations to obtain these weights, and those equations are indeed defined with the goal of maximizing network throughput ([30, § 3.8.3]). The equations use several approximations to simplify this goal: (1) they only consider exit circuits, although a proposal does exist to incorporate traffic over other circuit types [23]; (2) they assume guard relays are selected for each circuit using only bandwidth, although each client in fact selects a small number of guards for all its circuits and weights that selection by how long a guard has have possessed the Guard flag; and (3) they assume all clients will put the same amount of traffic load on the network.

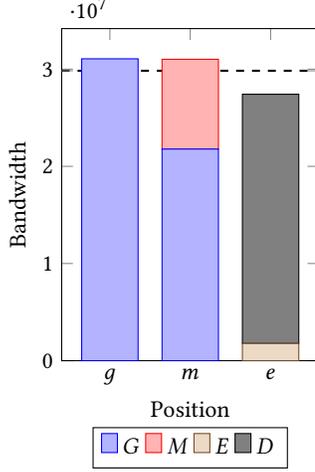
Using these approximations, we can estimate the bandwidth available in each position in order to perform load balancing across them. Let G, M, E , and D be the sum of relay bandwidth in classes $\mathcal{G}, \mathcal{M}, \mathcal{E}$, and \mathcal{D} , respectively. For $x \in \{g, m, e\}$ and $y \in \{\mathcal{G}, \mathcal{M}, \mathcal{E}, \mathcal{D}\}$, let y_x denote the bandwidth y weighted for position x , e.g. $D_g = D \cdot W_{gD}$. We require that all weights for a given class sum to one: $\sum_x W_{xy} = 1$. Then we denote by B_x the estimated bandwidth available in position x , which we define as follows:

$$\begin{aligned} B_g &= G_g + D_g \\ B_m &= G_m + M_m + E_m + D_m \\ B_e &= E_e + D_e. \end{aligned}$$

Because all traffic needs to pass through each position, the minimum value among these creates a “bottleneck” that limits the total expected network throughput. For a similar reason, given the total network bandwidth, $T = G + M + E + D$, the maximum network throughput is limited to $T/3$. This limit can be achieved if the bandwidths allocated to each position can be made equal. Combining two equations guaranteeing these equalities with the equations requiring the W_{xy} sum to one yields six equations for eight unknowns, and so there may be multiple satisfying solutions. Moreover, sometimes the position bandwidths cannot be made equal, for example when there is relatively scarce bandwidth available for a given position. Tor computes positional weights by considering various scarcity conditions and then choosing weights that attempt to maximize the minimum bandwidth allocated to the positions. Its weight equations appear in Appendix A.

Figure 1 illustrates an example of how Tor’s current positional weights allocate its bandwidth across circuit positions. It shows the fraction of each relay class allocated to each position for the consensus from 2021-12-31 23:00. In this network state, (which falls into case 3(a)iiB—see Section 3), the positional allocations cannot be balanced because the total bandwidth available for the exit position is less than a third of the total bandwidth (i.e. $E + D < T/3$). Observe that there are alternative allocations that maximize the minimum positional bandwidth—some of the G bandwidth can be moved from the guard to the middle position or vice versa without reducing either one below the bandwidth in the exit position. As we will show, this flexibility can be used to achieve goals secondary to throughput maximization.

Figure 1: Tor’s positional bandwidth allocation for 2021-12-31 23:00. The dashed line indicates the $T/3$ throughput limit.



3 DEFICIENCIES OF POSITIONAL WEIGHTS

The way that Tor computes positional weights in some cases fails to achieve the stated goal of maximizing network throughput, even accepting their acknowledged approximations as necessary simplifications. Moreover, the stated goal does not uniquely define the weights, and no criteria to choose specific weights are either given in the specification or apparent from their logic. To describe these deficiencies, again let G , M , E , and D be the aggregate bandwidth of \mathcal{G} , \mathcal{M} , \mathcal{E} , and \mathcal{D} , respectively, and let $T = G + M + E + D$ be the total bandwidth in the network.

3.1 Tor’s current weights

We express the current Tor weights as a sequence of 13 mutually exclusive cases covering the possible network states according to the relationships among G , M , E , and D . The case division and names largely follow those in the Tor specifications [30]. However, we do further divide some cases so that each case contains a full set of weights and to characterize the cases that fail to maximize throughput. In this section, we present the six cases that we use in the paper; as we will show, the first two exhibit a problematic discontinuity, the following two give rise to errors, and the final two appear in recent consensuses. Appendix A includes these and the other seven cases. For each case, we present the positional weights in tables where the entry in row $y \in \{\mathcal{G}, \mathcal{M}, \mathcal{E}, \mathcal{D}\}$ and column W_{xy} , $x \in \{g, m, e\}$ is the weight W_{xy} . Empty cells indicate weights of zero, and \oplus indicates the logical exclusive OR. The conditions defining each case are given before each weight table.

Case 1: $E \geq T/3$, $G \geq T/3$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------|---------------------|--------------------|
| \mathcal{G} | $\frac{G+E+M}{3G}$ | $\frac{2G-E-M}{3G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{2E-G-M}{3E}$ | $\frac{E+G+M}{3E}$ |
| \mathcal{D} | $\frac{1}{3}$ | $\frac{1}{3}$ | $\frac{1}{3}$ |

Case 2(b)i: $E < T/3$, $G < T/3$, $\min(G, E) + D \geq \max(G, E)$, $G \geq M$, $E \geq G - M$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------------|--------------------------|--------------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{G-M}{E}$ | $\frac{E-G+M}{E}$ |
| \mathcal{D} | $\frac{2D+2E-4G+2M}{6D}$ | $\frac{2D+2E-4G+2M}{6D}$ | $\frac{2D+2E-4G+2M}{6D}$ |

Case 2(b)iiiA: $E < T/3$, $G < T/3$, $\min(G, E) + D \geq \max(G, E)$, $G < M \vee E < G - M$, $M > T/3$, $T/3 - E \leq D$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------------|----------|-----------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{2D+2E-G-M}{3D}$ | | $\frac{D-2E+G+M}{3D}$ |

Case 2(b)iiiB: $E < T/3$, $G < T/3$, $\min(G, E) + D \geq \max(G, E)$, $G < M \vee E < G - M$, $M > T/3$, $T/3 - E > D$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------------|----------|-----------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{2D+2E-G-M}{3D}$ | | $\frac{D-2E+G+M}{3D}$ |

Case 3(a)iiB: $E < T/3 \oplus G < T/3$, $\min(G, E) + D < T/3$, $G \geq E$, $G \geq M$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------|------------------|----------|
| \mathcal{G} | $\frac{G+M}{2G}$ | $\frac{G-M}{2G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

Case 3(b)ii: $E < T/3 \oplus G < T/3$, $\min(G, E) + D \geq T/3$, $G \geq E$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------------|------------------------|-----------------------|
| \mathcal{G} | $\frac{G+M}{2G}$ | $\frac{G-M}{2G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{2D+2E-G-M}{6D}$ | $\frac{2D+2E-G-M}{6D}$ | $\frac{D-2E+G+M}{3D}$ |

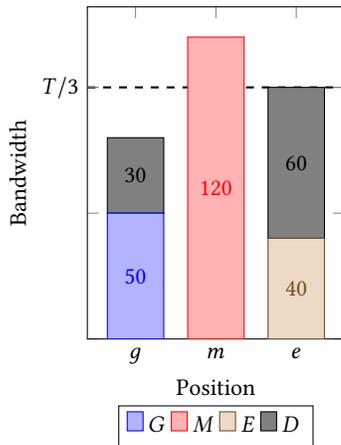
3.2 Problems with Tor’s current weights

We observe that Cases 2(b)iiiA and 2(b)iiiB fail to properly load-balance Tor, while in all other cases load balancing appears to be achieved. Cases 2(b)iiiA and 2(b)iiiB have the same weight formulas and appear as one case in the Tor specification. We divide these into two cases based on different errors that can occur with these.

The first error is in case 2(b)iiiA. In this case, the minimum bandwidth will occur in the guard position. The D bandwidth is used suboptimally; enough is used in the exit position to make the total exit bandwidth allocation $T/3$ ($W_{e\mathcal{D}} = T/3 - E$), but assigning the remainder to the guard position leaves its allocation at less than $T/3$. Fig. 2 illustrates an example of the bandwidth allocated in each position, showing a bottleneck throughput of 80 while the optimal would balance \mathcal{D} better and provide a throughput of 90.

The second error we observe occurs in Case 2(b)iiiB. This case fails to yield valid weights. $W_{g\mathcal{D}}$ will be less than 0 because $G+M = T - (D+E) > 2T/3 > 2(D+E)$, where both inequalities hold because $D+E < T/3$ in this case. Consequently, $W_{e\mathcal{D}}$ will be greater than 1.

Figure 2: Positional bandwidth allocation under Tor’s current weights with $G = 50$, $M = 120$, $E = 40$, and $D = 90$



Tor clients apply a range check on the computed weights to ensure they occur in $[0, 1]$. When the weights fail that check, the clients default to using $W_{xy} = 1$ for all x and y , which can yield very poor load balancing.

As an example of the problem, let $G = 90$, $M = 120$, $E = 40$, and $D = 50$, which falls into this case. Because the positional weighting is so unbalanced, different relay classes hit their capacity at different traffic loads. When client traffic reaches the smallest amount at which some class uses its total capacity, some additional client demand is no longer met by the network. In this example, that occurs when traffic reaches $1/(1/90 + 1/140 + 1/300) \approx 46.3$ because at that point the bandwidth of the \mathcal{D} relays becomes fully used. This limitation occurs much earlier than the optimal network throughput of 90 under load-balanced positional weights.

An additional deficiency of Tor’s current weights is that they have no stated goal other than maximizing network throughput. As a result, there are some cases with multiple throughput-optimizing allocations (e.g. Fig. 1), but one is not chosen from them in any principled way. This fact represents a missed opportunity to choose the weights that achieve some secondary goal subject to throughput maximization. It also makes the choices potentially inconsistent, yielding weights that might be quite different even for close network states that happen to fall in different cases.

4 ATTACKS FROM INDUCED CASE CHANGES

As noted above, under Tor’s current positional weighting algorithm exactly one case of relay category bandwidth distribution fails to assign positional weights to relays (Case 2(b)iiiB: see Sec. 3 and Appendix A.) And, an inspection of the Tor relay selection code current at the time of writing shows that when the weighting algorithm fails to assign positional weights, those weights are all set to 1 [32]. Thus, in this case, when selecting relays for a given circuit position, assuming a relay has a flag to permit being in that position, clients will rely just on the bandwidth assigned to that relay by the consensus and will not ensure that the relay’s bandwidth is optimally allocated across positions.

This deficiency enables attacks in which the adversary reduces network throughput by adding or modifying its own relays to put the network in the problematic 2(b)iiiB case. Of course, the Tor network could end up in this case without malicious action and consequently could experience a significant reduction in throughput. However, we consider attacks in which malicious relays are added or modified to put the network in this case. We note that an adversary may also be able to accomplish an intentional case shift via denial-of-service [12] instead of running malicious relays, by getting some relays to be dropped from the consensus or to lose the Guard flag. An adversary might also attempt to move the network state to Case 2(b)iiiA, which yields valid weights that fail to maximize minimum bandwidth, but we focus on the other attack because it yields particularly suboptimal load balancing.

As we show in Section 4.1, every Tor network state could be arbitrarily close (in bandwidth) to Case 2(b)iiiA, and small changes in relay bandwidths can lead to a large reduction in network throughput. In Section 4.2, we analyze past consensuses to estimate how vulnerable Tor has in fact been to such attacks.

Abrupt shifting of case can also create significant churn of network resource utilization, even if there is optimal weighting within the cases. Relays within a given class may shift from being mostly used in one position to mostly in another because of a case change. This could be especially problematic for the \mathcal{G} and \mathcal{D} relays because client guard choices persist after their initial selection. We explore the implications of this time inconsistency in Section 4.3.

4.1 Risk of case-change attacks

Before looking at actual consensus values that have occurred and the cases they fall into, we consider what it would take in principle for an attacker to shift from any case to another. If the adversary can shift to either of non-load-balanced cases from any other case with trivial effort, then they can induce a weighting failure (either failure to weight at all or failure to minimize the bottleneck).

Theorem 4.1 shows that all cases are arbitrarily close to one another. Its proof appears in Appendix F.

