Analyzing and Enhancing Routing Protocols for Friend-to-Friend Overlays

Extended Abstract of the Dissertation

Stefanie Roos
Chair of Privacy and Data Security
TU Dresden
stefanie.roos@tu-dresden.de

February 2016

In the last decades, digital communication has become an integral part of our life. At the same time, large-scale Internet surveillance through governmental and commercial parties has emerged as a serious threat to user privacy. Two of the most prominent examples illustrating the drastic extent of this threat are the unveiling of the NSA's global surveillance program [1] and Facebook's accidental publication of private data on a large scale [2]. The threat presented by accidental data loss is immediate, as normally uninvolved and uninformed parties can abuse the private information for undesired user profiling and criminal purposes such as burglary [3]. In contrast, the impact of global surveillance on users' everyday life is less palpable but all the more dangerous for society as a whole. Large-scale surveillance opens the door for global censorship, in particular the repression of minorities and inconvenient opinions. At worst, companies and governmental parties can abuse their knowledge of and power over enormous amounts of information to manipulate public opinions. In this manner, they can possibly even dictate the results of elections or similar important political decisions. Even without the actual execution of this capability, the knowledge of surveillance on its own calls forth self-censorship [4] and radical personality changes [5]. As a consequence, the current large-scale Internet surveillance drastically undermines freedom of speech, an essential human right and the foundation of modern democracy.

A key issue enabling large-scale surveillance is the convergence towards monopolization of Internet services, so that a handful of companies control the majority of Internet traffic [6]. By restricting the majority of the communication to a small number of parties, large-scale censorship merely requires the cooperation of these companies by sabotage, blackmail, or the application of legal (but potentially immoral) means. Particular examples of the latter are the National Security Letters of the United States [7] and the recent data retention laws ('Vorratsdatenspeicherung') [8] in Germany. The centralized nature of today's Internet service thus enables global surveillance and manipulation at a (relatively, if adversary is in a position of power) low cost. Hence, decentralization of services is the first step in the prevention of global surveillance.

Decentralized systems are either based on distributed servers, e.g., the social network Diaspora [9], or a completely decentralized P2P overlay, e.g., the file-sharing system BitTorrent [10]. While decentralized servers merely distribute the service on several, seemingly independent service providers, P2P networks consists of end devices of everyday users participating as both servers and consumers. In this manner, P2P overlays completely avoid the use of dedicated servers.

However, simply removing the central point of trust only mitigates global surveillance but fails to prevent powerful adversaries from obtaining large amounts of data. Adversaries with sufficient resources, including governmental organizations as well as large companies, are able to observe and possibly manipulate a large fraction of the communication by controlling a large number of the participating servers or end devices, referred to as *nodes*. As a consequence, infiltrating the system with a large number of artificially created nodes controlled by a single entity, commonly referred to as *Sybils*, is generally easy, as for instance recent attacks on the Tor system show [11].

In addition to being vulnerable to Sybil attacks by computationally powerful adversaries, distributed systems introduce additional vulnerabilities by enabling everyday users to directly influence the provided service. Being both consumers and service providers, users have a higher impact on the system and in particular more opportunities to cause injury to the system, intentionally or unintentionally. For example,

some vulnerabilities in a subroutine of the anonymous P2P system Freenet allow a single user to undermine the complete service without the need of Sybils or unusual computational power [12]. Furthermore, end devices reveal their network addresses, e.g., IPs, to many not necessarily trustworthy participants. As a consequence, the need to establish connections with untrusted strangers and thus reveal private information or at least the participation in the system to such parties induces numerous vulnerabilities and limits the desire of users to participate in the system. So, merely decentralizing the service does not entail privacy-preserving and censorship-resistant communication.

