Anonymity Online for Everyone: What Is Missing for Zero-Effort Privacy on the Internet?

Dominik Herrmann, Jens Lindemann, Ephraim Zimmer, Hannes Federrath

Anonymity Online for Everyone: What is missing for zero-effort privacy on the Internet?

Dominik Herrmann[1], Jens Lindemann[2], Ephraim Zimmer[2], and Hannes Federrath[2]

[1] University of Siegen, Institute for Information Systems, Germany
herrmann@wiwi.uni-siegen.de
[2] University of Hamburg, Computer Science Department, Germany
{firstname.lastname}@informatik.uni-hamburg.de

Abstract. Privacy is difficult to protect on the Internet, because surveillance is ubiquitous. Researchers have conceived many different countermeasures. However, these solutions have so far failed to find widespread adoption due to poor performance and usability. What is missing is an Internet access that offers a decent level of privacy for average users out of the box. In this paper, we survey suitable lightweight anonymity solutions and present avenues for future research so that Internet service providers can offer anonymity online without compromising performance and usability, i. e. an effortless solution for customers.

1 Introduction

Surveillance is commonplace on the Internet. Not only intelligence services, but also corporate service providers are interested in the activities of users. Ubiquitous connectivity and pervasive data collection are increasing the size of our digital footprint. Mobile devices like smartphones, tablets, and other wearable devices allow more detailed profiling of user behaviour and preferences, resulting in severe infringements of privacy. The right to informational self-determination is becoming more and more difficult to enforce.

Policy makers are struggling to keep up with the fast-paced development. In the future, legislation may become unable to protect the right to informational self-determination, and users may completely lose control of their private data. Some citizens have already taken matters into their own hands. They use privacy tools as a means of self-defence.

Unfortunately, the currently available solutions for self-defence have not seen widespread adoption so far. According to the Tor Metrics project [37], the number of average Tor clients that connect to the Tor network per day from Germany is about 200,000 (based on data obtained in the last three months of 2015). Only a tiny fraction of the estimated 69 million [42] German Internet users route their traffic over the Tor network on a daily basis. Although the adoption of Tor has increased since the Oxford Internet Institute carried out a similar analysis on a global level in 2013 [15]), Tor and other anonymisation systems like JAP are still far away from mainstream.

We believe that adoption of anonymity online will only increase if it causes (virtually) zero effort. Following the principle of privacy by design, anonymous Internet should be a available "out of the box" and not require any involvement of users. This paper surveys existing research on lightweight anonymity solutions that may be useful to turn this vision into a reality. Moreover, we summarise the most important challenges that should be addressed in future research.

The paper is organised as follows. In Sect. 2 we review the most important self-defence techniques that are available today, before we survey open issues and avenues for future research in Sect. 3. We conclude in Sect. 4.

2 State of the Art of Anonymity Online

In the following we review the most important tools that are used in practice at the moment.

Researchers have proposed a number of self-defence techniques. These efforts have resulted in tools such as Tor [27] and JAP [5] that encrypt and route the traffic of their users over multiple network nodes, providing relationship anonymity. Thus, the identity of users (their IP address, to be precise) is hidden from the destinations (for instance, the webservers) as well as observers on the network.

However, many users are switching from desktop computers to mobile and wearable devices, where the installation of client-side proxies or browser extensions is difficult or impossible. There *are* Tor clients for both Android and iOS, but their functionality is limited due to restrictions enforced by the operating systems. Apps like OnionBrowser for iOS [30] and Orbot for Android [31] can only anonymise their own traffic. Neither iOS nor Android offer users a straightforward way to route the traffic of all apps through an anonymisation network. Some apps try to work around these limitation, for instance by running a local proxy in the background or by intercepting traffic of other apps, which is usually only possible if the device has been rooted. However, these workarounds are hardly suitable for non-experts.

Recently, privacy activists have come up with dedicated network devices that can be plugged into the home network. Popular efforts are Anonabox [4], InvizBox [22], Safeplug [35], and eBlocker [8]. Typically, these designs consist of low-cost hardware running a Tor client. While such deployments improve usability, they also increase the attack surface (cf. Anonabox, which suffered from severe vulnerabilities [16]). Moreover, the utility of anonymising home network routers is limited. Users only benefit from them when they are at home, but not when they are on the road with their mobile devices. Furthermore, performance issues with the underlying anonymisation techniques are not addressed. Another open problem is the automatic and robust filtering of additional identifying information such as cookies, HTTP referrer, fingerprinting, or HTTP POST parameters (cf. Sect. 3.4). Nevertheless, anonymising home routers are in demand, as evidenced by successful crowd funding campaigns on Indiegogo and Kickstarter, some of them raising more than 100,000 Euro [24].

