

Towards a Privacy-Preserving Hybrid Radio Network: Design and Open Challenges

Mirco Schönfeld, Martin Werner, Claudia Linnhoff-Popien

Mobile and Distributed Systems Group
Ludwig-Maximilians-Universität in Munich
80538 Munich, Germany

Email: mirco.schoenfeld@ifi.lmu.de, martin.werner@ifi.lmu.de, linnhoff@ifi.lmu.de

Alexander Erk

Institut für Rundfunktechnik GmbH
Floriensmühlstr 60,
80939 Munich, Germany
Email: erk@irt.de

Abstract—A large-scale system hundreds of millions of people encounter every day is radio. While this “system” has a tremendous reach it also is technologically outdated. Technological constraints offer great protection of listeners’ privacy but prevent radio stations from implementing modern business models like personalization and mining valuable user information, at the same time. This paper describes a fully distributed system that aims at overcoming current technological constraints in interconnecting radio stations and their listeners while retaining a comparable protection of sensitive user data. Namely, a peer-to-peer architecture with integrated data-mining-capabilities employing differential privacy is proposed. The system will offer personalizable radio programs to listeners and it will enable radio stations to gather valuable information about their listeners. Furthermore, this paper points out key challenges in deploying, bootstrapping and maintaining such a distributed system.

I. INTRODUCTION

In Europe, there is a large-scale system 630 million people encounter every day – around 84% of the population. In Germany, there are still approximately 29 million people getting in touch with this system – between 6 am and 8 am alone. The large-scale system with such an impressive reach is called radio. In competition to modern music streaming services, radio retained an almost constant popularity across all age-groups. Even young people under the age of 30 spend over 2 hours listening to radio, every day.

From a technological point of view, however, radio retained a state of being touchingly out-of-step. Unlike modern music streaming or Internet services radio offers no interactivity, no personalization, and no feedback or return channel to service providers, whatsoever. Instead, technological constraints induce complete decoupling of service providers and listeners inhibiting radio stations from implementing features like personalization. Also, any mining of user data is prevented due to the unavailability of feedback or engagement mechanisms in the core service. Instead, radio stations rely on a complicated and expensive procedure to gather knowledge about their listeners. On the other hand, because this knowledge is always based on statistical assumptions about the entirety of the listeners and no profile of any single listener exists anywhere in the system, the current service design enforces strongest possible protection of listeners’ privacy.

Meanwhile, Internet services and especially music services identified personalization as a key feature for a tailored service

experience. Personalized music services offer a radio-like experience with integrated feedback mechanisms. By providing feedback to single songs users feed their personal profiles with valuable information about music taste and related context information. Based on profile information the services calculate different recommendations. Spotify, Deezer, last.fm, Napster, and Tidal are only a few of these on-demand music services, and, Spotify alone reported six million paying subscribers in 2014 [1].

Radio stations trying to prevent listeners from drifting away to streaming services could enrich their current service spectrum with personalization mechanisms. But, the scenario in which every radio station sets up their own streaming platform suggests itself only at first glance. Consequently, this would mean to disunite radio into stand-alone music services where each service would offer its own technical and user interfaces impeding bundled integration into consumer electronics or car radios, for example, and, thereby, exposing every single station to direct competition with streaming providers with a considerable group of convinced users.

In contrast, this paper describes the concept of a radio backbone network as a geographically distributed and resilient peer-to-peer (P2P) network interconnecting radio devices all over the world enabling them to aggregate and share information on available radio services, to deliver personalized content to listeners, to establish an interactive feedback channel for user engagement, and to integrate special data mining functionality that takes care of user’s privacy.

The proposed system supports a variety of devices serving as network nodes ranging from DAB- and IP-enabled smartphones or car radios with special interaction constraints to headless information forwarding nodes. At the same time, the proposed system offers benefits for both listeners and radio stations. While listeners are given the ability to interact with radio stations via a feedback channel and to personalize one radio station’s music stream, the privacy of their personal data and music preferences is of important relevance. Meanwhile, radio stations are able to gather more precise knowledge about their listeners in near real-time and to offer personalization features for their content. Furthermore, both listeners and radio stations are enabled to explore new program formats with integrated interactive features.

The remainder of the paper is organized as follows. The following Section II will raise some high-level requirements on




Fig. 1. Item Running- and Toggle Bits that can be exploited for exchanging content

the proposed network design. Section III will then present the vision of a radio network together with insights from research related to the key aspects of the formulated requirements. Some aspects remain unsolved, though, and, therefore, have to be considered for further research as Section IV will show. Section V concludes the paper.

