Free and Open
COMMUNICATIONS
https://foci.community on the Internet

Rethinking Realistic Adversaries for Anonymous Communication Systems

Kevin Gallagher
NOVA LINCS, NOVA School of Science and Technology

Diogo Barradas
University of Waterloo

Nuno Santos
INESC-ID
Instituto Superior Técnico

Free and Open Communications on the Internet 2023(1), 81-87

Rethinking Realistic Adversaries for Anonymous Communication Systems

Kevin Gallagher
NOVA LINCS, NOVA School of Science and Technology

Diogo Barradas
University of Waterloo

Nuno Santos
INESC-ID / Instituto Superior Técnico, Universidade de Lisboa

Abstract

Few anonymity works have deviated from the traditional global passive or global active adversaries, or describe adversaries that have justified limitations on their abilities to see or interfere with traffic. As such, anonymity systems that consider only these adversaries may miss opportunities to make informed trade-offs about security and utility. In this work we motivate the need for new adversaries against anonymous communication systems, and present some early work towards constructing novel practical adversaries in anonymity literature. Specifically, we discuss adversary limitations and expansions to adversary ontologies that could better model real world adversaries.

1 Introduction

Privacy and anonymity allow the average person to express their opinions or desires and contribute to society without fear of retaliation, as well as explore their identity and experiment with new opinions without repercussions [45]. However, the design of TCP/IP and other Internet protocols don't allow for privacy or anonymity; all packets that traverse the Internet are stamped with metadata, and current Internet monetization schemes incentivize websites, Internet Service Providers (ISPs) and others to invade the privacy of their users [13, 15, 47, 51].

Fortunately, previous research [16, 46] has focused on creating anonymity systems that break the linkability between a sender and their messages. Each of these systems provides anonymity with respect to a considered adversary that has capabilities and limitations. As discussed more in Section 2, existing research tends to use a standard set of adversaries.

However, typical adversaries considered in anonymous communication literature may be too strong compared to real adversary capabilities, given the tendency to consider i) the *strongest possible* adversaries (e.g., as in Dissent [12]), or ii) a moderate adversary that passes only heuristic examination (e.g., as in Tor [18]). This leads to the lack of consideration of anonymity systems research that frustrate weaker, more realistic adversaries.

Inspired by the work of Aumann and Lindell [4] in the area of Secure Multiparty Computation and Jaggard et. al [26] in the area of anonymity, we discuss the realistic and unrealistic aspects of the standard adversaries considered in anonymity research and argue for the creation of new, more realistic adversaries. We then discuss different potential expansions of adversary models based on limitations derived from real surveillance programs.

2 Background

Due to the design of the Internet, users must rely on overlay networks to gain anonymity properties. New anonymity systems are analyzed with respect to certain adversaries, who are assumed to be capable of observing and interacting with Internet traffic in different amounts.

In his seminal work on anonymity [10], Chaum considered a very strong adversary model. Any party could see any messages, and could artificially inject, modify, or remove traffic. Though this adversary model did not become standard, it did eventually converge on one model: the *global active adversary*, a single adversary that can see all messages that pass through a network, and can arbitrarily modify, remove, or inject traffic. This differs from Chaum's adversary in that the adversary is considered to be one entity, and not every entity in the protocol is considered to potentially be a global active adversary.

The global active adversary usually focuses on using their active attacks to deanonymize participants by linking communications to the sender. In traffic modification attacks, the adversary can use their ability to modify traffic to i) insert a timing signature, making the traffic more susceptible to correlation attacks, or ii) delay or drop packets to ensure that only one individual is interacting with the system, trivially deanonymizing them. Notable examples of works that consider a global active adversary include Loopix [40], Nym [17], and more.

Chaum's system was later revealed to be weak to an attack called the N-1 attack [48]. However, this adversary may be considered rather unrealistic, as the ability to modify, remove, and inject traffic in real time at any point of the network may defy many geopolitical and economic realities. Thus, works began using different adversaries, such as the *global passive adversary*. This adversary is also assumed to be global, but unlike the global active adversary, it cannot interact with traffic, making it capable of only performing computations and observing the network. Some works that consider a global passive adversary include Riffle [32], cMix [9], Groove [6], and more.

