

Opportunistic Social Networks for Academia

James Leedham¹, Thanassis Tiropanis²

*School of Electronics and Computer Science, University of Southampton
Southampton, UK*

¹jrl@ecs.soton.ac.uk

²tt2@ecs.soton.ac.uk

Abstract—This paper begins by describing the behavioural shift in online interactivity and how it has empowered the growth of online social networks and emergent forms of social awareness. It then discusses the notion of decentralised, peer-to-peer social networks and how this has been largely motivated by ideals that include giving users greater control over the data that they share in order to protect their privacy. At the same time it discusses the motivation for opportunistic contact that takes place when people are in geographical proximity (e.g. in the same office or attending the same event) and how those opportunities seem to establish the right context in which one can share and interpret ideas. Based on the hypothesis that the value of opportunistic contact in decentralised, peer-to-peer networks is more than the sum of its parts, we describe the deployment and evaluation of an application that lets users build opportunistic social networks, making use of decentralised, peer-to-peer infrastructures to ensure that users are in control of their information and leveraging geographic proximity for the exchange of ideas at the right place at the right time. The evaluation of this approach in an academic environment establishes opportunities and challenges for opportunistic social networks and a clear potential of this approach in terms of trust, context, and filtration to increase and improve social awareness.

Index Terms—Online Social Networks, Privacy, Opportunistic Networks, Peer-to-peer Networking, Decentralisation

I. INTRODUCTION

The sociological implications and technological design principles of the Web 2.0 paradigm have been the subject of study for some time [1][2]. Some of the highest profiling sites to encourage significant user participation are those that facilitate social interaction, more commonly referred to as Online Social Networks (OSN) [3]. While these sites continue to enjoy significant popularity, a growing number of researchers have begun to question the degree of security and integrity they demonstrate while conserving and displaying private and personally identifiable information [4][5]. Such authors have similarly drawn attention to the naivety shown by

many users when exercising control over their private data [6], as well as the lack of social trust and context afforded to users due to heavily invested third parties and the abundance of information presented upon entering the network. The apparent scalability of centrally managed OSNs is equally questionable owing to the costs of infrastructure and down time [7]. Some believe the answers to these concerns may lie within the ongoing research and development of new methods for facilitating social interaction over digital spaces, many of which are supported by notions of decentralisation [8] and opportunistic availability [9][10]. This paper examines each of these concepts with regard to their motivations, technical challenges and the results of experiments that have been carried out in order to measure their feasibility. It is important to note that although both notions are distinct, they are not mutually exclusive, and the potential within their overlap is of particular relevance. This paper will also present the findings of a new analysis, the data for which has followed as a result of a recent software trial. The relative significance of this analysis is discussed before several key areas for future work are recommended.

II. MOTIVATIONS

One of the earliest publications to emerge within this area of study appeared in 2008 at the first conference for online social networks [11], and was expanded upon by the same authors the following year [12]. In both papers they highlight that the vast majority of OSN users allow their profiles, as well as their contact lists, to be publicly visible, and suggest that such action is more likely due to negligence than intention. They also demonstrate how certain social networks have leaked user information through their Application Programming Interfaces (API).

Other investigators have emphasised the degree to which an OSN may change its Terms of Service (ToS), allowing it to modify the way in which it handles user data, as well as its functionality and, subsequently, its user experience [13].

With these issues in mind, some have discussed the facility for user-enacted encryption of data within existing OSNs [14]. However, such methods are not without vulnerability [15] and therefore authorities have instead begun to explore the viability of a Peer-to-Peer (P2P) infrastructure to decentralise OSNs and give users greater control over the data they own and share, as well as that which they create and contribute [16][17].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Copyright © 2011 Electronics and Computer Science, University of Southampton

III. PEER-TO-PEER SOCIAL NETWORKS

A. Arguments In Favour

The authors mentioned already advocate how freedom of speech may be uniquely realised within P2P OSNs, as there would exist no single authority with the power to censor, though they admit that this could also be a disadvantage. They go on to assert that the P2P paradigm scales better than its client-server equivalent, as a growing user base naturally brings additional infrastructure, and that such systems may be utilised without an Internet connection, effectively allowing the operation of local area social networks. Separate researchers have also posited that P2P systems may offer unique adaptation to variable circumstances through automatic self-organisation, fault-tolerance and load balancing [18]. Some of these academics have similarly argued that since social networking applications are inherently P2P in nature, as users contribute and consume content produced by one another, a P2P paradigm would most naturally suit their implementation. They also state that the use of P2P OSNs will represent a step towards ubiquitous social networking.