II. DESIGN CONSIDERATIONS

Radio listening habits are heterogeneous with respect to time of day, location and receiving devices. While some people prefer listening to radio in their car others prefer their smartphone or Internet-enabled stereo system. To offer a familiar yet seamless reception experience the envisioned radio network has to reflect this heterogeneity of devices. Consequently, the network's nodes are diverse types of devices offering a variety of interaction concepts ranging from voice-controlled- over touchscreen-enabled- to no interaction at all. Further, it is assumed that devices are able to establish IP-based Internet connections. Additionally, some devices will also be able to receive Digital Audio Broadcast (DAB) signals – the transmission-technique for digital delivery of radio content. Of course, this well-established technique should be integrated into the envisioned radio network. In fact, the network should function as an aggregator collecting and disseminating information on locally available DAB streams. This would allow mobile devices that are both IP- and DAB-enabled to seamlessly switch to DAB-reception where available conserving Internet bandwidth and increasing audio quality.

Considering a hybrid radio network blending reception via DAB- and IP-streams impacts the requirements on the envisioned radio network as well as possible client-side exploitations. On the one hand, one has to bear in mind that while interconnected devices are geographically well distributed the availability of reception via DAB is not. In Germany alone, only few stations exist that are available throughout the whole country. The average reach of a single station further decreases taking Europe or even other continents into account, as well. This leads to the requirement of knowledge being disseminated to relevant geographic regions, only. On the other hand, combining both technologies in single devices allows for privacy-preserving distributed personalization. That is, because DAB allows for transmission of additional meta-information describing the current and future content. For example, DAB standards explicitly ask for a so called DynamicLabel+ feature.

DynamicLabel+ allows the broadcaster to send textual information about the ongoing program that can be presented to listeners. Most importantly though, the DynamicLabel+ feature also provides a simple, yet powerful tool to indicate segments in the linear audio stream and let broadcasters assign metadata (e.g. Artist, Title, Composer) to relevant segments in the linear broadcast. Figure 1 depicts the so called item-running- and item-toggle bit as central elements for segmentation. It can be seen how content is played out regularly as indicated by the item running bit keeping its state. Songs by artists such as Rolling Stones, Elton John and ACDC follow each other and sometimes advertisements are played in between. The item toggle bit indicates exact start- and end-points of each element providing important timing information. Nevertheless, it might happen that an item is interrupted by an urgent announcement. This is indicated by the item running bit changing its state (to zero) for the duration of the interruption. The bit switches again when regular content continues. However, this exact timing information for content elements allows client-side personalization engines to accurately replace single songs with better suited content elements. The alternative content can be received directly from the radio stations streaming servers, from another node in the radio network distributing bandwidth usage evenly over nodes, or from radio history by some kind of time shift buffer either in the local device or in another nodes' buffer in the radio network.

This is an important aspect considering the amount of people listening to radio. For example, Germany's most popular radio station records an average of 1.4 million listeners per hour between 6am and 6pm. Supplying half of those listeners with a stream of alternative music content alone would be a considerable cost factor for one single radio station pushing the idea of a distributed network consisting of equal nodes. The distinction between radio station nodes and listener's client nodes would be only with respect to the usage of the network – either by primarily publishing information or by primarily using this information. However, each client could publish useful information, as well – for example for peer-to-peer personalization and each radio station could act as a data consumer – for example while surveying the clients listening to the current radio program. But, abstaining from a fixed or static network structure leads to a network that is a highly dynamic formation with a considerable amount of nodes entering and leaving, constantly. End-user devices establishing the network will only be available as long as the user actually listens to the

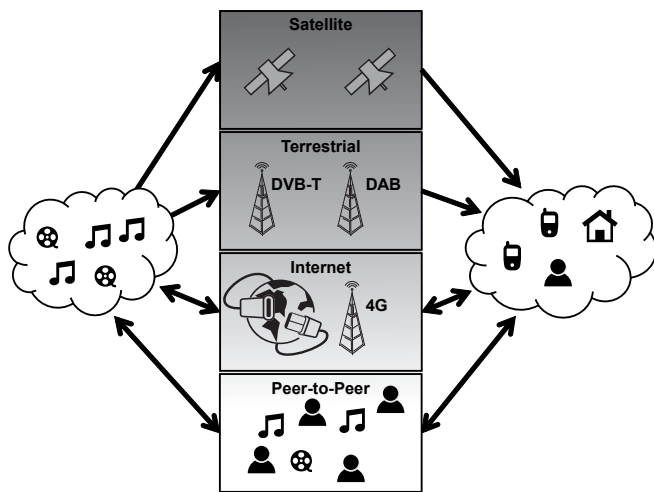


Fig. 2. Current and Future Radio Network Architecture.

radio, after all. This is of important relevance for dissemination of information with narrow geographic relevance as well as for streaming alternative content for personalization among nodes.