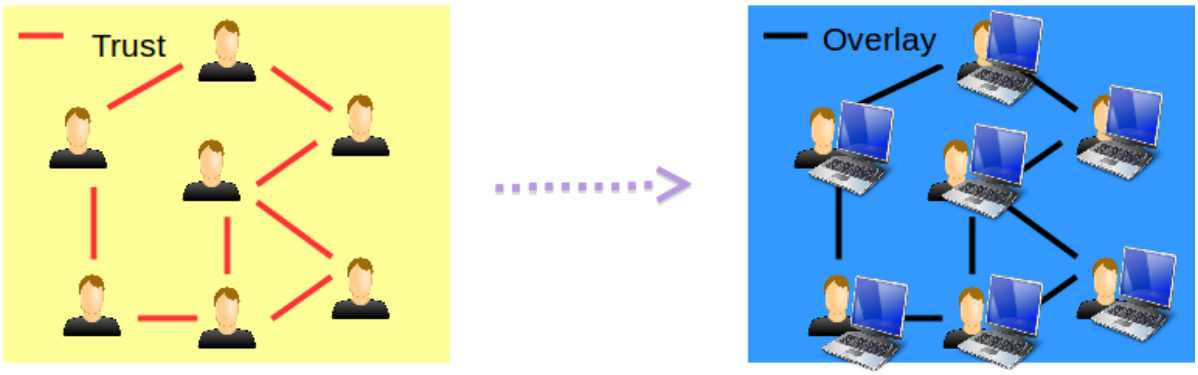

Figure 1: Concept Friend-to-Friend overlay: Overlay connections corresponds to trust relations

Friend-to-Friend (F2F) overlays or F2F networks, also called *Darknets* and first suggested in [13], are a highly promising P2P-based approach to overcome the above concerns. Note that the decisive factor governing the impact of Sybil attacks are the number of connections to honest devices, not the number of fake identities. Hence, rather than impeding the creation of fake identities, F2F overlays raise the cost for establishing connections to honest devices. For this purpose, they restrict direct communication to mutually trusted parties. In other words, the connections in a F2F overlay correspond to real-world trust relationships, as illustrated in Figure 1. In order to surveil or sabotage the F2F overlay, an adversary has to establish trust relationships through the use of social engineering. We assume that establishing trust relationships is costly, at least in comparison to establishing fake identities in an automated fashion. Thus, if the algorithms of the overlay are secure in the sense that they do not permit nodes with a low connectivity to sabotage the complete system, F2F overlays limit the impact of Sybil attacks. In addition to limiting the impact of Sybil attacks, F2F overlays abolish the need to directly communicate with untrusted strangers, thus providing membership-concealment towards untrusted participants.

By restricting the communication to trusted parties, F2F overlays offer a promising communication substrate for privacy-preserving applications such as email, instant-messaging (between trusted parties as well as between anonymous untrusted parties), file-sharing, social networking, and publish-subscribe. Three exemplary use cases for F2F overlays are the following, illustrated in Figure 2:

a) Alice is a whistle blower in an oppressive regime and wants to contact the journalist Bob while hiding the fact that she is the sender of the message from Bob and all other participants.

b) Bob has published information about a health issue he is suffering from under a pseudonym. Alice, who is unsure if she is affected by the same condition, wants to contact Bob anonymously for additional information knowing Bob's pseudonym only.

c) Bob has published information regarding the organization of a demonstration without revealing his identity. The information is stored by Claire, who does not want to reveal that she stores the information due to her fear of persecution. Alice wants to access the information without revealing her identity.

So, F2F overlays need to provide anonymous messaging between two untrusted parties and anonymous sharing of content. On a more abstract level, node and content discovery are the two essential required functionalities. In order to attract a sufficient amount of users, efficient realization of the two functionalities is essential, as is robustness to failures and resistance to attacks on the availability.

A number of deployed F2F overlays are available for use, e.g., Turtle [13], OneSwarm [14], GNUnet [15], and Freenet [16]. However, those overlays are poorly analyzed and exhibit high latency, bad quality of service

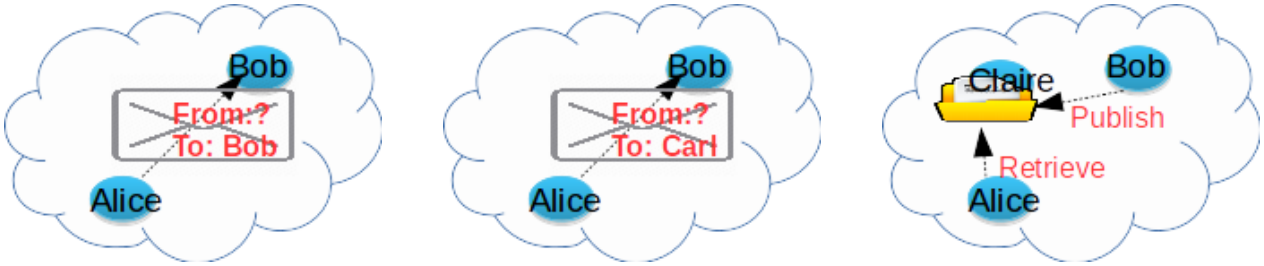

(a) (b) (c)