THEOREM 4.1. *For all positional-weight cases it is possible for G , M , and E to be within some ϵ of $T/3$ and $D \leq \epsilon$.*

Therefore, regardless of starting case, it is possible that the adversary needs to add less than 1% to the total relay bandwidth to move from any case to either of the cases with weighting failures. While this need not result in much loss in throughput, the following example shows that large throughput loss can indeed result from a small change in relay bandwidths.

Suppose that the network is such that $G = 98.75$, $M = 103$, $E = 0$, and $D = 98.25$. This would put the network in (load-balanced) Case 2(a)i. In this case, all the G bandwidth is allocated to the guard position ($W_{g\mathcal{G}} = 1$), all the M bandwidth is allocated to middle position ($W_{m\mathcal{M}} = 1$), and all the D and E bandwidth is allocated to the exit position ($W_{e\mathcal{E}} = 1$ and $W_{e\mathcal{D}} = 1$). The maximum throughput in this network is 98.25, where the minimum allocated bandwidth is in the exit position.

If the adversary added just one relay of bandwidth 1 (i.e. $1/300$ of the total bandwidth) to the \mathcal{D} class in this network, yielding $D = 99.25$, then the weight case would shift to Case 2(b)iiiB, which yields invalid weights because $W_{g\mathcal{D}} = -0.01$. Thus, clients would

use 1 for all positional weights when selecting paths. Because the D relays can be selected in the most positions and have nearly the same bandwidth as the G and M relays, they constitute a throughput bottleneck. When the traffic load reaches $1/(1/(D + G) + (1/(D + M + G + E)) + (1/(D + E))) = 54.21$, the D bandwidth is consumed, limiting all future traffic because only those relays are selected in the exit position. Thus, the adversary reduces the throughput of the network from 98.25 to 54.21, which is a 44.8% reduction, by adding less than 1% of additional bandwidth.

4.2 Case-change attack analysis of past consensus

The worst-case possibilities need not occur in practice, however. Therefore, we analyze past Tor consensus to determine how much case-shifting attacks could plausibly affect performance.

We first look at positional weight cases that have occurred in Tor consensus from 2009–2021. The case that fails to assign positional weights, 2(b)iiiB, does appear in one consensus in 2010. The case with correct but suboptimal positional weights, 2(b)iiiA, appears in 119 consensus in that time period, all in 2010–2011. It is not clear if in those years the Tor positional weights had the same or similar deficiencies to the current weights, as the positional weights have changed and other errors fixed over that time. Since 2015, only cases 3(a)iiB and 3(b)ii have appeared. Moreover, 3(a)iiB was the case in 99% of consensus during 12/2015–12/2021. Relative numbers varied somewhat between those two cases, but, e.g., this was also the distribution during the last nine months of our data (3/21–12/21).

We next turn to look at case changes an adversary can induce and their impacts on network performance starting from distributions of relay categories that have occurred in the consensus and ending in case 2(b)iiiB. We further consider the plausibility of attack scenarios from the perspective of adversary overhead, concerns about detectability, and adversary goals and incentives. Degrading performance to hurt network usefulness can be an end goal by itself. As noted in the introduction, however, it is also potentially a means to increasing self-censorship as well as increasing the quality of network surveillance observations. Our code for processing consensus distributions and sorting them into cases, as well as for analyzing the attacks discussed below is available [14].

4.2.1 Case-change attacks from adding bandwidth. A positional-weight case shift can reduce network throughput. Adding significantly to the relays in the network does not require the adversary to attack others' resources. It might still be noticed as an attack if a significant increase in network relays occurs close to some real-world events, though those can also often cause a rise in volunteering of relays so might not stand out.

For the focus of this paper, an important consideration is if the bandwidth the adversary must add to the network to induce such a shift increases capacity in a way that more than offsets the case change. If so, then this might not be an effective attack strategy. Indeed, as set out in Table 1 we find that across consensus during the five year period of 12/16–12/21, to induce moving to case 2(b)iiiB the median added bandwidth (relative to the particular consensus total bandwidth) necessary is 87%, and median relative throughput actually *increases* by 18%. Numbers do vary. So, e.g., during the last

year (12/20–12/21) for which we had data, median relative added bandwidth was 77% and did result in a median relative reduction of throughput of 0.5% (rounded to 1% in the table) and a maximum relative reduction of 24%. So, even when an adversary does manage to induce a small change in throughput, they need to contribute significant relay bandwidth to do so, over three-fourths of the existing network bandwidth. Note, however, that if the added bandwidth distribution were optimally weighted, instead of decreasing by 1% the median throughput would increase by over 100%.

4.2.2 Case-change attacks by shifting bandwidth. It is also possible for the adversary to induce a case change by simply shifting some bandwidth from one class to another, without adding or removing any relays or bandwidth. This might not even be noticed as an attack, especially if it is primarily a result of causing the adversary's own relays to lose the Guard or Exit flag temporarily and/or disrupting the relays of others (or even disrupting just their communication with Directory Authorities) just long enough to do so. Whether noticed or not, it is also potentially less overhead or effort for the adversary versus adding. And clients that already have selected malicious guards will continue to do so even if those lose the Guard flag, as long as the relay is in the consensus. So the adversary will minimally disrupt their ability to observe traffic. Further, inducing a case change only when it is strategic for them to reduce network performance, e.g., to discourage casual users, could increase the expected value of correlated observations of remaining network activity. As noted in the introduction, resource degradation has been used before to both censor use and improve the value of surveillance observations.

But how effective is such a shifting of relay class? Unlike for the adding-bandwidth attack, for the shifting-bandwidth attack, Table 1 shows a median relative reduction of throughput of 44% over the last five years. And while the median shift in relative bandwidth required is a substantial 40%, recall that this only requires a temporary change of relay class (flag) not a change in contributed capacity. Further, the median shift for the last year is 28% but with median reductions in throughput during that year comparable to the five year period.

Perhaps most importantly, throughput reduction is not due to a reduction of bottleneck-position capacity. When using optimal weighting, rather than the uniform weighting if in case 2(b)iiiB, instead of 44% there is no relative reduction of throughput in the median. Under the current weighting algorithm it is thus possible to have a significant negative impact on Tor network performance simply by inducing a case change—not just in theory but for most of the consensus distributions for the last several years. It is thus of more than theoretical import to have a weighting algorithm that is not subject to such attacks.

4.3 D-shifting attacks

The above attacks involve either adding capacity to various relay classes or shifting relays from one class to another. But D relays can serve as guard, exit, or both (though not for the same circuit). In this section, we explore *D-shifting* attacks, in which the adversary induces a positional-weight case change that causes a significant fraction of the D bandwidth to shift from being used in one position to another (without any abrupt addition of relays or shifting of the

| Relative values | | BW added/shifted | | | Throughput change | | | Tput change: wts optimal | | |
|-----------------|-------------|------------------|-----|------|-------------------|------|------|--------------------------|------|------|
| Attack | Period | min | med | max | min | med | max | min | med | max |
| Adding | 12/16–12/21 | 25% | 87% | 113% | 81% | 18% | -24% | 295% | 145% | 49% |
| | 12/20–12/21 | 40% | 77% | 97% | 30% | -1% | -24% | 173% | 102% | 49% |
| Shifting | 12/16–12/21 | 17% | 40% | 40% | -42% | -44% | -48% | 0% | 0% | -47% |
| | 12/20–12/21 | 28% | 28% | 34% | -42% | -45% | -48% | 0% | 0% | -47% |

Table 1: Throughput reduction attacks from adding relay bandwidth or shifting bandwidth between relay classes. ‘Relative’ numbers are relative to the bandwidth (resp. throughput) of the current consensus. Bandwidth percentages are relative to the total network bandwidth, cumulative across all classes.

class existing relays are in). These attacks can have a large impact on overall network performance, overhead, etc.

Also, suddenly adding significant numbers of relays to the network or significantly adding capacity to existing relays calls attention to those relays and makes them subject to scrutiny. Similarly, even just adding or removing guard or exit flags for significant numbers of relays suddenly might not be wise for an adversary wanting to minimize the risk of detection. *D*-shifting is relatively easier to perform while the resulting attacks are harder to distinguish from normal network dynamics or attribute to particular relays.

As we have noted, algorithm cases can be arbitrarily close to each other when *D* is a very small fraction of *T*. But can an adversary cause discontinuous case changes where the allocations are not close? First we explore whether an adversary can cause the fraction of *D* bandwidth used in the guard position to go from zero to a non-trivial portion of *T*. This is problematic even if most *D* relays are honest. For example, the adversary could make this change short-lived, and so there will only be a small number of clients who choose the relays during this period, and they will be thus less anonymous if their guard is identified (and guard discovery attacks exist [15, 20]). Even without owning or bridging those guards to find the specific client IP, this can also be combined with attacks by middle relays to fingerprint the activity of the clients who obtained guards at that time [11, 18].

Further, if most adversary-owned *D* bandwidth is devoted to exits except possibly during the short-lived shifting attack, then the adversary will be able to increase success rate of adversary-guard to adversary-exit correlation (or fingerprinted-guard to adversary-exit correlation) without changing the fraction of its contributions to the network, either by increasing its own contributions or by attacking and diminishing others. And as noted, the correlated traffic can be disproportionately that of clients targeted for having started to use the network at a specific time.

The following is an example demonstrating that such an attack is possible with only a small amount of adversary bandwidth. Let *T* be the total relay bandwidth. Given $\epsilon \in (0, 1/6]$, let $G = E = D = (1 - \epsilon)T/3$, and let $M = \epsilon T$. This network falls in case 2(b)i (which has occurred in many Tor consensuses, though none in the last decade). The weight W_{gD} in this case is equal to $(T/3 - G)/D = \epsilon T/(3D)$.

Suppose the adversary controls the *M* relays and that it decreases their bandwidth by ϵT , to zero. The resulting network falls in case 1 because (1) $E = T/3$ and (2) $G = T/3$. The weight W_{gD} in this case is $1/3$, and so the adversary has caused the total amount of *D* bandwidth used in the guard position to go from $\epsilon T/3$ to $(1 - \epsilon)T/9$,

which by varying ϵ can be made arbitrarily close to a change from 0 (i.e. no *D* bandwidth used in the guard position) to $T/9$ (i.e., $1/3$ of the *D* bandwidth used in the guard position). Note especially that this affects a potentially large number of clients selecting guards because *D* is a large fraction of the total bandwidth (nearly one-ninth) and is one-third of all bandwidth used in the guard position.

This attack causes a change in weighting-algorithm case. As already noted above, by causing a *D*-shift at strategic times an adversary can, disproportionately to its resources and with reduced risk of detection, facilitate observations of behavior by clients-of-interest. Another impact, however, is to disproportionately degrade performance for clients that select guards during the shift. Middles and exits are newly selected for each circuit, but a guard is kept for many months once selected. For our example, before the *D*-shift and after it has ended, $W_{eD} = (3D - 2M)/3D = 1 - 2\epsilon/(1 - \epsilon)$, which is at least $3/5$, depending on ϵ . The smaller ϵ the more any *D* relay is weighted to be selected as exit and the less of its bandwidth will be available for those clients using it as a guard, which should almost entirely be those who selected it during the *D*-shift. If the *D*-shift was part of a censorship event, note that in the past censoring adversaries have used degrading rather than outright blocking of secure communications such as over Tor [1].

Another complementary attack would be to shift temporarily to case 2(b)i from a distribution that has regularly been in case 1. This would cause a temporary competition for bandwidth among all clients that have *D*-guards. Case 1 is even rarer than Case 2(b)i in the overall consensus history. But the common cases could change at any time due to external events or changes in network-design parameters.

5 NEW POSITIONAL WEIGHTS FOR TOR

5.1 Stopgap fixes to current shortcomings

We start by identifying quick fixes to the shortcomings, identified in Section 3, in Tor’s current positional weights. These fixes are simple tweaks to the weights in the affected cases, but they are not part of a complete bandwidth allocation that is derived from clearly articulated design goals.

As noted in Section 3, Case 2(b)iiiA may lead to a smaller bottleneck than is necessary, and Case 2(b)iiiB may produce invalid weights. We propose changing, in both cases,

$$W_{gD} = \frac{2D + 2E - G - M}{3D}$$

$$W_{e\mathcal{D}} = \frac{D - 2E + G + M}{3D}$$

to

$$W'_{g\mathcal{D}} = \frac{E - G + D}{2D}$$

$$W'_{e\mathcal{D}} = \frac{D + G - E}{2D}.$$

The resulting weights would maximize network throughput in these cases.