Meanwhile, nodes leaving such a fully decentralized network will prevent private data from persist inside the network due to the lack of a central instance collecting and saving user profile information. Hence, the envisioned network structure supports strong privacy demands by design. But, to enable radio stations to gather precise knowledge about their listeners particular data mining techniques have to be integrated which will be highlighted in Chapter III.

In summary, the envisioned radio network will be a fully distributed network consisting of equal nodes, an information-centric network maintaining geographically relevant information, a content-delivery network enabling peer-to-peer personalization, and an overlay network allowing for data mining functionality. Furthermore, DAB reception will be preferred where available to save bandwidth still allowing for personalizing content or for buffering. The following section will provide insights into related research concerned with the mentioned aspects.

III. SYSTEM DESCRIPTION

Figure 2 depicts the classical radio and television network system as well as its current developments from satellite distribution over terrestrial distribution using DVB-T and DAB towards Internet distribution over mobile networks such as LTE. Additionally, we show the envisioned radio network based on peer-to-peer support and personalization as the last step of this innovation process. One of the key feature of the envisioned radio network enabled by peer-to-peer functionality is personalization. Consequently, an Information-centric network (ICN) architecture appears to be suitable. Information-centric networks are an overlay structure for the current Internet tackling upcoming challenges regarding content distribution and increasing client mobility. A key characteristic of an ICN is decoupling sender and receiver with respect to direct communication channels through transparent in-network caching relaxing the requirements in trust and security of

direct communication. This is achieved by employing named data objects enabling requests to directly address desired information instead of the information's memory location as in today's Internet Protocol [2], [3], [4].

The advantages of employing an ICN-overlay for the envisioned radio network are twofold. First, content elements used for personalizing a radio stream can be accessed directly using their name and are automatically cached inside the network. Second, information on the availability of local radio streams can be accessed directly using a name encoding a geographic region. Both types differ in their temporal and geographical importance, though. Since most radio stations are known only inside certain geographic regions information about their availability via DAB is relevant in that region, only. But, since this information barely changes it is expected to be relevant for a long time. Content elements, on the other hand, may be of particular relevance to a certain geographic region, too, especially considering a small radio station offering a certain element as an alternative to the current live-content, but, this increase in geographic relevance may only last a short time. Likewise, certain content elements like latest charts hits may have both geographical and temporal importance above average.

However, personalization of a radio stream can be performed regardless of the source of the radio stream itself – may it be received via HTTP-streaming or via DAB. In any case, a client-side personalization engine observing an upcoming content element being disliked by the user will request the radio station for a list of allowed alternative elements. This list has to be maintained by the radio station itself since radio programs usually comply to elaborate and complex rule sets defining, for example, that songs with certain characteristics (beats per minute, language, genre, etc.) may only be played once within a certain timespan, and, thereby, uniquely shaping the brand of a radio station. However, the client-side personalization engine receiving the list of allowed alternatives is now able to request the network for a chosen alternative using the reported content identifier. The alternative content is received from another node in the network and precisely superposes the unwanted element in the HTTP- or DAB-stream inducing the impression of a personalized radio stream.

Both the description of information characteristics and the personalization workflow above contain several important aspects to be covered in depth in this section. First, the ICN will be built upon a P2P network mainly consisting of intermittently connected, sporadically available nodes sharing information with varying geographical and temporal relevance. This has implications on information dissemination, content delivery and in-network caching of content with respect to network resilience and information availability. Further, analyzing which content elements of the original stream are liked or disliked by which target group, which content elements are being preferred most, or at what time of the day certain elements are replaced frequently constitutes valuable knowledge about a radio station's listeners. Hence, the network needs to integrate data mining functionality to analyze preferences regarding content and to analyze the structure of connected listeners.

In the following, special network monitoring and measurement techniques are introduced that allow information dis-

semination and in-network caching to adapt changing network structure. Descriptions of envisioned data mining functionality will conclude this section.