For many Internet users the main privacy concern is behavioural profiling, i. e. being tracked by third parties [28]. Over the past five years, a number of countermeasures have been released, either taking the form of browser add-ons or of stand-alone tools.

Ghostery [12] is a popular browser add-on for self-defence that blocks widespread advertising and tracking technologies embedded on websites. It is based on ad-hoc filtering techniques and a database. The embedded tracking snippets do not get blocked before the user approves, which allows for manual inspection and ensures compatibility. As a consequence, the user can balance the trade-off between anonymity, functionality, and usability.

A different approach at achieving the same goal has been studied and developed by the Electronic Frontier Foundation, resulting in a browser add-on called Privacy Badger [10]. The add-on prevents trackers from following user activities over several different websites and browsing sessions by keeping track of different tracking sources and content which is automatically loaded while surfing. Communication with servers is blocked as soon as their already loaded content appears to be used for tracking. The user gets visual feedback via green, yellow, or red slides. Furthermore, Privacy Badger maintains a so-called cookie-blocking yellowlist to identify and allow tracking activities, which are actually important for certain functionalities, while filtering of third-party cookies and referrers is still enforced.

Torres et al. have developed the browser extension FP-Block [38] to counter fingerprinting-based tracking. The authors motivate their work with the observation that Tor fails to meet the needs of ordinary users due to its poor usability. They argue that intra-domain tracking is an acceptable and sometimes even useful feature of web services, in contrast to cross-domain tracking, which reveals user behaviour and preferences not only to one specific website, but potentially globally over large portions of the Internet. Therefore, FP-Block creates and manages several different spoofed web identities as well as a set of fingerprintable browser characteristics, and presents those unrelated identities to different web domains.

ShareMeNot [34] focuses on preventing social media sites from tracking user behaviour and visited websites. To this end cookies relating to "share" widgets, which are embedded and automatically loaded in many websites, are automatically stripped from the HTTP traffic. The transit of these cookies is only allowed if a user wants to interact with them.

Solving the problem of being tracked or profiled by companies via advertisements, referrers, or other techniques that get activated during user clicks of browsing sessions is likewise targeted by the browser add-ons TrackMeNot [20] and AdNauseam [2]. In contrast to the aforementioned countermeasures, these tools utilise a dummy traffic approach to hide the attributable behaviour and preference of users, which is valuable information for tracking or advertising companies. TrackMeNot obfuscates requests to web search engines. It sends fake queries generated from lists of popular search queries as well as RSS feeds of popular news websites and then clicks on some of the search results to mimic typical

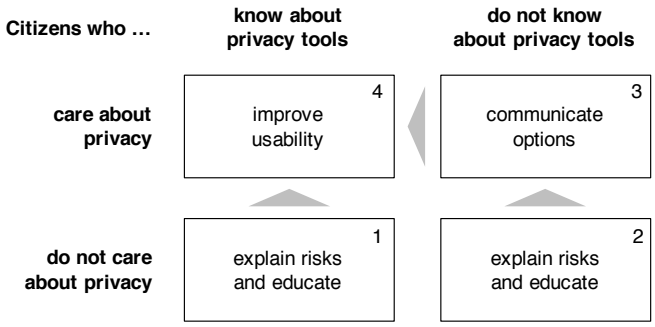

Fig. 1. Four approaches to increase the use of online anonymity solutions among different groups of citizens

user behaviour. AdNauseam automatically activates all advertisements blocked by an ad-blocker on all websites visited by a user, creating the impression that the user actually clicked these ads.

The demand for the mentioned add-ons seems to be high, although they can only be used with some browsers. However, a large part of the online activities nowadays include audio or video streaming as well as interactions on social networks that take place via dedicated (mobile) applications instead of a web browser. Furthermore, the wide range of browser-based countermeasures shows the many possibilities of being tracked and profiled by third parties. New tracking approaches, focusing on other environments like mobile platforms, will certainly be developed in the near future.

