Poster: Anonymity Trilemma — Strong Anonymity, Low Bandwidth Overhead, Low Latency — Choose Two.

Debajyoti Das
Purdue University, USA
das48@purdue.edu

Sebastian Meiser
University College London, UK
s.meiser@ucl.ac.uk

Esfandiar Mohammadi
ETH Zurich, Switzerland
mohammadi@inf.ethz.ch

Aniket Kate
Purdue University, USA
aniket@purdue.edu

ABSTRACT

Millions of users from all over the world employ anonymous communication networks, such as Tor [1], to protect their privacy over the Internet. The design choice made by the Tor network to keep the latency and bandwidth overheads small has made it highly attractive to its geographically diverse user-base. However, over the last decade, the academic literature [2]–[8] has demonstrated Tor's vulnerability to a variety of traffic correlation attacks. In fact, Tor also has been successfully attacked in practice [9].

It is widely accepted that low-latency low-bandwidth overhead of anonymous communication (AC) protocols, such as Tor [10], can only provide a weak form of anonymity [11]. In the anonymity literature, several AC protocols were able to overcome this security barrier to provide a stronger anonymity guarantee (cryptographic indistinguishability-based anonymity [12], [13]) by either increasing the latency overhead or the bandwidth overhead. In particular, high-latency approaches (such as threshold mix networks [14]) can ensure strong anonymity by introducing significant communication delays for users messages, while high-bandwidth approaches (such as Dining Cryptographers network [15] and its extensions [16]–[18]) can provide strong anonymity by adding copious noise (or dummy) messages.

There have been a few efforts to propose hybrid approaches [19]–[23] that try to provide anonymity by simultaneously introducing latency and bandwidth overhead. However, it is not clear how to balance such system parameters to ensure strong anonymity while preserving practical performance.

In general, in the last 35 years a significant amount of research efforts have been put towards constructing novel AC protocols, deploying them, and attacking real-world AC networks. However, unlike other security fields such as cryptography, our understanding regarding the fundamental limits and requirements of AC protocols remains limited. This work takes some important steps towards answering fundamental question associated with anonymous communication. "Can we prove that strong anonymity cannot be achieved without introducing large latency or bandwidth overhead? When we wish to introduce the latency and bandwidth overheads simultaneously, do we know the overhead range values that still fall short at providing stronger anonymity?"

In our work, we investigate the fundamental constraints of anonymous communication (AC) protocols. We analyze the relationship between bandwidth overhead, latency overhead, and sender anonymity or recipient anonymity against a global passive (network-level) adversary. We confirm the trilemma that an AC protocol can only achieve two out of the following three properties: strong anonymity (i.e., anonymity up to a negligible chance), low bandwidth overhead, and low latency overhead.

We further study anonymity against a stronger global passive adversary that can additionally passively compromise some of the AC protocol nodes. For both adversary classes, we analyze two different user distributions (i.e., distributions that determine at which time or rate users of the AC protocol send messages): (i) synchronized user distributions, where users globally synchronize their messages, and (ii) unsynchronized user distributions, where each user locally decides when to send his messages independent of other users.

We derive as a necessary constraint a trade-off between bandwidth and latency overhead whose violation make it impossible for an AC protocol to achieve *strong anonymity*, i.e., anonymity up to a negligible (in a security parameter η) chance of failure. For any AC protocol where only a fraction of $\beta \in [0, 1]$ users send noise messages per communication round, and where messages can only remain in the network for $\ell \geq 0$ communication rounds, we find that against a global network-level adversary no protocol can achieve strong anonymity if $2\beta\ell < 1 - 1/\text{poly}(\eta)$ even when all the protocol parties are honest. In the case where a strictly stronger adversary additionally passively compromises c (out of K) protocol parties, we show that strong anonymity is impossible if $2(\ell - \mathsf{c})\beta < 1 - 1/\text{poly}(\eta)$ (for $\mathsf{c} < \ell$), or $2\beta\ell < 1 - 1/\text{poly}(\eta)$ and $\ell \in \mathcal{O}(1)$ (for $\mathsf{c} \geq \ell$).