B. Arguments Against

One of the most noteworthy arguments for continuing development of traditional OSNs is that centrally hosted sites can offer better exposure for those wishing to reconnect with forgotten social communities or contacts, such as childhood friends. In addition, there exist attitudes towards privacy that are in opposition to those discussed already, as exemplified by OSN users who freely choose to publicise their online profiles, leveraging the disclosure of such sites in order to appear to the widest audience possible. There are also commercial benefits to OSNs which are easier to realise when employed by a centralised authority. Regardless of these factors, there remains a growing volume of research into P2P OSNs and a number of associated challenges have been identified.

C. Data Storage

One solution involves client-only storage, by which only an individual's own machine holds their personal information [13]. Despite being a low-cost solution, this would result in poor availability. Some have suggested Cloud storage but say that it is too expensive. Most researchers uphold methods for distributed storage whereby user data is replicated over numerous client machines [16], and there exist several ways in which this may be implemented so as to maintain the security of private data [19]. One approach involves the use of Public/Private Key Infrastructure (PKI), allowing user data to be stored securely with any host [20][21]. However, a more socially informed distribution mechanism, such as storing data purely with friends, would permit finer-grained access control [22]. For example, a friend may be allowed to store an individual's user data without encryption since they are trustworthy and may therefore be given complex instructions detailing to whom they may or may not serve all or part of that data to. It is comparably difficult to instruct an untrusted host serving key-encrypted data to do more than permit all or nothing access. An extension of the socially informed distribution model might be to store data of varying levels of sensitivity with friends of friends, and potentially their friends

as well. The drawback to using distributed storage in any of these manners is its high level of redundancy, requiring end users to sacrifice more storage space than they themselves use.

D. Content Dissemination

Ideally any P2P OSN should employ an efficient content dissemination strategy such that it may appear to exhibit the same level of synchronicity as a conventional, server-oriented network might. Some investigators have developed and tested a number of strategies designed to maximise the value of the data exchanged between clients connecting either consciously or opportunistically [23]. They focus on how clients who regularly connect to different parts of a network may best share data so as to propagate only the most recent version of a given piece of information. Such researchers specifically observe how strategies that consider either the geographical areas most visited by a client, or the friends with whom that client most frequently connects, may offer varying degrees of efficiency for different environments.

E. Performance – Indexing, Addressing and Search

Other academics argue that a more significant concern is performance [22]. A decentralised P2P architecture requires a sophisticated system for indexing data and clients, such that individuals are able to find one another and explicitly request whichever piece of data they desire. Most P2P OSNs that have already been developed utilise a Distributed Hash Table (DHT) to accomplish this [24][25][26], in a similar fashion to the BitTorrent protocol's distributed tracker. While this solves the problem of indexing, an associated issue remains; as clients leave and join the network, decisions may be made to replicate their data to other hosts in order to maintain an appropriate level of availability, potentially resulting in excessive consumption of bandwidth. This is referred to as a high rate of churn, and mechanisms are required to keep this as low as possible. Another concern is how to search for clients within a P2P network, as DHT-structured networks only permit exact match queries. For example, a user's request to see John Smith's profile will work as 'John Smith' will hash to a value that exists within the DHT, but a user cannot request the profiles of every person whose last name begins with the letter 'S'. However, one DHT variant has been developed which can offer support for lexicographic key ordering and range queries [27] [28].

Other areas of research include how to enable third party applications, such as those seen on existing, centralised OSNs, as well as how to implement effective and timely content revocation over a network that may be empowered only by ad-hoc, opportunistic propagation [29].