3 Avenues for Future Research

The reasons for the low adoption of tools that provide anonymity online are manyfold. Figure 1 presents four typical user groups and what could be done so that they become anonymous online. Some citizens just do not care for privacy online (Groups 1 and 2). Others do care for privacy, but are not aware of the available anonymisation tools (Group 3). Citizens in Group 4 are privacy-conscious *and* aware of anonymisation tools, but do not use them. This irrational behaviour is due to psychological effects such as immediate gratification and hyperbolic discounting [1]. The violation of privacy is a risk that is difficult to grasp and the benefits of using an anonymisation network unfold in the distant future. The impact on comfort, on the other hand, is very tangible: First, there is some initial effort, i. e. users have to install (and also correctly configure) a dedicated client software or browser extension. Second, there is a loss of comfort, because they have to use a restricted web browser that breaks some web sites and surfing on the Internet is slower than before.

Therefore, educating users and advertising the available anonymity solutions are insufficient. First, a certain fraction of users in Groups 1–3 cannot be reached by these efforts anyway, and second, eventually those who are susceptible to these

efforts will eventually end up in Group 4, being frustrated by the limitations of existing anonymity solutions. Ordinary users are unlikely to accept dedicated anonymity tools until their usage becomes virtually effortless and comfortable.

In our view, effortless anonymity that does not compromise comfort is possible by moving towards network layer anonymisation and relaxing the attacker model.

Moving the anonymisation functionality from the transport layer (TCP overlay networks like Tor) *to the network layer* (IP-level anonymisation) decreases its overhead and avoids undesirable effects of TCP congestion control. In addition, this allows users of mobile devices to connect to the anonymisation service without any additional client software, because all of their IP traffic could be transparently routed over the anonymisation network. However, feasibility and security of this approach are unknown so far. Note that the current version of Tor also supports transparent forwarding of traffic, but it is restricted to TCP connections and DNS queries [39].

The second step towards zero-effort privacy protection consists in *relaxing the attacker model*, as proposed by Hsiao et al. [21]. Not all users need the protection offered by Tor and JAP. Some are willing to trust their local Internet Service Provider (ISP). They only want their identity – which can be revealed both by their IP address as well as other tracking mechanisms, such as cookies – to be concealed from the visited websites and third parties (such as ad networks). This more limited attacker model would make it possible to delegate anonymisation to the ISP, as proposed by Raghavan et al. [32]. Ideally, users would not have to install any client-side software or set up additional hardware at home.

Thus, zero-effort anonymity trades off perfect anonymity for better usability, i. e. our objective is to raise the overall level of anonymity that is available out of the box for the majority of users. Achieving this goal involves a number of challenges, which we will describe in the following.

3.1 Relationship Anonymity

The classical approach to provide anonymity for users consists in relaying traffic through multiple mixes or using onion routing. Both techniques achieve relationship anonymity between senders and receivers. Unfortunately, the currently available anonymity services introduce a relatively large overhead, both in terms of message size as well as latency. Moreover, the available bandwidth is limited. The effective performance cannot compete with typical DSL and cable broadband access, as has been shown to be the case for Tor by Fabian et al. [11]. Thus, the existing anonymity services are not suitable for applications that require near real-time transmission of messages (e. g. IP telephony) or a high bandwidth (e. g. video streaming or downloads).

The main goal is therefore to improve the efficiency of mixing and onion routing. This goal can be achieved in two ways: The first approach tries to tailor anonymity solutions to a specific application or protocol, as is the case for the EncDNS [18] service that relays encrypted DNS messages via DNS. The

second approach consists in understanding and tuning the relationship between the underlying transport and the characteristics of the overlay network [40].

One possible direction consists in moving from application layer (such as JAP, which mainly anonymises HTTP traffic) and transport layer (such as Tor, which anonymises TCP traffic) overlays to *network layer* overlays in order to decrease the overhead. This would also eliminate performance issues introduced by the combination of TCP congestion control and multiplexing in the overlay such as *head-of-line blocking* and *cross-connection interference* [23,33,3].

LAP [21] is an example for this approach. It aims to achieve user anonymity by obscuring source and destination addresses via changes in path establishment during routing. HORNET, which was proposed by Chen et al. [7], is another concept aiming to provide anonymity "as an in-network service to all users". Dovetail [36] also strives to attain network-level security. Unlike LAP and HORNET, it fully replaces IP (while HORNET can build upon IP Segment Routing, it can also be used with replacement protocols). Both HORNET and Dovetail do not require a relaxation of the classical threat model and aim to protect users against all adversaries, including global adversaries and the ISP.

Ideally, all traffic should be routed through an anonymity network without the user having to manually configure the applications or the operating system to use the anonymity service. A system achieving this goal can be implemented in different ways.