We also assess the practical impact of our results by analyzing prominent AC protocols and depicting to which extent those satisfy our necessary constraints (Table I summarizes bounds on the bandwidth β and latency overhead ℓ, in the sense of this work). Our constraints mark an area on a 2D graph (see Figure 1) with latency overhead (x-axis) versus bandwidth overhead (y-axis) where strong anonymity is impossible. Our impossibility results naturally only offer nec-

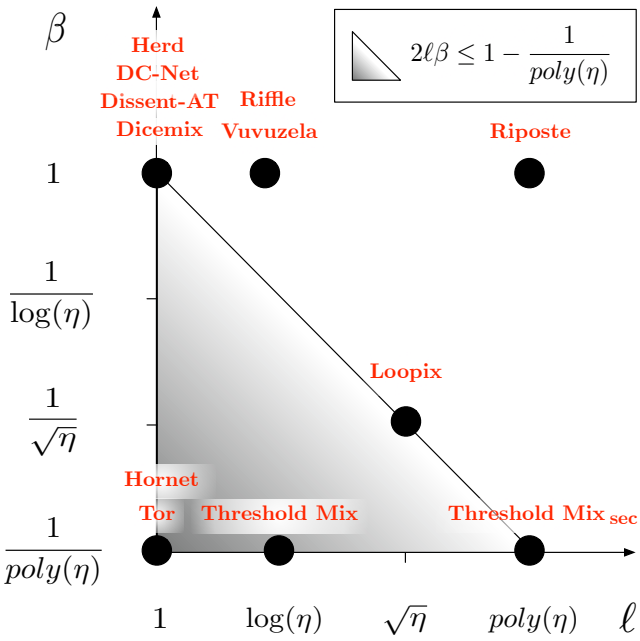

Fig. 1. Asymptotic latency overhead (ℓ) and bandwidth overhead (β) together with the "area of impossibility" where $2\ell\beta \leq 1 - \epsilon(\eta)$. We portray protocols as dots depending on their choices for ℓ and β. This graph assumes N is ca. $poly(\eta)$, the number of nodes K is ca. $\log \eta$, for security parameter η. The threshold for Threshold Mix $T = 1$ and for Threshold Mix$_{\text{sec}}$ $T = \mathsf{N} = poly(\eta)$. In the graph, both the axes are approximately in logarithmic scale. (For a more accurate visual representation we refer the readers to [24].)

TABLE I
Latency vs. bandwidth vs. strong anonymity of AC protocols, with the number of protocol-nodes K, number of clients N, and message-threshold T, expected latency ℓ' per node, dummy-message rate β.

Protocol	Latency	Bandwidth	Strong Anonymity
Tor	$\theta(1)$	$\theta(1/\mathsf{N})$	impossible
Hornet	$\theta(1)$	$\theta(1/\mathsf{N})$	impossible
Herd	$\theta(1)$	$\theta(\mathsf{N}/\mathsf{N})$	possible
Riposte	$\theta(N)$	$\theta(\mathsf{N}/\mathsf{N})$	possible
Vuvuzula	$\theta(\mathsf{K})$	$\theta(\mathsf{N}/\mathsf{N})$	possible
Riffle	$\theta(\mathsf{K})$	$\theta(\mathsf{N}/\mathsf{N})$	possible
Threshold mix	$\theta(T\mathsf{K})$	$\theta(1/\mathsf{N})$	impossible*
Loopix	$\theta(\sqrt{K}\ell')$	$\theta(\beta)$	possible
DC-Net	$\theta(1)$	$\theta(\mathsf{N}/\mathsf{N})$	possible
Dissent-AT	$\theta(1)$	$\theta(\mathsf{N}/\mathsf{N})$	possible
DiceMix	$\theta(1)$	$\theta(\mathsf{N}/\mathsf{N})$	possible

* if T in $o(poly(\eta))$

essary constraints for anonymity, but *not* sufficient conditions for the AC protocol. However, these necessary constraints for sender and recipient anonymity are crucial for understanding bi-directional anonymous communication. In fact, we find that several AC protocols in the literature are asymptotically close to the suggested constraints. Moreover, designers of new AC protocols can use our necessary constraints as guidelines for avoiding bad combinations of latency and bandwidth-overhead.

References

[1] "The Tor Project," https://www.torproject.org/, accessed in Nov 2017.

[2] A. Johnson, C. Wacek, R. Jansen, M. Sherr, and P. Syverson, "Users get routed: Traffic correlation on tor by realistic adversaries," in *Proc. ACM SIGSAC conference on Computer & communications security 2013*, 2013, pp. 337–348.

[3] L. Øverlier and P. F. Syverson, "Locating Hidden Servers," in *Proc. 27th IEEE Symposium on Security and Privacy*, 2006, pp. 100–114.