Ironically, global passive adversaries suffer from simultaneously being too weak and too strong to be realistic. The restriction of not being able to interact with traffic makes the global passive adversary too weak; most network-level adversaries have the ability to interact with at least some of the network flows they can observe. Being global is also unrealistic, as economic, geo-political, and legal restrictions make it difficult for one adversary the observe the entirety of the world's Internet infrastructure. Thus, a need arose for a more imperfect but realistic adversary.

3 Imperfect Adversaries

In the original onion routing paper [42], Reed et. al introduce the idea of an adversary that can arbitrarily interact with the packets on only a faction of the network. They also allowed for the adversary to participate in the anonymity network itself, posing as an honest participant. This adversary was then adopted by Tor [18], a large scale low-latency anonymity network with wide adoption.

Although this adversary is a lot closer to a realistic adversary, it remains abstract. Concrete details could be vital in creating an anonymity network that avoids deanonymization by predicting where the adversary is spying and avoiding that portion of the network.

Towards this end, in 2004 Feamster and Dingledine [20] explored the impact of node diversity on adversary abilities to intercept paths of anonymous communication. They find that multiple nodes are in the same Autonomous Systems (ASes) and that single ASes are often able to see both the entry and exit communication path at the same time. In 2009, Edman and Syverson [19] demonstrated that these results were underestimations, and that a large number of Tor circuits begin in a small set of ASes and end in another small set of ASes. They then propose a new heuristic for country-level diversity in Tor circuits.

These works contain a few drawbacks that limit their applicability. For example, Feamster and Dingledine's work assumes that separate ASes won't collude, however the Snowden leaks demonstrate that they do in practice. More, the work assumes that law enforcement agencies would not be willing to face accountability for illegally accessing data. However, recent events [1, 39, 49] demonstrate that law enforcement frequently violate this assumption.

Jaggard et. al [26] took these previous concepts and made them more concrete by building an ontology to express adversary capabilities. They then build *The Man*, an adversary that is capable of observing any independent group of portions of Internet infrastructure (ASes, IXPs, etc.) or Tor relay families. They also construct another adversary through the observation of Mutual Legal Assistance Treaties (MLATs) and access to undersea cables. Though this work begins to address some more subtle nuances of adversaries, it does not focus much on the limitations that adversaries may have in the wild. To this end, we discuss a potential expansion of the ontology presented in [26] and potential adversary limitations that have so-far gone under-researched.

3.1 Spheres of Influence

We argue that to expand the realism of the adversary model presented in [26], we must introduce the notion of *spheres of influence*, described below. More, adversary models should reflect that expanding a sphere of influence comes with a *cost*, typically in time, money, computation, storage, political capital, or all of the above. The specifics of these limitations can be derived from the reported constraints of surveillance states based on documents, news publications, and interviews given by whistleblowers, or data exposed via open source intelligence (osint) operations. We plan to extend this ontology in future work.

We argue for the distinction between two different spheres of influence: *direct influence*, in which an adversary can directly analyze or alter the traffic flowing under their observation, and *indirect influence*, in which the adversary gains access to information through a third party. To motivate this differentiation, we look to the current known practices of one example surveillance state: the United States of America. The National Security Agency (NSA) has multiple programs to collect information about US and foreign citizens. These programs include Xkeyscore [22], a database application used to query information collected through other programs, and PRISM [23], a data collection program for internet data given by third parties. These programs unveil the existence of an *indirect influence* of the NSA over infrastructure owned by third parties. Though Xkeyscore can query data, it relies on databases already being built and populated. Similarly, the PRISM program seems to be a one-way data flow: companies send information that they collect through their operations with users, but may or may not perform active attacks on them.

