Algorithmic Advances in Anonymous Communication over Networks

Giulia Fanti and Pramod Viswanath
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign

***Abstract*—Over the years, researchers have thought extensively about how to communicate anonymously over a network. In this position paper, we first provide an overview of the current research landscape—including a discussion of common communication and anonymity models—and examples of prominent work in this space. In the second half of the paper, we present our thoughts on future directions of interest in this area. In short, we notice an increasing amount of work in the anonymity literature that tackles nearly omnipotent adversaries, which presumably represent the NSA or a similarly powerful surveillance agency. Such strong adversarial models can only be defeated through heavy-handed techniques that may be difficult to scale in practice. We therefore present two guidelines that have guided our own algorithmic problem formulation in this space. These guidelines suggest alternative research directions that (1) exploit surveillance agencies' weaknesses at the physical layer, and (2) consider the many weaker, but still relevant, adversaries who contribute to global censorship and surveillance.**

I. Introduction

In a free society, people have the right to consume and distribute information without being surveilled or censored [1]. However, the typical Internet user reveals intimate personal details to search engines, ISPs, and surveillance agencies, simply by using the web. Countries that rely heavily on electronic surveillance may observe a chilling effect among Internet users, in which people become less willing to share or discuss controversial ideas for fear of targeting [2]. More dramatically, electronic surveillance in countries under authoritarian rule regularly leads to the punishment of individuals who espouse controversial ideas [3], [4].

The research community has studied many techniques for allowing people to communicate freely in the face of various adversaries, while providing provable anonymity guarantees. This body of work took off in the 1980s and remains an active research area as surveillance capabilities continue to grow.

In this position paper, we present a brief overview of the anonymous communication research landscape. Although there exist reviews of related topics [5], this paper is meant to be a brief introduction, highlighting open problems and our own views on this space. Notably, our emphasis is not necessarily on confidentiality; some of the discussed techniques do not explicitly use encryption at all.

We partition the space according to the most common communication models and definitions of anonymity, as illustrated in Figure 1. For communication models, we distinguish between point-to-point communication and one-to-many communication. *Point-to-point* communication refers to an individual who wishes to send a message to a single recipient, such as email. *One-to-many* communication instead assumes the individual wishes to send a message to as many people as possible, as in microblogging applications like Twitter. Regarding anonymity, we distinguish between systems that provide sender anonymity, receiver anonymity, and sender-receiver unlinkability. *Sender anonymity* refers to a system that prevents the adversary from learning which user sent a given message. *Receiver anonymity* prevents the adversary from learning who received a message. *Sender-receiver unlinkability* prevents the adversary from learning if two parties are communicating at any given time. Darkened boxes in Figure 1 indicate categories that are inherently true. For instance, one-to-many communication trivially gives receiver anonymity, so we do not consider this a research area.[1] Many systems provide one or more anonymity properties, depending on the adversary's capabilities.

	Point-to-Point Communication	One-to-Many Communication
Sender anonymity	**Onion Routing** Tor, Tarzan **DC Networks** Herbivore, Dissent	**DC Networks** Herbivore, Dissent **Statistical Obfuscation** Fanti *et al.*, Luo *et al.* **Secret Sharing** Riposte
Receiver Anonymity	**Broadcast** BitMessage, DC Networks **Message drop** Dolev and Ostrovsky	*Inherent*
Sender-Receiver Unlinkability	**Mix Networks** P5, Mixmaster, Vuvuzela **Onion Routing** Tor, Tarzan **DC Networks** Herbivore, Dissent	*Inherent*

Fig. 1: Research landscape for anonymous communication. Existing work can be categorized by the type of anonymity it provides, and the communication model it considers. Dark grey boxes are inherently encompassed by a different category. The examples in this chart are non-exhaustive.

Adversaries. An adversary can be either local or global. In this context, a *local* adversary monitors a single network link

[1] Notice that one-to-many communication algorithms do not trivially solve the point-to-point communication problem, even when content is encrypted with the intended recipient's key. They are bandwidth-intensive and cannot reasonably scale to support a full communication system.

with the goal of thwarting anonymous communication protocols. For example, in the point-to-point scenario with sender anonymity, the adversary could be modeled as the recipient of the message, or an eavesdropper that monitors the recipient's traffic. A *global* adversary instead monitors multiple network taps. For example, a police agency monitoring Facebook would count as a global adversary.[2] Both adversarial models have been of interest over the years. However, in light of recent revelations about governmental surveillance [6], [7], it is becoming increasingly important to develop countermeasures against global adversaries. The precise capabilities of a global adversary depend on the problem setting.

