Padding-only defenses add delay in Tor

Ethan Witwer, James Holland, Nicholas Hopper
{witwe004,holla556,hoppernj}@umn.edu
University of Minnesota

ABSTRACT

Website fingerprinting is an attack that uses size and timing characteristics of encrypted downloads to identify targeted websites. Since this can defeat the privacy goals of anonymity networks such as Tor, many algorithms to defend against this attack in Tor have been proposed in the literature. These algorithms typically consist of some combination of the injection of dummy "padding" packets with the delay of actual packets to disrupt timing patterns. For usability reasons, Tor is intended to provide low latency; as such, many authors focus on padding-only defenses in the belief that they are "zero-delay." We demonstrate through Shadow simulations that by increasing queue lengths, padding-only defenses add delay when deployed network-wide, so they should not be considered "zero-delay." We further argue that future defenses should also be evaluated using network-wide deployment simulations.

CCS CONCEPTS

• **Security and privacy** → Pseudonymity, anonymity and untraceability;

KEYWORDS

website fingerprinting; padding; delay; simulation

1 INTRODUCTION

Tor [9] is a low-latency anonymity network and web browser used daily by millions of users to evade state and corporate surveillance and censorship. Tor redirects connections through multiple intermediaries using layered encryption so that no single entity knows both the source and destination of a connection. This hides which sites a user visits and any contents requested from those sites.

However, Tor has been shown to be vulnerable to several traffic analysis attacks, including *Website Fingerprinting* (WF) attacks [15]. WF attacks use information about the timing, sequence, and volume of packets sent between a client and the Tor network to detect whether a client is downloading a targeted, or "monitored," website. These attacks have been found to be highly effective against Tor, identifying targeted websites with accuracy as high as 99% [6] under some conditions.

This has led to many proposed defenses that modify the characteristics of a connection in order to confuse a WF attacker. A canonical "high-overhead" WF defense is BuFlo [10]: BuFlo sets a constant traffic rate for every connection; in every fixed-length time slot, a Tor node sends one packet on each connection. If multiple packets are pending on a connection, some must be *delayed* to the next slot, and if no packets are pending, an encrypted *padding* packet is sent instead. By

making all downloads have roughly the same sequence, timing, and volume characteristics, BuFlo greatly reduces the accuracy of WF attacks, but it has a very high cost in terms of the additional bandwidth overhead and latency incurred.

Because of the resource constraints of the Tor network and its focus on providing low latency, the literature contains many proposals for WF defense schemes that have lower costs than BuFlo; we briefly describe several of these in Section 2. Typically, the cost of these defenses is measured either by *trace simulation*, in which the traces of several page downloads over Tor are captured and a simulator is used to add or delay packets in the trace according to the defense; or by implementing the defense as a *pluggable transport* [1] that encapsulates the connection between a single Tor client and relay in a defended connection. Researchers then compute the bandwidth overhead by comparing the number of bytes transmitted when downloading a site with and without the WF defense, and if the defense involves adding packet delays, the latency overhead is computed by comparing the overall time to download a site with and without the defense.

For padding-only defenses, this evaluation will not add any latency, and so such defenses are sometimes referred to as "zero-delay." However, if such defenses were to be implemented on a network-wide scale, the added packets sent as padding would necessarily consume resources that could otherwise be used to send non-padding packets to other clients. If enough padding is added, such a defense may actually delay connections *more* than a defense that uses less padding but sometimes delays packets. Thus, we contend that it is important when evaluating WF defenses to account for the effects of deploying to the full Tor network.

In this paper, we use the Shadow [17] network simulator to evaluate the effect of deploying three padding-only defenses – REB [25], Spring, and Interspace [28] – on a network-wide basis. By measuring the progress of downloads over time compared to results from the same network without padding, and comparing the download time for files to the number of padding packets injected per download, we show that padding-only defenses cause delay and that more padding causes additional delay. This illustrates that previous methodologies for measuring delay overhead give an incomplete picture of the costs of WF defense techniques.