Adversaries can expand their *indirect influence* through data sharing with other adversaries, such as between the NSA and other members of the Five Eyes [38]. Other methods include the expansion of their partners. It is worth noting that each of these expansions has a cost, both in money and

time. We note that even in the case of warrants through so-called "rubber stamp" courts or "kangaroo courts", the legal processes of submitting a case and waiting on decisions does incur a monetary and time cost, and does not guarantee success. For example, of the 44,269 requests made to the Foreign Intelligence Surveillance Court (FISC) between 1979 and 2022, 2,068 required modifications, increasing cost and delay [8]. Unfortunately, the average amount of time and cost of putting a request through the FISC is unknown, though during our continuation of this work we will attempt to quantify this through Freedom of Information Act requests to include in our models.

There also exist programs for active attacks performed by the NSA. An example of these are the attacks performed under the Turbulence program [21], which uses malware and other cyberwarfare techniques to collect data. The success of such an attack would put an infected machine under the *direct influence* of an adversary, allowing for potential deanonymization using IP address identification for compromised clients or allowing for the injection of traffic patterns for later identification by compromised portions of networks. However, these attacks may have a lower probability of success, or require more cost and overhead, including, hopefully, more legal scrutiny.

With this new distinction included, we see that no adversary is ever fully passive nor active. Adversaries may need to rely on indirect sources to confirm the existance of an injected watermark, for example. More, it is worth noting that adversaries also inherit some limitations of their third-party partners.

3.2 Limitations of Adversaries

The existing literature has overlooked many practical limitations in modeling realistic adversaries against anonymous communication systems. Next, we discuss a non-comprehensive yet significant set of important limitations and some of their implications.

Dealing with lack of coverage. Several traffic confirmation attacks on anonymity networks assume that an adversary has *full coverage*, meaning they have access to points-of-presence in the locations where user traffic enters and exits the anonymity network [34,35,37]. While these attacks are highly accurate and suggest that confirmation attacks pose a realistic threat, they fail to acknowledge the difficulties in achieving such coverage in practice. The cumbersome negotiations and legal processes that adversaries may need to undertake to establish widespread international cooperation agreements for achieving increased coverage (see Section 3.1) are often ignored. More, the research community has proposed multiple defenses that deliberately route anonymity networks' traffic (e.g., Tor) away from adversary-controlled points-of-presence, despite an adversary's potentially broad coverage. These defenses include evading specific autonomous systems (ASes) [2,7,36,50] or entire geographical regions [19,31,33].

Crucially, measuring the potential effects of imprecise claims regarding the accuracy of traffic confirmation attacks in *partial coverage* settings remains unaddressed. Tthe occurrence of false positives in such attacks may lead to the erroneous conclusion that two flows are correlated, even when the true corresponding matching network flow was never observed by the adversary. It is paramount to assess the implications of these attacks in real-world scenarios, where nation-state adversaries may lack coverage but still rely and act upon the results of (potentially inaccurate) traffic confirmation attacks to retaliate against citizens.

Dealing with time-bounded observations. An orthogonal dimension to the spatial coverage of communications is the temporal capability of an adversary to perform observations. Prior research by Wright et al. [53] established that an adversary with infinite time for observation could increase the precision of identifying communication parties to near certainty. Similarly, Danezis and Serjantov [14] demonstrate a statistical disclosure attack that observes each communication round and uses these observations to calculate a probability that a given client is communicating with a given server. However, in real-world scenarios, an adversary's observation time is inherently limited. In 2003 work by Kedogan et. al [30] measured the effectiveness of observation based attacks against mix networks and determined that the number of observations needed were related to the batch size, number of users, and the number of peer-partners. The first comprehensive study assessing the vulnerability of Tor users to traffic correlation attacks over time was conducted by Johnson et al. [27]. This research showed that if an adversary controls both the guard and exit nodes for a period of six months, the likelihood of successful deanonymization can rise up to 80%. Nevertheless, assessing to what extent realistic adversaries can perform such long-term observations remains an open question. For instance, the extensive storage requirements for logging prolonged observations of communications may constitute a serious obstacle to launching such attacks. Thus, the extent to which real-world adversaries can conduct such sustained observations is still uncertain.

Dealing with mutual distrust. Another important limitation of an adversary that relies on the collusion of multiple parties to deanonymize users' traffic is the potential issue of mutual distrust amongst the colluding parts. Indeed, although a global adversary may have access to points of presence distributed around the globe due to international cooperation agreements, it is unclear whether all parties would agree to provide timely and accurate data to enable deanonymization attempts.