On the one hand, the anonymity client software could be integrated into the user's router. This has the advantage of not requiring active support by the user's ISP (other than the ISP allowing their customers to use their own routers). Users are also not required to trust their ISPs, as providers would only see traffic flowing from a customer to the anonymisation network, but do not get to know the true destination and contents of the message. The disadvantage of a router-based solution is that it will be typically more expensive and less convenient for users due to the initial setup effort.

On the other hand, the responsibility for routing a user's traffic through an anonymisation network could also be delegated to the ISP. Such a design would require absolutely no effort from users. However, the provider has to explicitly support and offer this service, as such a change cannot be implemented by the user. Another problem is that this model requires users to trust their ISP, as the provider would be aware of the contents of all (not otherwise encrypted) messages as well as their source and destination.

3.2 Privacy-preserving Assignment of IP Addresses

Relaying traffic over multiple mixes or onion routers hides the source IP address of a user from destination servers as well as from other observers on the network, i. e. it obscures the relationship between sources and destinations. Many ordinary users do not demand this rather high level of protection. Their main concern is that destination servers can track their (parallel or consecutive) activities in order to create behavioural profiles. So far, tracking mainly happens on the

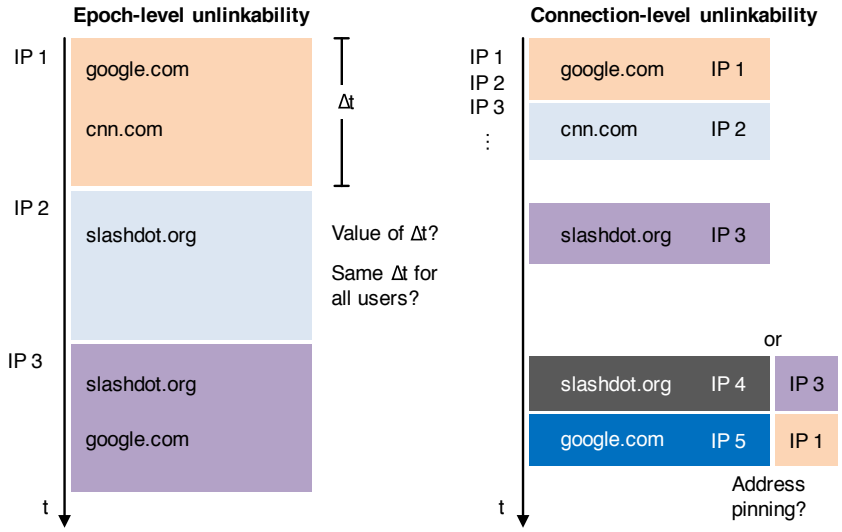

Fig. 2. Address assignment schemes for epoch- and connection-level unlinkability

application layer, for instance via cookies (see Sect. 3.4), and not on the network layer.

However, in principle ad networks could also rely on the source IP addresses to track individual users over time, especially if these addresses change only rarely [6,43]. We expect IP-based tracking to become more relevant in the future due to the rising adoption of IPv6. In the following, we will outline lightweight approaches that specifically prevent IP-based tracking. These approaches are of interest for users who only want their activities to be *unlinkable* rather than hiding their source IP address altogether.

We believe that this could be achieved at very low cost if ISPs changed their address assignment schemes. Today, ISPs assign a single IP address (or IPv6 prefix) to each customer for a certain period of time. Typical addresses change after 24 hours or whenever the broadband router goes offline and online again. In the following we describe two design alternatives for privacy-preserving address assignment and the challenges that have to be overcome (cf. also [17]). Figure 2 shows a side-by-side comparison.

First, ISPs could assign new IP addresses (or prefixes) much more frequently, i. e. resulting in epochs with a duration Δt of a few seconds or minutes. Implemented on its own, this approach provides *epoch-level unlinkability* of IP addresses, which limits the amount of information available for profiling. Obviously, Δt should be as short as possible, because all IP packets sent by a customer within an epoch originate from the same source IP address. Parallel activities and everything within an epoch can still be linked. Apart from determining suitable values for Δt, future research has to look into the feasibility of potential traffic analysis attacks: When a distinctive activity (such as browsing

a rather unpopular website) or a set of activities span multiple epochs, it may be possible to link these events – and potentially all other activities within these epochs. This could be prevented by source address pinning, i. e. re-using the same source IP address that was used for sending traffic to a destination (identified by its IP address or domain name) for the first time. However, address pinning considerably complicates address space management for the ISP.