A. Adaptive information dissemination and in-network caching

As the radio network will primarily interconnect radio reception devices the network structure will be reflected directly by peoples' listening habits: peaks in the morning are expected as well as lows during night hours. Nevertheless, information regarding availability of DAB services nearby and caches of content elements for personalization have to be retained inside the remaining network. Hence, a key feature of the envisioned network will be the ability to self-surveil its viability and adapt dissemination and caching strategies, accordingly, yielding a resilient structure with respect to high fluctuation of constituting nodes. Recent analysis regarding the resilience of Botnets as a special manifestation of P2P-networks provides valuable insights [5]. Also, Bitcoin serves as an example regarding the distribution of nodes across autonomous systems on the network level as well as the nodes' geographic distribution enabling a fast and robust information propagation throughout the network [6], [7].

A popular data structure in distributed computing and Information-centric networking are Bloom filters [8]. A Bloom filter allows to describe sets of elements with a small, constant-size array of bits. In order to compact the set's description and reduce memory consumption or bandwidth, Bloom filters suffer from false positives. Nevertheless, in distributed computing and Information-centric networking, Bloom filters are employed to support caching, routing, forwarding of information [9], [10], [11], [12]. Often, Bloom filters are also employed for monitoring a network. As the structure of the envisioned radio network is essential for retaining information inside the network, monitoring is a central challenge in this scenario. Here, monitoring the network with respect to geographical deployment of nodes is of interest.

Therefore, surveillance of the structure of the surrounding network will rely on a geo-aware variant of OVFSF-coded Bloom filters [13] that is sent by a random node to its neighbouring nodes, periodically. Instead of using OVFSF-codes to insert structured data into the filter, we will incorporate Geohash [14]. Geohash is a latitude/longitude geocode system in which nearby places are most oftenly represented by strings with common prefixes. Incorporating such a Geohash instead of an OVFSF-code tree enables a node receiving such a network surveillance message to insert itself at its specific geolocation instead of inserting itself at an abstract partition of an OVFSF-code as proposed in [13]. An additional geohash-sensitive decaying parameter will ensure that this surveillance poll is bound to a geohash with a certain precision, hence, a limited geographic region around the initiating node. If the encoded bound is reached the surveillance Bloom filter will be returned to the initiating node enabling this node to estimate the number of nodes at each neighbouring geohash-coded location within a certain region. Further, any node receiving or forwarding such surveillance messages is able to keep its own estimation and to adjust this estimation according to frequency and encoded information of incoming messages. Also, the aforementioned geohash-sensitive decaying parameter can be adjusted to the

estimated distribution of nodes resulting, for example, in forwarding a surveillance message to a wider separated region if necessary.

This network surveillance poll will be used to estimate node density in one node's immediate surrounding and to trigger a backup of geographically relevant information if node density impends to underrun a defined threshold. The idea is to prevent information loss caused by a listener turning off his device serving as the last node inside a geographic region. Therefore, information worth protecting has to be copied to regions with higher node density, at the right time, ensuring its retention by a sufficient number of nodes. Such proactive caching strategies sensitive to geographic distribution of mobile nodes have yet to be found for ICNs since most research proposes protocols for wireless P2P-networks [15], [16], [17], [18]. It has to be investigated whether it is feasible to encode the desired backup location into the name of the data object since this results in maintaining several names for single data objects. A different approach could focus on a lower network layer exploiting, for example, link-local-addresses in IPv6 address space using a geo-coded Bloom filter addressing nodes at certain geographic locations. As a side effect, hash collisions in Bloom filters would result in information being forwarded to other than desired locations possibly supporting prevention of information loss.

B. Data mining

For the envisioned radio network the set of features for the user and the broadcaster are of great importance. If the network fails in providing state-of-the-art application level services to both parties, it is of no interest, whether it is fault-tolerant, scalable or anyhow special in comparison to the common centralized approach. Therefore, this section elaborates on distributed data mining with in-network processing and the consequences for users and broadcasters in scalability and privacy.

Classically, radio programs are managed by experts from the field, intuition, and polls. The broadcaster selects the content and evaluates the appropriateness of the selection for his intended audience by producing a loop back from the audience to the broadcaster. A very classic format for this is given by surveys or polls. In these situations, a specific set of questions is given out to the public via the Internet or phone calls or to the audience via the broadcasting medium. Then, the answers to these inquiries are collected, summarized and analyzed in order to form a decision based on this information.