Figure 2: Exemplary use cases for F2F overlays: a) messaging with anonymous sender, b) messaging with two anonymous parties using a pseudonym for the receiver, c) anonymous content sharing storing the published content at an otherwise uninvolved user; anonymity here refers to hiding one's identity from communication partners as well as all remaining nodes in the overlay

as well as various security issues and privacy leaks [12, 17, 18]. In contrast, there are mainly academic projects, e.g., MCON [17] and XVine [19], which share similar goals. Both MCON and XVine achieve sufficient robustness to failures as well as certain malicious behavior. However, anonymization of communicating untrusted parties is not explicitly considered in both approaches. Furthermore, the communication complexity of messaging and content sharing scales with $\mathcal{O}(\log^2 n)$ for an overlay with n participants while the shortest path in the overlay between any two participants scales with $\mathcal{O}(\log n)$. Thus, neither MCON nor XVine are asymptotically optimal, i.e., their asymptotic communication complexity exceeds the asymptotic lower bound on the communication complexity given by the shortest path between nodes. In addition, the trade-off between the stabilization complexity when nodes join and leave the overlay and the communication complexity of messaging over an extended period of time has not been considered. In summary, the current state-of-the-art leaves room for improvement both with regard to conceptual evaluation and the performance of the algorithms.

In this thesis, we aim to design a F2F overlay that realizes the above functionalities based on censorship-resistant and efficient protocols. We mainly aim to evaluate the underlying concepts of both state-of-the-art approaches and our own designs rather than providing a ready-to-use prototype. For this purpose, we first perform an in-depth evaluation of the state-of-the-art. We show that the existing approaches are inherently unable to simultaneously achieve both efficiency and censorship-resistance. In the second part of the thesis, we design algorithms for communication in F2F overlays. Our design is based upon greedy embeddings, which assign coordinates to nodes in order to facilitate efficient node discovery. However, greedy embeddings have been designed for static environments without malicious parties and the need to hide the identity of participating nodes. Thus, we increase the robustness to failures as well as the resistance to denial-of-service attacks by considering alternative routes and applying multiple diverse greedy embeddings in parallel. Furthermore, we modify the embedding and messaging algorithm to rely on anonymous addresses rather than identifying coordinates. In order to provide content sharing, we establish a virtual overlay on top of the greedy embedding. In this manner, we realize the desired functionalities using efficient and resilient protocols.

We evaluate our algorithms predominately through theoretical analysis, combined with a simplified simulation study. More precisely, we show that the messaging and the stabilization complexity scales logarithmically with the number of participating nodes n. Content sharing is slightly more costly, requiring polylog complexity. Furthermore, we prove that our anonymous addresses indeed always provide sender and receiver anonymity, preventing the attacker from uniquely identifying sender and receiver even if it controls neighbors of both parties. Our extensive simulation study indicates that our algorithms allow faster node discovery and stabilization in the presence of joining and departing nodes than the state-of-the-art approaches. Moreover, the overhead for content sharing only slightly exceeds the overhead of the best state-of-the-art algorithm, while VOUTE requires considerably less stabilization overhead. Last, the robustness to failures and the resistance to various denial-of-service attacks is improved. In summary, our initial design satisfies our requirements and thus presents a promising approach for a real-world implementation.

The contributions presented in this thesis resulted in publications at several top conferences and journals, including INFOCOM [21, 22, 23], PETs [24], and ToMPECs [25]. In order to provide a more extensive overview of this thesis, we now give a concise but informal summary of our requirements and contributions.

1 Requirements

Providing both anonymous messaging as well as content sharing in a F2F overlay is a challenging problem. A F2F overlay has to satisfy many partially conflicting requirements regarding i) *scalability and efficiency*, ii) *robustness and censorship-resistance*, and iii) *anonymity and membership-concealment*. In the following, we motivate each of these three aspects.

Efficiency and Scalability: Efficiency and scalablity are key requirements for any large-scale communication system, because people are unlikely to participate if the system does not provide the expected quality-of-service. In particular, efficient communication implies low latencies, fast stabilization after topology changes, and a low communication overhead measured in the number of messages exchanged for messaging, content sharing, and stabilization after node joins and departures. Scalability implies that the communication complexity increases slowly with the number of participants n, indicating that an approach provides efficient service for large user groups, which we might be unable to simulate or observe in real-world systems. As current F2F overlays have been shown to be slow and unreliable [17], increasing the performance of F2F overlays while still achieving the certain protection goals is the main focus of this thesis.