As an extension to the first scheme, ISPs could provide multiple randomly selected addresses or IPv6 prefixes from their address pool to a customer in each epoch. This would enable customers to make use of multiple source IP addresses at the same time, either one for each destination address or – in the extreme case – one for each TCP connection or even IP packet. This *connection-level unlinkability* would prevent websites and ad networks from linking requests based on IP addresses. However, assigning multiple addresses at once requires considerable changes to the network stack on routers of both the ISP and the customer, as well as to protocols such as DHCP and PPP. Furthermore, it is critical to evaluate the effects of changing source addresses in practice. Changing the source IP too often may cause compatibility problems: For security reasons some web applications terminate HTTP sessions (maintained using cookies) if the source IP address changes during a session. Such issues could be resolved by the aforementioned source address pinning as well. However, enabling address pinning for all destinations by default is undesirable, because it would allow the webservers of ad networks (whose web servers have constant destination addresses) to link requests of the same user.

Apart from the already mentioned assignment variants and strategies for address re-use, there is another more fundamental design choice regarding the role of the ISP: Do we request the ISP merely to provide all the means for preserving privacy, or do we trust the ISP so much that we delegate the actual task of preserving privacy to the ISP? In the first case, the role of the ISP is solely to assign a sufficient amount of addresses to the customer's router and the decision which address is used for an outgoing packet is made on the customer's premises. This offers tech-savvy customers control and transparency, but is technically challenging due to necessary changes to the authentication protocols between the router and the ISP. The second case would be much easier to implement. In fact, it could be implemented by the ISP on its own in a transparent fashion, for instance via carrier-grade network address translation (cf. Raghavan et al. [32]).

3.3 Privacy by Obfuscation

User privacy can also be ensured by obfuscating *activities* rather than identities. This is possible by generating plausible dummy requests and sending them along with the real requests of a user. If the fake requests resemble normal user activity closely enough, it is impossible for the server to distinguish them from the real requests initiated by the user.

TrackMeNot implements this approach to obfuscate queries to search engines. The more recently presented browser add-on AdNauseam aims to make user

profiling for targeted advertising impossible by automatically "clicking" on every ad encountered during websurfing (cf. Sect. 2). Zhao et al. proposed a similar concept for DNS, referred to as "range queries" [44]. However, range queries have been shown to not obfuscate a user's surfing behaviour sufficiently, because the automatically generated dummy requests are not plausible [19]. Generating plausible dummy traffic is still an open problem.

Dummy traffic can either be generated by a software running on a user's client computer (or router) or by the ISP. Generating it on the premises of a user makes deployment cumbersome, not only for users but also for vendors: The dummy generator would need access to (individual or aggregated) traffic profiles of other users in order to create plausible dummies. Therefore, at first sight, delegating the generation of dummy traffic to the ISP is an intriguing idea. After all, the ISP has access to the traffic of all its customers. However, ISPs may be unwilling to offer such a feature, because they risk being sued by service providers that suffer from their dummy traffic.

Whether or not dummy traffic is a suitable tool to provide privacy is debatable. It challenges the dominant business model on the Internet that relies on revenues from advertisements. Moreover, it is a question of morality: Is it acceptable that individuals harm other parties and the environment in pursuit of their own privacy interests? After all, dummy traffic increases the burden on both the servers, which the dummy requests are being sent to, as well as the network infrastructure in general. Eventually, operators would have to invest in additional servers to respond to the dummy requests, which are of no use apart from cloaking user activities. In consequence, energy consumption and thus carbon emissions rise: If all search queries to Google were to be obfuscated by sending an extra ten dummy queries along with each genuine query, this would lead to an additional use of energy equivalent to an emission of about 1.85 million tonnes of carbon dioxide per year (assuming 1.2 trillion queries per year [14], 0.0003 kWh per query [13] and 515.735 grammes of carbon emissions per kWh generated in the US [41]).

3.4 Application Layer Issues

So far, we have discussed techniques to protect the identity of users on the network level (i. e. their IP address). On its own, none of these techniques is sufficient to protect anonymity in practice, because it is rendered ineffective if identifying pieces of information are present on other layers of the network stack, for instance on the transport (e. g. TCP timestamps [25]), session (e. g. TLS client certificates), or application layers. In this section we focus on the application layer.

A well-known example of this are HTTP cookies, whose purpose is to be able to re-identify a user upon subsequent visits to a website. However, re-identification of web users does not necessarily rely on explicit identifiers, but can also be performed in a more subtle way, such as via browser fingerprinting [9,29].