IV. OPPORTUNISTIC SOCIAL NETWORKS

This area is strongly linked to decentralised OSNs and as such a number of scholars have begun to investigate their combined potential within certain environments.

A. Conference Environments

In 2007 several researchers encouraged 28 attendees of an ACM conference in New York to help examine the potential of opportunistic, short-term social networks, and ultimately

confirmed the importance of central nodes with high popularity [30]. The researchers also noted the efficiency of using social neighbours to exchange messages during relatively small windows of opportunity. The experiment involved participants using a Bluetooth-enabled device running software, which indicated when a friend, or friend of a friend, was nearby. Initially participants were asked to select their immediate friends from a list of all attendees. During the three-day event each device logged its awareness of neighbouring handsets and any user actions, such as adding a new friend. After the experiment the researchers analysed each person's social graph over time with respect to where they travelled and who they were physically collocated with. They repeated the experiment at the same conference the following year and recorded similar results [31]. Both experiments confirmed the importance of central nodes, and the efficiency of socially informed data propagation. However, the authors admit that the conference setting is well suited to such behaviour as it provides geographic bounding and its attendees typically have pre-existing social relationships.

B. Office Environments

In 2009 a similar study took place within the University of Milano, involving a survey of almost 300 computer science faculty members and students, followed by a trace recording using 49 mobile devices [9]. 11,895 instances of opportunistic contact were recorded and analysed, leading to the conclusions that strangers can be just as useful, and sometimes more so, than friends in forwarding messages. The initial survey revealed a strong desire among faculty members for a ubiquitous solution to enable the receiving of notices about upcoming meetings or events. Students desired similar ways of receiving notices from friends. 80% specified that they see colleagues at least once a day, providing strong potential for reliable, opportunistic forwarding. The trace recording lasted 19 days, during which time the 49 recorders were carried by staff and PhD students working in a variety of offices and laboratories. The conclusion that strangers can sometimes be more useful than friends in forwarding messages is due to the fact strangers, by their definition, are more likely to visit areas that friends do not, and thus they provide vital links to broader regions of the social network.

Other researchers have also shown mathematically that socially informed content dissemination strategies can offer improved robustness and performance [32]. However, the following are open challenges that have either not yet been considered by this community, or been proposed as future work.

V. OPEN CHALLENGES

Although P2P networks have been motivated by concerns regarding the privacy and security of centralised OSNs, the P2P paradigm itself does not guarantee safety. Some solutions have been proposed, such as PKI, but other issues, such as the secure encryption of network traffic, have yet to be properly considered within this context [8].

Research on content dissemination also includes a number of areas for expansion. Most notably, the examples discussed here do not consider the storage limitations of clients moving

through a network. It may be wise to consider scenarios in which clients use small, limited-capability devices such as smartphones, where they may only have the ability to carry a small number of messages.

Perhaps most notably there appears to be only limited study of the sociological implications of P2P social networks, whether formed opportunistically or otherwise. Upon making this observation, the authors of this paper designed and conducted a short software trial in order to learn whether a social network made up of clients who connect opportunistically over P2P infrastructure can be used to improve awareness and collaborative efforts within a research-driven environment.

VI. OPPORTUNISTIC SOCIAL NETWORKS FOR ACADEMIA

Specifically the authors of this paper wished to know whether enabling and encouraging colleagues to opportunistically share resources, such as journal articles and conference papers, over a given P2P OSN, would result in those individuals gaining:

- A greater awareness of the work being done by colleagues,
- A greater awareness of the development within their department and institution,
- A desire to collaborate more with their colleagues,
- A stronger sense of community within their research group.

The initial hypothesis was that these affordances would become evident and be attributable to factors such as:

- Trust, in shared resources being of genuine quality and usefulness, as they come recommended by impartial colleagues rather than invested organisations or aggregation services,
- Context, of shared resources being relevant to the group's interests,
- Timing, of resources arriving during working hours only, at times when researchers are available to read them,
- Availability, of peers for discussion, as individuals may see when colleagues are online, and from whom shared resources originate, and
- Filtering, of lower quality resources by way of a 5 star rating system.

