DefenestraTor: Throwing Out Windows in Tor

Mashael AlSabah[1], Kevin Bauer[1,*], Ian Goldberg[1], Dirk Grunwald[2],
Damon McCoy[3], Stefan Savage[3], and Geoffrey M. Voelker[3]

[1] University of Waterloo
{malsabah,k4bauer,iang}@cs.uwaterloo.ca
[2] University of Colorado
grunwald@colorado.edu
[3] University of California, San Diego
{dlmccoy,savage,voelker}@cs.ucsd.edu

Abstract. Tor is one of the most widely used privacy enhancing technologies for achieving online anonymity and resisting censorship. While conventional wisdom dictates that the level of anonymity offered by Tor increases as its user base grows, the most significant obstacle to Tor adoption continues to be its slow performance. We seek to enhance Tor's performance by offering techniques to control congestion and improve flow control, thereby reducing unnecessary delays.

To reduce congestion, we first evaluate small fixed-size circuit windows and a dynamic circuit window that adaptively re-sizes in response to perceived congestion. While these solutions improve web page response times and require modification only to exit routers, they generally offer poor flow control and slower downloads relative to Tor's current design. To improve flow control while reducing congestion, we implement *N23*, an ATM-style per-link algorithm that allows Tor routers to explicitly cap their queue lengths and signal congestion via back-pressure. Our results show that N23 offers better congestion and flow control, resulting in improved web page response times and faster page loads compared to Tor's current design and other window-based approaches. We also argue that our proposals do not enable any new attacks on Tor users' privacy.

1 Introduction

Tor [10] is a distributed circuit-switching overlay network consisting of over two-thousand volunteer-run *Tor routers* operating around the world. Tor clients achieve anonymity by source-routing their traffic through three Tor routers using onion routing [14].

Context. Conventional wisdom dictates that the level of anonymity provided by Tor increases as its user base grows [8]. Another important, but often overlooked, benefit of a larger user base is that it reduces suspicion placed on users simply because they use Tor. Today, there are an estimated 150 to 250 thousand daily Tor users [20]. However, this estimate has not increased significantly since 2008.

---

[*] Work was done at University of Colorado and University of California, San Diego.

One of the most significant road blocks to Tor adoption is its excessively high and variable delays, which inhibit interactive applications such as web browsing.

Many prior studies have diagnosed a variety of causes of this high latency (see Dingledine and Murdoch [11] for a concise summary). Most of these studies have noted that the queuing delays often dominate the network latencies of routing packets through the three routers. These high queuing delays are, in part, caused by bandwidth bottlenecks that exist along a client's chosen circuit. As high-bandwidth routers forward traffic to lower-bandwidth downstream routers, the high-bandwidth router may be able to read data faster than it can write it. Because Tor currently has no explicit signaling mechanism to notify senders of this congestion, packets must be queued along the circuit, introducing potentially long and unnecessary delays for clients. While recent proposals seek to re-engineer Tor's transport design, in part, to improve its ability to handle congestion [18,29,36], these proposals face significant deployment challenges.

Improving Congestion and Flow Control. To reduce the delays introduced by uncontrolled congestion in Tor, we design, implement, and evaluate two classes of congestion and flow control. First, we leverage Tor's existing end-to-end window-based flow control framework and evaluate the performance benefits of using small fixed-size circuit windows, reducing the amount of data in flight that may contribute to congestion. We also design and implement a dynamic window resizing algorithm that uses increases in end-to-end circuit round-trip time as an implicit signal of incipient congestion. Similar solutions are being considered for adoption in Tor to help relieve congestion [6], and we offer a critical analysis to help inform the discussion. Window-based solutions are appealing, since they require modifications only to exit routers.

Second, we offer a fresh approach to congestion and flow control inspired by standard techniques from Asynchronous Transfer Mode (ATM) networks. We implement a per-link credit-based flow control algorithm called N23 [19] that allows Tor routers to explicitly bound their queues and signal congestion via back-pressure, reducing unnecessary delays and memory consumption. While N23 offers these benefits over the window-based approaches, its road to deployment may be slower, as it may require all routers along a circuit to upgrade.

Evaluation. We conduct a holistic experimental performance evaluation of the proposed algorithms using the ModelNet network emulation platform [35] with realistic traffic models. We show that the window-based approaches offer up to 65% faster web page response times relative to Tor's current design. However, they offer poor flow control, causing bandwidth under-utilization and ultimately resulting in poor download time. In contrast, our N23 experiments show that delay-sensitive web clients experience up to 65% faster web page responses and a 32% decrease in web page load times compared to Tor's current design.

2   Tor Background

The Tor network is a decentralized circuit-switching overlay consisting of volunteer-run Tor routers hosted around the world. Tor offers anonymity to clients by
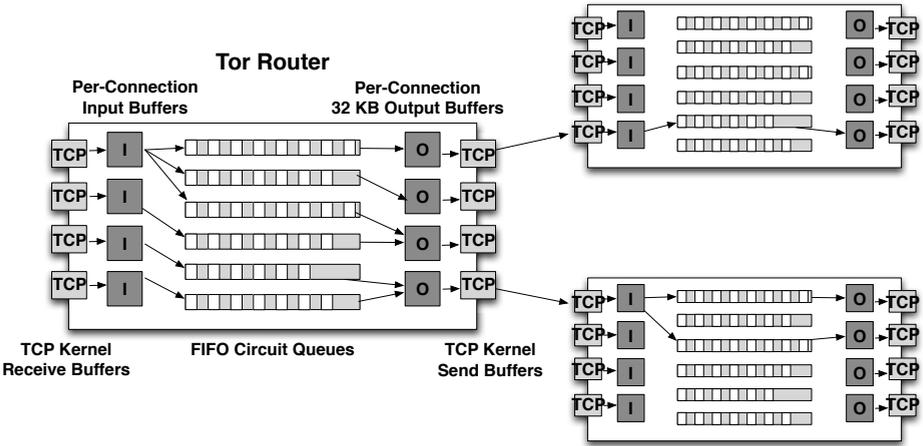
Fig. 1. A Tor router's queuing architecture

employing a layered encryption scheme [14] with three Tor routers. All data is sent in fixed-sized 512-byte units called cells. In general, the client selects routers to use on a circuit taking into account their bandwidth capacities, in order to balance the traffic load over the available router bandwidth. The first router on a circuit (called an "entry guard") is chosen carefully to reduce the threat of profiling and the predecessor attack [38]. Upon receiving a cell, the router removes its layer of encryption and forwards the cell to the next router on the circuit. Once the final (exit) router in the circuit removes its layer of encryption, the client's traffic is forwarded to the destination. A prior study found that the majority of Tor traffic by connection is interactive HTTP [21], and most of this traffic volume flows from the destination to the client. More details about Tor can be found in its design document [10] and its evolving protocol specification [9].

3   Tor's Approach to Congestion and Flow Control

Since the Tor network consists of volunteer-run routers from across the world, these routers have varying and often limited amounts of bandwidth available to relay Tor traffic. Consequently, as clients choose their circuits, some routers have large amounts of bandwidth to offer, while others may be bandwidth bottlenecks. In order for Tor to offer the highest degree of performance possible, it is necessary to have effective mechanisms in place to ensure steady flow control, while also detecting and controlling congestion. In this section, we discuss the many features that directly or indirectly impact congestion and flow control in Tor.