Organization. In Sections II and III, we discuss sender anonymity, in the point-to-point and one-to-many problems, respectively. Section IV considers receiver anonymity, and Section V considers sender-receiver unlinkability, both in the point-to-point setting. Finally, we provide a list of technical challenges and open questions in Section VI, and discuss our own thoughts regarding the most interesting and important directions moving forward.

II. Sender Anonymity, Point-to-Point Communication

Sender-anonymous point-to-point communication is commonly used by hosts connecting to web servers. In the simplest version of this problem, a sender wishes to communicate with a receiver without the receiver learning the sender's identity. This formulation models the (local) adversary as the message recipient, and it is effectively solved in practice by proxy servers. However, if the proxy itself is adversarial, or if the adversary is monitoring the link from the sender, this approach does not provide sender anonymity.

The problem becomes more difficult under global adversaries. A typical assumption is that the adversary can monitor a subset of network traffic. Under this adversarial model, the most common privacy primitive is *onion routing*. Onion networks route messages through a circuit of "onion router" relays before reaching the final destination. Messages are encrypted with the selected onion routers' keys, so any router can only see the preceding and subsequent hop in the circuit.

The best-known example of onion routing is Tor [8], which has about 2 million directly connecting clients as of December 2015 [9]. Tor provides sender anonymity—as well as sender-receiver unlinkability—under a global adversary that monitors a subset of links (i.e., it should not tap both entry and exit links). However, if the adversary monitors the first link in an onion routing circuit, sender anonymity is broken. This is not necessarily problematic, because message contents are encrypted, and unlinkability is preserved; however, this toy example raises important point: anonymous point-to-point communication can exist without sender (or receiver) anonymity—in this context, *unlinkability* may be a better anonymity metric. Nonetheless, sender anonymity can be useful for masking a user's involvement in a network or transaction.

If an adversary has slightly stronger capabilities, other attacks are possible, which primarily break sender-receiver unlinkability; for instance, if an adversary observes both the entry and exit node of an onion route, timing attacks are possible [10]. A sub-body of work has emerged around identifying and mitigating attacks on the Tor network in particular [11], [12], [13], [14]. Due to Tor's widespread adoption, this research area remains both active and relevant.

Another class of sender-anonymous algorithms is based on anonymous broadcast algorithms with encrypted messages. Even though broadcast algorithms are not typically associated with point-to-point communication, the sender can encrypt with the intended recipient's key. Since only the intended recipient can decrypt, the channel behaves as a low-capacity broadcast point-to-point channel. Because of this, anonymous broadcast algorithms like dining cryptographer networks (DC nets) [15] can be used for sender-anonymous, point-to-point communication. We discuss some of these algorithms in greater detail in Section III. Finally, XOR-trees are algorithmically similar to DC nets, but tailored for point-to-point anonymous communication [16]. XOR trees use distributed linear coding over modulated data sequences to reduce communication overhead compared to broadcasting messages.

Although the main approach to sender anonymity in point-to-point systems is onion routing, the research community has produced a number of other systems. These include privacy-preserving filesharing systems like Freenet [17] and Free Haven [18], and various messaging systems that exploit different privacy primitives: Tarzan and I2P, for instance, route messages through Chaumian mixes [19], [20]. Crowds uses a version of onion routing with random path selection [21]. Hordes builds atop Crowds, harnessing multicast addressing to provide stronger receiver anonymity [22] Herbivore [23] divides the network into anonymizing cliques, and each user disseminates messages using a DC net within a clique.

Overall, this research area has matured; a notable exception is studies on Tor, due to its popularity and utility. Recent work in this space (correctly) focuses explicitly on sender-receiver unlinkability rather than pure sender anonymity [24].

III. Sender Anonymity, One-to-Many Communication

One-to-many communication typically takes one of two forms: broadcast communication or multicast transmission. Existing anonymity research has primarily focused on the former problem, but we discuss both.

Broadcast communication. Anonymous broadcast communication is useful in tools like anonymous message boards. However, message boards and other broadcast messaging systems are susceptible to spam, limiting their practical use.