Robustness and Censorship-Resistance: As F2F overlays are dynamic systems with malicious or malfunctioning nodes, they have to be able to provide reliable communication despite topology changes, in particular node departures, and attacks. Here, we define robustness as the ability of a system to function despite failures of individual nodes. Similarly, we define censorship-resistance as the ability of system to function despite malicious nodes aiming to censor communication. The success ratio of the routing, i.e., the fraction of paths between source and destination in the same component that can be found, should decrease *gracefully* with the number of failed nodes or the number of edges between honest and malicious nodes. We aim to improve the robustness to failures and the resistance to censorship in comparison to state-of-the-art approaches.

Anonymity and Membership-concealment: As stated above, our goal is to prevent an adversary from identifying participants without establishing a direct overlay connection. As a consequence, we even demand that the anonymized topology of the social graph is not revealed, as individuals can be identified from anonymized graphs [20]. In addition, an adversary should never be able to uniquely identify the sender or receiver of a message.

2 Contributions

Overall, our contributions regarding the design of F2F overlays can be classified into three parts:

1. *Preparation*: Review of state-of-the-art, attainment of realistic data sets for our simulation study, and methodology development
2. *State-of-the-Art Evaluation*: In-depth evaluation of existing concepts that have not been sufficiently analyzed in existing literature
3. *Own Design and Evaluation*: Design and evaluation of algorithms for a novel F2F overlay concept based on the insights of the previous parts

Together, these parts present a thorough analysis of F2F overlays culminating in the design of a novel promising concept for messaging and content sharing in such overlays.

Preparation: We review the state-of-the-art and categorize the existing approaches into i) unstructured overlays, ii) virtual overlays, and iii) (network) embeddings. While unstructured approaches cannot achieve the required efficiency, it remains unclear if the virtual overlays and embeddings are suitable approaches. Thus, we conclude that we require a detailed analysis of these approaches.

An appropriate evaluation consists of a rigorous scientific methodology applied to realistic (user) models. In our evaluation, we require i) social graphs as models for the connections in the F2F overlay and ii) join and departure pattern of users in a F2F overlay in order to model the dynamics in the user group. While we

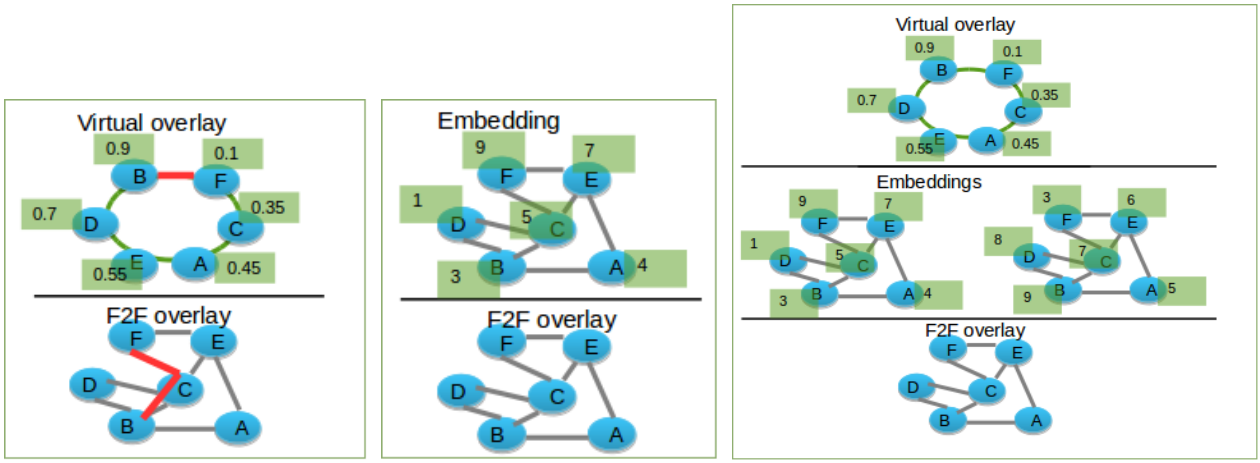

(a) Virtual Overlays (b) Embeddings (c) VOUTE

Figure 3: F2F overlays approaches: a) Virtual overlays such as [19, 17] establish a routable structure and connecting neighboring nodes in that structure through *tunnels* in the F2F overlay, e.g. red tunnel between B and F, b) network embeddings such as the Freenet swapping algorithm [16] assign coordinates that facilitate routing by forwarding to neighbors with coordinates close to destination, c) VOUTE [21] combines the two by facilitating messaging between nodes through multiple network embeddings and enabling content sharing in a virtual overlay leveraging the messaging protocol in the embeddings for communication between virtual neighbors