3.1   Congestion and Flow Control Mechanisms

Pairwise TCP.   All packets sent between Tor routers are guaranteed to be delivered reliably and in-order by using TCP transport. As a result of using TCP,

communications between routers can be protected with TLS link encryption. However, several circuits may be multiplexed over the same TCP connections, which could result in an unfair application of TCP's congestion control [29].

Tiered Output Buffers. Each Tor router's internal queuing architecture is illustrated in Figure 1. When a Tor router receives a cell on one of its TCP connections, the cell is first copied from the connection's receive kernel buffer into an application-layer input buffer to be decrypted. Next, the cell is pushed onto a FIFO *circuit queue* for the cell's respective circuit. For each outgoing TCP connection, a FIFO *output buffer* is maintained. The output buffer has a fixed size of 32 KiB, while the circuit queue has no explicit bound, but the circuit window size restricts how many cells may be in flight (described below). Since multiple circuits are often multiplexed over the same TCP connection, when there is space available in the outgoing connection's respective output buffer, the router must choose which circuits' cells to copy onto the output buffer. Initially, cells were chosen by round-robin selection across circuits. Recently, circuit prioritization has been proposed to give burstier circuits that likely correspond to interactive traffic priority over long-lived, bulk circuits [34].

Circuit and Stream Windows. Tor uses two layers of end-to-end window-based flow control between the exit router and the client to ensure steady flow control. First, a *circuit window* restricts how many cells may be in flight per circuit. By default, Tor uses a fixed 500 KiB (1000 cell) circuit window. For every 50 KiB (100 cells) received, an acknowledgment cell called a SENDME is sent, informing the sender that they may forward another 100 cells to the receiver[1].

Within each circuit window is a *stream window* of 250 KiB (500 cells) to provide flow control (or fairness) within a circuit. The receiver replies with a stream-level SENDME for every 25 KiB (50 cells) received. On receiving a stream-level SENDME, the sender may forward another 50 cells.

Both the stream-level and circuit-level windows are relatively large and static. To illustrate how this can degrade performance, consider the following scenario. Suppose a client downloads files through a circuit consisting of 10 MiB/s entry and exit routers and a 128 KiB/s middle router. Since the exit router can read data from the destination server faster than it can write it to its outgoing connection with the middle router, and the reli-

Fig. 2. The exit router's circuit queue delays for a 300 KiB download

able TCP semantics preclude routers from dropping cells to signal congestion, the exit router must buffer up to one full circuit window (500 KiB) worth of cells. Furthermore, as shown in Figure 2, these cells often sit idly for several seconds while the buffer is slowly emptied as SENDME cells are received. Since cells may
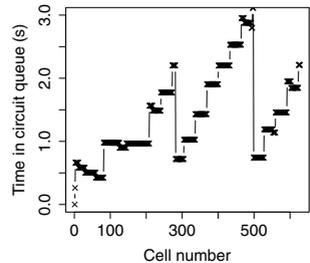
---

[1] Due to a bug, clients running Tor 0.0.0–0.2.1.19 erroneously reply with circuit-level SENDME cells after receiving 101 cells (rather than 100 cells).

travel down a circuit in large groups of up to 500 KiB followed by periods of silence while the exit router waits for SENDME replies, Tor's window-based flow control does not always keep a steady flow of cells in flight.

Token Bucket Rate Limiting. In order to allow routers to set limits on the amount of bandwidth they wish to devote to transiting Tor traffic, Tor offers token bucket rate limiting. Briefly, a router starts with a fixed amount of tokens, and decrements their token count as cells are sent or received. When the router's token count reaches zero, the router must wait to send or receive until the tokens are refilled. To reduce Tor's CPU utilization, tokens are refilled only once per second. However, it has been previously observed that refilling the tokens so infrequently contributes in part to Tor's overall delays [5].

3.2   Alternate Proposals to Reduce Congestion

There have been several recent proposals aimed specifically at reducing Tor's congestion. First, Tor has incorporated adaptive circuit-building timeouts that measure the time it takes to build a circuit, and eliminate circuits that take an excessively long time to construct [4]. The intuition is that circuits that build slowly are highly congested, and would in turn offer the user poor performance. While this approach likely improves the users' quality of service in some cases, it does not help to relieve congestion that may occur at one or more of the routers on a circuit *after* the circuit has been constructed.

In addition, user-level rate limiting has been proposed to throttle over-active or bulk downloading users. Here, the idea is to reduce the overall bandwidth consumption by bulk downloaders by using per-connection token bucket rate limiting at the entry guard. Early experiments indicate faster downloads for small file downloaders (the majority of Tor users), while harming bulk downloaders [7].

Finally, incentive schemes [17,24] have been proposed to reward users for operating fast Tor routers by offering them prioritized service. These proposals seek to reduce congestion and improve performance by increasing the bandwidth available for relaying Tor users' traffic.

4   Improving Tor's Congestion and Flow Control

Our primary goal is to improve Tor's performance, specifically by better understanding and improving Tor's congestion and flow control. We consider two broad classes of solutions. First, we wish to understand how much improvement is possible simply by adjusting Tor's existing end-to-end window-based flow control mechanisms to reduce the amount of data in flight, and thereby mitigate congestion. We also evaluate an end-to-end congestion control technique that enables exit Tor routers to infer incipient congestion by regarding increases in end-to-end round-trip time as a congestion signal. Second, we consider a fresh approach to congestion and flow control in Tor, eliminating Tor's end-to-end window-based flow control entirely, and replacing it with ATM-style, per-link flow control that caps routers' queue lengths and applies back-pressure to upstream routers to signal congestion.

4.1   Improving Tor's Existing End-to-End Flow Control

We first consider whether adjusting Tor's current window-based flow control can offer significant performance improvements. Keeping Tor's window-based mechanisms is appealing, as solutions based on Tor's existing flow control framework may be deployed immediately, requiring modifications only to the exit routers, not clients or non-exit routers.

Small Fixed-size Circuit Windows. The smallest circuit window size possible without requiring both senders and receivers to upgrade is 50 KiB (100 cells, or one circuit-level `SENDME` interval). We evaluate how fixed 50 KiB circuit windows impact clients' performance[2].

Dynamic Circuit Windows. It has been shown that protocols that use a small, fixed end-to-end window may achieve suboptimal throughput [28]. To avoid a potential loss in throughput that could result from an under-sized window, we next consider an algorithm that initially starts with a small, fixed circuit-window and dynamically increases the window size in response to positive end-to-end latency feedback. Inspired by latency-informed congestion control techniques for IP networks [3,37], we propose an algorithm that uses increases in perceived end-to-end circuit round-trip time (RTT) as a signal of incipient congestion.

The algorithm works as follows. Initially, each circuit's window size starts at 100 cells. First, the sender calculates the circuit's end-to-end RTT using the circuit-level `SENDME` cells, maintaining the minimum RTT (rtt_{min}) and maximum RTT (rtt_{max}) observed for each circuit. We note that rtt_{min} is an approximation of the base RTT, where there is little or no congestion on the circuit. Next, since RTT feedback is available for every 100 cells[3], the circuit window size is adjusted quickly using an additive increase, multiplicative decrease (AIMD) window scaling mechanism based on whether the current RTT measurement (rtt) is less than the threshold T, defined in Equation 1. This threshold defines the circuit's tolerance to perceived congestion.

$$T = (1 - \alpha) \times rtt_{min} + \alpha \times rtt_{max} \tag{1}$$

Choosing a small α value ensures that the threshold is close to the base RTT, and any increases beyond the threshold implies the presence of congestion along the circuit[4]. For each RTT measurement (*e.g.*, each received circuit-level `SENDME`), the circuit window size (in cells) is adjusted according to Equation 2.

$$new\_window(rtt) = \begin{cases} old\_window + 100 & \text{if } rtt \leq T \\ \lfloor old\_window/2 \rfloor & \text{otherwise} \end{cases} \tag{2}$$

---

[2] Due to the aforementioned bug, in practice, the window size should be 101 cells.