For instance, consider a scenario where country X aims to deanonymize a traffic flow, but country Y is cognizant that one of their own overseas agents was using the anonymity system during the timeframe in question. In such a situation, nation Y may become suspicious that state X intends to deanonymize nation Y's own agent's flow and either a) periodically refuse

to grant access to all of their data, or b) provide tampered data to impede correlation efforts at specific times. This scenario alludes to (and is further exacerbated by) the needs of an ideal global adversary – to achieve full coverage, nations will need to cooperate with other nations perceived as enemies.

To the best of our knowledge, the issue of mutual distrust amongst colluding entities wishing to attack anonymity networks has received no attention in the literature, despite its potential impact on the accuracy of traffic confirmation attacks' results. As it stands, the collaboration between colluding entities is perceived as a binary decision, but the lack of transparency (or verification) about the data collected and exchanged between colluding parties opens the floor for discussion about new obstacles and complicating factors during traffic confirmation attacks.

Dealing with concurrent active adversaries. Current attacks on anonymity networks typically assume the existence of a single distributed adversary that benefits from the collusion of multiple entities across the world. However, geopolitical relationships between different countries can be tense and may prevent full and honest cooperation between them. Thus, it may be more realistic to expect the emergence of concurrent, active, and partial coverage-enabled global adversaries with conflicting goals and interests, such as the countries composing the Five Eyes alliance [38] and the Shanghai Cooperation Organisation [3]. This concern was also previously shared by Johnson et al. [28], who considered the existence of adversaries with conflicting goals when modelling trust in anonymous communications. In their work, Johnson et al. assumed these adversaries could not only host their own network nodes for aiding deanonymization efforts, but also simultaneously compromise the same network nodes (even those controlled by each other). Johnson et al. also hint at the concept of an adversary-learning adversary, who can potentially observe a user's communication patterns and infer which adversary the user is trying to resist.

We argue that the rise of multiple concurrent adversaries may lead to interference with each other's ability to successfully perform traffic confirmation attacks in certain conditions. Consider attacks on anonymity networks that rely on traffic manipulations to introduce watermarks or other distinguishing features that aid in traffic confirmation efforts [24, 43, 44]. Since this kind of attack requires a high degree of accuracy in the timing and sequencing of network traffic, any form of network interference or perturbation can disrupt the effectiveness of traffic confirmation attacks. However, the presence of multiple adversaries that simultaneously attempt to introduce watermarks in network traffic (or normalize traffic to eliminate possible watermarks on flows crossing the network links they are able to observe) can wreak havoc in the predictable patterns of network traffic, effectively stopping adversaries from recognizing their own watermarks. Thus, adversaries may require more complex and resilient traffic confirmation techniques to operate in the presence of other concurrent active adversaries. This underscores the importance of understanding the impact of network conditions [11, 29] towards developing defenses against traffic analysis attacks.

Dealing with user data leakage and data protection laws. Traffic confirmation attacks are typically carried out under the assumption that a set of probes will share the necessary data to correlate traffic flows with the help of a correlator node. However, the correlator node will have access to all the data shared by each probe, even if the probes have removed PII or contents of the flows they collected. As a result, a correlator node may be able to leverage traffic analysis to infer specific information about the non-target users observed by that probe, such as identifying the websites they visited [41] or whom they spoke to in instant-messaging applications [5]. This problem is aggravated by the fact that nations colluding to increase their chances of correlating traffic may not place trust in a central correlation entity outside their control.

In addition, it is plausible that ISPs may need to abide by region-specific data protection laws. While these ISPs may still collaborate in traffic confirmation efforts, they may require increased privacy protections to keep the data about users away from the prying eyes of an honest-but-curious correlator. Thus, we envision that correlator nodes may be limited to the use of privacy-preserving computation schemes, like those based in multiparty computation [25, 52]. These schemes are known to be slow, adding to the already high cost and combinatorial nature of traffic confirmation used to attack anonymity networks, and making such schemes less attractive to use in practice.