Our contributions. We make the following contributions:

- We give the first full-network simulations of padding defense deployments. These show the importance of evaluating additional defenses in this setting.
- We propose a new methodology for WF defense evaluation, time-to-nth-byte. Tracking download progress over time gives a more accurate and complete depiction of the delay incurred by a WF defense mechanism.

- We show that "zero-delay" padding defenses cause delay and should not be over-prioritized in comparison with timing-based defenses.

2 BACKGROUND

Tor. The Tor network [9] is made up of several thousand volunteer-operated *relays* distributed globally, which provide service to roughly one million concurrent users at any time [18]. To connect to a website using Tor, a client constructs a three-hop circuit consisting of a *guard* relay, a *middle* relay, and an *exit* relay. The circuit uses layered encryption so that the guard can only see that it forwards from the client to the middle relay, the middle relay can only see that it forwards from a guard to an exit relay, and the exit relay can see the destination traffic but not the client identity. All traffic between the client and relays is encapsulated into 512-byte *cells*.

Website fingerprinting attacks. While neither an ISP on the network path between the client and the guard nor the guard relay itself can see the *contents* or ultimate *destination* of Tor cells, they can observe the timing, direction, and volume of cells sent in each direction on this connection, which are fairly consistent for any given website. WF attacks use statistical and machine-learning techniques to build "fingerprints" of these sequences for a set of targeted websites. A series of results [6, 13, 14, 24, 27, 30, 31, 33, 35] has shown with increasing accuracy that with no defenses in place to modify these sequences, WF attacks can identify visits to a targeted web page with over 99% accuracy in some settings.

Defenses. To defend against such attacks, it is necessary to modify the sequence of cells observed by the attacker. One way to do this is to inject extra "dummy cells" into either the connection between the client and the guard relay or the tunnel between the client and the middle relay (depending on whether the guard is considered a potential WF attacker or not). Based on the results of Shmatikov and Wang [32], and Juarez *et al.* [21], the Tor project has implemented a circuit padding framework [2] that can deploy stochastic state machines that adaptively pad a connection to fill unlikely gaps between cells. Matthews, Sirinam, and Wright used this framework to implement the "Random Extend Burst" (REB) defense mechanism [25], and Pulls used genetic algorithms to find the locally optimal padding machines Spring and Interspace [28]. Several other defenses [4, 5, 11] also use padding to disrupt fingerprints, but none have been implemented in the circuit padding framework.

In addition to padding, a defense could elect to delay the transmission of some actual cells to disrupt fingerprints, in an attempt to make a website's fingerprint match a different sequence, such as constant-rate traffic [7, 8], A "decoy trace" [36], a common but evolving traffic rate [16, 23], or an adversarially generated pattern [12, 26]. These defenses often have more easily-stated security arguments, but because they induce latency by sometimes delaying cells, they are seen as unacceptable in the low-latency context of Tor.

Attack and Defense Evaluation. Juarez *et al* [20] argued in 2014 that research on WF attacks was unfairly privileging the *attacker* by restricting to an unrealistic setting. These advantages include the use of "closed world" assumption in which only a fixed number of websites could be visited, the "single tab browsing" assumption, the use of website frontpages rather than subpages, and the lack of consideration of the effects of network conditions and retraining. Subsequent works have relaxed many of these conditions, though recently Cherubin, Jansen, and Troncoso [3] showed that modern techniques still degrade quickly with the size of the "monitored" set.

In contrast, Wang [34] and Pulls and Dahlberg [29] have argued that WF *defenses* should be evaluated in the most optimistic setting for attacks, assuming a single monitored page and even an oracle that can eliminate false positives, because preventing these attacks also protects against weaker attacks.

Since few defense mechanisms can provide such strong protection, we advocate for a more realistic evaluation of the *deployment cost* of WF defenses.

3 EXPERIMENTS AND RESULTS

3.1 Experimental Setup

The main tool we used to collect data was Shadow, a network simulator designed to enable realistic and reproducible experiments with Tor. Shadow allows configuration of a network *graph*, which specifies nodes (network-connected hosts), how they are connected, and the processes that they should run. We sought to create a network graph resembling the live Tor network, with relays running the Tor process, and clients and servers running Tor and other processes to send and receive traffic over Tor.