[3] Similar to the 50 KiB windows, `SENDME` cells may be available after 101 cells.

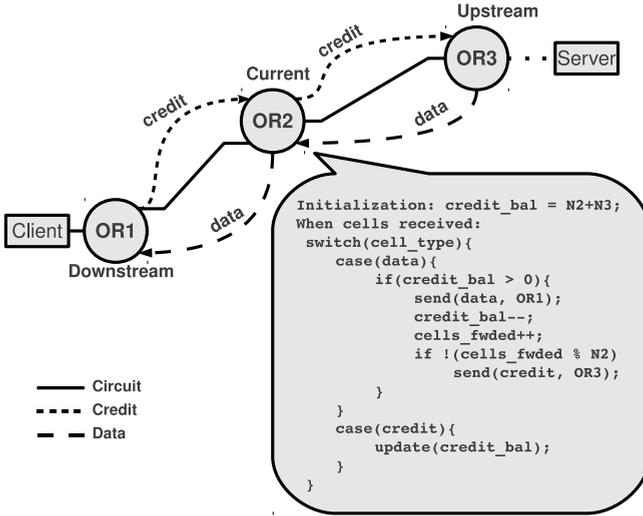[4] For our experiments, we use $\alpha = 0.25$.

Fig. 3. N23 credit-based flow control in Tor

Finally, we explicitly cap the minimum and maximum circuit window sizes at 100 and 1000 cells, respectively[5]. Note that for long-lived circuits, rtt_{min} may become increasingly small and rtt_{max} may grow very large. In practice, these values should decay over time, for example, using an exponentially weighted moving average of each respective parameter.

4.2   ATM-Style Congestion and Flow Control for Tor

Because Tor's flow control works at the circuit's edges—the client and the exit router—we seek to improve performance by implementing per-link flow control to ensure a steady flow of cells while reducing congestion at the intermediate routers. Implementing per-link flow control in Tor resembles the problem of link-by-link flow control (LLFC) in ATM networks. While the goals of Tor and ATM are certainly different, there are many similarities. Both networks are connection-oriented, in the sense that before applications can send or receive data, virtual circuits are constructed across multiple routers or switches, and both have fixed-sized cells. Furthermore, it has been shown that ATM's credit-based flow control approaches, such as the N23 scheme, eliminate cell loss due to buffer overflows [16], a feature that makes such approaches similar to Tor, where no packets may be dropped to signal congestion.

N23 Flow Control for Tor. Figure 3 depicts the N23 scheme that we integrated into Tor, and it works as follows. First, when a circuit is built, each router along the circuit is assigned an initial *credit balance* of $N2 + N3$ cells,

---

[5] Note that a selfish Tor client could attempt to increase their circuit window by preemptively acknowledging data segments before they are actually received. Prior work in mitigating similar behavior in selfish TCP receivers may be applied here [30,32].

where $N2$ and $N3$ are system parameters. $N2$ cells is the available steady state buffering per circuit, $N3$ cells is the allowed excess buffering, and the circuit's queue length is strictly upper bounded by $N2 + N3$ cells. In general, $N2$ is fixed at the system's configuration time, but $N3$ may change over a circuit's lifetime.

When a router forwards a cell, it decrements its credit balance by one for that cell's circuit. Each router stops forwarding cells if its credit balance reaches zero. Thus, routers' circuit queues are upper bounded by $N2 + N3$ cells, and congestion is indicated to upstream routers through this back-pressure. Next, for every $N2$ cells forwarded, the downstream router sends a *flow control cell* to the upstream router that contains credit information reflecting its available circuit queue space. On receiving a flow control cell, the upstream router updates the circuit's credit balance and may forward cells only if the credit balance is greater than zero.

Adaptive Buffer Sizes and Congestion Control. The algorithm as described assumes a static $N3$. We also developed an adaptive algorithm that reduces the $N3$ value when there is downstream congestion, which is detected by monitoring the delay that cells experience in the connection's output buffer. When the congestion subsides, $N3$ can increase again. The value of $N3$ is updated periodically and is bounded by a minimum and a maximum value (100 and 500 cells, respectively).

Advantages. The N23 algorithm has two important advantages over Tor's current flow control. First, the size of the circuit queue is explicitly capped, and guaranteed to be no more than $N2 + N3$ cells. This also ensures steady flow control, as routers typically have cells available to forward. Tor's current flow control algorithm allows the circuit queue of a circuit's intermediate routers to grow up to one circuit window in size, which not only wastes memory, but also results in unnecessary delays due to congestion. In contrast, for typical parameter values ($N3 = 500$ and $N2 = 10$), N23 ensures a strict circuit queue bound of 510 cells, while these queues currently can grow up to 1000 cells in length.

The second advantage is that adaptive N3 reacts to congestion within a single link RTT. When congestion occurs at a router, the preceding router in the circuit will run out of credit and must stop forwarding until it gets a flow control cell.

5   Experiments and Results

To empirically demonstrate the efficacy of our proposed improvements, we offer a whole-network evaluation of our congestion and flow control algorithms using the ModelNet network emulation platform [35]. Briefly, ModelNet enables the experimenter to specify realistic network topologies annotated with bandwidth, delay and other link properties, and run real code on the emulated network[6].

Our evaluation focuses on performance metrics that are particularly important to the end-user's quality of service. First, we measure *time-to-first-byte*, which is how long the user must wait from the time they issue a request for

---

[6] More details about our experimental environment can be found in Bauer *et al.* [2].
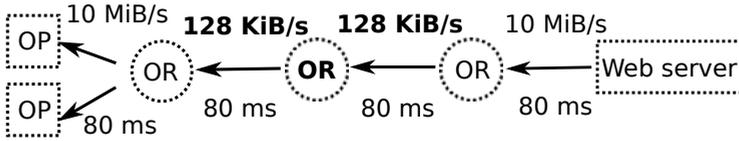
Fig. 4. A simple topology with a middle router bandwidth bottleneck

data until they receive the first byte (or until a web client starts to see content load on their current web page). The time-to-first-byte is two end-to-end circuit RTTs: one RTT to connect to the destination web server, and a second RTT to issue a request for data (*e.g.*, HTTP `GET`) and receive the first byte of data in response[7]. Second, we measure *overall download time* (including time-to-first-byte), which is how long the user must wait for their web page to load. For all experiments, we use the latest development branch of the Tor source code (version `0.2.3.0-alpha-dev`)[8].

5.1   Small-Scale Analysis

Setup.   We emulate the topology depicted in Figure 4 on ModelNet where two Tor clients compete for service on the same set of routers with a bandwidth bottleneck at the middle router[9]. One client downloads 300 KiB, which roughly corresponds to the size of an average web page [27]. The second client, a bulk downloader, fetches 5 MiB. Both clients pause for a random amount of time between one and three seconds, and repeat their downloads. Each experiment concludes after the web client completes 200 downloads. Both clients use the `wget` web browser and the destination runs the `lighthttpd` web server.