Anonymous broadcast messaging has been most studied in context of the dining cryptographers' problem. Dining cryptographer networks, or DC nets, allow a user to broadcast

[2]We treat adversaries with partial eavesdropping capabilities as global. For instance, a police agency might create fake accounts in a social network, or tap some fraction of network links. This adversary only has a sampled view of network activity, but we still call it 'global' because the attacks feasible for such an adversary are stronger than those feasible for local adversaries.

a message without being identifiable as the source of the message; they are robust against collusions of users in the network [15]. However, DC nets are very communication-intensive, and they are also sensitive to misbehaving or malicious players in the system. Much work has therefore studied how to make DC nets practically deal with malicious adversaries and reduce communication overhead [15], [25], [23], [26], [27].

Recently, [28] addressed a similar problem of anonymously writing to a public message board. Instead of relying on DC nets, Riposte uses techniques from private information retrieval to store multiple corrupted message copies on distributed servers. This approach prevents (sufficiently small) subsets of colluding servers cannot determine authorship.

Today, there are no practical systems built on DC nets, and bulletin boards in general are waning in popularity compared to two decades ago. Although this research topic remains unsolved in a practical sense, demand for pure anonymous broadcast tools may be less compelling than other tools.

Multicast Communication. An alternative application of one-to-many communication is microblogging, or network-constrained multicast communication. Unlike broadcast messaging, microblogging spreads content only to a user's contacts on an underlying contact graph (e.g., a social graph). If a user approves a message by 'liking' or 'retweeting' it, the message gets passed to the friend's friends, and so forth. In this way, content may eventually spread to the entire graph, but social filtering prevents the propagation of spam messages.

Various anonymous microblogging applications have seen significant growth in the last two years [29], but very little research has been done on the topic. Popular messaging applications in this space include Whisper [30] and Yik Yak [29]. This area is fairly new; it is not clear yet whether these applications have staying power. However, Yik Yak has seen widespread adoption, particularly by school-age users [29].

Research is needed because existing anonymous messaging services do not provide real anonymity. Most platforms store both messages and authorship information on centralized servers, which makes them vulnerable to government subpoenas, hacking, or direct company access, for the purpose of identifying the author of a message. Moreover, the spreading protocol used by existing platforms places users at risk of deanonymization against adversaries with side information, as proved in recent advances [31], [32], [33], [34], [35], [36]. Thus, even distributed architectures do not solve the problem.

In response, there has been some work on designing spreading protocols that prevent a message author from being identified by various global adversaries. For example, in [37], we consider a *snapshot* adversary, which uses side channels to infer whether a node received the message. This could represent a state-level adversary that attends a Twitter-organized protest; it implicitly learns who received the protest advertisement, but not the associated metadata. The main result of this work is a protocol called *adaptive diffusion*, which exhibits strong anonymity guarantees theoretically and in simulation. We also considered a *spy-based* adversary, which corrupts some fraction of nodes through bribery or coercion; these spy nodes pass along metadata like message timestamps and relay IDs [38]. A spy-based adversary could represent a government agency participating in social media to study users, for instance.

In a different approach to the problem, Luo *et al.* use a game-theoretic framework to design optimal infection and estimation strategies [39]. This takes into account the government's ability to adapt its estimation strategy as needed.

It is still unclear if and how this new research area will develop, but the topic is very timely, and there is plenty of space for innovation, both practical and theoretical. For example, nobody has considered the effect of malicious adversaries (i.e. adversaries that do not follow protocol) on the anonymity guarantees of the system. Another key point is that the protocols in [37], [38], [39] all assume a privacy-preserving ecosystem. That is, they assume the lower-level operations of the distributed communication network do not leak authorship information. However, running a privacy-preserving, distributed communication network is not a solved problem. Issues like privacy-preserving user presence [40] and contact discovery [41] are active research areas, and must be solved in order to build a fully-functioning system.

IV. Receiver Anonymity, Point-to-Point Communication

Very few anonymous communication systems offer only receiver anonymity. Receiver anonymity typically appears in conjunction with sender anonymity, whereas sender anonymity can appear without receiver anonymity. Nonetheless, there are a few systems that provide receiver anonymity.