We used TorNetTools [19] for this purpose. TorNetTools uses Tor Project metrics [22] to create *models* of the Tor network, which can be scaled down to account for resource constraints, specifically memory limitations, while remaining representative of the composition of the network. However, this scaling process involves random sampling, which can introduce error that might affect the conclusions of an experiment. Similarly, each simulation of a model involves random sampling, which could also introduce error.

To account for this, we generated multiple models and ran several simulations on each to produce our results. Using Tor Project metrics from March 2022, we generated two models at 0.5% of the scale of the live Tor network at that time and two models at 0.75% scale. We ran 5 simulations of each model with the default Tor configuration, which does not include any WF defenses, for a total of 20 simulations. We also compiled Tor with the padding-only defenses REB [25], Spring, and Interspace [28], running a total of 20 simulations with each defense.

The models generated by TorNetTools consist of Tor relays; clients and servers that generate realistic network traffic and send it over Tor; and 100 *benchmarking clients*, which repeatedly perform 50 KiB, 1 MiB, and 5 MiB downloads throughout a simulation. We used these benchmarking clients to measure bandwidth overhead, delay, and failure rate. We begin by looking at bandwidth overhead, the typical measurement used when reporting the cost of a padding-only defense, with standard Tor, REB, Spring, Interspace.
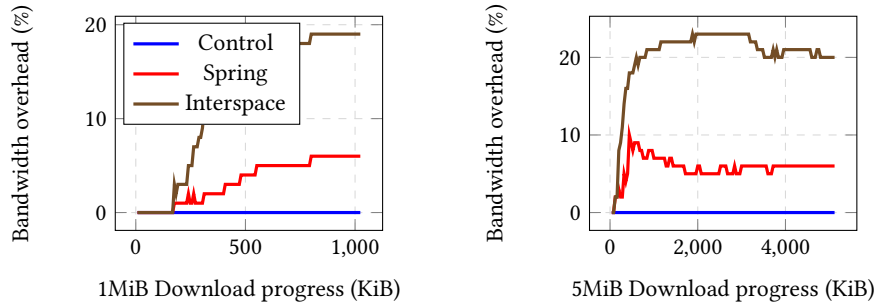
Figure 1: Bandwidth overhead: Median of receive bandwidth overhead with standard Tor, Spring, and Interspace
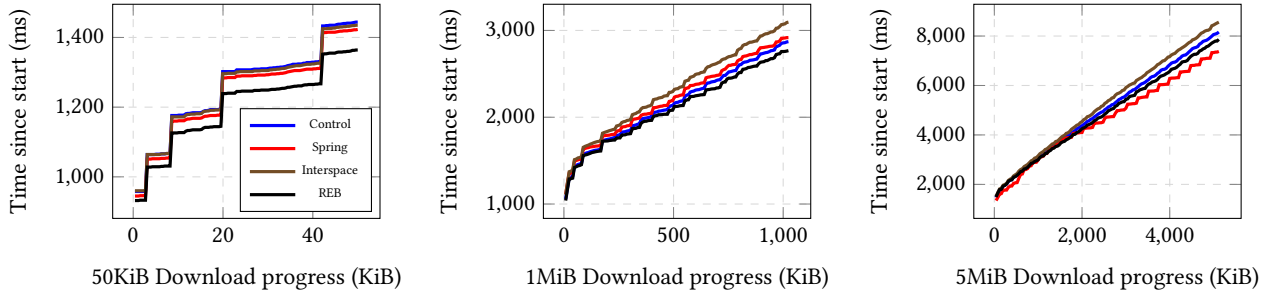