End-to-end Window-based Solutions.   Figure 5(a) shows that the time-to-first-byte for a typical web client using stock Tor is 4.5 seconds at the median, which is unacceptably high for delay-sensitive, interactive web users who must incur this delay for each web request. In addition, stock Tor's circuit queues fluctuate in length, growing up to 250 cells long, and remaining long for many seconds, indicating queuing delays, as shown in Figure 6(a). Reducing the circuit window size to 50 KiB (*e.g.*, one circuit `SENDME` interval) offers a median time-to-first-byte of less than 1.5 seconds, and dynamic windows offer a median time-to-first-byte of two seconds. In Figure 5(b), we see that the web client's download time is influenced by the high time-to-first-byte, and is roughly 40% faster with 50 KiB and dynamic windows relative to stock Tor. Also, the circuit queues are smaller with the 50 KiB and dynamic windows (see Figures 6(b) and 6(c)).

---

[7] Note that there is a proposal being considered to eliminate one of these RTTs [13].

[8] In our evaluation, we refer to unmodified Tor version `0.2.3.0-alpha-dev` as *stock Tor*, 50 KiB (100 cell) fixed windows as *50 KiB window*, the dynamic window scaling algorithm as *dynamic window*, and the N23 algorithm as *N23*.

[9] Note that a 128 KiB/s router corresponds to the 65th percentile of routers ranked by observed bandwidth, as reported by the directory authorities. Thus, it is likely to be chosen fairly often by clients. Also, prior work [26] found the median round-trip time between live Tor routers to be about 80 ms.
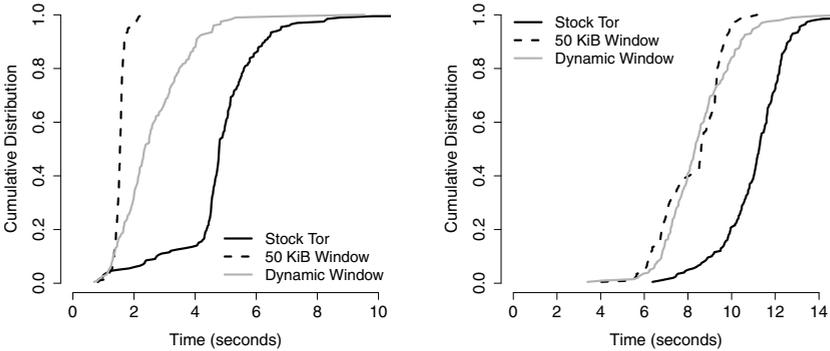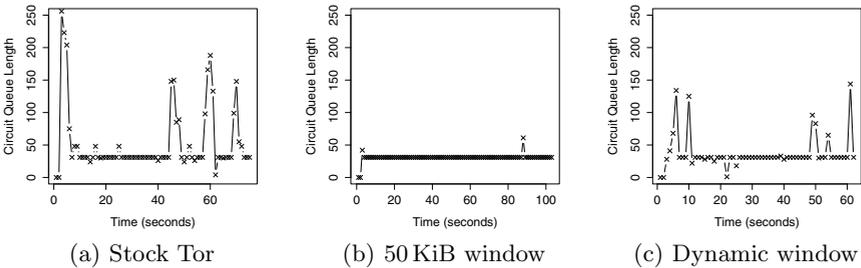
(a) Web client's time-to-first byte

(b) Web client's download time
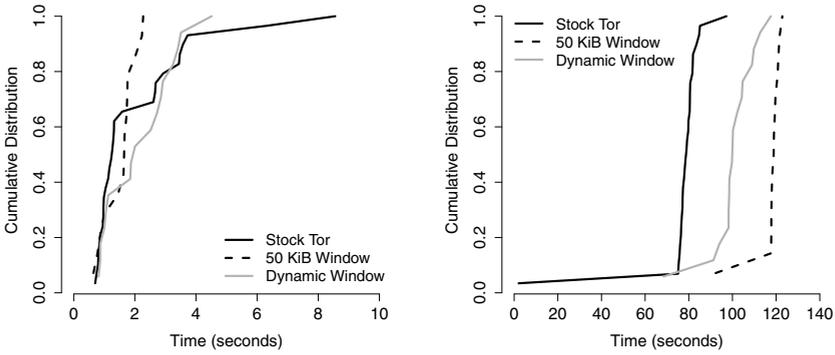
(c) Bulk client's time-to-first-byte

(d) Bulk client's download time

Fig. 5. Performance comparisons for window approaches in a bottleneck topology

(a) Stock Tor

(b) 50 KiB window

(c) Dynamic window

Fig. 6. Bulk client's circuit queues at the exit router over the course of a download

The bulk client experiences significantly less time-to-first-byte delays (in Figure 5(c)) than the web client using stock Tor. This highlights an inherent un-fairness during congestion: web clients' traffic is queued behind the bulk traffic

(a) Web client's download time   (b) Bulk client's download time
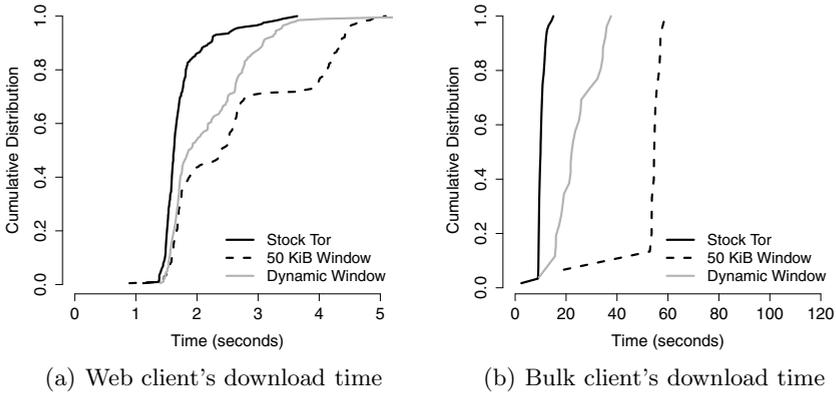
Fig. 7. Download time comparisons for windows in a non-bottleneck network

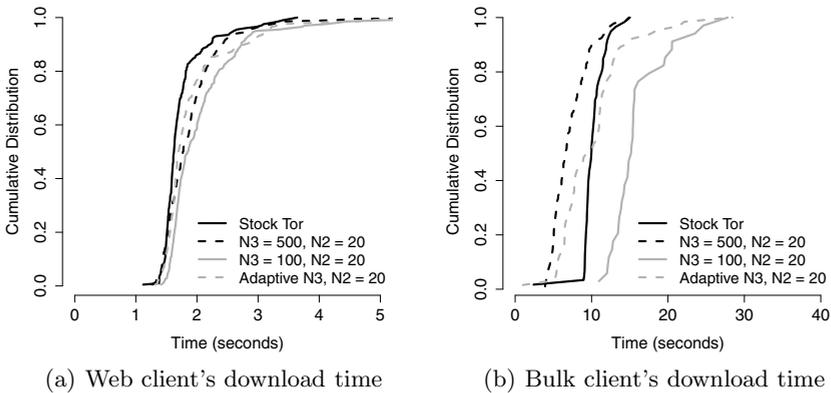(a) Web client's download time   (b) Bulk client's download time

Fig. 8. Download time comparisons for Tor and N23 in a non-bottleneck network

and, consequently, delay-sensitive clients must wait longer than delay-insensitive bulk downloaders to receive their first byte of data. Using a small or dynamic window reduces this unfairness, since the bound on the number of unacknowledged cells allowed to be in flight is lower.