found suitable existing data sets for social graphs, we performed a measurement study in Freenet to obtain the required join and departure patterns. The measurement study has been published in PETs 2014 [24].

Our methodology for evaluating F2F overlays both theoretical and by simulation leverages the standard methodology for evaluating distributed systems. However, we adapt the methodology for F2F overlays and extend it slightly in order to address certain issues such as the trade-off analysis of stabilization and node discovery complexity. These extensions are of general interest and are partially published in [26].

State-of-the-Art Evaluation: Our review of the state-of-the-art reveals that virtual overlays and embeddings have not been sufficiently evaluated to determine if they fulfill our requirements. Hence, we perform an in-depth analysis based on our previously developed methodology.

First, we show that virtual overlays offer the desired functionalities but fail to fulfill all requirements simultaneously. More precisely, virtual overlays aim to achieve a routable structure by assigning random coordinates to nodes. Neighboring nodes in the routable structure then have to establish *tunnels*, paths in the F2F overlay, to communicate indirectly. The concept is illustrated in Figure 3a. We show that maintaining tunnels of a sufficiently short length requires exceedingly high stabilization complexity in the presence of joining and departing nodes. The result has been published in INFOCOM 2015 [23].

Second, we show that embeddings in their current form cannot combine efficiency and resistance to censorship. Embeddings assign coordinates to nodes in order to facilitate messaging between non-neighboring nodes. More precisely, the embedding coordinates take the role of the receiver address in the message. In addition, content sharing is realized by assigning coordinates to content and storing content on nodes with similar coordinates. We illustrate the concept of embeddings in Figure 3b. In previous work, we showed that purely local round-based embeddings into Euclidean space fail to achieve the necessary efficiency [27, 22, 25]. Thus, we focus on greedy tree-based embeddings. Here, the coordinates are assigned based on a rooted spanning tree of the underlying social graph of the F2F overlay. However, spanning trees do not allow for efficient censorship-resistant content sharing without requiring costly stabilization schemes. The results have been partially published in [28, 29].

In addition, we execute a thorough code analysis of Freenet, the most widely used F2F overlay. In the course, we discover two vulnerabilities of the former implementation. We subsequently design alternative protocols. The improved protocols and their evaluation are published in PETs 2014 [24] and NetSys 2015 [30]. Furthermore, they have been integrated in the current implementation of Freenet.

Thus, our analysis shows that the current concepts are inherently unable to satisfy our requirements.

Own Design and Evaluation: In order to achieve all our requirements, we design and evaluate Virtual Overlays Using Tree Embeddings (VOUTE). Our design relies on tree-based embeddings for messaging between nodes. For content sharing, we establish an additional virtual overlay using the coordinates of the embeddings for communication rather than tunnels. Figure 3c exemplary illustrates the main idea of our design. We modified the standard greedy embedding scheme in order to achieve all our requirements, as detailed in the following.

The scalability and efficiency of messaging and content sharing follows from the efficiency of greedy embedding and virtual overlays. Greedy embeddings allow for efficient stabilization, and removing the need for tunnels thus reduces the stabilization complexity of virtual overlays. We prove that the overall communication complexity increases at most polylog with the number of participants. Furthermore, we perform a simulation study to show that we indeed achieve an improved performance in comparison to the state-of-the-art.

Robustness and censorship-resistance require several modifications due to the vulnerability of the tree structure underlying the embeddings. For this purpose, we use multiple embeddings, allow backtracking if a message cannot be forwarded anymore, and optionally include additional nodes in the communication. In our theoretical analysis, we prove that the changes i) do not considerably reduce the scalability, and ii) indeed increase robustness and censorship-resistance in comparison to the unmodified scheme. In our simulation study, we quantify the improvement in comparison to both the unmodified scheme and state-of-the-art approaches.