4 Conclusions

In this work we outlined the common adversaries assumed in anonymity works and discussed the limitations adversaries have that are not reflected in the academic literature. We then introduced the concept of an adversary's *direct* and *indirect spheres of influence*, and argued for the need to expand the ontology presented in [26]. In future work we will expand this ontology with these concepts using information about real world mass surveillance programs.

Acknowledgments

We thank the anonymous reviewers for their comments and insightful feedback. This work was partially supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) under project UIDB/50021/2020, NOVA LINCS (UIDB/04516/2020) with the financial support of FCT.IP, the SmartRetail project (ref. C6632206063-00466847) financed by IAPMEI, and by NSERC under grant DGECR-2023-00037.

References

[1] Trevor Aaronson. Court ruling shows how fbi abused nsa mass surveillance. https://theintercept.com/2019/10/10/fbi-nsa-mass-surveillance-abuse/, Oct 2019.

[2] Masoud Akhoondi, Curtis Yu, and Harsha V. Madhyastha. Lastor: A low-latency as-aware tor client. In *2012 IEEE Symposium on Security and Privacy*, pages 476–490, 2012.

[3] Stephen Aris. The Shanghai Cooperation Organisation:'Tackling the three evils'. A regional response to non-traditional security challenges or an anti-Western bloc? *Europe-Asia Studies*, 61(3):457–482, 2009.

[4] Yonatan Aumann and Yehuda Lindell. Security against covert adversaries: Efficient protocols for realistic adversaries. In Salil P. Vadhan, editor, *Theory of Cryptography*, pages 137–156, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.

[5] Alireza Bahramali, Ramin Soltani, Amir Houmansadr, Dennis Goeckel, and Don Towsley. Practical Traffic Analysis Attacks on Secure Messaging Applications. page arXiv:2005.00508, May 2020.

[6] Ludovic Barman, Moshe Kol, David Lazar, Yossi Gilad, and Nickolai Zeldovich. Groove: Flexible Metadata-Private messaging. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 735–750, Carlsbad, CA, July 2022. USENIX Association.

[7] Armon Barton and Matthew Wright. Denasa: Destination-naive as-awareness in anonymous communications. *Proceedings on Privacy Enhancing Technologies*, 2016(4).

[8] Electronic Privacy Information Center. Foreign intelligence surveillance act court orders 1979-2022. https://epic.org/foreign-intelligence-surveillance-court-fisc/fisa-stats/, 2023.

[9] David Chaum, Debajyoti Das, Farid Javani, Aniket Kate, Anna Krasnova, Joeri De Ruiter, and Alan T. Sherman. cmix: Mixing with minimal real-time asymmetric cryptographic operations. In Dieter Gollmann, Atsuko Miyaji, and Hiroaki Kikuchi, editors, *Applied Cryptography and Network Security*, pages 557–578, Cham, 2017. Springer International Publishing.

[10] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, feb 1981.

[11] Giovanni Cherubin, Rob Jansen, and Carmela Troncoso. Online website fingerprinting: Evaluating website fingerprinting attacks on tor in the real world. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 753–770, Boston, MA, August 2022. USENIX Association.

[12] Henry Corrigan-Gibbs and Bryan Ford. Dissent: Accountable anonymous group messaging. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, CCS '10, page 340–350, New York, NY, USA, 2010. Association for Computing Machinery.

[13] Levente Csikor, Himanshu Singh, Min Suk Kang, and Dinil Mon Divakaran. Privacy of dns-over-https: Requiem for a dream? In *2021 IEEE European Symposium on Security and Privacy (EuroSP)*, pages 252–271, 2021.

[14] George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In Jessica Fridrich, editor, *Information Hiding*, pages 293–308, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

[15] Ha Dao and Kensuke Fukuda. Alternative to third-party cookies: Investigating persistent pii leakage-based web tracking. In *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '21, page 223–229, New York, NY, USA, 2021. Association for Computing Machinery.

[16] Debajyoti Das, Sebastian Meiser, Esfandiar Mohammadi, and Aniket Kate. Anonymity trilemma: Strong anonymity, low bandwidth overhead, low latency - choose two. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 108–126, 2018.