5.2 Preliminaries to a more general approach

5.2.1 Metagoals. Our approach to defining bandwidth allocations involves defining an ordered list of goals and then determining which allocation satisfies those goals as fully as possible. In particular, we attempt to satisfy each goal as much as possible without reducing the extent to which any higher priority goal is satisfied. This typically involves case analysis conditioned on the values of G , M , E , and D . To facilitate this approach, we state our goals using “maximize” or “minimize”.

We identify the following metagoals for developing goal lists.

Security The goals should promote Tor security

Performance The goals should promote Tor performance

Uniqueness The goals should lead to a unique weight solution (if they do not, extend the goal list to resolve any ambiguity)

Continuity The goals should produce a solution in which all the allocations C_p are continuous functions of D , E , M , and G

Promoting security and performance is very broad. This might be pursued in a wide variety of ways, using either metrics or general heuristics. Theoretical and experimental evaluation of the bandwidth allocations induced by goals may provide insight into side effects of pursuing any particular set of goals.

5.2.2 The V_C measure of expected advantage. Some of our goals use metrics that capture the expected advantage an adversary might obtain by adding bandwidth to a relay class. We motivate and define these before defining our goal sets.

We consider an adversary who will add a small amount of bandwidth, relative to the sizes of each existing bandwidth class, to some bandwidth class C . This adversary hopes that the bandwidth it contributes increases its chances of being used in a sensitive position—*i.e.*, as a guard or exit—in Tor circuits. Being used in one of these positions allows the adversary to observe the Tor client or the destination, respectively, while being used in both of these positions allows correlation of the source and destination. Adding bandwidth controlled by the adversary may increase its power; we want to measure (and minimize) any *extra* advantage the adversary would gain by adding to a class C instead of a different class C' .

We define a quantity V_C to capture the advantage, in the sense just discussed, that the adversary gains by adding bandwidth to class C . In the case of adding bandwidth to class \mathcal{G} , the guard position is the only sensitive position in which this bandwidth could be used. When a client chooses a guard, a relay from \mathcal{G} is chosen with probability G_g/B_g . If a relay from \mathcal{G} is chosen, the probability that it is controlled by an adversary who adds amount a of bandwidth to G is $\frac{a}{G+a}$; for relatively small values of a , we approximate this as $1/G$. We thus view the adversary’s advantage, in the sense we are considering, as $(G_g/B_g)(1/G)$. Similarly, for

adding bandwidth to \mathcal{E} and being selected for the exit position, we view the adversary’s advantage as $(E_e/B_e)(1/E)$. For adding bandwidth to \mathcal{D} , we consider both sensitive positions, and we add $(D_g/B_g)(1/D)$ and $(D_e/B_e)(1/D)$. This gives the following definitions:

$$V_G = \frac{G_g}{B_g} \frac{1}{G} \tag{1}$$

$$V_E = \frac{E_e}{B_e} \frac{1}{E} \tag{2}$$

$$V_D = \left(\frac{D_g}{B_g} + \frac{D_e}{B_e} \right) \frac{1}{D} \tag{3}$$

In constructing our goal lists, we seek to minimize the maximum value of any V_C , *i.e.*, to minimize the expected advantage (under these metrics) for an adversary who might add a little additional bandwidth in the hope of increasing its chance of being chosen in a sensitive position (guard or exit). Minimizing the maximum V_C value promotes security, in particular with respect to an adversary who might add bandwidth to the network in order to be more likely to be selected in a sensitive position.

5.3 Alt1 and Alt2 goals and allocations

We next present two candidate sets of goals for positional weighting, the Alt1 and Alt2 goals, and derive weights that achieve them.

5.3.1 Alt1. We state the Alt1 goals before discussing their motivation. They are, in order starting with the most important,

- (1) Maximize the bandwidth of the minimum-bandwidth position
- (2) Maximize the amount of class- D bandwidth put into position e , *i.e.*, maximize D_e
- (3) Minimize the maximum V_C value
- (4) Minimize the second-largest V_C value
- (5) Maximize the bandwidth in the position with the second-smallest bandwidth

Goal 1 promotes Tor performance; it seeks to ensure that the bottleneck position has as much bandwidth as possible. One implication of this goal is that, if it is possible to balance the positions by allocating $T/3$ bandwidth to each, that should be done.

Goal 2 puts as much bandwidth from \mathcal{D} relays into the exit position as is possible. This goal, which will not be part of the Alt2 goals, might be adopted if relays in \mathcal{D} (which have the Guard and Exit flags) are viewed as more trustworthy than relays in \mathcal{E} (which have the Exit flag but have not had sufficient uptime or bandwidth to earn the Guard flag). The goal then promotes security through the use of these more trustworthy relays in the sensitive exit position. By contrast, the difference between relays in \mathcal{D} and in \mathcal{G} is in how they set their exit policies. We do not view that as suggesting different levels of trustworthiness, and so we do not consider an analogous goal for the guard position. This goal introduces some asymmetry into the goals and thus provides a point of contrast with the Alt2 goals.

Goal 3 minimizes the advantage, as measured by the V_C values discussed above, that an attacker could gain by adding a little extra bandwidth to the network. Goal 4 minimizes this advantage with respect to the relay class that gives the second-largest advantage.

Goal 5 seeks to give the position with the second-least amount of bandwidth as much bandwidth as possible. Because this position is not the bottleneck position (covered by Goal 1), we put this goal below all the others in this list. Balancing the positions automatically satisfies this goal as well as Goal 1.

We present in Section 5.3.3 the bandwidth allocations that are induced by these goals. We will refer to this as the Alt1 allocation. Theorem 5.1 shows that those weights satisfy the Alt1 goals. It also shows that these goals fully determine the weights, i.e., that they are the unique solution. We do not show continuity but believe that it holds as a consequence of our principled goals-based approach.

THEOREM 5.1. *The case analysis in Section 5.3.3 satisfies the following properties:*

- (1) For each network state (G, M, E, D) , the case analysis defines positional weights $\{W_{xy} | x \in \{g, m, e\}, y \in \{\mathcal{G}, \mathcal{M}, \mathcal{E}, \mathcal{D}\}\}$ (the Alt1 weights) with $\sum_x W_{xy} = 1$ for every $y \in \{\mathcal{G}, \mathcal{M}, \mathcal{E}, \mathcal{D}\}$.
- (2) For each network state (G, M, E, D) , the bandwidth allocation (the Alt1 allocation) corresponding to the Alt1 weights satisfies the Alt1 ordered goals to the greatest extent possible. The Alt1 allocation from class y to position x is given by the Alt1 weight W_{xy} times the total bandwidth in the y relays.
- (3) Any other bandwidth allocation is strictly worse at satisfying the ordered Alt1 goals than the Alt1 allocation is.

We leave the proof of Theorem 5.1 to Appendix B

5.3.2 Alt2 goals. We turn now to the goals that define our Alt2 approach. We start by stating the ordered list of goals.

- (1) Maximize the bandwidth of the minimum-bandwidth position
- (2) Minimize the maximum V_C value
- (3) Maximize the bandwidth in the position with the second-smallest bandwidth

This is a sublist of the Alt1 ordered list of goals that in particular does not prefer the \mathcal{D} relays for the exit position. Note that it does not have any goal that imposes asymmetry. We present weights in Appendix C that Theorem 5.2 proves satisfy these goals. We do not prove uniqueness or continuity of these weights but believe that they hold.

THEOREM 5.2. *The case analysis in Appendix C satisfies the following properties:*

- (1) For each network state (G, M, E, D) , the case analysis defines positional weights $\{W_{xy} | x \in \{g, m, e\}, y \in \{\mathcal{G}, \mathcal{M}, \mathcal{E}, \mathcal{D}\}\}$ (the Alt2 weights) with $\sum_x W_{xy} = 1$ for every $y \in \{\mathcal{G}, \mathcal{M}, \mathcal{E}, \mathcal{D}\}$.
- (2) For each network state (G, M, E, D) , the bandwidth allocation (the Alt2 allocation) corresponding to the Alt2 weights satisfies the Alt2 ordered goals to the greatest extent possible. The Alt2 allocation from class y to position x is given by the Alt2 weight W_{xy} times the total bandwidth in the y relays.

A proof of Thm. 5.2 appears in Appendix D.

5.3.3 Alt1 case analysis. We now present a case analysis, based on the values (G, M, E, D) and $T = G + M + E + D$, that gives the C_p and V_C values for any network state (G, M, E, D) . We start with some observations.

If $M > T/3$ (all subcases of Case 1), then we cannot balance the positions because all M bandwidth must be put into position m . If $M \leq T/3$ (Case 2) and $E + D < T/3$ (Case 2a) or $G + D < T/3$ (Case 2b), then we cannot balance the position because we can neither allocate more than $E + D$ to e nor allocate more than $G + D$ to g . (If $M \leq T/3$, then at least one of these sums must be at least $T/3$, otherwise we would have $T = G + M + E + D \leq G + D + M + E + D < T/3 + T/3 + T/3$.)

If $M \leq T/3$, $E + D \geq T/3$, and $G + D \geq T/3$, then we can always balance the positions (subcases of Case 2c). Within these subcases, the values of V_C are used when $D > T/3$ and $D + G > 2T/3$ (in Cases 2(c)iiB.I and 2(c)iiB.II). These conditions ensure that e can be allocated $T/3$ from D and that, between the G and remaining D bandwidth, there is more than enough bandwidth to allocate $T/3$ to g . This then raises the question of how much D (and thus G) should be allocated to g and how much should be allocated to m ; this question is answered by considering V_G and V_D .

Here, we partition the space of (G, M, E, D) tuples using conditions on their constituent values. For each case, we present the non-zero weights W_{xy} for that case and indicate which lemma in Appendix B proves that these are the unique weights induced by the Alt1 goals for this case. Note that the proofs are given in terms of bandwidth allocations, but the weights can be derived by dividing by the bandwidth of the appropriate class.

(1) $M > T/3$

(a) $G \geq M$ As shown in Lem. B.1, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------|------------------|----------|
| \mathcal{G} | $\frac{G+M}{2G}$ | $\frac{G-M}{2G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

(b) $G < M$

(i) $E \geq M$ As shown in Lem. B.2, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|------------------|------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{E-M}{2E}$ | $\frac{E+M}{2E}$ |
| \mathcal{D} | 1 | | |

(ii) $E < M$

(A).(I) $(G \geq E) \wedge (G \geq E + D)$ As shown in Lem. B.3, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|----------|----------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

(A).(II) $(G \geq E) \wedge (G < E + D)$ As shown in Lem. B.4, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------|----------|--------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{D+E-G}{2D}$ | | $\frac{D+G-E}{2D}$ |

(B).(I) $(G < E) \wedge (E \geq G + D)$ As shown in Lem. B.5, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|----------|----------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | 1 | | |

(B).(II) $(G < E) \wedge (E < G + D)$ As shown in Lem. B.6, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------|----------|--------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{D+E-G}{2D}$ | | $\frac{D-E+G}{2D}$ |

(2) $M \leq T/3$

(a) $E + D < T/3$ As shown in Lem. B.7, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------|------------------|----------|
| \mathcal{G} | $\frac{G+M}{2G}$ | $\frac{G-M}{2G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

(b) $G + D < T/3$ As shown in Lem. B.8, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|------------------|------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{E-M}{2E}$ | $\frac{E+M}{2E}$ |
| \mathcal{D} | 1 | | |

(c) $(E + D \geq T/3) \wedge (G + D \geq T/3)$

(i) $D \leq T/3$

(A) $G < T/3$ As shown in Lem. B.12, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|-------------------|---------------------------|-----------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $1 - \frac{2T-3D-3G}{3E}$ | $\frac{2T-3D-3G}{3E}$ |
| \mathcal{D} | $\frac{T-3G}{3D}$ | | $1 - \frac{T-3G}{3D}$ |

(B) $G \geq T/3$ As shown in Lem. B.9, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------------|-----------------------|-------------------|
| \mathcal{G} | $\frac{T}{3G}$ | $1 - \frac{T}{3G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $1 - \frac{T-3D}{3E}$ | $\frac{T-3D}{3E}$ |
| \mathcal{D} | | | 1 |

(ii) $D > T/3$

(A) $G + D \leq 2T/3$ As shown in Lem. B.12, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|-------------------|---------------------------|-----------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $1 - \frac{2T-3D-3G}{3E}$ | $\frac{2T-3D-3G}{3E}$ |
| \mathcal{D} | $\frac{T-3G}{3D}$ | | $1 - \frac{T-3G}{3D}$ |

(B).(I) $G + D > 2T/3, D \geq G$ As shown in Lem. B.10, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------------------------------------------|---------------------------------------------------|----------------|
| \mathcal{G} | $\frac{2T}{3(D+G)}$ | $1 - \frac{2T}{3(D+G)}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | 1 | |
| \mathcal{D} | $\frac{T}{3D} - \frac{2T}{3D} \frac{G}{G+D}$ | $1 - \frac{2T}{3D} + \frac{2T}{3D} \frac{G}{G+D}$ | $\frac{T}{3D}$ |

(B).(II) $G + D > 2T/3, D < G$ As shown in Lem. B.11, the weights for this case are:

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------------|--------------------|----------------|
| \mathcal{G} | $\frac{T}{3G}$ | $1 - \frac{T}{3G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | 1 | |
| \mathcal{D} | | $1 - \frac{T}{3D}$ | $\frac{T}{3D}$ |

5.4 Comparison of alternative allocations

We now compare results from the three approaches we consider—Tor’s current approach, Alt1, and Alt2—under different potential network states. We use the tuple (G, M, E, D) as the network state.