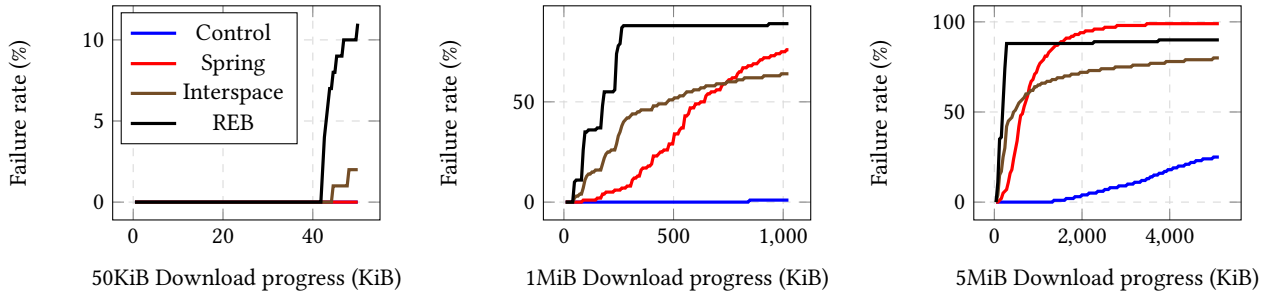**Figure 2: Delay: Time-to-byte data with standard Tor, Spring, Interspace, and REB**

Figure 3: Failure rate: Percentage of attempted downloads that failed with standard Tor, REB, Spring, and Interspace

3.2 Bandwidth overhead

To visualize bandwidth overhead, we make a distinction between bytes of *content* received during a download and bytes of *padding* received. This allows us to examine the bytes of padding that have been received so far at any point in a download as a percentage of the bytes of content received. This is often referred to as *receive* bandwidth overhead; similarly, *send* bandwidth overhead is the padding sent as a percentage of the bytes of content sent, and total overhead takes both bytes sent and bytes received into account.

During each simulation, we measured the receive bandwidth overhead at various times throughout every download made by the benchmark clients. After filtering out the partial results of failed downloads, we calculated the median at each time for the three download sizes with standard Tor (the *control*), REB, Spring, and Interspace. We omit the results for REB since its median receive bandwidth overhead was 0% for every download size. We also exclude all 50 KiB results as overhead was 0% for the control and with each defense; the data for the 1 MiB and 5 MiB download sizes is shown in Figure 1.

For both download sizes, we observed lower values than those originally reported by the authors of the three defenses (for REB, 83% [25]; for Spring, 89%; and for Interspace, 88% [28]). This is likely due to two factors: first, the different traffic patterns induced by single-file downloads as opposed to the web page downloads used as a basis for comparison in previous work will naturally lead to different patterns of adaptive cover traffic; and second, we observed a large fraction of failed downloads apparently caused by errors in the Tor circuit padding framework code, which we explore in the next section.

3.3 Delay

To measure delay, we recorded a timestamp and number of bytes each time data was received by a benchmark client. This allowed us to determine the time taken to reach any given byte count; that is, to examine the progress of each download over time. We aggregated these results over all simulations to obtain the median time to reach a number of progress points for the control and three defenses. In Figure 2, we compare the results for the 50 KiB, 1 MiB and 5 MiB download sizes.

It may be noted that the 50 KiB results are quite counter-intuitive: the control actually had the highest median latency to every byte, and the median time to reach the 50 KiB mark for REB specifically was 5.5% shorter than that of the control. The results for 5 MiB downloads are similar: although Interspace appears to have caused some additional delay, REB and Spring had shorter median times to 5 MiB than the control.

As with the bandwidth overhead data in Figure 1, we note that the time-to-byte data in Figure 2 does not include the partial results of failed downloads. Since REB had a very high failure rate but little padding overhead, it is likely that download failures kept the total load on the network below that of the control throughout each simulation, allowing downloads that did succeed to complete more quickly.

Similarly, although Spring and Interspace involved more padding overhead, they both had very high failure rates for the 1 MiB and 5 MiB download sizes. As different benchmark clients performed downloads of all three sizes in parallel, it is likely that these download failures reduced the load on the network enough to affect the timing of successful downloads, including 50 KiB downloads.

To further illustrate this, we note that the time-to-byte data for 5 MiB downloads would seem to indicate that Spring had a median time to 5 MiB that was 9.6% lower than that of the control. However, Spring had a 99% failure rate for the 5 MiB download size, whereas the control had a 25% failure rate, so this data is based on a much smaller pool of downloads that likely occurred when there was less load on the network.

