

On Building Onion Routing into Future Internet Architectures

Daniele Asoni, Chen Chen, David Barrera, Adrian Perrig

► **To cite this version:**

Daniele Asoni, Chen Chen, David Barrera, Adrian Perrig. On Building Onion Routing into Future Internet Architectures. International Workshop on Open Problems in Network Security (iNetSec), Oct 2015, Zurich, Switzerland. pp.71-81, 10.1007/978-3-319-39028-4_6 . hal-01445794

HAL Id: hal-01445794

<https://hal.inria.fr/hal-01445794>

Submitted on 25 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



On Building Onion Routing into Future Internet Architectures

Daniele E. Asoni, Chen Chen, David Barrera, and Adrian Perrig

Network Security Group
Department of Computer Science
ETH Zürich

Abstract. User privacy on the Internet has become a pressing concern in recent years largely due to the revelations of large scale network surveillance programs. Research initiatives around future Internet architectures (FIAs) offer a unique opportunity to integrate privacy protection measures into the architecture of the network itself. In this paper, we survey the main design challenges of network layer onion routing protocols in FIAs. We empirically investigate the requirements and trade-offs of different design choices. Our goal is to identify promising research directions and incentivize further exploration of the field.

1 Introduction

Recent revelations about global-scale pervasive surveillance [13] programs have demonstrated that Internet users' privacy is severely threatened. These revelations suggest massive amounts of private traffic, including web browsing activities, location information, and personal communications are being harvested in bulk by domestic and foreign intelligence agencies. In response to these threats, an increasing number of privacy-concerned users are resorting to anonymity tools and services. The state-of-the-art solutions today are onion routing systems (most notably Tor [11]), which try to strike a balance between privacy and performance, enabling low-latency anonymous communication suitable for typical Internet activities (e.g., web browsing and instant messaging). Many of these systems work on top of the Internet as overlay networks: they rely on a number of servers, typically provided and run by volunteers, which anonymize user traffic by relaying it across the network a number of times. While these systems are gaining popularity, the active number of users still represents only a small fraction of the entire Internet population, partly because of these systems' limitations in terms of latency and scalability.

In recent years, to overcome the performance and scalability limitations of traditional anonymity systems, researchers have explored a new approach: building anonymity systems directly *into* the network architecture [20, 28, 10, 23, 6]. Instead of relying on volunteer-run servers, these proposals require Internet routers to perform traffic anonymization in addition to their typical packet forwarding operations. Current research in Future Internet Architectures (FIAs) [17, 33, 34] gives the opportunity to concretely plan for and evaluate this paradigm shift.

Research on network-layer anonymity systems is still in its infancy. Only a small fraction of the design space has been explored so far, and many of the important challenges and design decisions in the field have not been analyzed in detail. This paper

aims to help fill this gap by identifying the main problems that arise when designing such systems, and by analyzing the trade-offs brought by those design choices.

The remainder is organized as follows. Section 2 gives background about traditional anonymity systems, FIAs, and recent network-layer anonymity systems based on FIAs. In Section 3, we discuss the necessary considerations when defining a threat model for onion routing systems at the network layer, showing how performance requirements and the network topology itself bound the privacy guarantees that such systems can provide. In Section 4 we present a set of design challenges and propose possible solutions as well as potential research directions. We conclude in Section 5.

2 Background and Related Work

Network-layer anonymity systems are usually an adaptation of traditional overlay anonymity systems. Thus, many of the fundamental concepts remain the same in both types of systems. For this reason we begin by presenting the traditional systems, in particular focusing on mix networks and onion routing. We then describe relevant FIA-based proposals, and finally present recent research on network-layer anonymity systems.