The experiment involved a total of 15 faculty members, PhD students and technical staff working within the University of Southampton's School of Electronics and Computer Science. The significant technical abilities belonging to participants were not considered to be of detriment to the experiment given its purpose. The 15 volunteers were initially surveyed to find out how they rated their current levels of awareness regarding work done by colleagues, the department as a whole, and the school. Certain existing friendships were also exposed and used to divide participants into three groups of five. Each group was engineered to include some individuals who already knew one another, and were therefore strong, central nodes, as well as some lesser-known members.

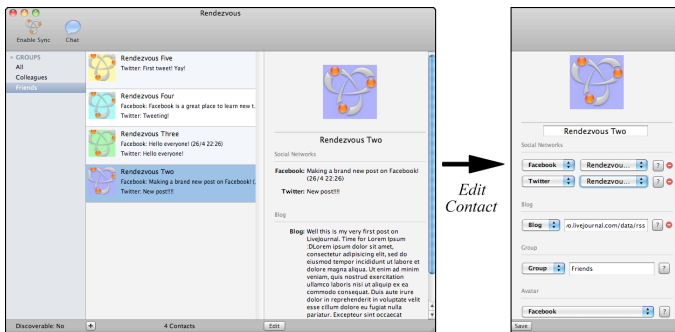


Figure 1 Rendezvous while editing contact information

A. Control Group

The first group acted a control group and its members were asked to simply record the details of each interaction with a colleague that resulted in learning about a new resource, such as an article or paper.

B. 'Rendezvous' Group

Participants in the second group were each given a copy of a custom-written piece of software codenamed 'Rendezvous'. This had been developed prior to the experiment by the authors of this paper and is designed to offer users a contact-oriented view of their social networks by aggregating data from popular providers such as Facebook and Twitter, as well as blogging sites. Rendezvous uses Apple's 'Bonjour' implementation for zero-configuration service discovery in order to find other instances of itself over the local area network. Doing this allows it to provide an Instant Messaging (IM) feature, as well as identify peers with mutual friends and consequently recommend new ways of connecting to such contacts. For example:

Assume Jim and John are friends; both work in the same office, and are both friends with Carly. All three are using Facebook and Twitter, but while Jim is following Carly's status updates via both social networks, John is only following her on Twitter. However, when Jim and John are both running Rendezvous, Jim will be given the opportunity to share with John the knowledge that Carly uses Facebook. If he takes this opportunity, John's client will automatically accept the new data and begin following Carly on Facebook, though John is free to cancel this if he wishes. In addition, Jim may choose to make the process ubiquitous by configuring his client to automatically propagate known social aliases to trusted friends without asking him each time.

Operating in this fashion, Rendezvous enables the provision of 'enhanced presence' within digital social networks by making its users aware of the availability of nearby peers, and potentially revealing the mentalities of those peers. This is another area of relevant research and one that actively considers the sociological and psychological effects of social networking [32][33].

By using Rendezvous, the second group of participants had a potential advantage over the control group, as they could learn about, and share, resources by posting status updates, blog posts, or sending instant messages to peers.

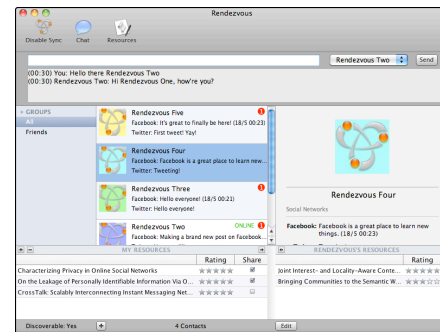


Figure 2 Rendezvous Pro showing the effects of network synchronisation, instant messaging and resource sharing

C. 'Rendezvous Pro' Group

Members of the third group were each given a copy of 'Rendezvous Pro', an extension to Rendezvous, which had been developed exclusively for use during this trial. Rendezvous Pro included the same features as Rendezvous with the additional ability to store and forward resources, such as research material. Users could add their own, or favourite, publications to their client, give each a rating up to five stars, and browse the resources being advertised by other peers on the network. Upon discovering a new article being recommended by a colleague, users could download it to their machines and begin advertising it to their other contacts.