First, if the network state is $(1, 4, 1, 1)$,¹ then Tor’s current approach does not produce valid weights (yielding a negative value for W_{gD}). In this state, Alt1 and Alt1 produce the same results; these are illustrated in Figure. 3. This illustrates that all of the class- M bandwidth is allocated to m ; this is more than $T/3$ (horizontal dashed line). The positions g and e each get the same amount of bandwidth. We note that the stopgap fix identified in Section 5.1 produces the same bandwidth allocation as Alt1 and Alt2.

Second, we consider the network state $(50, 120, 40, 90)$. As noted in Section 3, Tor’s current bandwidth weights allocate too little to the g position in this network state, leaving a smaller bottleneck than is necessary. The stopgap fix suggested in Section 5.1, Alt1, and Alt2 all produce the same bandwidth allocation for this network state. That allocation is shown in Figure 4.

The network state $(1, 1, 1, 1)$ is one instance in which all three approaches produce different results. Figure 5 shows the three different allocations. Tor’s current allocation (top) spreads D evenly across all positions, with each other class going to its “natural” position. Alt1 (middle) uses the amount of D that is required to bring g to $T/3$ and the puts the rest of D into e . This determines all bandwidth allocations before considering V_C values. Alt2 (bottom) puts some of each non- M class into m to reduce the V_C values.

¹Note that, in all approaches we consider, the case covering a network state and the bandwidth weights W_{xy} are unchanged when the values of (G, M, E, D) are rescaled to $(k \cdot G, k \cdot M, k \cdot E, k \cdot D)$ for any $k > 0$. Here, for example, we could consider $(50, 200, 50, 50)$ instead.

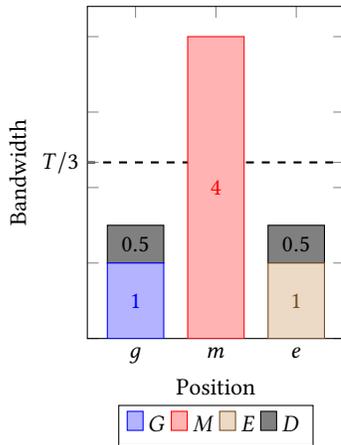


Figure 3: Bandwidth allocation given by both Alt1 and Alt2 for $(G, M, E, D) = (1, 4, 1, 1)$. Tor’s current algorithm does not produce valid weights in this case; our stopgap fix produces these weights as well.

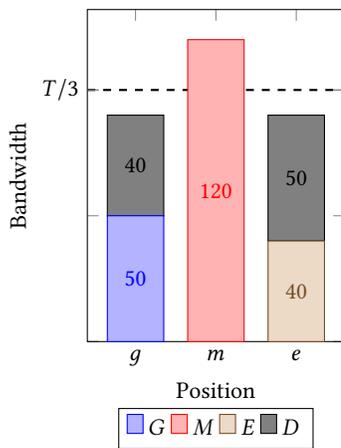


Figure 4: Bandwidth allocation given by Alt1, Alt2, and our stopgap corrections for $(G, M, E, D) = (50, 120, 40, 90)$. Compare with Figure 2, which shows the allocation given by Tor’s current algorithm.

Most network states that have appeared in the Tor network fall into Case 3(b)ii (in the current Tor case breakdown). The following proposition shows that all three of the approaches we consider produce the same bandwidth allocation; we defer its proof to Appendix E.

PROPOSITION 5.3. *If (G, M, E, D) is a network state with $E + D < T/3$, $G \geq T/3$, and $G \geq M$ then the bandwidth allocation produced by Tor’s rules, Alt1, and Alt2 is always given by*

| C | g | m | e |
|-----|-----------------|-----------------|-----|
| G | $\frac{G+M}{2}$ | $\frac{G-M}{2}$ | |
| M | | M | |
| E | | | E |
| D | | | D |

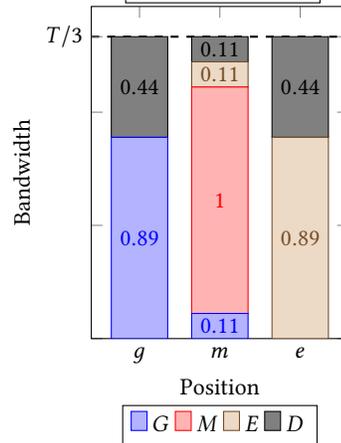
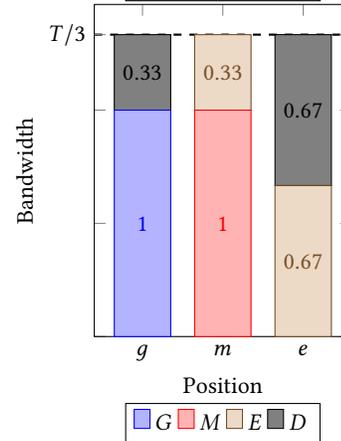
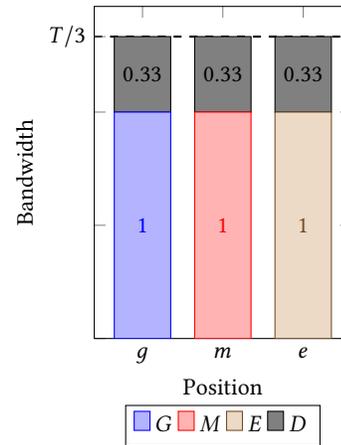


Figure 5: Bandwidth allocations given by (top to bottom) Tor’s current approach, Alt1, and Alt2 for network state $(G, M, E, D) = (1, 1, 1, 1)$.

Occasionally, network states have appeared in the Tor network that fall into Case 3(a)iiB (in the current Tor case breakdown). By contrast to those states that fall into Case 3(b)ii, this case produces

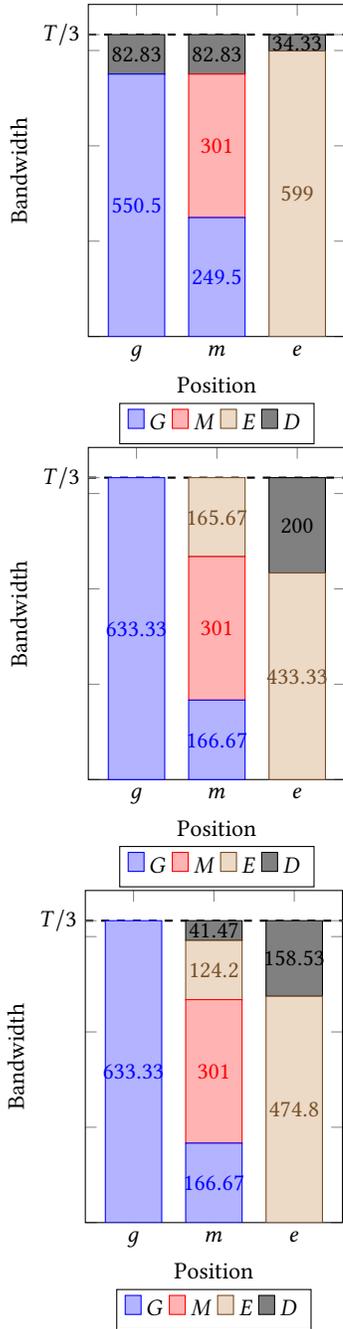


Figure 6: Bandwidth allocations produced by Tor’s current rules (top), Alt1 (middle), and Alt2 (bottom) as applied to the network state (800, 301, 599, 200).

different bandwidth allocations for all three approaches. An example network state covered in this case is (800, 301, 599, 200). Figure 6 shows the differing allocations for this network state from the three different sets of positional weights.

6 DISCUSSION

Recommendations. In April 2022, we disclosed to the Tor Project the deficiencies in the position weights that we have identified. Tor developers had already noticed that the applicable weights case sometimes produces invalid weights [21]. However, they did not identify exactly when the case occurs, which we characterize as case 2(b)iiiB. Moreover, they did not produce a solution to the problem.

Our immediate recommendation is for Tor to adopt one of the alternatives we introduce. Because position weights are defined in the consensus and must be agreed upon by at least half of the Directory Authorities, updating them should be done through the creation of a new consensus version (currently at 32 [29]), which will then start to be used when the majority supports it. Each of the alternatives has advantages over Tor’s current positional weighting, but between them we recommend Alt1. As noted above, Alt1 maximizes the amount of class- D bandwidth in the exit position, after maximizing the bandwidth of the minimum-bandwidth position. Numerous papers have noted that it is less of a persistent overhead commitment for adversaries to conduct attacks if their relays do not require the Guard flag, and that this can be a factor in the likelihood of attempting an attack or of it succeeding [10, 11, 16]. Since D relays must have the Guard flag, Alt1 will guarantee that the highest proportion of exit traffic possible will go through Guard-flagged relays (consistent with maximizing minimal positional-bandwidth). While this means that an adversary willing to put in the persistent overhead will at that point need to have to compromise less overall bandwidth versus Alt2, that is only because the exits without the stability and minimum capacity to qualify as guards are chosen less frequently. In other words, an adversary willing to persistently contribute substantial resources to the Tor network can do better, which is a long-recognized reality about Tor.

Future compatibility and limitations. Tor Proposal 265 [23] outlines another approach to improving position weighting. In that proposal, “overhead” parameters are used to incorporate traffic from sources other than exit circuits into the balancing equations. Parameters $O_g, O_m \in [0, 1]$ indicate the fraction of guard and middle bandwidth, respectively, that is consumed by various types of non-exit traffic (e.g. onion-service traffic). To simplify the derivation of the load balancing equations, the D class is simply treated like the E class. However, the overhead notion used in the proposal is compatible with allowing the D relays to be used in the guard position, which gives more flexibility to achieve the goals targeted by the positional weights. Moreover, the weights we have proposed can generalize to incorporate the proposed overhead parameters.

This generalization can be accomplished by including the overhead factor in the goal sets. For the primary goal of maximizing the minimum-bandwidth position, the guard bandwidth and middle bandwidth need to be adjusted to $(1 - O_g)B_g$ and $(1 - O_m)B_m$, respectively. The resulting case analysis depends on the key value of $T/(1/(1 - O_g) + 1/(1 - O_m) + 1)$, instead of $T/3$, because it is an upper bound on exit-traffic throughput.

Our approach does have some limitations that may need to be addressed in the longer term. The solutions we identify are somewhat complicated, potentially making future design changes more

difficult to create and analyze. Simpler designs, such as only allowing a relay to be used in a single circuit position, may compensate for their decreased efficiency by being easier to modify for future needs. At the same time, our solutions are based on a simplified traffic model. Only exit traffic is considered (as discussed regarding Proposal 265), only expected throughput is considered, and no feedback on actual load and performance is taken into account. More sophisticated designs that use more complicated or dynamic traffic modeling [34] may obviate the use of the types of load equations used in this work.

Other onion-routing systems. Our analysis is somewhat specific to the Tor design. However, the issues Tor faces do arise in general for onion routing. Other networks based on onion-routing principles include Loki [13], I2P [27], and Orchid [3]. All these networks allow traffic to exit to an arbitrary destination, and so they face similar challenges as Tor when selecting routing paths. An adversary observing entry and exit nodes can perform traffic correlation attacks [16], serving as an exit node exposes the operator to abuse complaints, and bandwidth is a scarce resource and should be used efficiently. Indeed, several of these networks have implemented some of Tor’s solutions to these problems, including a limited set of exit nodes in Loki, designated entry nodes (or “fast peers”) in I2P, and weighted path selection in Orchid (although on the basis of a cryptocurrency “stake”). Any attempt to improve the throughput of these systems by balancing resource usage across circuit positions, as Tor has, would benefit from the problems and solutions we have identified.