2.1 Traditional Anonymity Systems

The first system designed to enable anonymous communication over the Internet was proposed by Chaum in 1981 [4]. The main idea in this system is as follows: an end-host (the *source*) wishing to communicate anonymously with another end-host (the *destination*) chooses a sequence of servers (generically called *nodes*) that will relay the traffic. We call this sequence a *path*. Additionally, the source encrypts each message it sends multiple times in such a way that every node on the path will be able to remove one layer of encryption, until the final node (or the destination) obtains the original message. This technique is called onion encryption. Since messages look different (as a result of encryption or decryption) before and after being processed by a node, and under the assumption that many users will send messages through the system, it is non-trivial for the adversary to trace messages and thus to de-anonymize the communicating parties.

Chaum's design also includes batching and mixing of messages at every node, which increase the difficulty for an adversary to trace those messages across the network. For this reason Chaum's system, and others that are based on the same principles [18, 8, 9], are called mix networks or mix-nets. These systems typically do on-the-fly key establishment for every message using the long-term public keys of the nodes on the path. Key negotiations, together with batching and mixing, make mix-nets very slow, and thus suitable only for latency-tolerant applications like email.

The other important category of anonymity systems, which derives from mix-nets, is that of onion routing systems. The main examples are Tor [11], I2P [32] and JonDonym [22, 14]. These systems also use onion encryption, but they typically do not perform mixing or batching to avoid the performance penalty. They also create *circuits* (also called tunnels or sessions), i.e., they establish shared keys with each node on the path, and then use these keys to send many messages/packets over the same path. The

overall speed of onion routing systems means that, unlike mix-nets, they can be used for applications like web browsing and instant messaging. The drawback of these systems is that they provide weaker security guarantees, which are typically expressed by considering threat models with more limited adversaries.

Likely due to their low latency, (circuit-based) onion routing systems are the most used anonymous networks today. Furthermore, the network-layer anonymity systems described in Section 2.3 are also mostly based on onion routing.

2.2 Future Internet Architectures

In response to the problems that the current Internet is facing, a number of research initiatives were started with the goal of defining new network architectures for the next-generation Internet [26]. As research in re-defining the Internet is still ongoing, there is an opportunity to integrate support for privacy-enhancing technologies into the network architecture itself.

Some of the new architectures that have been proposed already include features which can be leveraged by anonymity networks (though the reason for their inclusion in the design lies strictly in networking aspects). In particular, some of these FIAs grant the end hosts a certain degree of control over the path that their traffic takes to traverse the network [31, 17, 34]. Control, or at least knowledge of the path, is typically offered at the granularity of Autonomous Systems (ASes) or Internet Service Providers (ISPs). Control and visibility of network paths is a fundamental property leveraged by network-layer onion routing systems, with the main realization being that intermediate ISPs and/or ASes can act as nodes to perform traffic anonymization. Assuming that the ISPs and ASes have public cryptographic keys that can be obtained and verified by the source, it is even possible for the source to negotiate keys with the nodes on the path to perform cryptographic operations on packets, (e.g., onion encryption [6]).

2.3 Network-Layer Anonymity Systems

The most practical and most used anonymity systems today are onion routing systems. However, application-layer overlay networks, on which today's onion routing systems are based, have inherent performance limitations. First, since each hop can traverse the entire network, the total propagation delay can be high. Second, the end hosts' network stacks add substantial processing and queuing delay [12]. Finally, compared to ISP infrastructure, volunteer-run nodes typically offer only low to medium throughput [30].

Recent works have proposed to address the performance limitation of onion routing systems by building anonymity systems into the network layer [20, 28, 10, 23, 6]. LAP [20] and Dovetail [28] (so-called *lightweight systems*) hide end hosts' network locations by concealing routing information. However, in these two protocols, packets remain unchanged as they traverse the network, making both schemes vulnerable to trivial packet matching attacks. In contrast, Tor instead of IP [23] and HORNET [6] advocate performing onion routing at the network layer, where Internet Service Providers (ISPs) assume the role of onion relays and support per-hop high-speed onion encryption/decryption.

We note that an important difference (and limitation) of network-layer anonymity system compared to overlay systems is that in the former the nodes typically only forward traffic to adjacent nodes. In overlay anonymity networks, on the other hand, it is assumed that each node can communicate with any other node. In Section 3 we show the limitations that this difference entails.