The 1 MiB download size data comes closer to reflecting our expectations: it indicates that Spring and Interspace did add delay, although REB remained similar to the control. At the 25% mark,
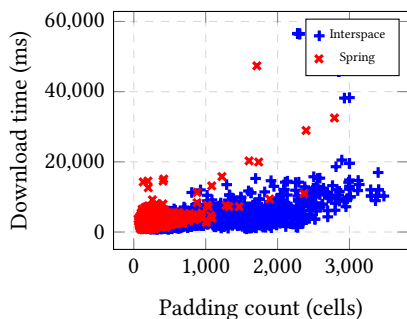
Figure 4: Download time vs padding count (1 MiB)

Spring had a median latency overhead of 2.7% and Interspace of 6.3%; at completion, Spring and Interspace had median latency overheads of 1.6% and 7.8%, respectively. REB finished slightly below the control, but had the highest median failure rate (89%), which supports our conclusions about download failures. Figure 4 also shows this trend, with download time increasing as a function of padding cells, but with an R^2 of only 0.21 for Interspace and 0.28 for Spring.

Although the precise mechanisms by which these failures occurred are unclear, we found that downloads were more likely to fail as the total number of padding cells sent and received increased (except REB, since it had negligible padding overhead) as seen in Figure 5. That is, most of the downloads which did succeed involved a relatively low number of padding
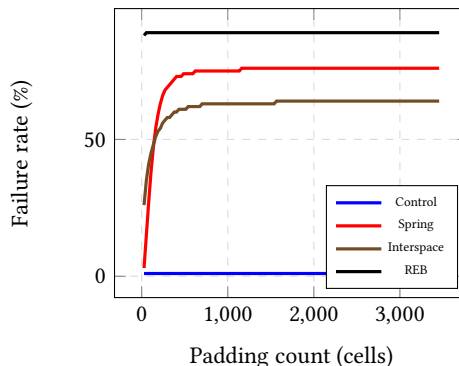
Figure 5: Failure rate vs padding count (1 MiB)

cells. This means that, in the absence of failures, the delay incurred by each defense should be expected to be higher than what we report.

Ultimately, even though we can't conclude that these defenses would result in a similar failure rate on the live Tor network, we suspect that latency overhead would be greater than what we observed in practice. Both cases represent potentially significant impacts on the usability of the Tor network, suggesting that rigorous evaluation of all potential usability factors is necessary before the deployment of any defense.

4 CONCLUSIONS AND FUTURE WORK

Due to the prevalence of download failures in our simulations, it is difficult to draw precise conclusions about the relationship between padding overhead and delay, but it is clear that padding-based defenses place additional stress on the Tor network, leading to nonzero added delays. This is especially true since download *failures* will either lead to higher delays as users reload pages and resources that fail to download, or simply lead to more severe user frustrations than those caused by the increased latency WF defense designers hope to avoid.

Our results highlight the fact that WF defenses require more complete evaluation than simple trace-based simulation or single-edge deployments, because these evaluations fail to capture the interaction between multiple defended circuits in the Tor network. They also suggest that the existing circuit padding framework may not be ready for deployment on the live network. Finally, in light of this need for further evaluation, it may be desirable to consider allowing cell delays, since adding padding already incurs latency, and cell delays may in fact reduce the resource stress caused by WF defenses.

Thus, an interesting avenue for future work is to explore how the circuit padding framework can be modified to allow delays, and to implement more recent low-overhead defenses such as FRONT [11], RegulaTor [16] and Surakav [12]. This will also allow more complete evaluations of these defenses. Similarly, considering our observation that TorNetTools measurements lead to different bandwidth overheads than previous workloads, another interesting avenue for future work is to explore the overhead of WF defenses on more realistic workloads. Time-to-byte measurements can help to normalize comparisons across workloads, but other metrics based

on time-to-event for various browser events may be useful as well.

REFERENCES

[1] 2016. Tor at the Heart: Bridges and Pluggable Transports. https://blog.torproject.org/tor-heart-bridges-and-pluggable-transports. (2016).

[2] 2020. Circuit Padding Developer Documentation. https://github.com/torproject/tor/blob/main/doc/HACKING/CircuitPaddingDevelopment.md. (2020).