However, Figure 5(d) indicates that the bulk client's download takes significantly longer to complete with 50 KiB windows relative to stock Tor. Thus, 50 KiB windows enhance performance for web clients at the cost of slower downloads for bulk clients. The bulk clients experience slower downloads because they keep less data in flight and, consequently, must incur additional round-trip time delays to complete the download. Dynamic windows offer a middle-ground solution, as they ameliorate this limitation by offering an improvement in download time for web clients while penalizing bulk clients less than small windows, but bulk clients are still penalized relative to stock Tor's performance.

We next consider the same topology shown in Figure 4, except we replace the bottleneck middle router with a 10 MiB/s router. In such a topology, congestion is minimal, as evidenced by a median time-to-first-byte of 0.75 s for both the web and bulk clients (regardless of the window size). However, because the 50 KiB and dynamic windows generally keep less data in flight, these solutions offer slower downloads relative to stock Tor, as shown in Figures 7(a) and 7(b). We also found that manipulating Tor's circuit windows in combination with circuit-level prioritization offers even more improvement for the web client, while not further harming the bulk client's performance. These results are in Appendix A.

Despite the improvements in time-to-first-byte in the presence of bandwidth bottlenecks, we find that smaller circuit windows may under-utilize the available bandwidth[10] and the dynamic window scaling algorithm is unable to adjust the window size fast enough, as it receives congestion feedback infrequently (only every 100 cells). Also, even in the non-bottleneck topology, the 50 KiB window web client's time-to-first-byte is higher than the optimal delay from two circuit RTTs, which is 0.64 s. Lastly, 50 KiB windows offer worse flow control than Tor's current design, since only 50 KiB can be in flight, and the exit router must wait for a full circuit RTT until more data can be read and sent down the circuit.

Based on these drawbacks, we conclude that in order to achieve an improvement in both time-to-first-byte and download speed, it is necessary to re-design Tor's fundamental congestion and flow control mechanisms. We next offer an evaluation of per-link congestion and flow control for Tor.

Per-link Congestion and Flow Control. We first implemented N23 with fixed values of N2 and N3 (*static N23*) and then with N3 values that react to network feedback (*adaptive N3*). We disabled Tor's window-based flow control, so that exit routers ignored SENDMEs they received from clients. We discuss the results of adaptive N3 with our large-scale experiments. In this section, we present the results of N23 for both the bottleneck and non-bottleneck topologies.

For the non-bottleneck topology, we see in Figure 8(b) that N23 provides a substantial improvement in download time for the 5 MiB downloads compared to stock Tor only for higher values of N3 — 500 cells, comparable to stock Tor's stream window size. The graph shows that there is a 25% decrease in delay for 50% of the bulk downloads when N23 is used. Since the maximum throughput is bounded by W/RTT, where W is the link's TCP window size and RTT is the link's round-trip time, and since N23's per-link RTT is significantly smaller than a stock Tor's complete circuit RTT, throughput is increased when N23 is used. This improvement suggests that in non-bottleneck scenarios, bulk traffic data cells are
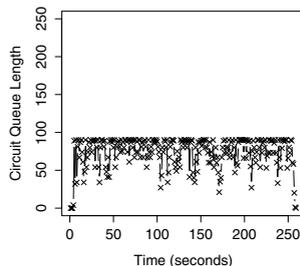

Fig. 9. Circuit queue length with bottleneck: N3 = 70, N2 = 20

---

[10] We note that bandwidth under-utilization may only be a problem if there is not sufficient demand from Tor clients to fully consume the network's bandwidth.
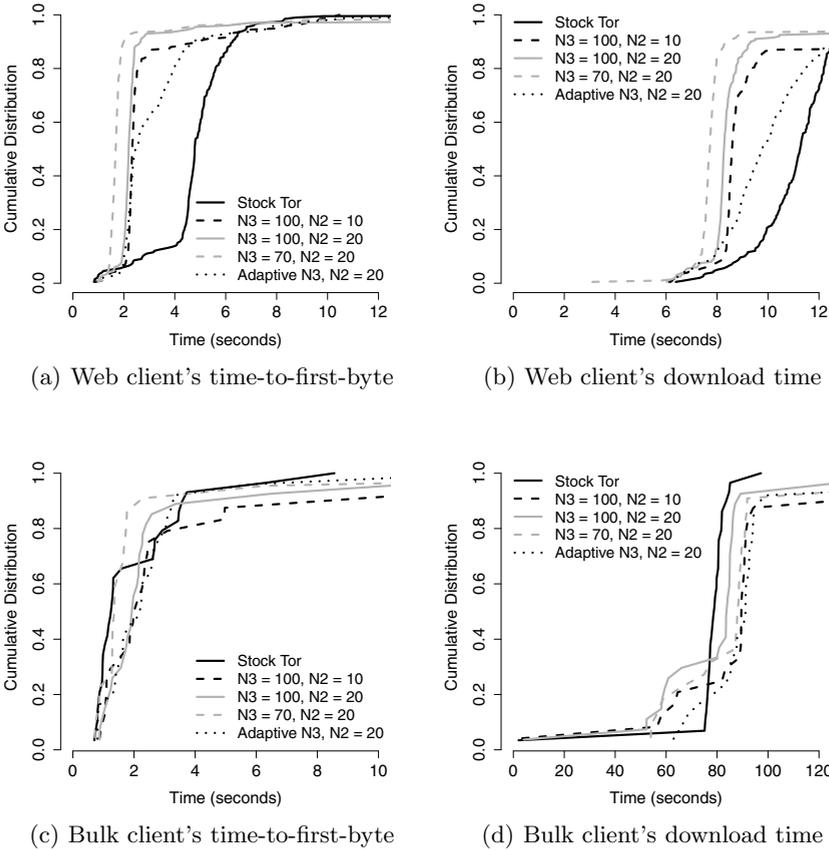
(a) Web client's time-to-first-byte

(b) Web client's download time

(c) Bulk client's time-to-first-byte

(d) Bulk client's download time

Fig. 10. Performance comparisons for Tor and N23 in a bottleneck topology

unnecessarily slowed down by Tor's flow control at the circuit's edges. For web traffic, Tor's current flow control and N23 have similar performance for fixed and adaptive N3, as shown in Figure 8(a). Also, the median time-to-first-byte is the same for the web and bulk clients at 0.75 s.

For bottleneck scenarios, Figures 10(a) and 10(b) show that smaller values of N3 improve both the download time and time-to-first-byte for the bursty web traffic. For example, the web browsing client experiences a 20% decrease in download time for 80% of the requests when N23 is used. Also, the web client's time-to-first-byte is only two seconds for 90% of the requests, whereas for the stock Tor client, 80% of web requests take more than four seconds to receive the first byte. Figure 9 shows that the circuit queue length is upper bounded by $N2 + N3 = 90$ cells.