We need to modify embeddings in order to provide anonymity and membership-concealment. Using the coordinates of nodes as addresses allows to identify the receivers of messages. Furthermore, the coordinates reveal essential information about the structure of the social graph, which can be used to identify participants. In order to prevent the identification of nodes, we first modify the nature of the assigned coordinates such that guessing a node's coordinate is no longer efficiently possible. Second, we propose a protocol for encrypting coordinates such that the encrypted coordinates represent anonymous addresses. Nodes can thus communicate without requiring the actual receiver coordinates. We show that the modification indeed achieves sender and receiver anonymity at the price of a slightly increased local computation complexity while maintaining the same communication complexity.

In this manner, VOUTE achieves all our requirements on a conceptual level. The results have been accepted for publication in INFOCOM 2016 [21].

In addition to the results presented in the appendix, we apply our novel evaluation techniques in the areas of large-scale discovery services [31, 32, 33], Botnets [34], P2P live-streaming [35], and content-centric networking [29].

In summary, we provide the foundation for a usable privacy-preserving communication system in order to protect freedom of speech. We identify vulnerabilities of the existing approaches, offer immediate solutions for an increased security in the Freenet system, and develop a promising conceptual approach. In the process, we introduce novel widely applicable methods for the evaluation of distributed systems. Our theory-based evaluation shows that our approach offers the required functionalities and opens the door for the development of a real-world implementation within Freenet, a popular privacy-preserving communication system.

References

[1] Nsa files timeline. `http://www.theguardian.com/world/2013/jun/23/edward-snowden-nsa-files-timeline`. Accessed: 2015-10-05.

[2] Facebook data loss. `http://www.reuters.com/article/2013/06/21/us-facebook-security-idUSBRE95K18Y20130621`. Accessed: 2015-10-05.

[3] How burglars use facebook. `http://www.ibtimes.com/how-burglars-use-facebook-target-vacationing-homeowners-1341325`. Accessed: 2015-10-05.

[4] Dawinder S Sidhu. The chilling effect of government surveillance programs on the use of the internet by muslim-americans. *University of Maryland Law Journal of Race, Religion, Gender and Class*, 7, 2007.

[5] We live in public. `http://www.theguardian.com/film/2009/nov/04/josh-harris-we-live-public`. Accessed: 2015-10-05.

[6] Top 10 traffic hogs. http://www.statista.com/chart/1620/top-10-traffic-hogs/. Accessed: 2015-10-05.

[7] National security letters. https://www.eff.org/de/issues/national-security-letters. Accessed: 2015-12-18.

[8] Vorratsdatenspeicherung. http://www.spiegel.de/netzwelt/netzpolitik/vorratsdatenspeicherung-die-wichtigsten-texte-zum-comeback-der-vds-a-1068480.html. Accessed: 2015-12-18.

[9] Diaspora. https://joindiaspora.com/. Accessed: 2015-10-05.

[10] Bittorrent. https://joindiaspora.com/. Accessed: 2015-10-05.

[11] Tor traffic confirmation attack. https://blog.torproject.org/blog/tor-security-advisory-relay-early-traffic-confirmation-attack. Accessed: 2015-10-05.

[12] Nathan S Evans, Chris GauthierDickey, and Christian Grothoff. Routing in the dark: Pitch black. In *ACSAC*, 2007.

[13] Bogdan Popescu. Safe and private data sharing with turtle: friends team-up and beat the system (transcript of discussion). In *Security Protocols*, 2006.

[14] Tomas Isdal, Michael Piatek, Arvind Krishnamurthy, and Thomas Anderson. Privacy-preserving p2p data sharing with oneswarm. In *ACM SIGCOMM Computer Communication Review*, volume 40, 2010.

[15] Nathan S Evans and Christian Grothoff. R5n: Randomized recursive routing for restricted-route networks. In *NSS*, 2011.

[16] Ian Clarke, Oskar Sandberg, Matthew Toseland, and Vilhelm Verendel. Private communication through a network of trusted connections: The dark freenet. *Network*, 2010.

[17] Eugene Vasserman, Rob Jansen, James Tyra, Nicholas Hopper, and Yongdae Kim. Membership-concealing overlay networks. In *CCS*, 2009.

[18] Swagatika Prusty, Brian Neil Levine, and Marc Liberatore. Forensic investigation of the oneswarm anonymous filesharing system. In *CCS*, 2011.

[19] Prateek Mittal, Matthew Caesar, and Nikita Borisov. X-vine: Secure and pseudonymous routing in dhts using social networks. In *NDSS*, 2012.