7 RELATED WORK

Rochet and Pereira [25] use and modify the Tor position weights for their “waterfilling” algorithm for path selection. They propose per-relay position weights that do not change the total bandwidth each relay class allocates to each position under the existing W_{xy} equations. The per-relay values put more weight on lower-bandwidth guard and exit relays to require an adversary to compromise more of them to perform traffic analysis attacks. This waterfilling approach is fully compatible with our suggested improvements to Tor’s position weights.

Indeed, they complement well our recommendation. Alt1 ensures that as much of D as possible is allocated for use in the exit position. But of all the relay weight that is allocated for exit usage, waterfilling will spread the selection across more relays. So the adversary will have to make the persistent network contribution noted in Section 6 to compromise as much exit bandwidth as without Alt1 weighting priorities. At the same time the adversary advantage of having to compromise less bandwidth (if more persistently), also noted in Section 6, is diminished by waterfilling in that more separate relays need to be compromised.

A waterfilling variation is also proposed that changes the position weights to make the amount of bandwidth allocated to the guard and exit positions as similar as possible while maximizing the bottleneck bandwidth. Weights are only derived under this goal for one network condition case. This goal improves the security of waterfilling path selection under their target threat model. This change would not be compatible with the position weights we have suggested, but it would constitute another possible set of goals

from which one could derive provably optimal position weights that cover all network cases.

Rochet et al. [26] propose the CLAPS scheme for Tor path selection which takes into account client location. In CLAPS, each relay may allocate its bandwidth across the circuit positions differently. Like Tor’s current per-class positional weights, the CLAPS per-relay weights are computed to optimize network throughput by maximizing the bandwidth in the bottleneck position. CLAPS is thus similar to our proposed weights in that it changes the way that relay bandwidth is allocated across positions, but it does so to allow location-aware path selection and constitutes a much bigger change from the way the Tor network currently operates.

Several works have looked at how to estimate Tor relay bandwidths [5, 6, 33]. This problem is complementary to the question of how to allocate that bandwidth across clients, circuits, and relays. Our work benefits from improved relay bandwidth estimates and makes better use of the measured bandwidths.

Our focus in this paper is on onion routing, and especially the positional weighting in Tor. Anonymous-communication protocols that are not based on onion routing typically do not have Tor’s combination of semi-dedicated node positions and positional weighting. Deployed internet mix networks have traditionally fallen into one of two possibilities. Message-based systems such as Mixminion [4] used free routes, in which any node can appear in any position. JonDonym [2, 17] supported web connections through a mix cascade, in which all communication followed the same path of nodes in the same order, though there were multiple cascades to choose from.

The more recent Nym system [19] is based on Loopix [24], which uses a stratified topology [9]. This topology requires communication be sent through nodes arranged in fixed layers (positions) where layer order is fixed for all communicants, but originators can choose any node within each layer. Loopix thus shares with Tor the random selection among nodes (relays) that have restricted positional order. Unlike Tor (and more like cascades), in all Loopix communication, nodes must always be used in the same position. Also unlike Tor, Loopix does not appear to use weighting for positions or for nodes within them.

8 CONCLUSION

We identify cases in which the positional weighting system in Tor fails to produce valid weights or fails to properly load balance across positions. We describe how these failures could allow an attacker to reduce Tor’s performance by inducing case changes by adding relays or changing their class. Our analysis of past Tor consensus shows that these attacks could have reduced the throughput of the network by 45% over a correctly load-balanced network. We then propose a new, systematic method to derive positional weights and propose two goal sets according to that method. We derive new sets of weights, prove that they satisfy these goal sets, and give examples of how they would change the weights from the current system. Tor could use our results to quickly fix the main deficiencies of its positional weights as well as adopt a better approach long-term.

ACKNOWLEDGMENTS

This work was supported by the Office of Naval Research.

REFERENCES

[1] Simurgh Aryan, Homa Aryan, and J Alex Halderman. 2013. Internet censorship in Iran: A first look. In *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*.

[2] Oliver Berthold, Hannes Federrath, and Stefan Köpsell. 2000. Web MIXes: A system for anonymous and unobservable Internet access. In *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*.

[3] Jake S. Cannell, Justin Sheek, Jay Freeman, Greg Hazel, Jennifer Rodriguez-Mueller, Brian J. Fox Eric Hou, and Steven Waterhouse. 2019. *Orchid: A Decentralized Network Routing Market (Version 2.0)*. Technical Report. <https://www.orchid.com/whitepaper/english.pdf>

[4] George Danezis, Roger Dingledine, and Nick Mathewson. 2003. Mixminion: Design of a Type III Anonymous Remailer Protocol. In *Proceedings, 2003 IEEE Symposium on Security and Privacy*. IEEE Computer Society, Berkeley, CA, 2–15.

[5] Hussein Darir, Geir Dullerud, and Nikita Borisov. 2023. ProbFlow: Using Probabilistic Programming in Anonymous Communication Networks. In *Proceedings of the 30th Annual Network and Distributed System Security Symposium (NDSS 2023)*. Internet Society.

[6] Hussein Darir, Hussein Sibai, Chin-Yu Cheng, Nikita Borisov, Geir E. Dullerud, and Sayan Mitra. 2022. MLEFlow: Learning from History to Improve Load Balancing in Tor. *Proc. Priv. Enhancing Technol.* 2022, 1 (2022), 75–104.

[7] Roger Dingledine and Nick Mathewson. 2006. Anonymity Loves Company: Usability and the Network Effect. In *Fifth Workshop on the Economics of Information Security (WEIS 2006)*, Ross Anderson (Ed.).

[8] Roger Dingledine and Nick Mathewson. Accessed February 2023. Tor Path Specification. <https://gitlab.torproject.org/tpo/core/torspec/-/blob/main/path-spec.txt>.

[9] Roger Dingledine, Vitaly Shmatikov, and Paul Syverson. 2005. Synchronous Batching: From Cascades to Free Routes. In *Privacy Enhancing Technologies: 4th International Workshop, PET 2004*, David Martin and Andrei Serjantov (Eds.). Springer-Verlag, LNCS 3424, 186–206.

[10] Tariq Elahi, Kevin Bauer, Mashaal AlSabah, Roger Dingledine, and Ian Goldberg. 2012. Changing of the Guards: A Framework for Understanding and Improving Entry Guard Selection in Tor. In *ACM Workshop on Privacy in the Electronic Society (WPES '12)*. ACM, Raleigh, North Carolina, USA.

[11] Aaron D. Jaggard and Paul Syverson. 2017. Onions in the Crosshairs: When The Man really is out to get you. In *ACM Workshop on Privacy in the Electronic Society (WPES '17)*. ACM, Dallas, Texas, USA.

[12] Rob Jansen, Tavish Vaidya, and Micah Sherr. 2019. Point Break: A Study of Bandwidth Denial-of-Service Attacks against Tor. In *USENIX security symposium*.

[13] Kee Jefferys, Simon Harman, Johnathan Ross, and Paul McLean. 2018. *Loki: Private transactions, decentralised communication (Version 3)*. Technical Report. https://loki.network/wp-content/uploads/2018/10/LokiWhitepaperV3_1.pdf

[14] Aaron Johnson, Aaron D. Jaggard, and Paul Syverson. 2023. Throwing Your Weight Around: Fixing Tor’s Positional Weighting. <https://github.com/ajohnson-nrl/tor-positional-weights>. Software artifact.

[15] Aaron Johnson, Rob Jansen, Aaron D. Jaggard, Joan Feigenbaum, and Paul Syverson. 2017. Avoiding The Man on the Wire: Improving Tor’s Security with Trust-Aware Path Selection. In *Proceedings of the 24th Network and Distributed System Security Symposium (NDSS 2017)*. Internet Society.

[16] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. 2013. Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (Berlin, Germany) (CCS '13)*. ACM, New York, NY, USA, 337–348. <https://doi.org/10.1145/2508859.2516651>

[17] JonDonym 2008. JonDonym – the internet anonymisation service. <https://www.jondos.de/en/>. Commercial version of the Java Anon Proxy (JAP).

[18] Marc Juarez, Rob Jansen, Rafael Galvez, Tariq Elahi, Claudia Diaz, and Matthew Wright. 02017. Poster: Fingerprinting Hidden Service Circuits from a Tor Middle Relay. In *IEEE Symposium on Security and Privacy*.

[19] nym Accessed February 2023. Nym: The Next Generation of Privacy Infrastructure. <https://nymtech.net/>.

[20] Lennart Oldenburg, Gunes Acar, and Claudia Diaz. 2022. From “Onion Not Found” to Guard Discovery. *Proceedings on Privacy Enhancing Technologies* 2022, 1 (2022).

[21] pastly. 2016. can’t create valid case 2b3 consens weight calculation. <https://gitlab.torproject.org/tpo/core/torspec/-/issues/2>. Accessed February 2023.

[22] Mike Perry. 2009. Torflow: Tor network analysis. *Proc. 2nd HotPETS* (2009).

[23] Mike Perry. 2016. Tor Proposal 265. <https://gitweb.torproject.org/torspec.git/tree/proposals/265-load-balancing-with-overhead.txt>. Accessed February 2023.

[24] Ania M. Piotrowska, Jamie Hayes, Tariq Elahi, Sebastian Meiser, and George Danezis. 2017. The Loopix Anonymity System. In *26th USENIX Security Symposium (USENIX Security 17)*.

[25] Florentin Rochet and Olivier Pereira. 2017. Waterfilling: Balancing the Tor network with maximum diversity. *Proceedings on Privacy Enhancing Technologies* 2017, 2 (2017), 4–22.

[26] Florentin Rochet, Ryan Wails, Aaron Johnson, Prateek Mittal, and Olivier Pereira. 2020. CLAPS: Client-Location-Aware Path Selection in Tor. In *Proceedings of the 27th ACM Conference on Computer and Communications Security (CCS '20)*. ACM.

[27] Lars Schimmer. 2009. Peer profiling and selection in the I2P anonymous network. In *PetCon 2009.1*. Technische Universität Dresden, 59–70.

[28] Paul Syverson. 2011. A Peel of Onion. In *Proceedings of 2011 Annual Computer Security Applications Conference (ACSAC '11)*, Orlando, FL, USA.

[29] Tor. Accessed February 2023. Consensus Health. <https://consensus-health.torproject.org/>.

[30] Tor. Accessed February 2023. Tor directory protocol, version 3. <https://gitlab.torproject.org/tpo/core/torspec/-/blob/main/dir-spec.txt>.

[31] Tor. Accessed February 2023. Tor Metrics. <https://metrics.torproject.org/>.

[32] Tor. Accessed February 2023. Tor node selection code. https://gitlab.com/torproject/tor/-/blob/master/src/feature/modelist/node_select.c.

[33] Matthew Traudt, Rob Jansen, and Aaron Johnson. 2021. FlashFlow: A secure speed test for Tor. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 381–391.

[34] Tao Wang, Kevin Bauer, Clara Forero, and Ian Goldberg. 2012. Congestion-aware path selection for Tor. In *Financial Cryptography and Data Security (FC 2012)*.

A CURRENT TOR WEIGHTS

We express the current Tor weights as a sequence of mutually exclusive cases covering the possible network states. The case division largely follows the one that appears in the Tor specifications [30]. However, we do further divide some cases so that each case contains a full set of weights and to characterize the cases that fail to maximize throughput. The positional weights in both cases are as follows, where column x and row y contains W_{xy} :

(1) $E \geq T/3 \wedge G \geq T/3$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------|---------------------|--------------------|
| \mathcal{G} | $\frac{G+E+M}{3G}$ | $\frac{2G-E-M}{3G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{2E-G-M}{3E}$ | $\frac{E+G+M}{3E}$ |
| \mathcal{D} | $\frac{1}{3}$ | $\frac{1}{3}$ | $\frac{1}{3}$ |

(2) $E < T/3 \wedge G < T/3$. Let $R = \min(G, E)$ and $S = \max(G, E)$.

- (a) $R + D < S$
- (i) $G > E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|----------|----------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

(ii) $G \leq E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|----------|----------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | 1 | | |

(b) $R + D \geq S$

- (i) $G \geq M \wedge E \geq G - M$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------------|--------------------------|--------------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{G-M}{E}$ | $\frac{E-G+M}{E}$ |
| \mathcal{D} | $\frac{2D+2E-4G+2M}{6D}$ | $\frac{2D+2E-4G+2M}{6D}$ | $\frac{2D+2E-4G+2M}{6D}$ |

(ii) $(G < M \vee E < G - M) \wedge (M \leq T/3)$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|-----------------------|-----------------------|-----------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{D+E-2G+M}{3D}$ | $\frac{D-2M+G+E}{3D}$ | $\frac{D-2E+G+M}{3D}$ |

 (iii) $(G < M \vee E < G - M) \wedge (M > T/3)$. This case is implied by the simpler condition of $M > T/3$ given the earlier requirement that $G < T/3$.