Most receiver-anonymous systems use some form of broadcast transmission. Broadcast communication inherently provides receiver anonymity; this explains the darkened gray boxes in Figure 1. For example, BitMessage is a P2P, point-to-point messaging system; it routes messages by encrypting them and broadcasting them [42]. Only the intended recipient is able to decrypt each message, but the adversary cannot learn *which* user was able to decrypt the message. Notice that BitMessage does not provide sender anonymity.

A similar approach can be used in other broadcast systems, such as DC networks [15], with encryption to ensure point-to-point confidentiality. By design, this approach also provides sender anonymity. A conceptually similar approach is taken by XOR trees, which use shared randomness between the receiver and the sender to achieve receiver anonymity [16].

V. Sender-Receiver Unlinkability, Point-to-Point Communication

The identities of communicating parties can often be as sensitive as their communication contents. As such, sender-receiver unlinkability is important in applications like email, telephony, and financial transactions [43], [44], [45]. Research in this space can roughly be divided into two main applications: anonymous messaging and cryptocurrencies.

Anonymous Messaging. Anonymous messaging with transaction unlinkability has been a popular research topic for

decades, starting with Chaum's seminal paper on email unlinkability [43]. This work introduces the notion of a *mix network*, which accumulates encrypted input packets, re-encodes them cryptographically, and forwards them to their final destination. This process introduces latency, as the mix requires multiple inputs, but in exchange it provides strong unlinkability.

Mix networks have been the topic of much work, both theoretical and practical. On the theoretical side, researchers have asked how to optimally add delays in order to satisfy latency constraints while ensuring the unlinkability of inputs and outputs, either with complete observer data [46] or partial timing side information [47], [48]. Other researchers have considered how to build distributed mix networks that are robust to colluding mix nodes [49].

On the systems side, a number of papers have studied how to use mix nodes to obfuscate the links between communicating parties—challenges include scalability and tolerating increasingly powerful adversaries. Examples include Free Haven [18], P5 [50], MixMinion [51], Tarzan [19], and Vuvuzela [24]. None of these systems have seen widespread adoption in practice, but they spurred discussion in the privacy community.

Overall, transaction anonymity in messaging applications is still an active research area, albeit somewhat less popular than a decade ago. Theoretical investigations of mix networks are receiving renewed attention, and the systems community continues to design new systems that reflect the changing landscape of modern surveillance.

Cryptocurrencies. Another important application space for sender-receiver unlinkability is cryptocurrencies. Anonymous, distributed currency transactions are sometimes viewed as a mechanism for reducing the power of centralized financial institutions, while protecting users' financial privacy. Even though financial transactions are not traditionally considered a communication problem, cryptocurrencies fundamentally rely on the flow of information, not physical goods—this construction introduces many similarities with messaging applications.

Cryptocurrencies use cryptographic verification to govern transactions and the "mining" of new currency. Although cryptocurrencies have been a point of discussion since (and preceding) Chaum [52], the most well-known work in this space is Bitcoin [53], which emerged in 2008. Bitcoin works by storing all transactions in a public ledger, called a *blockchain*. Each transaction must be accepted through distributed consensus—other Bitcoin nodes verify transactions by solving a computational puzzle and posting their results to the blockchain, and they are rewarded for their efforts with more bitcoins.

It is a popular misconception that Bitcoin is anonymous. However, Bitcoin transactions are displayed publicly on the blockchain, and participating parties are identified by their public keys—Bitcoin is actually *pseudonymous*. This has led to a number of deanonymization attacks; for example, the transaction graph between users reveals information [54]. In practice, people often use Bitcoin 'laundries', which anonymize the origin of a particular Bitcoin using Chaum mixes to cryptographically mix bitcoins from different users. However, even laundries may be susceptible to deanonymization attacks in practice [55].

This observation has spurred a great deal of research on improving the anonymity of Bitcoin. A prominent example is Zerocoin [45], which strengthens the cryptographic properties of Bitcoin to guarantee anonymity. However, since Bitcoins and altcoins are run on P2P networks, network properties can impact an adversary's ability to deanonymize a transaction [56], particularly in more recent lightweight Bitcoin implementations [57]. In light of this, interesting questions may include how to exploit properties of P2P networks (such as peer churn) in order to provide stronger anonymity guarantees.

Cryptocurrencies have become a popular research topic, sparked in large part by interest in Bitcoin, both academic and industrial. There are still many important, unanswered questions, both in a practical and theoretical sense [58].