[17] Claudia Diaz, Harry Halpin, and Aggelos Kiayias. The Nym network. https://nymte.ch/nym-whitepaper.pdf, 2021.

[18] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th Conference on USENIX Security Symposium - Volume 13*, SSYM'04, page 21, USA, 2004. USENIX Association.

[19] Matthew Edman and Paul Syverson. As-awareness in tor path selection. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, CCS '09, page 380–389, New York, NY, USA, 2009. Association for Computing Machinery.

[20] Nick Feamster and Roger Dingledine. Location diversity in anonymity networks. In *Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society*, WPES '04, page 66–76, New York, NY, USA, 2004. Association for Computing Machinery.

[21] Siobhan Gorman. Costly nsa initiative has a shaky takeoff. http://articles.baltimoresun.com/2007-02-11/news/0702110034_1_turbulence-cyberspace-nsa, 2007.

[22] Glenn Greenwald. XKeyscore: NSA tool collects' nearly everything a user does on the internet. *The Guardian*, 31, 2013.

[23] Glenn Greenwald and Ewen MacAskill. NSA Prism program taps into user data of Apple, Google and others. *The Guardian*, 7(6):1–43, 2013.

[24] Zhong Guan, Chang Liu, Gang Xiong, Zhen Li, and Gaopeng Gou. Flowtracker: Improved flow correlation attacks with denoising and contrastive learning. *Computers Security*, 125:103018, 2023.

[25] Marcella Hastings, Brett Hemenway, Daniel Noble, and Steve Zdancewic. Sok: General purpose compilers for secure multi-party computation. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1220–1237, 2019.

[26] Aaron D Jaggard, Aaron Johnson, Sarah Cortes, Paul Syverson, and Joan Feigenbaum. 20,000 in league under the sea: Anonymous communication, trust, mlats, and undersea cables. *Proceedings on Privacy Enhancing Technologies*, 2015(1):4–24, 2015.

[27] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. Users get routed: Traffic correlation on tor by realistic adversaries. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer Communications Security*, CCS '13, page 337–348, New York, NY, USA, 2013. Association for Computing Machinery.

[28] Aaron M. Johnson, Paul Syverson, Roger Dingledine, and Nick Mathewson. Trust-based anonymous communication: Adversary models and routing algorithms. In *Proceedings of the 18th ACM Conference on Computer and Communications Security*, CCS '11, page 175–186, New York, NY, USA, 2011. Association for Computing Machinery.

[29] Marc Juarez, Sadia Afroz, Gunes Acar, Claudia Diaz, and Rachel Greenstadt. A critical evaluation of website fingerprinting attacks. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, page 263–274, New York, NY, USA, 2014. Association for Computing Machinery.

[30] Dogan Kedogan, Dakshi Agrawal, and Stefan Penz. Limits of anonymity in open environments. In Fabien A. P. Petitcolas, editor, *Information Hiding*, pages 53–69, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

[31] Katharina Kohls, Kai Jansen, David Rupprecht, Thorsten Holz, and Christina Pöpper. On the challenges of geographical avoidance for tor. In *Proceedings of 26th Annual Network and Distributed System Security Symposium*, 2019.

[32] Albert Kwon, David Lazar, Srinivas Devadas, and Bryan Ford. Riffle: An efficient communication system with strong anonymity. *Proceedings on Privacy Enhancing Technologies*, 2016(2):115–134, 2016.

[33] Zhihao Li, Stephen Herwig, and Dave Levin. Detor: Provably avoiding geographic regions in tor. In *Proceedings of the 26th USENIX Conference on Security Symposium*, SEC'17, page 343–359, USA, 2017. USENIX Association.

[34] Milad Nasr, Alireza Bahramali, and Amir Houmansadr. Deepcorr: Strong flow correlation attacks on tor using deep learning. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18, page 1962–1976, New York, NY, USA, 2018. Association for Computing Machinery.

[35] Milad Nasr, Amir Houmansadr, and Arya Mazumdar. Compressive traffic analysis: A new paradigm for scalable traffic analysis. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, page 2053–2069, New York, NY, USA, 2017. Association for Computing Machinery.