Figure 1 shows how a user can browse and add contacts in Rendezvous. Similarly, Figure 2 indicates how the software reacts to local area network synchronisation with friendly peers. It also shows the instant messaging window in use (top) and the resource-sharing window (bottom). Note that the latter feature is exclusive to Rendezvous Pro.

By using Rendezvous Pro, members of the third group were afforded an even greater potential for increased interactivity with colleagues than those in groups one or two.

The experiment ran for approximately 15 days, with some members beginning and ending at different times due to unavoidable commitments elsewhere. As members in the first group kept manual recordings of their resource-sharing activities, the Rendezvous and Rendezvous Pro applications kept automated logs of theirs. A total of 238 instances of opportunistic contact were recorded, along with 115 and 163 occurrences of Facebook and Twitter statuses being updated respectively. In addition, 26 blog posts were updated, 53 hypertext links (URL) were followed and 117 instant messages were sent between the 10 participants using the custom-written software. Among those using Rendezvous Pro, a total of 15 resources were carried, and there were 25 instances of individuals observing a shared resource.

Following the trial, all participants were asked to complete another survey designed to discover how their levels of awareness regarding colleagues had changed, and how the tools used by groups two and three were of benefit. The results are detailed below.

D. Results

During the initial survey a number of questions were asked which were not repeated in the final survey, as they were designed to validate the motivations for this trial rather than graph sociological change. Answers to the first two questions

were as expected; on average participants travelled to other parts of the campus, distant from their usual place of work, every 15 days, and travelled to other universities or conferences approximately once per year. These figures might have been higher if lecturing staff had been included, as they tend to travel more than PhD students, who represented the majority of volunteers. Encouragingly, 76% indicated that they enjoy forming new social connections during conferences or staff meetings for communicative and collaborative purposes, and that during such an event they would connect with an average of 4 new individuals. Additionally, the authors of this paper wished to confirm that participants had sufficiently similar research areas to be of use to one another. Of those surveyed, 93% said that the people with whom they make social contact worked in the same, or similar, areas as they did. When asked about their levels of awareness, 88% said they felt aware of what their colleagues, and the research group as a whole, worked on and was interested in. However, only 40% felt aware of what the school worked on. Finally, participants were asked about the tools they use for communication and research. 88% indicated they were happy with tools for intra-office communication, with 86% using social networks and most citing Windows Live Messenger as another useful tool. However, only 35% were satisfied with existing tools for accessing and retrieving academic material, and even fewer, 25%, were happy with tools for sharing such material. These results particularly supported the motivations of the authors to develop and deploy the Rendezvous Pro software.

Results from the final survey revealed that among participants in the first group, 40% felt they had gained an increased awareness of what their colleagues were doing, while 20% felt more aware of the department's actions. However, on average, the control group did not have any discernable increase in awareness. Of those in the second group using Rendezvous, 60% felt they knew more about their colleagues while 40% knew more about the department as a whole. Generally those in the second group did not feel Rendezvous was an improvement on existing tools for sharing academic material, for reasons such as it not offering anything more than other social networking tools, for example TweetDeck, and the fact that not everyone is likely to post academic material to their Facebook or Twitter profile. However, as expected, the most promising results emerged from the third group, for which an average of 80% felt more aware of the work being done by colleagues, though only 40% felt more aware of the department's work. 100% agreed that Rendezvous Pro afforded them more trust in the appropriateness of shared resources than existing tools, and 80% agreed that the software provided a better context and mechanism for filtering resources than existing tools. Meanwhile, 60% agreed that being shown when colleagues were available for discussion was useful, though only 20% agreed that the timing of resource retrieval was relevant. In general, 80% of participants felt Rendezvous Pro was an improvement on existing tools for discovering and sharing published material, and although 60% were satisfied with the built-in rating system, some suggestions for improvement included using a distributed rating method and including citation counts.