As discussed in Section 2.2, we assume that the network architecture provides end hosts with information about the network and the paths, which is a fundamental requirement of many of these network-layer anonymity systems. However, it is worth noting that LAP [20] differs from the other proposed schemes in this respect, as it does not require path information to be known to the source. The reason LAP does not rely on this assumption is that its privacy guarantees are weak: the source has no control over whether its traffic is truly anonymized, and has to fully trust its ISP. We discuss anonymity guarantees in the next section.

3 Threat Model Considerations

Traditionally, high-latency mix systems [4] consider powerful Dolev-Yao adversaries (i.e., adversaries that control the entire network), and typically try to guarantee the highest degree of anonymity. Defining a threat model that is as clear for low-latency onion routing systems is generally more difficult. Low-latency schemes are unable to defend against Dolev-Yao adversaries, and almost any observation of the network increases the knowledge of the adversary, thereby affecting anonymity. This means that the definition of anonymity needs to be quantitative [2], and this is a challenging task as it requires an analysis of the actual network topology and the entities involved. For this reason even the most popular anonymity systems provide only some approximate notion of what an adversary is allowed to do, and what guarantees the system provides for its users [11]. For network-layer anonymity systems these challenges remain, but additional elements must be considered.

Performance constraints. First, it is important to note that performance is a primary goal for network-layer low-latency anonymity systems. This implies, for example, that performing cryptographic operations on the packet should not constitute a bottleneck that limits a node’s throughput. Some of the proposed lightweight schemes [20, 28] lower the anonymity guarantees significantly in order to achieve better performance and scalability (see Figure 1), by considering a weak threat model in which at most one node can be compromised. HORNET [6], an onion routing systems, tries to raise the bar, preserving user anonymity against more powerful adversaries, but has to sacrifice some performance. Unlike LAP and Dovetail, HORNET requires an initial circuit setup which involves expensive asymmetric cryptography operations, and data packets need to be fully onion encrypted/decrypted at each node. While the loss in performance is a clear downside of such a scheme, it appears to be unavoidable if the goal is to protect against stronger (and arguably more realistic) adversaries.

Topology constraints. Another important aspect of network-layer protocols is that they are constrained by the network topology and the business relations between network entities. This reduces the anonymity guarantees that such systems can provide [6], as

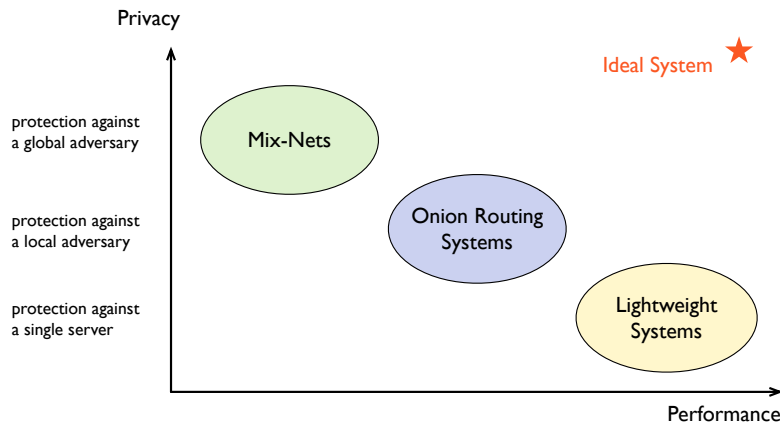


Fig. 1. Conceptual representation of the performance-privacy trade-off for anonymity systems, according to the categories described in Section 2 (figure adapted from Hsiao et al. [20]).

sources are not free to choose arbitrary paths traversing the network (a property explicitly enabled by overlay systems). For example, in the extreme case where the source and destination are in the same ISP, complete de-anonymization is trivial if the ISP is malicious. This shows one of the important challenges in the definition of the threat model: typically low-latency anonymity systems require that an adversary should not be able to observe both source and destination. However, for a variety of reasons a user may not trust its own ISP, so in such cases the system does not provide any guarantees.¹