To understand how N23 performs with different $N2$ values, we repeated the bottleneck experiments while varying that parameter. Although a higher value for $N2$ has the undesirable effect of enlarging the circuit buffer, it can be seen

in Figures 10(a) and 10(b) that when $N3$ is fixed at 100 cells, increasing $N2$ to 20 cells slightly improves both download time and time-to-first-byte. It can be observed from Figure 10(a) that time-to-first-byte is significantly improved by keeping a smaller $N3 = 70$ and a larger $N2 = 20$. Decreasing $N3$ to 70 cells makes up for the increase in the $N2$ zone of the buffer, which means we gain the benefits of less flow control overhead, and the benefits of a small buffer of $N2 + N3 = 90$ cells. While performance is improved for the web clients, the bulk client's time-to-first-byte is not affected greatly, as seen in Figure 10(c), but its downloads generally take longer to complete, as we see in Figure 10(d). In addition, adaptive N3 offers improved time-to-first-byte and download times for the web client, while slowing downloads for the bulk client. By N23 restricting the amount of data in flight, the bandwidth consumed by bulk clients is reduced, improving time-to-first-byte and download time for delay-sensitive web clients.

We also evaluate N23 in combination with circuit-level prioritization in a bottleneck topology. We observe that circuit-level prioritization with N23 offers no performance benefit over N23 alone. The full results are in Appendix B.

Finally, the bandwidth cost associated with the N23 scheme is relatively low. For instance, with $N2 = 10$, a flow control cell must be sent by each router on the circuit for every 10 data cells forwarded, which requires a 10% bandwidth overhead per router. For $N2 = 20$, a flow control cell is sent for every 20 data cells, which is only a 5% overhead per router. While this cost is higher than Tor's window-based flow control (*e.g.*, one stream-level SENDME for every 50 data cells and one circuit-level SENDME for every 100 data cells, resulting in a 3% overhead per circuit), the cost of N23 is nonetheless modest.

5.2   Larger-Scale Analysis

Setup. We next evaluate the window-based solutions and N23 with adaptive N3 in a more realistic network topology[11]. We deploy 20 Tor routers on a ModelNet topology whose bandwidths are assigned by sampling from the live Tor network. Each link's round-trip time is set to 80 ms. Next, to generate a traffic workload, we run 200 Tor clients. Of these, ten clients are bulk downloaders who fetch files between 1–5 MiB, pausing for up to two seconds between fetches. The remaining 190 clients are web clients, who download files between 100–500 KiB (typical web page sizes), pausing for up to 30 seconds between fetches. This proportion of bulk-to-non-bulk clients approximates the proportion observed on the live Tor network [21]. To isolate the improvements due to our proposals, circuit-level prioritization is disabled for this experiment.

Results. For web clients, Figure 11(a) shows that both the 50 KiB fixed and dynamic windows still offer improved time-to-first-byte. However, both algorithms perform worse than stock Tor in terms of overall download time, as shown in Figure 11(b). Note that this observation holds for both web clients and bulk downloaders. Because smaller windows provide less throughput than larger windows when there is no bottleneck, non-bottlenecked circuits are under-utilized.

---

[11] In this experiment, we only consider N23 with adaptive N3 because in practice, N23 should discover the right buffer size for the given network conditions.

(a) Web client's time-to-first-byte
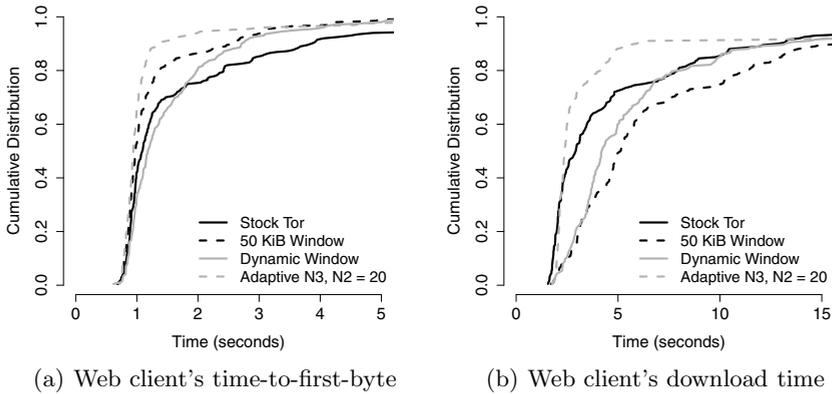
(b) Web client's download time

Fig. 11. Performance results for large-scale experiments

N23 with the adaptive N3 algorithm, in contrast, has the ability to react to congestion quickly by reducing routers' queue lengths, causing back pressure to build up. Consequently, our results indicate that N23 offers an improvement in both time-to-first-byte *and* overall download time for web clients, while bulk clients experience roughly the same performance as stock Tor.

These experiments again highlight the potential negative impact of 50 KiB and small dynamic windows, since even in a larger network with a realistic traffic load, smaller windows offer worse performance for typical delay-sensitive web requests relative to Tor's current window size. Thus, to achieve maximal improvements, we suggest that Tor adopt N23 congestion and flow control.

6   Discussion

Having empirically evaluated our proposed congestion and flow control approaches, we next discuss a variety of open issues.

6.1   Limitations of Experiments and Results

The results presented in Section 5 generally show an improvement in time-to-first-byte and download time with N23 flow control relative to end-to-end windows. However, these results were obtained in a testbed environment with a single, artificial traffic load; thus, an analysis of expected performance as the traffic load varies, on the live Tor network with real traffic loads, and with exhaustive N23 parameter configurations is future work.

6.2   Incremental Deployment

In order for our proposed congestion and flow control mechanisms to be practical and easily deployable on the live Tor network, it is important that any modifications to Tor's router infrastructure be incrementally deployable. Any solutions

based on Tor's existing window-based flow control require upgrades only to the exit routers; thus they can be slowly deployed as router operators upgrade. N23 may also be deployed incrementally, however, clients may not see substantial performance benefits until a large fraction of the routers have upgraded.

6.3   Anonymity Implications

A key question to answer is whether improving Tor's performance and reducing congestion enables any attack that was not previously possible. It is well known that Tor is vulnerable to congestion attacks wherein an attacker constructs circuits through a number of different routers, floods them with traffic, and observes if there is an increase in latency on a target circuit, which would indicate a shared router on both paths [22]. More recent work has suggested a solution that would mitigate bandwidth amplification variants of this attack, but not the shared router inference part of the attack [12]. We believe that by reducing congestion (and specifically, by bounding queue lengths), our proposed techniques may increase the difficulty of mounting congestion attacks.

However, if only a fraction of the routers upgrade to our proposals and if clients only choose routers that support the new flow control, then an adversary may be able to narrow down the set of potential routers that a client is using. Thus, it is important to deploy any new flow control technique after a large fraction of the network has upgraded. Such a deployment can be controlled by setting a flag in the authoritative directory servers' consensus document, indicating that it is safe for clients to use the new flow control.

Another well-studied class of attack is end-to-end traffic correlation. Such attacks endeavor to link a client with its destination when the entry and exit points are compromised, and these attacks have been shown to be highly accurate [1,23,25,31,33]. Reducing latency might improve this attack; however, Tor is already highly vulnerable, so there is little possibility for additional risk.