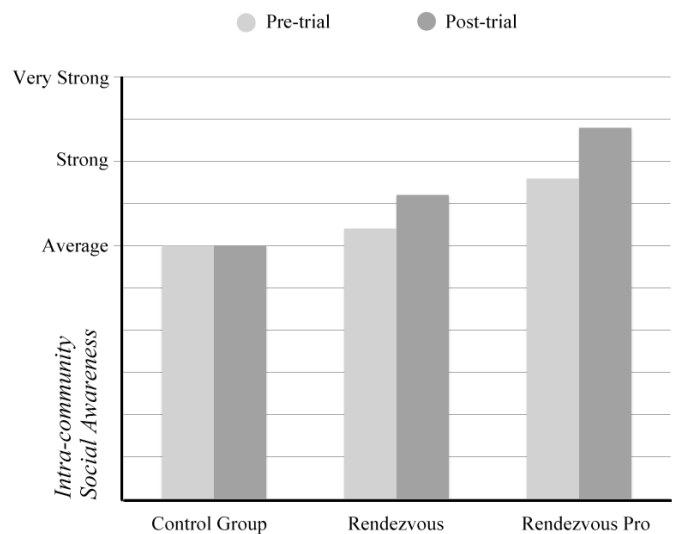


Figure 3 Graph showing the perceived sense of awareness among colleagues before and after the experiment

E. Analysis of Results

The degree to which participants travel between parts of the campus, while expected, is never the less low and might have been detrimental to the propagation of messages if numerous, geographically-separated academic networks had been tested. Many of the other figures provided by the initial survey supported the implementation of Rendezvous Pro and this trial, and more generally support the efforts of those developing other resource sharing tools. The levels of participant awareness indicated by the final survey for all three groups are in line with the initial hypothesis of the authors and are shown in figure 3. From this it is clear that a localised social network, empowered by opportunistic connections, has great potential for improving sociality within a given network, as both the second and third group show increased levels of awareness while the first group shows none. Note that although the percentages suggest a lesser increase for those in the second group, awareness improved almost as much as in the third group, as its participants reported a lower sense of awareness during their initial surveys. It also worth considering whether the increase shown by the third group may have been greater if its members had started with a similarly low level. Finally, it has been seen that the element of ‘timing’ of resources arriving, is of less significance to individuals than originally anticipated; this is particularly relevant to those studying delay tolerant networks.

VII. CONCLUSIONS

This paper initially presented the principal motivations for researching and developing decentralised opportunistic social networks, and highlighted the desire for increased security and ownership over private data, as well as context for sharing ideas, as chief among these. Evaluation of existing studies has indicated not only the feasibility of opportunistic social networking, but also that ‘strangers’ or ‘social neighbours’ can be effective at propagating data. Further results have identified the benefits of ‘socially-aware’ content dissemination strategies. The academic trial discussed previously has shown

empirically that decentralised social networks can be leveraged to afford increased awareness to their users, providing they exhibit characteristics such trust, context and a mechanism for filtering of data. The trial has also demonstrated that timing is of little significance to users, at least within the research environment considered by this study. These are the key findings.

FUTURE WORK

This experiment has shown the potential within given academic environments. However, there may exist value in observing the effects within less densely populated or geographically separated sub networks. In addition, the notion of a distributed rating system is of particular interest and if supported by those attending conferences as well as working in research laboratories, may enable scholars to promote their school's best efforts while off-site. This, however, raises the question of whether ratings become inherently ephemeral and proportional with respect to broader communities.

As high profile, centralised OSNs increasingly support cryptographic connection protocols (SSL), so too must decentralised networks incorporate methods for traffic encryption if they are to advertise improved security.