Trust assumptions. Traditional onion routing protocols typically use secure channels between nodes to protect against eavesdroppers. Tor, for instance, specifies TLS connections to transmit data between pairs of relays. At the network layer, an equivalent security mechanism would be some form of link encryption (e.g., IPsec [29]). In HORNET, the authors argue that link encryption may not be required under all network settings and for all adversaries. The reasoning is that if the threat model considers an adversary that can legally get access to the communication links, then the same adversary may be able to coerce ISPs into handing over all encryption keys, which would make link encryption pointless. However, should the threat model be different, it could be beneficial to use link encryption. One such case would be if multiple links between nodes, usually considered separate, all traverse a single Internet exchange Point (IXP) [21]. This example shows the importance of the threat model definition, which directly influences the design decisions.

¹ It is possible to mitigate these restrictions by adding redirection in the network (see Section 4.5).

4 Trade-offs and Challenges in the Design Space

We now analyze some of the main design questions and challenges that arise when building onion routing systems at the network layer. While most of the items below concern performance, some of these design decisions also depend on the chosen threat model (see Section 3) and expected security guarantees.

4.1 Stateful vs. Stateless Node Design

To process and forward a packet, onion routing relays require state including cryptographic keys and routing information. In Tor, each relay node maintains the state for all the circuits traversing the relay. Given the high throughput and the limited amount of high-speed memory on routers, a stateful relay-node design creates scalability problems. We estimate, based on packet traces of an edge router [3], that a router with 100 Gb/s would need around 20 GB of state if all flows traversing through that router were Tor flows.

The ever decreasing costs of memory hardware might overcome the problem in the long term, but we note that to achieve high performance, routers need specialized high-speed memory, whose cost is higher than system memory. Furthermore, stored state is always a challenge for parallelization: if multiple cores need to access the same state, there are synchronization problems, and caching becomes less effective; if multiple machines need to access the same state, the state needs to be replicated. Hence, to mitigate such scalability problems in a stateful design, ISPs must either equip routers with a large high-speed memory, or carefully conduct load-balancing to delegate the state across multiple routers.

The diametrically opposed design choice is a stateless relay design, where each packet carries the necessary state embedded in its header [20, 28, 6]. However, the stateless design requires large packet headers, reducing the effective throughput; it also requires cryptographic schemes that protect packet-carried state from tampering, and prevent it from leaking information about end hosts. This leads to more complex designs which require additional cryptographic operations and whose security is more difficult to analyze (cf. the simplicity of Tor [11], which is stateful, with more complex design of HORNET [6], which is stateless).

4.2 Transport Control

Tor guarantees lossless and in-order delivery of packets by using reliable transport between neighboring relays. Although this scheme (based on TLS) enables the detection of malicious packet modification, replay, or dropping, it introduces a high overhead due to additional processing and queuing² [16]. In comparison, protocols without reliable per-hop transport control [20, 28, 6] can reach lower latency and high throughput, but they consider different (usually weaker) threat models.

² Alternative transport designs have been studied to improve the performance of Tor in this respect [27, 1], but the fundamental problem of queuing remains.

Another disadvantage of having only end-to-end transport control is that the transport layer of the network stack is exposed to the remote endpoint. This can, for instance, allow a malicious destination to perform network stack fingerprinting (e.g., learn what operating system is running). Note, however, that regardless of the design, some part of the network stack is always exposed. For example, in Tor it is still possible to obtain information about the host at the application layer (e.g., HTTP).

4.3 Circuit Setup

Low-latency onion routing systems need to set up circuits, a process which typically involves asymmetric cryptographic operations. The established circuits are then used to forward traffic efficiently (i.e., using only symmetric key cryptography). Existing onion routing systems adopt one of two strategies for the circuit setup: telescopic setup and direct setup. The telescopic setup [11] consists in the extension of the circuit one hop at a time, so that the key exchange with the i -th node is completed before the $(i + 1)$ -th node is contacted. The telescopic setup guarantees forward secrecy, but imposes high communication latency ($O(N^2)$, where N is the number of nodes on the path). Furthermore, this type of setup requires that the underlying FIA allow traffic to flow in both direction on all links, which might not always be the case [19].