Finally, previous work has shown that network latency can be used as a side channel to infer a possible set of client locations [15]. By decreasing the variance in latency, we might expose more accurate RTT measurements, thus improving the effectiveness of this attack. However, reducing congestion does not enable a new attack, but rather may potentially increase the effectiveness of a known attack. To put this attack in perspective, Tor's design has already made many performance/anonymity trade-offs, and thus, we believe that our performance improvements outweigh any potential decrease in anonymity brought about by reducing the variance in latency.

7   Conclusion

We seek to improve Tor's performance by reducing unnecessary delays due to poor flow control and excessive queuing at intermediate routers. To this end, we have proposed two broad classes of congestion and flow control. First, we

tune Tor's existing circuit windows to effectively reduce the amount of data in flight. However, our experiments indicate that while window-based solutions do reduce queuing delays, they tend to suffer from poor flow control, under-utilizing the available bandwidth, and consequently, smaller windows provide slower downloads than unmodified Tor.

To solve this problem, we offer a fresh approach to congestion and flow control in Tor by designing, implementing, and experimentally evaluating a per-link congestion and flow control algorithm from ATM networks. Our experiments indicate that this approach offers the promise of faster web page response times and faster overall web page downloads.

References

1. Bauer, K., McCoy, D., Grunwald, D., Kohno, T., Sicker, D.: Low-resource routing attacks against Tor. In: Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2007), Washington, DC, USA (October 2007)
2. Bauer, K., Sherr, M., McCoy, D., Grunwald, D.: ExperimenTor: A testbed for safe and realistic Tor experimentation. Technical Report CACR 2011-12 (May 2011), `http://www.cacr.math.uwaterloo.ca/techreports/2011/cacr2011-12.pdf`
3. Brakmo, L.S., O'Malley, S.W., Peterson, L.L.: TCP Vegas: New techniques for congestion detection and avoidance. In: Proceedings of the Conference on Communications Architectures, Protocols and Applications, SIGCOMM 1994, pp. 24–35. ACM, New York (1994)
4. Chen, F., Perry, M.: Improving Tor path selection (July 2008), `https://gitweb.torproject.org/torspec.git/blob_plain/HEAD:/proposals/151-path-selection-improvements.txt`
5. Dhungel, P., Steiner, M., Rimac, I., Hilt, V., Ross, K.W.: Waiting for anonymity: Understanding delays in the Tor overlay. In: Peer-to-Peer Computin. IEEE, Los Alamitos (2010)
6. Dingledine, R.: Prop 168: Reduce default circuit window (August 2009), `https://gitweb.torproject.org/torspec.git/blob_plain/HEAD:/proposals/168-reduce-circwindow.txt`
7. Dingledine, R.: Research problem: adaptive throttling of Tor clients by entry guards (September 2010), `https://blog.torproject.org/blog/research-problem-adaptive-hrottling-tor-clients-entry-guards`

8. Dingledine, R., Mathewson, N.: Anonymity loves company: Usability and the network effect. In: Workshop on the Economics of Information Security (June 2006)

9. Dingledine, R., Mathewson, N.: Tor Protocol Specificiation (2010), https://gitweb.torproject.org/tor.git/blob_plain/HEAD:/doc/spec/tor-spec.txt

10. Dingledine, R., Mathewson, N., Syverson, P.: Tor: The second-generation onion router. In: Proceedings of the 13th USENIX Security Symposium (August 2004)

11. Dingledine, R., Murdoch, S.: Performance improvements on Tor or, why Tor is slow and what we're going to do about it (March 2009), http://www.torproject.org/press/presskit/2009-03-11-performance.pdf

12. Evans, N., Dingledine, R., Grothoff, C.: A practical congestion attack on Tor using long paths. In: Proceedings of the 18th USENIX Security Symposium (2009)

13. Goldberg, I.: Prop 174: Optimistic data for Tor: Server side, https://trac.torproject.org/projects/tor/ticket/1795

14. Goldschlag, D.M., Reed, M.G., Syverson, P.F.: Hiding routing information. In: Proceedings of Information Hiding: First International Workshop (May 1996)

15. Hopper, N., Vasserman, E.Y., Chan-Tin, E.: How much anonymity does network latency leak? In: Proceedings of ACM CCS (October 2007)

16. Jain, R.: Congestion control and traffic management in ATM networks: Recent advances and a survey. Computer Networks and ISDN Systems (1995)

17. Jansen, R., Hopper, N., Kim, Y.: Recruiting new Tor relays with BRAIDS. In: Proceedings of the ACM CCS (2010)

18. Kiraly, C., Bianchi, G., Cigno, R.L.: Solving performance issues in anonymiziation overlays with a L3 approach. University of Trento Information Engineering and Computer Science Department Technical Report DISI-08-041 (September 2008)

19. Kung, H.T., Blackwell, T., Chapman, A.: Credit-based flow control for ATM networks: credit update protocol, adaptive credit allocation and statistical multiplexing. SIGCOMM Comput. Commun. Rev. 24, 101–114 (1994)

20. Loesing, K.: Measuring the Tor network: Evaluation of client requests to the directories. Tor Project Technical Report (June 2009)

21. McCoy, D., Bauer, K., Grunwald, D., Kohno, T., Sicker, D.: Shining light in dark places: Understanding the Tor network. In: Proceedings of the 8th Privacy Enhancing Technologies Symposium (July 2008)

22. Murdoch, S.J., Danezis, G.: Low-cost traffic analysis of Tor. In: Proceedings of the 2005 IEEE Symposium on Security and Privacy. IEEE CS, Los Alamitos (2005)

23. Murdoch, S.J., Zieliński, P.: Sampled traffic analysis by Internet-exchange-level adversaries. In: Privacy Enhancing Technologies Workshop (June 2007)

24. Ngan, T.W.J., Dingledine, R., Wallach, D.S.: Building Incentives into Tor. In: Proceedings of Financial Cryptography (January 2010)

25. Øverlier, L., Syverson, P.: Locating hidden servers. In: Proceedings of the 2006 IEEE Symposium on Security and Privacy. IEEE CS, Los Alamitos (2006)

26. Pries, R., Yu, W., Graham, S., Fu, X.: On performance bottleneck of anonymous communication networks. In: Parallel and Distributed Processing (2008)

27. Ramachandran, S.: Web metrics: Size and number of resources, https://code.google.com/speed/articles/web-metrics.html

28. Rapier, C., Bennett, B.: High speed bulk data transfer using the SSH protocol. In: 15th Mardi Gras Conference on Distributed Applications. ACM, New York (2008)

29. Reardon, J., Goldberg, I.: Improving Tor using a TCP-over-DTLS tunnel. In: Proceedings of the 18th USENIX Security Symposium (August 2009)