However, application layer tracking does not apply to web browsing only. For instance, some BitTorrent clients disclose the IP address of the client's machine

in the handshake messages [26]. Anonymity is only maintained if all identifying pieces of information are filtered from transmitted messages. Automated filtering of identifying pieces of data is challenging for three reasons: First, the identifiers may take many different forms, which makes it difficult to locate them within the transferred data. Strategic adversaries could even set up a subtle side-channel to bypass the filter. Second, the filter has to replace the identifying pieces of data in a syntactically and semantically correct way so that the functionality of the application is not impaired. Third, the filter must be able to differentiate between unintentionally transferred identifiers and intentional linkage, for instance for the purpose of maintaining state with session cookies during online shopping.

Further, there is the question of placement of such a filter. In order to minimise effort for users, it would be desirable to place the filter at the ISP side. As in routing traffic through anonymisation networks (as described in Section 3.1), placing responsibility with the ISP requires users to trust their ISP. However, such a setup is ineffective here, because identifiers in encrypted streams (including very commonly used HTTPS connections) would pass the filter unaltered. Users would have to install a root certificate issued by their ISP in order to enable the ISP to filter identifiers in encrypted traffic. However, this comes at the cost of losing end-to-end security, which makes it a very poor trade-off.

As an alternative, the traffic could be filtered on the user's side, for instance directly on each of the user's devices. In contrast to deployment at the ISP, this solution is more cumbersome, because users would have to install plug-ins for their applications or – if this is not possible – explicitly configure them to send the traffic through the filtering tool. Effort for the user could be decreased by deploying the filter on the user's router, analogous to the concepts described in Section 3.1. Regardless of the concrete placement, users would still have to install a trusted root certificate on their devices for the inspection of encrypted connections. If this certificate was generated and stored only on the premises of the user, end-to-end security would be intact (under the assumption that the vendor of the filter cannot access the private key).

3.5 Policy Issues

In many countries, data retention legislation is in place to enable law enforcement to combat cyber crime. Typically, this kind of legislation requires ISPs to retain certain data for a set period of time, most importantly which IP address a customer was assigned at what time.

If the responsibility for anonymising a customer's traffic is transferred to the ISP, the ISP may be in conflict with such legislation, which would effectively prevent the implementation of ISP-based anonymisation services. Therefore, it needs to be ensured that the provider can still abide the law by retaining all information required. Depending on the interpretation of the law, this may require ISPs to store the outward-facing IP address(es), i. e. the one of the exit node, and all internal identifiers required to trace back the traffic to a particular customer. Efficient storage techniques will have to be employed to cope with the resulting amount of data.

3.6 Verifiability by End Users

A zero-effort anonymisation solution that is active out of the box should ideally not interfere with the user experience at all. However, good usability has its price: Users would not notice if their traffic was *not* routed through the anonymisation system, either due to unintentional system failures or malicious activity.

One approach to solving this problem consists in a website that checks what kind of information it can get about a user. While such an *anonymity self assessment* is simple to set up, it must be visited by users manually. Moreover, the site has to be trusted by users to show correct results and it might be difficult to determine whether the IP address the webserver sees actually belongs to the user's computer or to a server within the anonymisation network.

Another approach would be a client-side *watchdog software* that continuously checks whether IP anonymisation works properly. The watchdog, which would have to be installed by users on their devices, would automate the process of manually visiting a self-assessment website by contacting one or multiple (trustworthy) remote servers. In case of hardware-based anonymity systems, this could be implemented on the customer's network router that notifies its owner via a status indicator LED.

4 Conclusion

Since the extensive surveillance capabilities of security agencies have been revealed, many citizens have lost faith in technological solutions for privacy protection. As a result, some users have fallen into a state of apathy. They are unwilling to concern themselves with privacy tools for self-defence at all. On the other hand, users who want to protect their privacy have to make difficult decisions that impede both usability and performance, causing analysis paralysis.

In order to change the state of affairs, we believe that it is worthwhile to pursue the goal of designing zero-effort privacy solutions. In this paper we have reviewed existing ideas and presented a number of areas for future work that are based on relaxing the attacker model, network-level anonymization, privacy-preserving IP address assignment, and privacy via obfuscation of activities. Deploying them in practice would raise the overall level of anonymity available out of the box for the majority of users.