The direct setup [6] does not suffer from these drawbacks. In particular, the communication latency is linear in N , but it cannot provide forward secrecy. Whether this is acceptable or not depends on the threat model: if it is assumed that the adversary will not be able to compromise certain nodes in the future, then a direct setup can be used. Otherwise it is still possible to mitigate the impact of a future compromise by changing circuit setup keys regularly. It can be challenging, however, to distribute the new keys if they are changed frequently. We note that direct setup does not preclude the end hosts from establishing an end to end channel that guarantees forward secrecy for the contents of the communication.

4.4 Bootstrapping Anonymous Communication

A difficult and still open problem is how to bootstrap anonymous communications in future Internet architectures. A source needs certain information to set up circuits, which usually at least includes the address of the intended destination, of a path to that destination³, and of the public keys of the nodes on that path. The challenge lies in retrieving this information in an anonymity-preserving way, while still providing the scalability and low latency properties of the main anonymity protocol (see Section 3).

Overlay anonymity systems (for instance, Tor) have an advantage in this case, because any sequence of mixes/onion relays can be used to reach any destination. This allows a source to construct a circuit without risking potential de-anonymization, as the circuit should give the attacker no information about the intended destination. Once the circuit is established, the source can ask the last mix/onion relay to perform the destination lookup through the established circuit. For onion routing protocols at the network layer this scheme cannot be directly applied, as the circuits are constructed on network

³ A path is made of a sequence of nodes that can route traffic from the source to the destination.

paths which leak information about the source’s location and the intended destination. We point out that this problem affects also the lightweight anonymity systems LAP [20] and Dovetail [28].

The simplest solution would be to use an overlay system only for bootstrapping, and then switch to a network layer onion routing system for the actual anonymous communication. The drawback of such a scheme would be a higher delay before the circuit is established, which is a problem for usability. Furthermore, the most popular anonymity systems today support a few million users a day [30], but if network anonymity protocols are used by a more significant fraction of the Internet population there might be scalability issues [6], so this option’s feasibility requires further investigation.

Another possibility is to implement a broadcast mechanism which pushes all topology information (or a large part of it) to all users. Pathlet routing [17] might be used to achieve such a broadcasting system, as proposed by Sankey and Wright for Dovetail [28]. It is unclear whether such a system would scale to the size of the Internet, especially considering that public keys and certificates would also need to be broadcast. A more plausible scheme could be based on the idea of Federrath et al. [15], which combines a broadcasting mechanism for popular destinations with an overlay system for the rest.

For completeness we mention that there are systems that allow private information retrieval (PIR) [7], which means that a client is able to obtain information from a database without the database server knowing what piece of information is being accessed. Such systems could be used to solve certain parts of the problem, for example retrieval of the destinations address. However, such schemes impose high computational burden on servers, and while there have been some applications to anonymous communications [25], these do not appear to be well suited for our purposes.

4.5 Hybrids between Network-Layer and Overlay Anonymity Systems

To mitigate the problem of topology constraints (Section 3), Sankey and Wright [28] suggest that ISPs could deviate from the standard valley-freeness constraint⁴ to allow some redirection. It is unclear, however, whether such a scenario would be deployable in practice, since it requires ISPs to use resources to forward traffic that is neither originating from, nor destined to, one of their clients. A similar, but more radical approach is to assume that a number of end hosts could act as proxies, and thus add an additional global redirection that would eliminate the problem. This approach is an example of combining network layer anonymity protocols with overlay systems, albeit with the drawback of increased in communication latency.

It may be desirable to allow the user (or software acting on the user’s behalf) to select from a range of protocols for a specific connection. Such flexibility would allow clients to dynamically trade off performance and privacy as needed for each particular case and adversary. More research is needed to accurately classify protocols and connection types based on their performance and privacy guarantees.