VI. Moving Forward

Anonymous communication has been and continues to be an active research area. Based on the breadth and depth of censorship and surveillance today [3], [59], [6], relevant adversaries are best represented by global models. We find that many interesting open problems in the space of anonymous communication algorithms share a common thread: they take a nuanced view of what it means to be a "global" adversary. In particular, we note that while anonymity researchers often target NSA-like adversaries [60], [24], it may be worth instead considering a "guerilla warfare" approach to anonymity. We share two guidelines that have shaped our own formulation of research problems: (1) Don't try to beat the NSA at its own game; (2) Censorship is the enemy, not the NSA.

(1) Don't try to beat the NSA at its own game. As surveillance agencies grow progressively stronger, an increasing number of research papers assume a quasi-omnipotent adversary, possibly modeling the NSA [24], [60], [42]. In the face of such a strong adversary, resulting algorithms often rely on heavy-handed measures to avoid detection, such as continuously sending dummy traffic to prevent the adversary from detecting when a real message is sent, as in [24]. This may be the best one can do against a seemingly invincible adversary. Moreover, most research in this space assumes the adversary cannot decipher encrypted data, but recent revelations suggest that this assumption (which is the bedrock of most modern security systems) does not necessarily hold [61], [62]. Faced with what appears to be a losing battle, it may be worth spending some time to think carefully about which adversaries we want to take on, and what their capabilities are.

For example, suppose our adversary is a major surveillance agency (e.g., the NSA). Assume this adversary can wiretap every Internet link with at least one terminus in the United States and break standard encryption protocols. Such an adversary seems invincible, but notice that a surveillance agency is less likely to be omnipresent in networks that do *not* rely on Internet infrastructure—for example, agencies are not *physically* omnipresent. Local, self-organizing networks, such as mesh networks, may transform omnipotent adversaries into

weaker ones. Mesh networks face severe scaling challenges in theory [63] and in practice [64], but a hybrid of mesh and traditional infrastructure might avoid these issues. The larger point is that well-funded surveillance agencies are equipped to surveil telecommunications; if the goal is to prevent privacy invasions, it may be useful to consider alternative communication paradigms, focusing on physical-layer solutions.

(2) Censorship is the enemy, not the NSA. In many ways, the NSA is an easy target. It is conspicuous, and it directly impacts the lives of many anonymity researchers. However, many smaller adversaries worldwide contribute to global surveillance and censorship, often in significantly more extreme ways [4]. These (comparatively) weaker surveillance agencies exhibit very different behavior from that of the NSA, and therefore demand different, oftentimes more manageable, adversarial models. For instance, our own work on anonymous broadcast messaging [38] considers a spy-based adversarial model, inspired by the behavior of government agencies in countries like Iran or Thailand [65], [4]. Asking canonical anonymity questions under such alternative adversarial models is both of societal and technical interest.

Open Problems. Inspired by these guidelines, we identify a few anonymous communication problems of potential interest:

- **Mix networks with partial monitoring.** Mix networks have been used since the 1980s to provide sender-receiver unlinkability [43]. However, under continuous surveillance, it is difficult to provide strong unlinkability guarantees while satisfying low-latency constraints [46]—this observation is central to timing-based deanonymization attacks on Tor, for instance [10]. Alternatively, strong unlinkable guarantees are provided but only in very limited contexts of a single mix node and a very small number of sender-receiver pairs [66], [67], [68]. An important question is how to build mix networks that are robust to *partial* observation [47], [48]. For example, state-level adversaries might be able to monitor only a subset of inputs/outputs in a mix net, particularly if the mix node is located outside the adversary's reach. How should the mixnet and users behave to provide optimal sender-receiver unlinkability? Is there an optimal way to introduce delays in the mix node?
- **Anonymous routing in P2P networks.** A big problem in anonymous P2P networks is routing—particularly in point-to-point communications. A number of anonymous routing protocols have emerged, many of which rely on message flooding and/or expensive encryption operations [42], [69], [70], [71], [72]. These measures tend to be expensive, in terms of both communication and computation. However, these anonymous routing protocols may be unnecessarily inefficient—many of the adversaries we care about are not omnipresent in P2P networks. As in our work on anonymous broadcast messaging, a spy-based adversary is very reasonable for this problem, particularly if users are not all in the same country. Can we achieve more efficient routing if only a subset of users are malicious?
  A quick note on this problem space: While P2P networks have fallen out of fashion as a research topic, the reality is that they see tremendous traffic in practice [73] and underpin the basic anonymous transactional framework of Bitcoin [58]. Most censorship countermeasures today use either a distributed or decentralized architecture. As such, this space remains an important one and deserves a renewed look through the lens of anonymous communication.
- **Harnessing physical resources for anonymity.** The NSA may be capable of monitoring a majority of domestic Internet traffic. However, there are means of communication that do not rely on Internet infrastructure, including mesh networks. Through judicious use of these means, can we build a communication network that provides user anonymity, even against an adversary with complete network oversight? Years of research and experimental deployment suggest that mesh networks cannot gracefully scale to support millions of users [63], [64]. However, this does not mean that infrastructure-less networks cannot aid in the pursuit of anonymous communication. An interesting question is whether small-scale deployment of alternative communication networks (e.g., mobile mesh networks) can impact anonymity. For instance, consider a messaging system that transmits each message to another user's mobile phone via Bluetooth prior to accessing Internet infrastructure. Message authorship would be obfuscated up to a factor of the users' spatial mobility. What anonymity can this provide? Are there other ways to anonymize one's location prior to sending a message? Are there other types of networks one can harness to add diversity to Internet messaging?