[4] S. J. Murdoch and G. Danezis, "Low-cost traffic analysis of Tor," in *Proc. IEEE Symposium on Security and Privacy 2005*, 2005.

[5] K. S. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. C. Sicker, "Low-resource routing attacks against tor," in *Proc. 6th ACM Workshop on Privacy in the Electronic Society (WPES)*, 2007, pp. 11–20.

[6] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal, "RAPTOR: Routing attacks on privacy in Tor," in *Proc. 24th USENIX Security Symposium*, 2015.

[7] R. Jansen, F. Tschorsch, A. Johnson, and B. Scheuermann, "The sniper attack: Anonymously deanonymizing and disabling the Tor network," in *Proc. Network and Distributed Security Symposium - NDSS '14*, 2014.

[8] Y. Gilad and A. Herzberg, "Spying in the Dark: TCP and Tor Traffic Analysis," in *Proc. 12th Privacy Enhancing Technologies Symposium (PETS 2012)*, 2012.

[9] The Tor Blog, "One cell is enough to break Tor's anonymity," https://blog.torproject.org/blog/one-cell-enough, accessed Nov 2017.

[10] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The Second-Generation Onion Router," in *Proc. 13th USENIX Security Symposium (USENIX)*, 2004, pp. 303–320.

[11] S. Chakravarty, M. V. Barbera, G. Portokalidis, M. Polychronakis, and A. D. Keromytis, "On the effectiveness of traffic analysis against anonymity networks using flow records," in *Proc. 15th International Conference on Passive and Active Measurement*, 2014, pp. 247–257.

[12] N. Gelernter and A. Herzberg, "On the limits of provable anonymity," in *Proc. Workshop on Privacy in the Electronic Society (WPES 2013)*, 2013, pp. 225–236.

[13] A. Hevia and D. Micciancio, "An indistinguishability-based characterization of anonymous channels," in *Proc. Eighth International Symposium on Privacy Enhancing Technologies (PETS 2008)*, N. Borisov and I. Goldberg, Eds., 2008, pp. 24–43.

[14] A. Serjantov, R. Dingledine, and P. Syverson, "From a trickle to a flood: Active attacks on several mix types," in *5th Information Hiding Workshop (IH 2002)*, 2003, pp. 36–52.

[15] D. Chaum, "The dining cryptographers problem: Unconditional sender and recipient untraceability," *Journal of Cryptology*, vol. 1, no. 1, pp. 65–75, 1988.

[16] T. Ruffing, P. Moreno-Sanchez, and A. Kate, "P2P Mixing and Unlinkable Bitcoin Transactions," in *Proc. 25th Annual Network & Distributed System Security Symposium (NDSS)*, 2017.

[17] H. Corrigan-Gibbs and B. Ford, "Dissent: Accountable Anonymous Group Messaging," in *Proc. 17th ACM Conference on Computer and Communication Security (CCS)*, 2010, pp. 340–350.

[18] P. Golle and A. Juels, "Dining cryptographers revisited," in *Proc. of Eurocrypt 2004*, 2004.

[19] H. Corrigan-Gibbs, D. I. Wolinsky, and B. Ford, "Proactively Accountable Anonymous Messaging in Verdict," in *Proc. 22nd USENIX Security Symposium*, 2013, pp. 147–162.

[20] J. van den Hooff, D. Lazar, M. Zaharia, and N. Zeldovich, "Vuvuzela: Scalable private messaging resistant to traffic analysis," in *Proc. 25th ACM Symposium on Operating Systems Principles (SOSP 2015)*, 2015.

[21] A. Kwon, D. Lazar, S. Devadas, and B. Ford, "Riffle: An Efficient Communication System With Strong Anonymity," in *Proc. Privacy Enhancing Technologies Symposium (PETS 2016)*, 2016, pp. 115–134.

[22] D. I. Wolinsky, H. Corrigan-Gibbs, B. Ford, and A. Johnson, "Dissent in Numbers: Making Strong Anonymity Scale," in *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, 2012, pp. 179–182.

[23] A. Piotrowska, J. Hayes, T. Elahi, S. Meiser, and G. Danezis, "The loopix anonymity system," in *Proc. 26th USENIX Security Symposium*, 2017.

[24] Anonymity Trilemma Project Webpage, "Anonymity trilemma: Strong anonymity, low bandwidth overhead, low latency overhead—choose two," https://freedom.cs.purdue.edu/projects/anonymity/trilemma/.