⁴ Valley freeness is a routing property that derives from the fact that ISPs and ASes have an incentive to only forward traffic that either comes from, or is destined to, one of their customers. When this property holds, for example, no ISP reflects traffic arriving from outside the ISP back to where it originated.

4.6 Legal Issues and Deployment Incentives

To date, the research community has focused on the technical aspects of anonymity systems while largely neglecting the legal aspects and economical incentives for the entities who should deploy those systems. We argue that it is important to consider, at least at a high level, what legal obstacles network entities might face when trying to deploy anonymity systems, and what incentives these entities have for making such systems available to their customers. Indeed, even the most secure and high-performing system is pointless if it cannot be used, so it is worth considering these issues during technical protocol design.

ISPs may benefit from offering anonymity as a service to both their immediate subscribers and to subscribers of other ISPs, but ISPs must simultaneously comply with legal requirements (e.g., state-mandated data retention laws) to facilitate the investigation of criminal activity. A naive solution here is to build anonymity systems with so called “master keys” that ISPs or entities with judicial power can use to de-anonymize communicating parties. This approach, however, is prone to abuse which could lead to pervasive de-anonymization of all users.

A perhaps better-suited solution may be for ISPs to keep logs of users generating anonymous traffic, while not allowing the immediate de-anonymization of traffic. If legally compelled to de-anonymize traffic, the ISP can assist in recording anonymous traffic, possibly cooperating with neighboring ISPs. De-anonymization attempts can then be performed for specific points on the path.

The recently proposed cMix [5] system consists of a fixed sequence of mix nodes, and the authors suggest that each node could be run by a different nation: such a distributed system would guarantee that several legal domains must be involved and must agree on sharing information to de-anonymize single communications.

But while the technical community can find various options with different trade-offs, lawmakers should decide the correct balance between the right to privacy and free speech, as well clearly define the role played by law enforcement in cases where anonymous networks are used. Recent work [24] points out the need for further research on the legal questions that arise around anonymous communications. The sooner there is clarity on these matters, the sooner ISPs and other parties can decide whether to and how to invest into these technologies.

5 Conclusion

This paper has given an overview of the design considerations, trade-offs, and challenges in deploying onion routing anonymity systems on future Internet architectures. Recent research has shown that network-level anonymous networks are not only feasible in practice, but can provide better performance and better privacy to any application. As deployment of future Internet architectures gains momentum, we expect that these theoretical anonymous network proposals will begin to see real-world adoption. While more research is needed to further investigate these and other aspects, we hope the discussion herein will guide exploration of the design space, ultimately leading to more efficient and more secure anonymous networks.