30. Savage, S., Cardwell, N., Wetherall, D., Anderson, T.: TCP congestion control with a misbehaving receiver. SIGCOMM Comput. Commun. Rev. 29, 71–78 (1999)
31. Serjantov, A., Sewell, P.: Passive attack analysis for connection-based anonymity systems. In: Snekkenes, E., Gollmann, D. (eds.) ESORICS 2003. LNCS, vol. 2808, pp. 116–131. Springer, Heidelberg (2003)
32. Sherwood, R., Bhattacharjee, B., Braud, R.: Misbehaving TCP receivers can cause Internet-wide congestion collapse. In: Proceedings of the 12th ACM Conference on Computer and Communications Security (2005)
33. Shmatikov, V., Wang, M.H.: Timing analysis in low-latency mix networks: Attacks and defenses. In: Gollmann, D., Meier, J., Sabelfeld, A. (eds.) ESORICS 2006. LNCS, vol. 4189, pp. 18–33. Springer, Heidelberg (2006)
34. Tang, C., Goldberg, I.: An improved algorithm for Tor circuit scheduling. In: Proceedings of the 2010 ACM Conference on Computer and Communications Security. ACM, New York (2010)
35. Vahdat, A., Yocum, K., Walsh, K., Mahadevan, P., Kostić, D., Chase, J., Becker, D.: Scalability and accuracy in a large-scale network emulator. SIGOPS Oper. Syst. Rev. 36, 271–284 (2002)
36. Viecco, C.: UDP-OR: A fair onion transport. HotPETS (July 2008)
37. Wang, Z., Crowcroft, J.: Eliminating periodic packet losses in the 4.3-Tahoe BSD TCP congestion control algorithm. SIGCOMM Comput. Commun. Rev. 22, 9–16 (1992)
38. Wright, M.K., Adler, M., Levine, B.N., Shields, C.: The predecessor attack: An analysis of a threat to anonymous communications systems. ACM Trans. Inf. Syst. Secur. 7(4), 489–522 (2004)
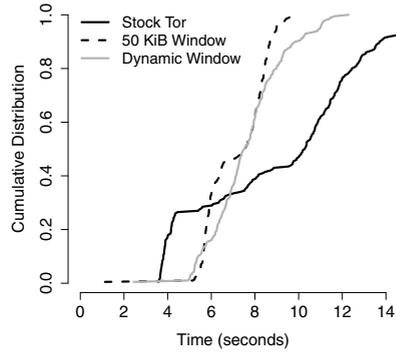
A   End-to-End Windows with Circuit Prioritization

Circuit-level prioritization has been proposed [34] to enable routers to process bursty circuits ahead of bulk circuits. In this appendix, we evaluate small and dynamic circuit windows in combination with circuit-level prioritization[12]. For the web client using stock Tor, the time-to-first-byte is reduced from 4.5 seconds (see Figure 5(a)) to 3 seconds, and the time-to-first-byte for 50 KiB and dynamic windows are roughly the same. However, as shown in Figure 12(a), roughly 25% of requests experience no significant improvement when using small or dynamic circuit windows. For these same requests, stock Tor's large window allows more data in flight without acknowledgment and, as shown in Figure 12(b), induces faster downloads (compared to Figure 5(b)). However, for the remaining 75%, small and dynamic windows offer faster downloads. The bulk client's time-to-first-byte and overall download times are not significantly altered by the circuit prioritization, as shown in Figures 12(c) and 12(d), relative to non-prioritized circuit scheduling (see Figures 5(c) and 5(d)). These observations are consistent with the claims made by Tang and Goldberg [34].
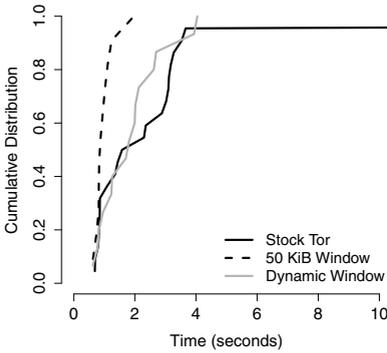
---

[12] For all prioritization experiments, we set `CircuitPriorityHalflifeMsec` to 30 seconds, the current value used on the live Tor network.
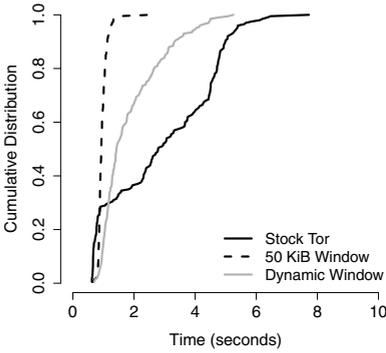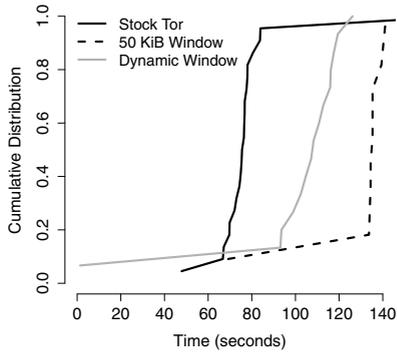
(a) Web client's time-to-first byte

(b) Web client's download time
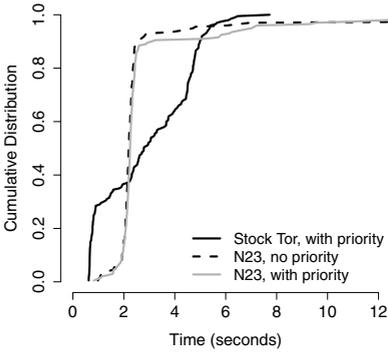

(c) Bulk client's time-to-first-byte
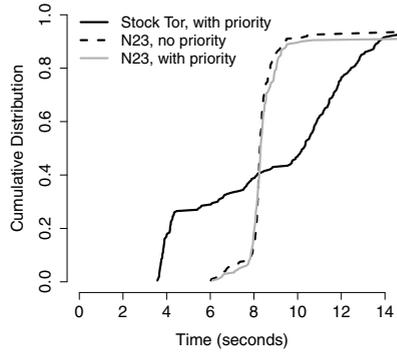

(d) Bulk client's download time

Fig. 12. Performance for window-based flow control with circuit prioritization

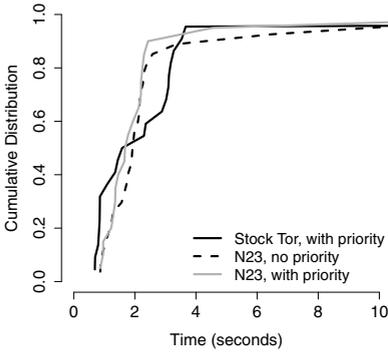# B   N23 with Circuit Prioritization

The circuit-level prioritization algorithm enhances the bursty clients' experience because it remembers how many cells each circuit has recently sent, and gives more priority to the circuits that have sent less. For stock Tor, this algorithm is useful since circuit queues can grow to 1000 cells, which means bulk-traffic circuits can grow large queues and are able to send continuously. However, with N23, circuit queue sizes are significantly smaller and are equal for both bulk and bursty clients. This allows both applications to have a fairer share of the bandwidth. Therefore, for N23, circuit-level prioritization does not provide any performance benefits. Figures 13(a)–13(d) depict the results of the performance of N23 in combination with circuit-level prioritization. Both time-to-first-byte and download times are unaffected by enabling prioritization.
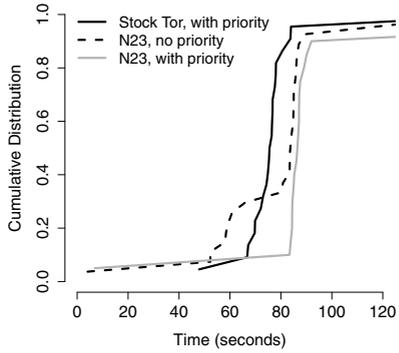
(a) Web client's time-to-first byte

(b) Web client's download time

(c) Bulk client's time-to-first-byte

(d) Bulk client's download time

Fig. 13. Performance for N23 ($N3 = 100$, $N2 = 20$) with circuit prioritization