Due to the ubiquitous availability of Internet services, it becomes more and more interesting to be able to perform such surveys online together with automated analysis. Though it would be possible to augment radio with a smartphone app and a cloud service providing this service, this is not a good idea for several reasons: It is not clear, how the tradeoff between user's privacy, survey credibility and anonymity can be resolved. Furthermore, the scalability of the system could be a great concern given the huge amount of listeners to radio programs world-wide.

To this challenge, distributed data mining and especially privacy-preserving data mining propose first solutions in which a network of nodes collaborates in order to create a consensus

about some outcome of distributed observations (e.g., a poll or sensor readings) thereby protecting privacy, credibility, and scalability for a slight amount of communication overhead. In many cases, these consensus protocols [19], [20] create the result of the function for each node such that no special node is needed in order to collect the input information and no node can infer about the answers of any network node unless an attacker controls an unrealistic amount of the network vicinity of a possible victim.

In many cases, the mean and the standard deviation of some value are distributedly agreed on. Still, protocols exist for calculating many different statistics like the sum, the minimum, the maximum, and others.

As privacy is a central concern for public services such as radio, it is noteworthy that distributed privacy provides a host of novel definitions relating privacy not to the amount of information being exposed but to the result of a collaborative algorithm and its sensitivity for specific changes of values of different users. The most important framework of this type is named differential privacy in which privacy is defined as the property of a function not changing more than a predefined threshold if the input values of a single user are varied. It is – in general – fairly easy to protect differential privacy by adding random noise to the data. However, this also degrades the usefulness of the data and might render the complete algorithms useless.

In this context, distributed algorithms are of special importance which gracefully handle erroneous data over time such that the errors inherent in early rounds of a consensus algorithm as well as the errors due to differential privacy are handled over time.

IV. OPEN CHALLENGES

For the envisioned distributed, fault-tolerant, resilient, scalable information network to become a reality, in our opinion, the following central challenges have to be addressed by research and industry.

a) Addressing of information: In an information-centric network, the central way of addressing information is given by metadata. Sometimes, however, the source of some information shall be identified in order to perform some specific advanced operation. Therefore, a mapping from information object metadata to physical addresses such as an IP address and vice versa should be possible without compromising the beneficial properties of the envisioned overlay network.

b) Language support for addressing in ICNs: It is not clear, how different languages and semantic concepts can be used in order to relate information elements which have a different address but similar content. In this area, text mining and relationships between words can be used. Alternatively, semantic technologies such as ontologies can be developed, published and agreed on. Still, both approaches have their specific drawbacks and limitations. Currently, we prefer the first approach over the second, since the development of a sound ontology and the public agreement including standardization is a quite tedious task.

c) Node Bootstrap: In order to have the desired properties in our network, each node has to be able to reach a representative set of nodes. After some time, it is in general easy to do so, though overlay networks can suffer from skewness in the probability of successful communication. Extra care has to be taken on nodes entering the network to enable such nodes to learn about nodes offering relevant information, targeted, ensuring the viability of the overall network.

d) Quality at a specific point in time: The most challenging question for privacy and differential privacy in practice is how to assess the quality of an outcome inside the network at a specific point in time not knowing the global consensus quality or the chosen thresholds for differential privacy without global communication.

e) Information dissemination: A central measure for the effectiveness of an overlay network is the performance of information dissemination among interested peers. Basically, a point-to-point communication protocol in which relay nodes do not store packets has a very small information dissemination performance: Several packets along the path between source and sink are generated for a single information communication. Contrarily, multicast and broadcast protocols have a very good information dissemination performance as long as the capacity is available. For practical reasons, however, tradeoffs have to be discussed in which it might be sensible to reduce the probability of forwarding information in gossip protocols or to somehow avoid loops in the communication reducing redundancy inside the network.

f) Malicious fake nodes: In a large-scale network enabling informed decisions in a business intelligence context, the credibility of the data at hand must be carefully engineered. Basically, producing lots of fake nodes is possible in all network protocols protecting privacy. For our case, however, it is a largely unaddressed challenge how to integrate Turing tests such as CAPTCHAs deeply into the network architecture in order to proof – at least from time to time – that the given data is actually related to an interacting human.

V. CONCLUSION

In this paper, we describe our vision of a large-scale distributed radio network offering personalization, data mining functionalities, and inherent protection of users' privacy. The network will be an Information-centric overlay built upon a resilient peer-to-peer core structure without any central entity. Hence, mining of user data is done on-demand initiating polls and surveys based on distributed consensus protocols. Further, personalization relies on blending DAB radio streams with alternative content received from caching peers in geographical proximity exploiting distinctive benefits of Information-centric networks enhanced with pro-active geo-aware caching strategies.