These questions are of key importance for society as surveillance techniques grow increasingly widespread and sophisticated. Answering them will require a nuanced understanding the many adversaries that present themselves in today's global landscape, as well as a broad understanding of the relation between the physical world and communication systems.

REFERENCES

[1] L. M. Elison and D. Nettik-Simmons, "Right of privacy," *Mont. L. Rev.*, vol. 48, p. 1, 1987.
[2] A. Marthews and C. Tucker, "Government surveillance and internet search behavior," *Available at SSRN 2412564*, 2014.
[3] P. Holley, "Post correspondent Jason Rezaian sentenced to prison term in iran," *The Washington Post*, November 2015.
[4] Agencies, "Thai man jailed for six years over facebook posts," *Al Jazeera*, 2016.
[5] N. Unger, S. Dechand, J. Bonneau, S. Fahl, H. Perl, I. Goldberg, and M. Smith, "Sok: Secure messaging," in *Security & Privacy*. IEEE, 2015.
[6] G. Greenwald, "NSA collecting phone records of millions of verizon customers daily," *The Guardian*, June 2013.
[7] "A glimpse at the packet sniffing that the NSA does," liveLeak.
[8] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," DTIC Document, Tech. Rep., 2004.
[9] "Tor metrics," http://metrics.torproject.org.
[10] B. Levine, M. Reiter, C. Wang, and M. Wright, "Timing attacks in low-latency mix systems," in *Financial cryptography*. Springer, 2004.
[11] N. S. Evans, R. Dingledine, and C. Grothoff, "A practical congestion attack on tor using long paths." in *USENIX Security Symposium*, 2009, pp. 33–50.
[12] L. Xin and W. Neng, "Design improvement for tor against low-cost traffic attack and low-resource routing attack," in *CMC*, vol. 3. IEEE, 2009.

[13] Z. Ling, J. Luo, W. Yu, X. Fu, D. Xuan, and W. Jia, "A new cell counter based attack against tor," in *CCS*. ACM, 2009, pp. 578–589.

[14] D. Arp, F. Yamaguchi, and K. Rieck, "Torben: A practical side-channel attack for deanonymizing tor communication," in *Information, Computer and Communications Security*. ACM, 2015, pp. 597–602.

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

[16] S. Dolev and R. Ostrovsky, "XOR-trees for efficient anonymous multicast and reception," *TISSEC*, vol. 3, no. 2, pp. 63–84, 2000.

[17] I. Clarke, O. Sandberg, B. Wiley, and T. Hong, "Freenet: A distributed anonymous information storage and retrieval system," in *Designing Privacy Enhancing Technologies*. Springer, 2001.

[18] R. Dingledine, M. Freedman, and D. Molnar, "The free haven project: Distributed anonymous storage service," in *Designing Privacy Enhancing Technologies*. Springer, 2001.

[19] M. Freedman and R. Morris, "Tarzan: A peer-to-peer anonymizing network layer," in *Proc. CCS*. ACM, 2002.

[20] M. Herrmann and C. Grothoff, "Privacy-implications of performance-based peer selection by onion-routers: a real-world case study using I2P," in *Privacy Enhancing Technologies*. Springer, 2011.