 (A) $T/3 - E \leq D$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------------|----------|-----------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{2D+2E-G-M}{3D}$ | | $\frac{D-2E+G+M}{3D}$ |

 (B) $T/3 - E > D$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------------|----------|-----------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{2D+2E-G-M}{3D}$ | | $\frac{D-2E+G+M}{3D}$ |

 (3) $E < T/3 \oplus G < T/3$. Let $S = \min(G, E)$.

 (a) $S + D < T/3$

 (i) $G < E$

 (A) $E < M$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|----------|----------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | 1 | | |

 (B) $E \geq M$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|------------------|------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{E-M}{2E}$ | $\frac{E+M}{2E}$ |
| \mathcal{D} | 1 | | |

 (ii) $G \geq E$

 (A) $G < M$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|----------|----------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

 (B) $G \geq M$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------|------------------|----------|
| \mathcal{G} | $\frac{G+M}{2G}$ | $\frac{G-M}{2G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

 (b) $S + D \geq T/3$

 (i) $G < E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|-----------------------|------------------------|------------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{E-M}{2E}$ | $\frac{E+M}{2E}$ |
| \mathcal{D} | $\frac{D-2G+E+M}{3D}$ | $\frac{2D+2G-E-M}{6D}$ | $\frac{2D+2G-E-M}{6D}$ |

 (ii) $G \geq E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------------|------------------------|-----------------------|
| \mathcal{G} | $\frac{G+M}{2G}$ | $\frac{G-M}{2G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{2D+2E-G-M}{6D}$ | $\frac{2D+2E-G-M}{6D}$ | $\frac{D-2E+G+M}{3D}$ |

B PROOF OF THEOREM 5.1

We prove Theorem 5.1 by case analysis. Each lemma below shows that, in one or more cases, the Alt1 goals are best satisfied by a unique bandwidth allocation and this bandwidth allocation is as given in the tables here. From these bandwidth-allocation tables, the weights in Section 5.3.3 can be derived. By construction, the cases presented in Section 5.3.3 partition the entire space. We note that Lemma B.12 applies to multiple cases in Section 5.3.3.

Each table shows the allocation of bandwidth from a class (rows) to a position (columns). The rightmost column shows the values V_C for the applicable classes. Even when these are not needed to determine bandwidth allocations, these values may be of interest. The bottom row shows the total bandwidth allocated to each position.

B.1 Subcases of Case 1 ($M > T/3$)

LEMMA B.1. *If $M > T/3$ and $G \geq M$, allocate bandwidth as:*

| C | g | m | e | V_C |
|--------------|-----------|-----------|-------|-----------|
| G | $(G+M)/2$ | $(G-M)/2$ | | $1/G$ |
| M | | M | | -- |
| E | | | E | $1/(E+D)$ |
| D | | | D | $1/(E+D)$ |
| Total | $(G+M)/2$ | $(G+M)/2$ | $E+D$ | |

PROOF. We must have $E+D < T/3 < M \leq G$. We will be able to give at least $T/3$ to both g and m , so e will be the smallest position; to maximize its value, we must put all D and all E into e . This satisfies the first and second goals. The division of G between g and m will be guided by the remaining goals.

We have $V_G = (G_g/B_g)(1/G) = 1/G$, $V_E = (E_e/B_e)(1/E) = 1/B_e = 1/(E+D)$, and $V_D = ((D_g/B_g)+(D_e/B_e))(1/D) = (D/B_e)(1/D) = 1/B_e = 1/(E+D)$. The bandwidth of e is determined in achieving

the first goal, and the value of G is not under our control, so the third goal does not influence the allocation. The fourth goal means we should balance g and m if possible; because $G \geq M$, we can do this, assigning $(G + M)/2$ to each $((G + M)/2)$ from G to g , all of M to m , and $(G - M)/2$ from G to m . \square

LEMMA B.2. *If $M > T/3$ and $E \geq M > G$, allocate bandwidth as:*

| C | g | m | e | V_C |
|-------|---------|-------------|-------------|-------------|
| G | G | | | $1/(G + D)$ |
| M | | M | | -- |
| E | | $(E - M)/2$ | $(E + M)/2$ | $1/E$ |
| D | D | | | $1/(G + D)$ |
| Total | $G + D$ | $(E + M)/2$ | $(E + M)/2$ | |

PROOF. We must have $G + D < T/3 < M \leq E$. We will be able to give at least $T/3$ to both e and m , so g will be the smallest position; to maximize its value, we must put all D and all G into g . This satisfies the first and (trivially) second goals. The division of E between e and m will be guided by the remaining goals.

We have $V_G = (G_g/B_g)(1/G) = 1/B_g = 1/(G + D)$, $V_E = (E_e/B_e)(1/E) = 1/E$, and $V_D = (D_g/B_g + D_e/B_e)(1/D) = (D/B_g)(1/D) = 1/B_g = 1/(G + D)$; at this point, we cannot change these values, so we move to the fourth goal. That dictates that we balance m and e , which requires that, after we allocate all of M to m , we allocate $(E + M)/2$ from E to e and $(E - M)/2$ of E to m . \square

LEMMA B.3. *If $M > T/3$, $M > G \geq E + D$, allocate bandwidth as:*

| C | g | m | e | V_C |
|-------|-----|-----|---------|-------------|
| G | G | | | $1/G$ |
| M | | M | | -- |
| E | | | E | $1/(E + D)$ |
| D | | | D | $1/(E + D)$ |
| Total | G | M | $E + D$ | |

PROOF. We cannot give more than $E + D < T/3$ to e . Assigning all of D to e and all of E to e is the only way to achieve this, and it does not force the other positions to be smaller than e , so we make this allocation. This addresses the first two goals.

We have $V_D = (D/B_e)(1/D) = 1/B_e = 1/(E + D)$, $V_E = (E/B_e)(1/E) = 1/B_e = 1/(E + D)$, and $V_G = (G_g/B_g)(1/G) = 1/G$; these values are independent of any remaining choices we might make, so we move to the fourth goal. Because $G < M$, this requires us to allocate all of G to g . \square

Lemmas B.4 and B.6 are similar; they could be viewed as a single case (the scarcer or E and G , combined with D , is at least as large as the more plentiful of E and G).

LEMMA B.4. *If $M > T/3$, $M > G \geq E$, and $E + D > G$, allocate bandwidth as:*

| C | g | m | e | V_C |
|-------|-----------------|-----|-----------------|-----------------|
| G | G | | | $2/(G + D + E)$ |
| M | | M | | -- |
| E | | | E | $2/(G + D + E)$ |
| D | $(D + E - G)/2$ | | $(D + G - E)/2$ | $2/(G + D + E)$ |
| Total | $(D + E + G)/2$ | M | $(D + E + G)/2$ | |

PROOF. We can balance e and g at $(1 - M)/2 = (G + D + E)/2$, achieving the first goal. Doing so completely determines the allocation as follows: allocate all E to e , all G to g , and all M to m . Because $E + D > G$, we divide D between e and g , assigning $(D + G - E)/2$ from D to e and $(D + E - G)/2$ from D to g . \square

For completeness, we calculate the V_C values for each class C . $V_G = (G_g/B_g)(1/G) = 1/B_g = 2/(G + D + E)$. $V_E = (E_e/B_e)(1/E) = 1/B_e = 2/(G + D + E)$. $V_D = ((D_g/B_g) + (D_e/B_e))(1/D) = ((D + E - G)/(D + E + G) + (D + G - E)/(D + E + G))(1/D) = 2/(D + E + G)$.

LEMMA B.5. *If $M > T/3$, $M > E > G$, and $E \geq G + D$, allocate bandwidth as:*

| C | g | m | e | V_C |
|-------|---------|-----|-----|-------------|
| G | G | | | $1/(G + D)$ |
| M | | M | | -- |
| E | | | E | $1/E$ |
| D | D | | | $1/(G + D)$ |
| Total | $G + D$ | M | E | |

PROOF. Because $G + D \leq E < M$, we must allocate all of G and D to g in order to achieve the first goal. The second goal is then irrelevant. We then turn to the third goal; we may divide E between e and m to achieve it as long as we allocate at least $G + D$ of E to e (to avoid making e smaller than g).

We have $V_G = (G_g/B_g)(1/G) = (G/(G + D))(1/G) = 1/(G + D)$, $V_D = ((D_g/B_g) + (D_e/B_e))(1/D) = (D/(G + D) + 0)(1/D) = (1/G + D)$, and $V_E = (E_e/B_e)(1/E) = (E_e/E_e)(1/E) = 1/E$. These are all independent of how we divide E between e and m , so we consider the fourth goal. In order to maximize e , we put all of E into e . \square

LEMMA B.6. *If $M > T/3$, $M > E > G$, and $E < G + D$, allocate bandwidth as:*

| C | g | m | e | V_C |
|-------|-----------------|-----|-----------------|-----------------|
| G | G | | | $2/(G + D + E)$ |
| M | | M | | -- |
| E | | | E | $2/(G + D + E)$ |
| D | $(D + E - G)/2$ | | $(D - E + G)/2$ | $2/(G + D + E)$ |
| Total | $(G + D + E)/2$ | M | $(G + D + E)/2$ | |

PROOF. The bandwidth allocated to at least one of g and e must be no bigger than $(T - M)/2 < T/3$; we will allocate this much bandwidth to each of those positions in order to satisfy the first goal, and that will determine the entire allocation.

Because $G < E < M$, we have $G < (T - M)/2$; because $E < G + D$, we have $E < (T - M)/2$. We thus put all G into g and all E into

e. We split D to balance these positions: put $(T - M - 2G)/2 = (D + E - G)/2 < D$ from D into g , and put the remaining $(2D - T + M + 2G)/2 = (D - E + G)/2$ from D into e . This results in $(G + D + E)/2$ bandwidth in g and $(G + D + E)/2$ bandwidth in e . \square

For completeness, we compute V_C for the various classes C . We have $V_G = (G_g/B_g)(1/G) = (G/B_g)(1/G) = 1/B_g = 2/(G + D + E)$. We have $V_D = ((D_g/B_g) + (D_e/B_e))(1/D) = ((D + E - G)/(G + D + E) + (D - E + G)/(G + D + E))(1/D) = 2/(G + D + E)$. We have $V_E = (E_e/B_e)(1/E) = (E/B_e)(1/E) = 2/(G + D + E)$.

B.2 Subcases of Case 2 when unable to balance

Here, $M \leq T/3$. If either $E + D < T/3$ or $G + D < T/3$, then we cannot balance the positions. Note that, if $M \leq T/3$, we cannot have both of those conditions simultaneously.

LEMMA B.7. *If $M \leq T/3$ and $E + D < T/3$, allocate bandwidth as:*

| C | g | m | e | V_C |
|--------------|-----------------|-----------------|---------|-----------------|
| G | $\frac{G+M}{2}$ | $\frac{G-M}{2}$ | | $\frac{1}{G}$ |
| M | | M | | -- |
| E | | | E | $\frac{1}{E+D}$ |
| D | | | D | $\frac{1}{E+D}$ |
| Total | $\frac{G+M}{2}$ | $\frac{G+M}{2}$ | $E + D$ | |

PROOF. Because we then have $G > T/3$ and $G + M > 2T/3$, we can ensure g and m are both allocated at least $T/3$, and e will be the smallest position. To satisfy the first goal, we allocate all of D and all of E to e ; this also satisfies the second goal. For the third goal, we compute the various V_C values. We have $V_G = (G_g/B_g)(1/G) = (B_g/B_g)(1/G) = 1/G$. We have $V_E = (E/(E+D))(1/E) = 1/(E+D)$. We have $V_D = (0 + D/(E + D))(1/D) = 1/(E + D)$. These are all independent of any allocation decisions we might make, so we proceed to the fourth goal. We will achieve that by balancing the allocations to g and m , which we can do because $G \geq M$. We allocate $(G + M)/2$ of G to g and the rest of G (i.e., $(G - M)/2$) to m . \square

LEMMA B.8. *If $M \leq T/3$ and $G+D < T/3$, then allocate bandwidth as:*

| C | g | m | e | V_C |
|--------------|---------|-----------------|-----------------|-----------------|
| G | G | | | $\frac{1}{G+D}$ |
| M | | M | | -- |
| E | | $\frac{E-M}{2}$ | $\frac{E+M}{2}$ | $\frac{1}{E}$ |
| D | D | | | $\frac{1}{G+D}$ |
| Total | $G + D$ | $\frac{E+M}{2}$ | $\frac{E+M}{2}$ | |

PROOF. Because $G + D < T/3$ and $E > T/3$, g will be the position with the smallest allocation. We allocate all of G and all of D to g ; this satisfies the first goal and leaves nothing to do for the second goal.