[3] 2022. Online Website Fingerprinting: Evaluating Website Fingerprinting Attacks on Tor in the Real World. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA. https://www.usenix.org/conference/usenixsecurity22/presentation/cherubin

[4] Ahmed Abusnaina, Rhongho Jang, Aminollah Khormali, DaeHun Nyang, and David Mohaisen. 2020. DFD: Adversarial learning-based approach to defend against Website fingerprinting. In *IEEE INFOCOM 2020 -IEEE Conference on Computer Communications*. 2459–2468.

[5] Khaled Al-Naami, Amir El-Ghamry, Md Shihabul Islam, Latifur Khan, Bhavani Thuraisingham, Kevin W. Hamlen, Mohammed Alrahmawy, and Magdi Z. Rashad. 2021. BiMorphing: A Bi-Directional Bursting Defense against Website Fingerprinting Attacks. *IEEE Transactions on Dependable and Secure Computing (TDSC)* 18, 2 (2021), 505–517.

[6] Sanjit Bhat, David Lu, Albert Kwon, and Srinivas Devadas. 2019. Var-CNN: A Data-Efficient Website Fingerprinting Attack Based on Deep Learning. *Proceedings on Privacy Enhancing Technologies* 4 (2019), 292–310.

[7] Xiang Cai, Rishab Nithyanand, Tao Wang, Rob Johnson, and Ian Goldberg. 2014. A systematic approach to developing and evaluating website fingerprinting defenses. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 227–238.

[8] Xiang Cai, Xin Cheng Zhang, Brijesh Joshi, and Rob Johnson. 2012. Touching from a distance: Website fingerprinting attacks and defenses. In *Proceedings of the 2012 ACM conference on Computer and communications security*. 605–616.

[9] Roger Dingledine, Nick Mathewson, and Paul F. Syverson. 2004. "Tor: The second-generation Onion router". In *USENIX Security Symposium*. 303–320.

[10] Kevin P. Dyer, Scott E. Coull, Thomas Ristenpart, and Thomas Shrimpton. 2012. Peek-a-Boo, I still see you: Why efficient traffic analysis countermeasures fail. In *IEEE Symposium on Security and Privacy (S&P)*. 332–346.

[11] Jiajun Gong and Tao Wang. 2020. Zero-delay lightweight defenses against website fingerprinting. In *29th USENIX Security Symposium (USENIX Security 20)*. 717–734.

[12] Jiajun Gong, Wuqi Zhang, Charles Zhang, and Tao Wang. 2022. Surakav: Generating Realistic Traces for a Strong Website Fingerprinting Defense. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 1525–1525.

[13] Jamie Hayes and George Danezis. 2016. k-fingerprinting: A robust scalable website fingerprinting technique. In *USENIX Security Symposium*. 1–17.

[14] Dominik Herrmann, Rolf Wendolsky, and Hannes Federrath. 2009. Website fingerprinting: attacking popular privacy enhancing technologies with the multinomial Naïve-Bayes classifier. In *ACM Workshop on Cloud Computing Security*. 31–42.

[15] Andrew Hintz. 2002. Fingerprinting websites using traffic analysis. In *International workshop on privacy enhancing technologies*. Springer, 171–178.

[16] James K. Holland and Nicholas Hopper. 2022. RegulaTor: A Straightforward Website Fingerprinting Defense. *Proc. Priv. Enhancing Technol.* 2022, 2 (2022), 344–362. https://doi.org/10.2478/popets-2022-0049

[17] Rob Jansen and Nicholas Hopper. 2012. Shadow: Running Tor in a Box for Accurate and Efficient Experimentation. In *19th Annual Network and Distributed System Security Symposium, NDSS 2012, San Diego, California, USA, February 5-8, 2012*. The Internet Society. https://www.ndss-symposium.org/ndss2012/shadow-running-tor-box-accurate-and-efficient-experimentation

[18] Rob Jansen and Aaron Johnson. 2016. Safely measuring tor. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 1553–1567.