[21] M. K. Reiter and A. D. Rubin, "Anonymous web transactions with crowds," *Communications of the ACM*, vol. 42, no. 2, pp. 32–48, 1999.

[22] C. Shields and B. N. Levine, "A protocol for anonymous communication over the internet," in *CCS*. ACM, 2000, pp. 33–42.

[23] S. Goel, M. Robson, M. Polte, and E. Sirer, "Herbivore: A scalable and efficient protocol for anonymous communication," Cornell University, Tech. Rep., 2003.

[24] J. Van Den Hooff, D. Lazar, M. Zaharia, and N. Zeldovich, "Vuvuzela: Scalable private messaging resistant to traffic analysis," in *SOSP*. ACM, 2015.

[25] H. Corrigan-Gibbs and B. Ford, "Dissent: accountable anonymous group messaging," in *CCS*. ACM, 2010.

[26] P. Golle and A. Juels, "Dining cryptographers revisited," in *Advances in Cryptology-Eurocrypt 2004*. Springer, 2004.

[27] L. von Ahn, A. Bortz, and N. Hopper, "K-anonymous message transmission," in *Proc. CCS*. ACM, 2003.

[28] H. Corrigan-Gibbs, "Riposte: An anonymous messaging system handling millions of users," 2014.

[29] "Yik yak," http://www.yikyakapp.com/.

[30] "Whisper," http://whisper.sh.

[31] D. Shah and T. Zaman, "Rumors in a network: Who's the culprit?" *Information Theory, IEEE Transactions on*, vol. 57, no. 8, pp. 5163–5181, Aug 2011.

[32] P. C. Pinto, P. Thiran, and M. Vetterli, "Locating the source of diffusion in large-scale networks," *Physical review letters*, vol. 109, no. 6, 2012.

[33] Z. Wang, W. Dong, W. Zhang, and C. Tan, "Rumor source detection with multiple observations: Fundamental limits and algorithms," in *ACM SIGMETRICS*, 2014.

[34] B. A. Prakash, J. Vreeken, and C. Faloutsos, "Spotting culprits in epidemics: How many and which ones?" in *ICDM*, vol. 12, 2012.

[35] V. Fioriti and M. Chinnici, "Predicting the sources of an outbreak with a spectral technique," *arXiv preprint arXiv:1211.2333*, 2012.

[36] W. Luo, W. Tay, and M. Leng, "How to identify an infection source with limited observations," 2013.

[37] G. Fanti, P. Kairouz, S. Oh, and P. Viswanath, "Spy vs. spy: Rumor source obfuscation," *SIGMETRICS Perform. Eval. Rev.*, 2015.

[38] G. Fanti, P. Kairouz, S. Oh, K. Ramchandran, and P. Viswanath, "Hiding the rumor source," *arXiv preprint arXiv:1509.02849*, 2015.

[39] W. Luo, W. P. Tay, and M. Leng, "Rumor spreading and source identification: A hide and seek game," *arXiv:1504.04796*, 2015.

[40] N. Borisov, G. Danezis, and I. Goldberg, "Dp5: A private presence service," *PETS*, 2015.

[41] M. Wachs, M. Schanzenbach, and C. Grothoff, "A censorship-resistant, privacy-enhancing and fully decentralized name system," in *Cryptology and Network Security*. Springer, 2014.

[42] "BitMessage," http://bitmessage.org.

[43] D. L. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, vol. 24, no. 2, 1981.

[44] T. B. Lee, "Here?s how phone metadata can reveal your affairs, abortions, and other secrets," *Washington Post*, 2013.

[45] I. Miers, C. Garman, M. Green, and A. D. Rubin, "Zerocoin: Anonymous distributed e-cash from bitcoin," in *Security & Privacy*. IEEE, 2013.

[46] P. Venkitasubramaniam and V. Anantharam, "On the anonymity of chaum mixes," in *ISIT*. IEEE, 2008, pp. 534–538.

[47] X. Gong, N. Kiyavash, and N. Borisov, "Fingerprinting websites using remote traffic analysis," in *CCS*. ACM, 2010, pp. 684–686.

[48] X. Gong and N. Kiyavash, "Timing side channels for traffic analysis," in *ICASSP*. IEEE, 2013, pp. 8697–8701.

[49] M. Jakobsson and A. Juels, "An optimally robust hybrid mix network," in *Principles of distributed computing*. ACM, 2001, pp. 284–292.