We consider the V_C values. We have $V_G = \frac{G}{G+D} \frac{1}{G} = \frac{1}{G+D}$ and $V_D = \left(\frac{D}{G+D} + 0\right) \frac{1}{D} = \frac{1}{G+D}$. For V_E , we have $E_e = B_e$, and thus $V_E = 1/E$. These are all independent of the allocation of E between e and m , so we turn to the fourth goal and balance these two positions by allocating $(E + M)/2$ from E to e and $(E - M)/2$ from E to m . \square

B.3 Subcases of Case 2 when able to balance

We now turn to cases where $M \leq T/3$ and we have both $E + D \geq T/3$ and $G + D \geq T/3$. When these conditions hold, we can balance all three positions.

LEMMA B.9. *If $M \leq T/3$, $G \geq T/3$, $D \leq T/3$, and $E + D \geq T/3$, allocate bandwidth as:*

| C | g | m | e | V_C |
|--------------|---------------|----------------------|------------------|--------------------------|
| G | $\frac{T}{3}$ | $G - \frac{T}{3}$ | | $\frac{1}{G}$ |
| M | | M | | -- |
| E | | $E - \frac{T-3D}{3}$ | $\frac{T-3D}{3}$ | $\frac{T-3D}{E \cdot T}$ |
| D | | | D | $\frac{3}{T}$ |
| Total | $T/3$ | $T/3$ | $T/3$ | |

PROOF. We will be able to balance all positions. To satisfy the second goal, we allocate all of D to e , and we allocate $T/3 - D$ from E to e . We allocate the rest of E (i.e., $E - (T - 3D)/3$) to m to avoid having more than $T/3$ in e . We must allocate $T/3$ from G to g , and we allocate the rest of G to m . \square

For completeness, we compute the various V_C values. We have $V_G = ((T/3)/(T/3))(1/G) = 1/G$. We have $V_E = \frac{T-3D}{3} \frac{1}{E} = \frac{T-3D}{T \cdot E}$. We have $V_D = (3D/T)(1/D) = 3/T$.

LEMMA B.10. *If $M \leq T/3$, $D > T/3$, $D + G > 2T/3$, and $D \geq G$, then allocate bandwidth as:*

| C | g | m | e | V_C |
|--------------|----------------------------------------------|---------------------------------------------------|-------|-----------------|
| G | $\frac{2T-G}{3} \frac{G}{D+G}$ | $G - \frac{2T-G}{3} \frac{G}{D+G}$ | | $\frac{2}{D+G}$ |
| M | | M | | -- |
| E | | E | | 0 |
| D | $\frac{T}{3} - \frac{2T-G}{3} \frac{G}{D+G}$ | $D - \frac{2T}{3} + \frac{2T-G}{3} \frac{G}{D+G}$ | $T/3$ | $\frac{2}{D+G}$ |
| Total | $T/3$ | $T/3$ | $T/3$ | |

PROOF. We will balance all three positions, satisfying the first goal. We allocate $T/3$ from D to e to satisfy the second goal. We now consider the various V_C values. $V_G = (G_g/B_g)(1/G) = \frac{3G_g}{T \cdot G}$. We have $V_E = (E_e/B_e)(1/E) = 0$. We have $V_D = ((D_e/B_e) + (D_g/B_g))(1/D) = (1 + (D_g/B_g))(1/D) = (1 + \frac{T/3 - G_g}{T/3}) \frac{1}{D} = \frac{2-3G_g}{D}$ (assuming $D_g + G_g = B_g = T/3$). We consider the value for G_g in $[0, \frac{T}{3}]$ that will minimize the larger of V_G and V_D . Taking $D_g + G_g = T/3$ (which is possible because $D + G \geq 2T/3$), we have that V_G increases with G_g and V_D decreases with $G_g = T/3 - D_g$. Solving $V_G = V_D$, we obtain $D_g = \frac{T}{3} - G_g = \frac{T}{3} - \frac{2T}{3} \frac{G}{D+G}$. For this value to be non-negative, we need $\frac{T}{3} \geq \frac{2T}{3} \frac{G}{D+G}$, which holds if, and only if, $D \geq G$ (assuming $D > 0$, which holds by assumption). Thus, as long as D is at least G and is also at least $\frac{T}{3} - \frac{2T}{3} \frac{G}{D+G}$ plus $D_e = T/3$, we may allocate bandwidth to obtain $V_G = V_D$; this holds by assumption. We allocate the remaining bandwidth from D to m . We then also have $G_g = \frac{2T}{3} \frac{G}{D+G}$. In turn, this gives $V_G = V_D = \frac{2}{D+G}$. \square

LEMMA B.11. *If $M \leq T/3$, $D > T/3$, $D + G > 2T/3$, and $D < G$, then allocate bandwidth as:*

| C | g | m | e | V_C |
|--------------|---------------|-------------------|-------|---------------|
| G | $\frac{T}{3}$ | $G - \frac{T}{3}$ | | $\frac{1}{G}$ |
| M | | M | | -- |
| E | | E | | 0 |
| D | | $D - \frac{T}{3}$ | $T/3$ | $\frac{1}{D}$ |
| Total | $T/3$ | $T/3$ | $T/3$ | |

PROOF. We will balance all three positions, satisfying the first goal. We allocate $T/3$ from D to e to satisfy the second goal. The analysis is as in the proof of Lemma B.10, but, because $D < G$, the value of G_g that makes $V_G = V_D$ is larger than $T/3$. In order to balance all positions while minimizing $V_D > V_G$, we allocate $T/3$ from G to g and the rest of G to m . We allocate the remaining $D - T/3$ from D to m as well. \square

For completeness, we note that $V_G = (G_g/B_g)(1/G) = 1/G$, and $V_D = (D_e/B_e + D_g/B_g)(1/D) = (0 + 1)(1/D) = 1/D$. Because $D > G$, we have $V_D > V_G$. Note that $G > D$ implies $D < T/2$, so $V_D > 2/T$.

LEMMA B.12. If $M \leq T/3$, $G < T/3$, and $T/3 \leq G + D \leq 2T/3$, allocate bandwidth as:

| C | g | m | e | V_C |
|--------------|-----------|----------------------------|------------------------|----------------------------------|
| G | G | | | $\frac{3}{T}$ |
| M | | M | | -- |
| E | | $E + D + G - \frac{2T}{3}$ | $\frac{2T}{3} - D - G$ | $\frac{2T - 3D - 3G}{T \cdot E}$ |
| D | $T/3 - G$ | | $D + G - \frac{T}{3}$ | $\frac{3}{T}$ |
| Total | $T/3$ | $T/3$ | $T/3$ | |

PROOF. We will balance the three positions. Because $G + D \leq 2T/3$, in order to allocate as much of D as possible to e , we must allocate all of $G < T/3$ to g . We then allocate $T/3 - G$ from D to g and the rest (i.e., $D + G - T/3$) of D to e . We allocate $2T/3 - D - G$ from E to e . We allocate the rest (i.e., $E + D + G - 2T/3$) of E to m . \square

For completeness, we compute the various V_C values. We have $V_G = 3G/G = 3/T$. We have $V_E = 3(2T/3 - D - G)/(T \cdot E) = (2T - 3D - 3G)/(T \cdot E)$. We have $V_D = ((T/3 - G)/(T/3) + (D + G - T/3)/(T/3))(1/D) = (1 - 3G/T + 3D/T + 3G/T - 1)/D = 3/T$.

C Alt2 WEIGHTS

Following are weights satisfying the Alt2 goals described in Section 5.3.2:

(1) $M > T/3$:

(a) $D \geq |G - E|$

(i) $G \geq E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------------|----------|------------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{D - (G - E)}{2D}$ | | $\frac{D + G - E}{2D}$ |

(ii) $G < E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|------------------------|----------|--------------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | $\frac{D + E - G}{2D}$ | | $\frac{D - (G - E)}{2D}$ |

(b) $D < |G - E|$

(i) $\max(G, E) \geq M$

(A) $G \geq E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------|--------------------|----------|
| \mathcal{G} | $\frac{G + M}{2G}$ | $\frac{G - M}{2G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

(B) $G < E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|--------------------|--------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{E - M}{2E}$ | $\frac{E + M}{2E}$ |
| \mathcal{D} | 1 | | |

(ii) $\max(G, E) < M$

(A) $G \geq E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|----------|----------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

(B) $G < E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|----------|----------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | 1 | | |

(2) $(M \leq T/3) \wedge (D + \min(G, E) < T/3)$:

(a) $G \geq E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|--------------------|--------------------|----------|
| \mathcal{G} | $\frac{G + M}{2G}$ | $\frac{G - M}{2G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | | 1 |
| \mathcal{D} | | | 1 |

(b) $G < E$

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|--------------------|--------------------|
| \mathcal{G} | 1 | | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{E - M}{2E}$ | $\frac{E + M}{2E}$ |
| \mathcal{D} | 1 | | |

(3) $(M \leq T/3) \wedge (D + \min(G, E) \geq T/3)$:

- (a) $D \geq |G - E|$
 Let $x = 2T/(3(T - M))$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------------------|----------|----------------------|
| \mathcal{G} | x | $1 - x$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $1 - x$ | x |
| \mathcal{D} | $\frac{T/3 - xG}{D}$ | $1 - x$ | $\frac{T/3 - xE}{D}$ |

- (b) $D < |G - E|$
 (i) $G \geq E$

Let $x = T/(3(E + D))$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|-----------------|---------------------|----------|
| \mathcal{G} | $\frac{T/3}{G}$ | $\frac{G - T/3}{G}$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $1 - x$ | x |
| \mathcal{D} | | $1 - x$ | x |

- (ii) $G < E$

Let $x = T/(3(G + D))$.

| y | W_{gy} | W_{my} | W_{ey} |
|---------------|----------|---------------------|-----------------|
| \mathcal{G} | x | $1 - x$ | |
| \mathcal{M} | | 1 | |
| \mathcal{E} | | $\frac{E - T/3}{E}$ | $\frac{T/3}{E}$ |
| \mathcal{D} | x | $1 - x$ | |

D PROOF OF THEOREM 5.2

We prove Theorem 5.2 by case analysis. Each lemma shows the Alt2 goals are satisfied by the bandwidth allocation given in their tables. In each bandwidth-allocation table, each entry can be derived from the analogous weight in Appendix C simply by multiplying it by the class bandwidth (e.g. the entry for class \mathcal{D} and position e is $W_{e\mathcal{D}}D$). By case inspection, we can see that the cases partition the entire space and that each case has T total bandwidth allocated. Therefore, to prove Theorem 5.2, we just need to show that the allocation satisfies the Alt2 goals, in order, to the greatest extent possible.

D.1 Subcases of Case 1

D.1.1 Case 1(a)i.

LEMMA D.1. *If $M > T/3$, $D \geq |G - E|$, and $G \geq E$, allocate bandwidth as:*

| | g | m | e |
|---------------|-------------------|-----|-----------------|
| \mathcal{G} | G | | |
| \mathcal{M} | | M | |
| \mathcal{E} | | | E |
| \mathcal{D} | $(D - (G - E))/2$ | | $(D + G - E)/2$ |

PROOF. To satisfy the first goal, we must maximize the minimum bandwidth allocated to the guard and exit positions. The reason is that $M > T/3$, and therefore the bandwidth allocated to the middle position will be more than $T/3$, causing at least one of the guard and exit positions to be allocated less than $T/3$.

This allocation puts as much possible bandwidth in the guard and exit positions. Only M is used in the middle position, and it can

only be used in that position. Moreover, this allocation results in the same amount of bandwidth in those two positions: $G/2 + E/2 + D/2$.

This allocation is the only one that satisfies the first goal. Moving any more bandwidth to the middle position would necessarily reduce the bandwidth allocated across the guard and exit positions, therefore reducing at least one of them and thus the minimum. Moving bandwidth from the guard to exit position (or vice versa) would reduce its bandwidth and thus the minimum. Therefore, the second goal and the third goal cannot be further optimized subject to the first goal. \square

D.1.2 Case 1(a)ii.