References

- [1] Mashael AlSabah and Ian Goldberg. “PCTCP: per-circuit TCP-over-IPsec transport for anonymous communication overlay networks”. In: *ACM CCS*. 2013.
- [2] Michael Backes and Aniket Kate. “AnoA: A Framework For Analyzing Anonymous Communication Protocols”. In: *IEEE CSF*. 2013.
- [3] CAIDA UCSD Anonymized Internet Traces 2014. http://www.caida.org/data/passive/passive_2014_dataset.xml. Retrieved on 2015.04.30.
- [4] David L. Chaum. “Untraceable electronic mail, return addresses, and digital pseudonyms”. In: *Communications of the ACM* 24.2 (1981).
- [5] David Chaum et al. *cMix: Anonymization by high-performance scalable mixing*. Tech. rep. 2016. URL: <https://eprint.iacr.org/2016/008>.
- [6] Chen Chen et al. “HORNET: High-speed Onion Routing at the Network Layer”. In: *ACM CCS*. 2015.
- [7] Benny Chor et al. “Private information retrieval”. In: *Journal of the ACM* 45.6 (1998).
- [8] George Danezis, Roger Dingledine, and Nick Mathewson. “Mixminion: Design of a Type III Anonymous Remailer Protocol”. In: *IEEE S&P*. 2003.
- [9] George Danezis and Ian Goldberg. “Sphinx: A Compact and Provably Secure Mix Format”. In: *IEEE S&P*. 2009.
- [10] Steven DiBenedetto et al. “ANDaNA : Anonymous Named Data Networking Application”. In: *NDSS*. 2011.
- [11] Roger Dingledine, Nick Mathewson, and Paul Syverson. “Tor: The Second-Generation Onion Router”. In: *USENIX Security*. 2004.
- [12] Roger Dingledine and Steven J. Murdoch. *Performance Improvements on Tor or, Why Tor is slow and what we’re going to do about it*. Tech. rep. The Tor Project, 2009. URL: <https://research.torproject.org/techreports/performance-2009-11-09.pdf>.
- [13] S. Farrell and H. Tschofenig. *Pervasive Monitoring Is an Attack*. IETF RFC 7258.
- [14] Hannes Federrath. “AN.ON – Privacy protection on the Internet”. In: *ERCIM News* 49 (2002), p. 11.
- [15] Hannes Federrath et al. “Privacy-preserving DNS: Analysis of broadcast, range queries and mix-based protection methods”. In: *ESORICS*. 2011.
- [16] John Geddes, Rob Jansen, and Nicholas Hopper. “IMUX: Managing Tor connections from two to infinity, and beyond”. In: *WPES*. 2014.
- [17] P. Brighten Godfrey et al. “Pathlet routing”. In: *ACM SIGCOMM* (2009).
- [18] Ceki Gülcü and Gene Tsudik. “Mixing Email with Babel”. In: *NDSS*. 1996.
- [19] Yihua He et al. “On routing asymmetry in the Internet”. In: *IEEE GLOBECOM*. 2005.
- [20] Hsu Chun Hsiao et al. “LAP: Lightweight anonymity and privacy”. In: *IEEE S&P*. 2012.
- [21] Aaron Johnson et al. “Users get routed: traffic correlation on Tor by realistic adversaries”. In: *ACM CCS*. 2013.
- [22] *JonDonym*. <https://anonymous-proxy-servers.net/>. Retrieved on 2016.02.24.
- [23] Vincent Liu et al. “Tor instead of IP”. In: *ACM HotNets*. 2011.

- [24] Tomas Minarik and Anna-Maria Osula. “Tor does not stink: Use and abuse of the Tor anonymity network from the perspective of law”. In: *Computer Law & Security Review* 32.1 (2016), pp. 111–127.
- [25] Prateek Mittal et al. “Scalable Anonymous Communication with Provable Security”. In: *USENIX HotSec* (2010).
- [26] Jianli Pan, Subharthi Paul, and Raj Jain. “A survey of the research on future internet architectures”. In: *IEEE Communications Magazine* 49.7 (2011), pp. 26–36.
- [27] Joel Reardon and Ian Goldberg. “Improving Tor using a TCP-over-DTLS Tunnel”. In: *USENIX Security*. 2009.
- [28] Jody Sankey and Matthew Wright. “Dovetail: Stronger anonymity in next-generation internet routing”. In: *PETS*. 2014.
- [29] Karen Seo and Stephen Kent. *Security Architecture for the Internet Protocol*. IETF RFC 4301. 2005.
- [30] *Tor Metrics*. <https://metrics.torproject.org>. Retrieved on 2015.05.13.
- [31] Xiaowei Yang, David Clark, and Arthur W Berger. “NIRA: a new inter-domain routing architecture”. In: *IEEE/ACM Transactions on Networking* (2007).
- [32] Bassam Zantout and Ramzi Haraty. “I2P data communication system”. In: *ICN*. 2011.
- [33] Lixia Zhang et al. “Named data networking”. In: *ACM SIGCOMM*. 2014.
- [34] Xin Zhang et al. “SCION: Scalability, Control, and Isolation on Next-Generation Networks”. In: *IEEE S&P*. 2011.