[20] Arvind Narayanan and Vitaly Shmatikov. De-anonymizing social networks. In *Security and Privacy*, 2009.

Publication List

[21] Stefanie Roos, Martin Beck, and Thorsten Strufe. Anonymous addresses for efficient and resilient routing in f2f overlays. In *INFOCOM*, 2016 (to appear).

[22] Stefanie Roos and Thorsten Strufe. A contribution to analyzing and enhancing darknet routing. In *INFOCOM*, 2013.

[23] Stefanie Roos and Thorsten Strufe. On the impossibility of efficient self-stabilization in virtual overlays with churn. In *INFOCOM*, 2015.

[24] Stefanie Roos, Benjamin Schiller, Stefan Hacker, and Thorsten Strufe. Measuring freenet in the wild: Censorship-resilience under observation. In *PETs*, 2014.

[25] Stefanie Roos and Thorsten Strufe. Dealing with dead ends-efficient routing in darknets. In *TOMPECS*, 2016 (to appear).

[26] Stefanie Roos, Giang T Nguyen, and Thorsten Strufe. Integrating churn into the formal analysis of routing algorithms. In *NetSys*, 2015.

[27] Stefanie Roos and Thorsten Strufe. Provable polylog routing for darknets. In *HotPost*, 2012.

[28] Andreas Hoefer, Stefanie Roos, and Thorsten Strufe. Greedy embedding, routing and content addressing for darknets. In *NetSys*, 2013.

[29] Stefanie Roos, Liang Wang, Thorsten Strufe, and Jussi Kangasharju. Enhancing compact routing in ccn with prefix embedding and topology-aware hashing. In *MobiArch*, 2014.

[30] Stefanie Roos, Florian Platzer, Jan-Michael Heller, and Thorsten Strufe. Inferring obfuscated values in freenet. In *NetSys*, 2015.

[31] Hani Salah, Stefanie Roos, and Thorsten Strufe. Diversity entails improvement: A new neighbour selection scheme for kademlia-type systems. In *IEEE P2P*, 2014.

[32] Daniel Germanus, Stefanie Roos, Thorsten Strufe, and Neeraj Suri. Mitigating eclipse attacks in peer-to-peer networks. In *IEEE CNS*, 2014.

[33] Stefanie Roos, Hani Salah, and Thorsten Strufe. Determining the hop count in kademlia-type systems. In *ICCCN*, 2015.

[34] Shankar Karuppayah, Stefanie Roos, Christian Rossow, Max Mühlhäuser, and Mathias Fischer. Zeusmilker: Circumventing the p2p zeus neighbor list restriction mechanism. In *ICDCS*, 2015.

[35] Giang Nyugen, Stefanie Roos, Thorsten Strufe, and Mathias Fischer. Rbcs: A resilient backbone construction scheme for hybrid peer-to-peer streaming. In *LCN*, 2015.

[36] Chris Biemann, Lachezar Krumov, Stefanie Roos, and Karsten Weihe. Network motifs are a powerful tool for semantic distinction. In *Towards a Theoretical Framework for Analyzing Complex Linguistic Networks*. Springer, 2016.

[37] Frederik Armknecht, Manuel Hauptmann, Stefanie Roos, and Thorsten Strufe. An additional protection layer for confidential osns posts. In *ICC*, 2014.

[38] Hani Salah, Stefanie Roos, and Thorsten Strufe. Characterizing graph-theoretic properties of a large-scale dht: Measurements vs. simulations. In *ISCC*, 2014.

[39] Patrick Welzel, Stefanie Roos, Andreas Höfer, and Thorsten Strufe. Darknetsim: a simulation framework for social overlays. In *Communications & Networking Simulation Symposium*, 2014.

[40] Chris Biemann, Stefanie Roos, and Karsten Weihe. Quantifying semantics using complex network analysis. In *COLING*, 2012.

[41] Benjamin Schiller, Stefanie Roos, Andreas Hoefer, and Thorsten Strufe. Attack resistant network embeddings for darknets. In *SRDSW*, 2011.

[42] Simone A Ludwig and Stefanie Roos. Prognosis of breast cancer using genetic programming. In *KES*. 2010.

[43] Simone A Ludwig, Stefanie Roos, Monique Frize, and Nicole Yu. Medical outcome prediction for intensive care unit patients. *IJCMAM*, 1(4), 2010.