LEMMA D.2. *If $M > T/3$, $D \geq |G - E|$, and $G < E$, allocate bandwidth as:*

| | g | m | e |
|---------------|-----------------|-----|-------------------|
| \mathcal{G} | G | | |
| \mathcal{M} | | M | |
| \mathcal{E} | | | E |
| \mathcal{D} | $(D + G - E)/2$ | | $(D - (G - E))/2$ |

PROOF. This case is symmetric with $G \geq E$, and the allocation has the same symmetry. \square

D.1.3 Case 1(b)iA.

LEMMA D.3. *If $M > T/3$, $D < |G - E|$, $\max(G, E) \geq M$, and $G \geq E$, allocate bandwidth as:*

| | g | m | e |
|---------------|-------------|-------------|-----|
| \mathcal{G} | $(G + M)/2$ | $(G - M)/2$ | |
| \mathcal{M} | | M | |
| \mathcal{E} | | | E |
| \mathcal{D} | | | D |

PROOF. First, observe that the case conditions imply that $D + E < T/3$ because otherwise there would exist the contradiction $T = D + E + G + M > T/3 + T/3 + T/3$. Therefore, this allocation achieves the minimum bandwidth in the exit position. That bandwidth is $D + E$, and no larger allocation is possible in that position, as only D and E can be used in the exit position. This allocation thus satisfies the first goal, and any allocation satisfying the first goal must put all D and E bandwidth in the exit position.

To consider the second goal, we observe that both D and E achieve the maximum V_C . Neither can be reduced without violating the first goal, and therefore the second goal is satisfied subject to the first goal.

To consider goal the third goal, we note that the guard and middle positions both have the second-smallest bandwidth allocation. All their bandwidth must come from G and M because D and E must be fully allocated to the exit position to satisfy the first goal. Thus, increasing the guard allocation could only come at the expense of the middle allocation and vice versa. Therefore, the second-smallest bandwidth is maximized, subject to the first goal, satisfying the third goal. \square

D.1.4 Case 1(b)iB.

LEMMA D.4. *If $M > T/3$, $D < |G - E|$, $\max(G, E) \geq M$, and $G < E$, allocate bandwidth as:*

| | g | m | e |
|---------------|-----|-------------|-------------|
| \mathcal{G} | G | | |
| \mathcal{M} | | M | |
| \mathcal{E} | | $(E - M)/2$ | $(E + M)/2$ |
| \mathcal{D} | D | | |

PROOF. This case is symmetric with the analogous case when $G \geq E$, and the allocation respects that symmetry. \square

D.1.5 Case 1(b)iiA.

LEMMA D.5. *If $M > T/3$, $D < |G - E|$, $\max(G, E) < M$, and $G \geq E$, allocate bandwidth as:*

| | g | m | e |
|---------------|-----|-----|-----|
| \mathcal{G} | G | | |
| \mathcal{M} | | M | |
| \mathcal{E} | | | E |
| \mathcal{D} | | | D |

PROOF. First, observe that the case conditions of $D + E < G$ and $M > T/3$ imply that $D + E < T/3$ because otherwise there would exist the contradiction $T = D + E + G + M > T/3 + T/3 + T/3$. Therefore, by allocating all D and E to the exit position, all G to the guard position, and all M to the middle position, this allocation achieves minimum bandwidth of $D + E$ in the exit position. No larger allocation is possible in that position, as only D and E can be used in the exit position. This allocation thus satisfies the first goal, and any allocation satisfying the first goal must put all D and E bandwidth in the exit position.

Second, to consider the second goal, we observe that both D and E achieve the maximum V_C . Neither can be reduced without violating the first goal, and therefore the second goal is satisfied subject to the first goal.

Third, to consider the third goal, we note that the guard position has the second-smallest bandwidth allocation. No more bandwidth can be allocated to it because because D must be fully allocated to the exit position to satisfy the first goal. Therefore, the second-smallest bandwidth is maximized, subject to the first goal, satisfying the third goal. \square

D.1.6 Case 1(b)iiB.

LEMMA D.6. *If $M > T/3$, $D < |G - E|$, $\max(G, E) < M$, and $G < E$, allocate bandwidth as:*

| | g | m | e |
|---------------|-----|-----|-----|
| \mathcal{G} | G | | |
| \mathcal{M} | | M | |
| \mathcal{E} | | | E |
| \mathcal{D} | D | | |

PROOF. This case is symmetric with the analogous case when $G \geq E$, and the allocation respects that symmetry. \square

D.2 Subcases of Case 2

D.2.1 Case 2a.

LEMMA D.7. *If $M \leq T/3$, $D + \min(G, E) < T/3$, and $G \geq E$, allocate bandwidth as:*

| | g | m | e |
|---------------|-------------|-------------|-----|
| \mathcal{G} | $(G + M)/2$ | $(G - M)/2$ | |
| \mathcal{M} | | M | |
| \mathcal{E} | | | E |
| \mathcal{D} | | | D |

PROOF. First, observe that this allocation puts $E + D < T/3$ bandwidth in the exit position, and it puts $(G + M)/2 = (T - E - D)/2 > (T - T/3)/2 = T/3$ in both the guard and middle positions. Therefore, the minimum allocated bandwidth is $E + D$ and is in the exit position. No more can be allocated to the exit position because only D and E can be allocated to it. Therefore, this allocation maximizes the minimum bandwidth, satisfying the first goal.

To consider the second goal, observe that the maximum V_C is obtained by both the D and E classes. This holds because they are fully allocated to the exit position, and less total bandwidth is allocated to it than to the guard position. This amount cannot be reduced because doing so would require moving some bandwidth from one of those classes away from the exit position, violating the first goal. Therefore, this allocation achieves the minimum maximum V_C , satisfying the second goal.

To consider the third goal, observe that the first goal has already required that all D and E be allocated to the exit position, and observe that the exit position has the minimum bandwidth. Moreover, observe that allocating some G bandwidth to the middle instead of just the guard position does not violate the first goal or the second goal as long as at least $D + E$ bandwidth of G is in the guard position. This allocation does so, as it puts $(G + M)/2 > T/3 > D + E$ bandwidth in both the guard and middle positions. The allocation then satisfies the third goal by allocating just enough G to the middle position that the guard and middle positions both have the second-smallest allocated bandwidth. \square

D.2.2 Case 2b.

LEMMA D.8. *If $M \leq T/3$, $D + \min(G, E) < T/3$, and $G < E$, allocate bandwidth as:*

| | g | m | e |
|---------------|-----|-------------|-------------|
| \mathcal{G} | G | | |
| \mathcal{M} | | M | |
| \mathcal{E} | | $(E - M)/2$ | $(E + M)/2$ |
| \mathcal{D} | D | | |

PROOF. This case is symmetric with the analogous case when $G \geq E$, and the allocation respects that symmetry. \square

D.3 Subcases of Case 3

D.3.1 Case 3a.

LEMMA D.9. *If $M \leq T/3$, $D + \min(G, E) \geq T/3$, and $D \geq |G - E|$, let $x = 2T/(3(T - M))$, and allocate bandwidth as:*

| | g | m | e |
|---------------|------------|----------|------------|
| \mathcal{G} | xG | $(1-x)G$ | |
| \mathcal{M} | | M | |
| \mathcal{E} | | $(1-x)E$ | xE |
| \mathcal{D} | $T/3 - xG$ | $(1-x)D$ | $T/3 - xE$ |

PROOF. This allocation puts the $T/3$ bandwidth in each position, and so the first and third goals are satisfied.

For the second goal, we observe that $V_G = V_E = V_D = 2/(T - M)$. There is no way to reduce any of these values while maintaining the bandwidths allocated across positions without increasing another of these values. Therefore, the second goal has been satisfied subject to the first goal, and this is the only allocation that can do so. \square

D.3.2 Case 3(b)i.

LEMMA D.10. *If $M \leq T/3$, $D + \min(G, E) \geq T/3$, $D < |G - E|$, and $G \geq E$, let $x = T/(3(E + D))$, and allocate bandwidth as:*

| | g | m | e |
|---------------|-------|-----------|------|
| \mathcal{G} | $T/3$ | $G - T/3$ | |
| \mathcal{M} | | M | |
| \mathcal{E} | | $(1-x)E$ | xE |
| \mathcal{D} | | $(1-x)D$ | xD |

PROOF. This allocation puts $T/3$ bandwidth in each position, and so the first and third goals are satisfied.

For the second goal, we observe that $V_G = 1/G$ and $V_E = V_D = 1/(E + D)$. We have that $G > E + D$ because $G \geq E$ and $D < |G - E|$ in this case. Therefore, $V_G < V_E = V_D$. It is not possible to reduce both V_E and V_D while satisfying the first goal because only D and E are (and can be) allocated to position e and D is not allocated to g (the other sensitive position). Therefore, this allocation minimizes the maximum V_C subject to the first goal, satisfying the second goal. \square

D.3.3 Case 3(b)ii.

LEMMA D.11. *If $M \leq T/3$, $D + \min(G, E) \geq T/3$, $D < |G - E|$, and $G < E$, let $x = T/(3(G + D))$, and allocate bandwidth as:*

| | g | m | e |
|---------------|------|-----------|-------|
| \mathcal{G} | xG | $(1-x)G$ | |
| \mathcal{M} | | M | |
| \mathcal{E} | | $E - T/3$ | $T/3$ |
| \mathcal{D} | xD | $(1-x)D$ | |

PROOF. This case is symmetric with the analogous case when $G \geq E$, and the allocation respects that symmetry. \square

E PROOF OF PROPOSITION 5.3

PROOF. This falls into Tor's Case 3(a)iiB: exactly one of E and G is less than $T/3$; for $S = E$, the smaller of G and E , we have $S + D < T/3$; we have $G > E + D \geq E$; and we have $G \geq M$. That case produces the claimed bandwidth allocation.

For Alt1: if $M > T/3$, then, because $G \geq M$, this would fall under Case 1a, which has the claimed allocation. If $M \leq T/3$ then, because $E + D < T/3$, this would fall under Case 2a, which also has the claimed allocation.

For Alt2: if $M > T/3$, then, because $G > E + D$, $G \geq M$, and $G > E$, this would fall under Case 1(b)iA, which has the claimed allocation. If $M \leq T/3$, then, because E is the smaller of G and E and $D + E < T/3$, and because $G \geq E$, this would fall under Case 2(a), which also has the claimed allocation. \square

F PROOF OF THEOREM 4.1

PROOF. Case 1. Let $G = E = M = T/3$, $D = 0$

Case 2(a)i. Let $G = T/3 - \epsilon/4$, $E = T/3 - \epsilon/2$, $M = T/3 + 5\epsilon/8$, $D = \epsilon/8$.

Case 2(a)ii. Let $G = T/3 - \epsilon/2$, $E = T/3 - \epsilon/4$, $M = T/3 + 5\epsilon/8$, $D = \epsilon/8$.

Case 2(b)i. Let $G = E = M = T/3 - \epsilon/4$, $D = 3\epsilon/4$.

Case 2(b)ii. Let $G = E = T/3 - \epsilon/2$, $M = T/3$, $D = \epsilon$.

Case 2(b)iiiA. Let $G = T/3 - \epsilon/2$, $E = T/3 - 3\epsilon/8$, $M = T/3 + \epsilon/4$, $D = 5\epsilon/8$.

Case 2(b)iiiB. Let $G = T/3 - \epsilon/2$, $E = T/3 - 3\epsilon/8$, $M = T/3 + 5\epsilon/8$, $D = \epsilon/4$.

Case 3(a)iA. Let $G = T/3 - \epsilon$, $E = T/3$, $M = T/3 + \epsilon/2$, $D = \epsilon/2$.

Case 3(a)iB. Let $G = T/3 - \epsilon$, $E = T/3 + \epsilon/2$, $M = T/3$, $D = \epsilon/2$.

Case 3(a)iiA. Let $G = T/3 + \epsilon/2$, $E = T/3 - \epsilon$, $M = T/3$, $D = \epsilon/2$.

Case 3(a)iiB. Let $G = T/3$, $E = T/3 - \epsilon$, $M = T/3 + \epsilon/2$, $D = \epsilon/2$.

Case 3(b)i. Let $G = T/3 - \epsilon$, $E = T/3$, $M = T/3$, $D = \epsilon$.

Case 3b.ii. Let $G = T/3$, $E = T/3 - \epsilon$, $M = T/3$, $D = \epsilon$. \square