References

1. Acquisti, A.: Privacy in electronic commerce and the economics of immediate gratification. In: Proceedings of the 5th ACM Conference on Electronic Commerce. pp. 21–29. EC '04, ACM, New York, NY, USA (2004)
2. AdNauseam (2015), https://dhowe.github.io/AdNauseam/
3. AlSabah, M., Goldberg, I.: PCTCP: per-circuit TCP-over-IPsec transport for anonymous communication overlay networks. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) Conference on Computer and Communications Security (CCS 2013). pp. 349–360. ACM (2013)

4. Anonabox (2015), https://www.anonabox.com
5. Berthold, O., Federrath, H., Köpsell, S.: Web mixes: A system for anonymous and unobservable internet access. In: Federrath, H. (ed.) Designing Privacy Enhancing Technologies, Lecture Notes in Computer Science, vol. 2009, pp. 115–129. Springer Berlin Heidelberg (2001)
6. Casado, M., Freedman, M.J.: Peering Through the Shroud: The Effect of Edge Opacity on IP-Based Client Identification. In: 4th Symposium on Networked Systems Design and Implementation (NSDI 2007), April 11–13, 2007, Cambridge, Massachusetts, USA. Proceedings. USENIX (2007)
7. Chen, C., Asoni, D.E., Barrera, D., Danezis, G., Perrig, A.: HORNET: high-speed onion routing at the network layer. In: Ray, I., Li, N., Kruegel, C. (eds.) Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12–16, 2015. pp. 1441–1454. ACM (2015)
8. eBlocker (2015), https://www.eblocker.com/
9. Eckersley, P.: How unique is your web browser? In: Atallah, M.J., Hopper, N.J. (eds.) Privacy Enhancing Technologies, 10th International Symposium, PETS. Proceedings. Lecture Notes in Computer Science, vol. 6205, pp. 1–18. Springer (2010)
10. Electronic Frontier Foundation: Privacy Badger (2015), https://www.eff.org/privacybadger
11. Fabian, B., Goertz, F., Kunz, S., Müller, S., Nitzsche, M.: Privately waiting – A usability analysis of the Tor anonymity network. In: Santana, M., Luftman, J.N., Vinze, A.S. (eds.) 16th Americas Conference on Information Systems, AMCIS 2010, Lima, Peru, August 12–15, 2010. Association for Information Systems (2010)
12. Ghostery (2015), https://www.ghostery.com/en/
13. Google Official Blog: Powering a google search (2009), https://googleblog.blogspot.de/2009/01/powering-google-search.html
14. Google Zeitgeist (2012), https://www.google.com/zeitgeist/2012/#the-world
15. Graham, M., De Sabbata, S.: Information Geographies at the Oxford Internet Institute The anonymous Internet (2015), http://geography.oii.ox.ac.uk/?page=tor
16. Greenberg, A.: Anonabox Recalls 350 "Privacy" Routers for Security Flaws (2015), http://www.wired.com/2015/04/anonabox-recall/
17. Herrmann, D., Arndt, C., Federrath, H.: IPv6 Prefix Alteration: An Opportunity to Improve Online Privacy. In: 1st Workshop on Privacy and Data Protection Technology (PDPT 2012), co-located with Amsterdam Privacy Conference (APC 2012), Amsterdam, Netherlands, October 7–10, 2012. Proceedings (2012), http://arxiv.org/abs/1211.4704
18. Herrmann, D., Fuchs, K., Lindemann, J., Federrath, H.: EncDNS: A Lightweight Privacy-Preserving Name Resolution Service. In: Kutylowski, M., Vaidya, J. (eds.) Computer Security – ESORICS 2014 – 19th European Symposium on Research in Computer Security, Wroclaw, Poland, September 7–11, 2014. Proceedings, Part I. Lecture Notes in Computer Science, vol. 8712, pp. 37–55. Springer (2014)
19. Herrmann, D., Maaß, M., Federrath, H.: Evaluating the Security of a DNS Query Obfuscation Scheme for Private Web Surfing. In: Cuppens-Boulahia, N., Cuppens, F., Jajodia, S., Kalam, A.A.E., Sans, T. (eds.) ICT Systems Security and Privacy Protection – 29th IFIP TC 11 International Conference, SEC 2014, Marrakech, Morocco, June 2–4, 2014. Proceedings. IFIP Advances in Information and Communication Technology, vol. 428, pp. 205–219. Springer (2014)
20. Howe, D.C., Nissenbaum, H.: TrackMeNot: Resisting surveillance in web search. In: Lessons from the Identity Trail: Anonymity, Privacy, and Identity in a Networked Society, pp. 417–436 (2009)