[50] R. Sherwood, B. Bhattacharjee, and A. Srinivasan, "p 5: A protocol for scalable anonymous communication," in *Security & Privacy*. IEEE, 2002.

[51] G. Danezis, R. Dingledine, and N. Mathewson, "Mixminion: Design of a type III anonymous remailer protocol," in *Security & Privacy*. IEEE, 2003, pp. 2–15.

[52] D. Chaum and S. Brands, "Minting electronic cash," *Spectrum, IEEE*, vol. 34, no. 2, pp. 30–34, 1997.

[53] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Consulted*, vol. 1, no. 2012, p. 28, 2008.

[54] P. Koshy, D. Koshy, and P. McDaniel, *An analysis of anonymity in bitcoin using P2P network traffic*. Springer, 2014.

[55] M. Möser, "Anonymity of bitcoin transactions," in *Münster Bitcoin Conference*, 2013.

[56] A. Biryukov and I. Pustogarov, "Bitcoin over tor isn't a good idea," *arXiv preprint arXiv:1410.6079*, 2014.

[57] A. Gervais, S. Capkun, G. O. Karame, and D. Gruber, "On the privacy provisions of bloom filters in lightweight bitcoin clients," in *Computer Security Applications Conference*. ACM, 2014.

[58] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten, "Research perspectives and challenges for bitcoin and cryptocurrencies (extended version)," Cryptology ePrint Archive, Report 2015/452, Tech. Rep., 2015.

[59] G. King, J. Pan, and M. E. Roberts, "How censorship in china allows government criticism but silences collective expression," *American Political Science Review*, 2013.

[60] H. Corrigan-Gibbs, D. Boneh, and D. Mazières, "Riposte: An anonymous messaging system handling millions of users," in *Security & Privacy*. IEEE, 2015.

[61] B. Schneier, "Did NSA put a secret backdoor in new encryption standard?" *URL: http://www. wired. com/politics/security/commentary/securitymatters/2007/11/securitymatters_1115 (2007-11-15)*, 2007.

[62] D. Adrian, K. Bhargavan, Z. Durumeric, P. Gaudry, M. Green, J. A. Halderman, N. Heninger, D. Springall, E. Thomé, L. Valenta, B. VanderSloot, E. Wustrow, S. Zanella-Béguelin, and P. Zimmermann, "Imperfect forward secrecy: How Diffie-Hellman fails in practice," in *CCS*, 2015.

[63] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *Information Theory, IEEE Transactions on*, 2000.

[64] P. De Filippi, "It?s time to take mesh networks seriously (and not just for the reasons you think)," *Wired*, 2014.

[65] J. Finkle, "RPT-iranian hackers use fake facebook accounts to spy on u.s., others," *Reuters*, 2014.

[66] S. Kadloor, N. Kiyavash, and P. Venkitasubramaniam, "Mitigating timing based information leakage in shared schedulers," in *INFOCOM*. IEEE, 2012.

[67] S. Kadloor, P. Venkitasubramaniam, and N. Kiyavash, "A statistical inference perspective on preventing timing analysis in networks."

[68] X. Gong and N. Kiyavash, "Quantifying the information leakage in timing side channels in deterministic work-conserving schedulers," *arXiv preprint arXiv:1403.1276*, 2014.

[69] B. Zhu, Z. Wan, M. S. Kankanhalli, F. Bao, and R. H. Deng, "Anonymous secure routing in mobile ad-hoc networks," in *Local Computer Networks*. IEEE, 2004, pp. 102–108.

[70] S. Seys and B. Preneel, "Arm: Anonymous routing protocol for mobile ad hoc networks," *International Journal of Wireless and Mobile Computing*, vol. 3, no. 3, pp. 145–155, 2009.

[71] A. Boukerche, K. El-Khatib, L. Xu, and L. Korba, "Sdar: a secure distributed anonymous routing protocol for wireless and mobile ad hoc networks," in *Local Computer Networks*. IEEE, 2004, pp. 618–624.

[72] L. Zhuang, F. Zhou, B. Y. Zhao, and A. Rowstron, "Cashmere: Resilient anonymous routing," in *NSDI*. USENIX, 2005.

[73] T. McDaniel, "P2P traffic predicted to double by 2014; online video to dominate," *Daily Tech*, 2010.