[36] Rishab Nithyanand, Oleksii Starov, Adva Zair, Phillipa Gill, and Michael Schapira. Measuring and mitigating AS-level adversaries against Tor. In *Proceedings of the 23rd Network and Distributed System Security Symposium*, 2016.

[37] Se Eun Oh, Taiji Yang, Nate Mathews, James K Holland, Mohammad Saidur Rahman, Nicholas Hopper, and Matthew Wright. Deepcoffea: Improved flow correlation attacks on tor via metric learning and amplification. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1915–1932, 2022.

[38] Christopher Parsons. Beyond privacy: Articulating the broader harms of pervasive mass surveillance. *Media and Communication*, 3(3):1–11, 2015.

[39] Person and Zeba Siddiqui. Fbi misused intelligence database in 278,000 searches, court says. https://www.reuters.com/world/us/fbi-misused-intelligence-database-278000-searches-court-says-2023-05-19/, May 2023.

[40] Ania M. Piotrowska, Jamie Hayes, Tariq Elahi, Sebastian Meiser, and George Danezis. The loopix anonymity system. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1199–1216, Vancouver, BC, August 2017. USENIX Association.

[41] Mohammad Saidur Rahman, Payap Sirinam, Nate Mathews, Kantha Girish Gangadhara, and Matthew Wright. Tik-tok: The utility of packet timing in website fingerprinting attacks. *Proceedings on Privacy Enhancing Technologies*, 2020(3), 2020.

[42] M.G. Reed, P.F. Syverson, and D.M. Goldschlag. Anonymous connections and onion routing. *IEEE Journal on Selected Areas in Communications*, 16(4):482–494, 1998.

[43] Fatemeh Rezaei and Amir Houmansadr. Tagit: Tagging network flows using blind fingerprints. *Proceedings on Privacy Enhancing Technologies*, 2017(4):290–307, 2017.

[44] Fatemeh Rezaei and Amir Houmansadr. Finn: Fingerprinting network flows using neural networks. In *Annual Computer Security Applications Conference*, ACSAC '21, page 1011–1024, New York, NY, USA, 2021. Association for Computing Machinery.

[45] Neil M Richards. The dangers of surveillance. *Harvard Law Review*, 126(7):1934–1965, 2013.

[46] Sajin Sasy and Ian Goldberg. SoK: Metadata-protecting communication systems. Cryptology ePrint Archive, Paper 2023/313, 2023. https://eprint.iacr.org/2023/313.

[47] Asuman Senol, Gunes Acar, Mathias Humbert, and Frederik Zuiderveen Borgesius. Leaky forms: A study of email and password exfiltration before form submission. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1813–1830, Boston, MA, August 2022. USENIX Association.

[48] Andrei Serjantov, Roger Dingledine, and Paul F. Syverson. From a trickle to a flood: Active attacks on several mix types. In *Revised Papers from the 5th International Workshop on Information Hiding*, IH '02, page 36–52, Berlin, Heidelberg, 2002. Springer-Verlag.

[49] Chris Strohm. Fbi searched data of millions of americans without warrants. https://www.bloomberg.com/news/articles/2022-04-29/fbi-searched-the-data-of-millions-of-americans-without-warrants, Apr 2022.

[50] Yixin Sun, Anne Edmundson, Nick Feamster, Mung Chiang, and Prateek Mittal. Counter-raptor: Safeguarding tor against active routing attacks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 977–992, 2017.

[51] Narseo Vallina-Rodriguez, Srikanth Sundaresan, Christian Kreibich, and Vern Paxson. Header enrichment or isp enrichment? emerging privacy threats in mobile networks. In *Proceedings of the 2015 ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization*, HotMiddlebox '15, page 25–30, New York, NY, USA, 2015. Association for Computing Machinery.

[52] Jean-Luc Watson, Sameer Wagh, and Raluca Ada Popa. Piranha: A GPU platform for secure computation. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 827–844, Boston, MA, August 2022. USENIX Association.

[53] Matthew K. Wright, Micah Adler, Brian Neil Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2002.