REFERENCES

- [1] I. F. of the Phonographic Industry (IFPI), "Digital music report," www.ifpi.org/digital-music-report.php, 2014.
- [2] A. Eriksson and B. Ohlman, "Scalable object-to-object communication over a dynamic global network," in *Future Network and Mobile Summit, 2010*, June 2010, pp. 1–8.

- [3] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," *Communications Magazine, IEEE*, vol. 50, no. 7, pp. 26–36, July 2012.
- [4] M. D'Ambrosio, C. Dannewitz, H. Karl, and V. Vercellone, "Mdht: A hierarchical name resolution service for information-centric networks," in *Proceedings of the ACM SIGCOMM Workshop on Information-centric Networking*, ser. ICN '11. New York, NY, USA: ACM, 2011, pp. 7–12.
- [5] C. Rossow, D. Andriessse, T. Werner, B. Stone-Gross, D. Plohmann, C. Dietrich, and H. Bos, "Sok: P2pwned - modeling and evaluating the resilience of peer-to-peer botnets," in *Security and Privacy (SP), 2013 IEEE Symposium on*, May 2013, pp. 97–111.
- [6] S. Feld, M. Schönfeld, and M. Werner, "Analyzing the deployment of bitcoin's p2p network under an as-level perspective," in *SPINS 2014, International Workshop on Secure Peer-to-Peer Intelligent Networks & Systems*, 2014, pp. 1121 – 1126.
- [7] C. Decker and R. Wattenhofer, "Information propagation in the bitcoin network," in *Peer-to-Peer Computing (P2P), 2013 IEEE Thirteenth International Conference on*, Sept 2013, pp. 1–10.
- [8] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Communications of the ACM*, vol. 13, no. 7, pp. 422–426, 1970.
- [9] S. Tarkoma, C. Rothenberg, and E. Lagerspetz, "Theory and practice of bloom filters for distributed systems," *Communications Surveys Tutorials, IEEE*, vol. 14, no. 1, pp. 131–155, First 2012.
- [10] M. Bari, S. Chowdhury, R. Ahmed, R. Boutaba, and B. Mathieu, "A survey of naming and routing in information-centric networks," *Communications Magazine, IEEE*, vol. 50, no. 12, pp. 44–53, December 2012.
- [11] G. Xylomenos, C. Ververidis, V. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. Katsaros, and G. Polyzos, "A survey of information-centric networking research," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 2, pp. 1024–1049, Second 2014.
- [12] H. Cai, P. Ge, and J. Wang, "Applications of bloom filters in peer-to-peer systems: Issues and questions," in *Networking, Architecture, and Storage, 2008. NAS '08. International Conference on*, June 2008, pp. 97–103.
- [13] M. Schönfeld and M. Werner, "Node wake-up via ovsf-coded bloom filters in wireless sensor networks," in *Proceedings of the 5th International Conference on Ad Hoc Networks (ADHOCNETS 2013)*, 2013.
- [14] G. Niemeyer, "Geohash," 2008.
- [15] G. Lai, S. Yang, and Y. Dehui, "Structuring peer-to-peer networks using temporal and semantic locality," in *Database Technology and Applications (DBTA), 2010 2nd International Workshop on*, Nov 2010, pp. 1–4.
- [16] H. Cai and J. Wang, "Exploiting geographical and temporal locality to boost search efficiency in peer-to-peer systems," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 17, no. 10, pp. 1189–1203, Oct 2006.
- [17] A. Kumar, J. Xu, and E. Zegura, "Efficient and scalable query routing for unstructured peer-to-peer networks," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 2, March 2005, pp. 1162–1173 vol. 2.
- [18] A. Esnault, N. Le Sommer, and F. Guidec, "A peer-to-peer overlay system for message delivery in wide intermittently-connected hybrid networks," in *Wired/Wireless Internet Communications*, ser. Lecture Notes in Computer Science, A. Mellouk, S. Fowler, S. Hoceini, and B. Daachi, Eds. Springer International Publishing, 2014, vol. 8458, pp. 200–213.
- [19] M. Schönfeld and M. Werner, "Distributed privacy-preserving mean estimation," in *PRISMS 2014 The 2nd International Conference on Privacy and Security in Mobile Systems (PRISMS 2014)*, 2014.
- [20] R. Wolff, "Local thresholding in general network graphs," *Knowledge and Data Engineering, IEEE Transactions on*, vol. 26, no. 4, pp. 916–928, April 2014.