21. Hsiao, H.C., Kim, T.J., Perrig, A., Yamada, A., Nelson, S.C., Gruteser, M., Meng, W.: LAP: Lightweight anonymity and privacy. In: Security and Privacy (SP), 2012 IEEE Symposium on. pp. 506–520. IEEE (2012)
22. InvizBox (2015), https://www.invizbox.io/
23. Karol, M., Hluchyj, M., Morgan, S.: Input versus output queueing on a space-division packet switch. IEEE Trans. on Communications 35(12), 1347–1356 (1987)
24. Kickstarter: InvizBox Go – Privacy Made Easy (2015), https://www.kickstarter.com/projects/683682172/invizbox-go
25. Kohno, T., Broido, A., Claffy, K.C.: Remote physical device fingerprintin. IEEE Trans. Dependable Sec. Comput. 2(2), 93–108 (2005)
26. Manils, P., Abdelberi, C., Blond, S.L., Kâafar, M.A., Castelluccia, C., Legout, A., Dabbous, W.: Compromising Tor anonymity exploiting P2P information leakage. CoRR abs/1004.1461 (2010), http://arxiv.org/abs/1004.1461
27. Mathewson, N., Syverson, P., Dingledine, R.: Tor: the second-generation onion router. In: the Proceedings of the 13th USENIX Security Symposium (2004)
28. Mayer, J.R., Mitchell, J.C.: Third-party web tracking: Policy and technology. In: IEEE Symposium on Security and Privacy, SP 2012, 21–23 May 2012, San Francisco, California, USA. pp. 413–427. IEEE Computer Society (2012)
29. Mowery, K., Shacham, H.: Pixel perfect: Fingerprinting canvas in HTML5. Proceedings of Web 2.0 Security and Privacy 2012 (2012)
30. OnionBrowser iOS App (2015), https://mike.tig.as/onionbrowser/
31. Orbot Android App (2015), https://guardianproject.info/apps/orbot/
32. Raghavan, B., Kohno, T., Snoeren, A., Wetherall, D.: Enlisting ISPs to Improve Online Privacy: IP Address Mixing by Default. In: Goldberg, I., Atallah, M. (eds.) Privacy Enhancing Technologies, Lecture Notes in Computer Science, vol. 5672, pp. 143–163. Springer Berlin Heidelberg (2009)
33. Reardon, J., Goldberg, I.: Improving Tor using a TCP-over-DTLS Tunnel. In: USENIX Security Symposium. pp. 119–134. USENIX Association (2009)
34. Roesner, F., Kohno, T., Wetherall, D.: Detecting and defending against third-party tracking on the web. In: Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation. pp. 12–12. NSDI'12, USENIX Association, Berkeley, CA, USA (2012)
35. Safeplug (2015), https://pogoplug.com/safeplug
36. Sankey, J., Wright, M.K.: Dovetail: Stronger anonymity in next-generation internet routing. In: Cristofaro, E.D., Murdoch, S.J. (eds.) Privacy Enhancing Technologies – 14th International Symposium, PETS 2014, Amsterdam, The Netherlands, July 16–18, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8555, pp. 283–303. Springer (2014)
37. TorMetrics: Top-10 countries by directly connecting users (2015), https://metrics.torproject.org/userstats-relay-table.html
38. Torres Ferreira, C., Jonker, H., Mauw, S.: FP-Block: usable web privacy by controlling browser fingerprinting. In: Computer Security – ESORICS 2015 – 20th European Symposium on Research in Computer Security, Vienna, Austria, September 21–25, 2015, Proceedings, Part II, Lecture Notes in Computer Science, vol. 9327, pp. 3–19. Springer (2015)
39. Tor Trac Wiki: Transparently Routing Traffic Through Tor (2015), https://trac.torproject.org/projects/tor/wiki/doc/TransparentProxy
40. Tschorsch, F., Scheuermann, B.: How (not) to build a transport layer for anonymity overlays. ACM SIGMETRICS Performance Evaluation Review 40(4), 101–106 (2013)

41. United States Environmental Protection Agency: Power profiler (2015), http://www2.epa.gov/energy/power-profiler
42. Worldbank Open Data (2015), http://data.worldbank.org/
43. Xie, Y., Yu, F., Achan, K., Gillum, E., Goldszmidt, M., Wobber, T.: How dynamic are IP addresses? In: Conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM 2007). Proceedings. pp. 301–312. ACM, New York, NY, USA (2007)
44. Zhao, F., Hori, Y., Sakurai, K.: Analysis of privacy disclosure in DNS query. In: 2007 International Conference on Multimedia and Ubiquitous Engineering (MUE 2007), 26–28 April 2007, Seoul, Korea. pp. 952–957. IEEE Computer Society (2007)