[19] Rob Jansen, Justin Tracey, and Ian Goldberg. 2021. Once is Never Enough: Foundations for Sound Statistical Inference in Tor Network Experimentation. In *30th USENIX Security Symposium (Sec)*. See also https://neverenough-sec2021.github.io.

[20] Marc Juarez, Sadia Afroz, Gunes Acar, Claudia Diaz, and Rachel Greenstadt. 2014. A critical evaluation of website fingerprinting attacks. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 263–274.

[21] Marc Juarez, Mohsen Imani, Mike Perry, Claudia Diaz, and Matthew Wright. 2016. Toward an efficient website fingerprinting defense. In *European Symposium on Research in Computer Security*. Springer, 27–46.

[22] Karsten Loesing, Steven J. Murdoch, and Roger Dingledine. 2010. A Case Study on Measuring Statistical Data in the Tor Anonymity Network. In *Proceedings of the Workshop on Ethics in Computer Security Research (WECSR 2010) (LNCS)*. Springer.

[23] David Lu, Sanjit Bhat, Albert Kwon, and Srinivas Devadas. 2018. Dynaflow: An efficient website fingerprinting defense based on dynamically-adjusting flows. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*. 109–113.

[24] Liming Lu, EC Chang, and MC Chan. 2010. Website fingerprinting and identification using ordered feature sequences. In *European Symposium on Research in Computer Security (ESORICS)*. Springer, 199–214.

[25] Nate Mathews, Payap Sirinam, and Matthew Wright. 2018. Understanding feature discovery in website fingerprinting attacks. In *2018 IEEE Western New York Image and Signal Processing Workshop (WNYISPW)*. IEEE, 1–5.

[26] Milad Nasr, Alireza Bahramali, and Amir Houmansadr. 2021. Defeating DNN-Based Traffic Analysis Systems in Real-Time With Blind Adversarial Perturbations. In *USENIX Security Symposium*. 2705–2722.

[27] Andriy Panchenko, Fabian Lanze, Andreas Zinnen, Martin Henze, Jan Pennekamp, Klaus Wehrle, and Thomas Engel. 2016. Website fingerprinting at Internet scale. In *Network & Distributed System Security Symposium (NDSS)*. 1–15.

[28] Tobias Pulls. 2020. Towards effective and efficient padding machines for tor. *arXiv preprint arXiv:2011.13471* (2020).

[29] Tobias Pulls and Rasmus Dahlberg. 2020. Website Fingerprinting with Website Oracles. *Proc. Priv. Enhancing Technol.* 2020, 1 (2020), 235–255.

[30] Mohammad Saidur Rahman, Payap Sirinam, Nate Mathews, Kantha Girish Gangadhara, and Matthew Wright. 2019. Tik-Tok: The utility of packet timing in website fingerprinting attacks. *arXiv preprint arXiv:1902.06421* (2019).

[31] Vera Rimmer, Davy Preuveneers, Marc Juárez, Tom van Goethem, and Wouter Joosen. 2018. Automated Website Fingerprinting through Deep Learning. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*. The Internet Society. http://wp.internetsociety.org/ndss/wp-content/uploads/sites/25/2018/02/ndss2018_03A-1_Rimmer

[32] Vitaly Shmatikov and Ming-Hsiu Wang. 2006. Timing analysis in low-latency mix networks: Attacks and defenses. In *European Symposium on Research in Computer Security*. Springer, 18–33.

[33] Payap Sirinam, Mohsen Imani, Marc Juarez, and Matthew Wright. 2018. Deep fingerprinting: Undermining website fingerprinting defenses with deep learning. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 1928–1943.

[34] Tao Wang. 2021. The one-page setting: A higher standard for evaluating website fingerprinting defenses. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. 2794–2806.

[35] Tao Wang and Ian Goldberg. 2013. Improved website fingerprinting on Tor. In *ACM Workshop on Privacy in the Electronic Society (WPES)*. 201–212.

[36] Tao Wang and Ian Goldberg. 2017. {Walkie-Talkie}: An Efficient Defense Against Passive Website Fingerprinting Attacks. In *26th USENIX Security Symposium (USENIX Security 17)*. 1375–1390.