REFERENCES

- [1] E. Hargittai, J. P. Robinson, P. DiMaggio, W. R. Neuman, "Social Implications of the Internet," *Annual Review of Sociology*, Vol. 27, (2001), pp. 307-336.
- [2] J. Thompson, "Is Education 1.0 Ready for Web 2.0 Students?," *Innovate: Journal of Online Education*, vol 3, issue 4, April/May 2007.
- [3] K. E. Skouby, L. Sorensen, "Next Generation Social Networks - What Users Want," *Centre for Communication, Media and IT, Aalborg University Copenhagen, Lautrupvang*, 2009.
- [4] C. Zhang, J. Sun, X. Zhu, Y. Fang, "Privacy and Security for Online Social Networks: Challenges and Opportunities," *Network IEEE*, vol. 24, no. 4, pp. 13-18, July-August 2010.
- [5] A. Acquisti, R. Gross, "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook," *Privacy Enhancing Technologies, Lecture Notes in Computer Science*, vol. 4258, pp. 36-58, 2006.
- [6] A. Acquisti, R. Gross, "Information Revelation and Privacy in Online Social Networks," in *ACM Workshop on Privacy in the Electronic Society (WPES)*, Alexandria, VA, 2005.
- [7] Q. Wang, Y. Idris, "Affordances of Facebook for learning," *International Journal of Continuing Engineering Education and Life Long Learning*, vol. 19, pp. 247-255, 2009.
- [8] A. Datta, D. Schiöberg, L. Vu, S. Buchegger, "PeerSoN: P2P Social Networking: Early Experiences and Insights," *Proceedings of the 2nd ACM EuroSys Workshop on Social Network Systems*, p.46-52, March 31-31, 2009, Nuremberg, Germany.
- [9] E. Pagani, G.P. Rossi, S. Gaito, "Opportunistic Forwarding in Workplaces," in *Proceedings of the 2nd ACM SIGCOMM Workshop on Online Social Networks (WOSN)*, Barcelona (Spain), Aug. 2009.
- [10] D. N. Tse, M. Grossglauser, "Mobility Increases the Capacity of Ad-hoc Wireless Networks," *IEEE/ACM Trans. Netw.* 10, 4 (Aug. 2002), 477-486.
- [11] B. Krishnamurthy, C. E. Wills, "Characterizing Privacy in Online Social Networks," *Proceedings of the 1st workshop on Online Social Networks*, August 18-18, 2008, Seattle, WA, USA.
- [12] B. Krishnamurthy, C. E. Wills, "On the Leakage of Personally Identifiable Information via Online Social Networks," *Proceedings of the 2nd ACM Workshop on Online Social Networks*, August 17-17, 2009, Barcelona, Spain.
- [13] A. Shakimov, A. Varshavsky, L. P. Cox, R. Cáceres, "Privacy, Cost, and Availability Tradeoffs in Decentralized OSNs," *Proceedings of the 2nd ACM Workshop on Online Social Networks*, August 17-17, 2009, Barcelona, Spain.
- [14] K. Tang, P. Francis, S. Guha, "NOYB: Privacy in Online Social Networks," in *Proceedings of the 1st ACM Workshop on Online Social Networks (WOSP '08)*. ACM, New York, NY, USA, 49-54.
- [15] M. M. Lucas, N. Borisov, "FlyByNight: Mitigating the Privacy Risks of Social Networking," in *Proceedings of the 7th ACM Workshop on Privacy in the Electronic Society (WPES '08)*. ACM, New York, NY, USA, 1-8.
- [16] A. Datta S. Buchegger, "A Case for P2P Infrastructure for Social Networks - Opportunities and Challenges," in *WONS 2009, 6th International Conference on Wireless On-demand Network Systems and Services*, Snowbird, Utah, USA, February 2009.
- [17] D. Schiöberg, "A Peer-to-Peer Infrastructure for Social Networks," *Diplom Thesis, TU Berlin, Berlin, Germany, December 17, 2008*.
- [18] B. Yang, H. Garcia-Molina, N. Daswani, "Open Problems in Data-Sharing Peer-to-Peer Systems," in *Proceedings of the 9th International Conference on Database Theory*, p.1-15, January 08-10, 2003.
- [19] R. Campbell, R. Hasan, L. Brumbaugh, W. Yurcik, Z. Anwar, "A Survey of Peer-to-Peer Storage Techniques for Distributed File Systems," in *ITCC '05: Proceedings of the International Conference on Information Technology: Coding and Computing (ITCC'05) - Volume II*, pages 205-213, Washington, DC, USA, 2005. *IEEE Computer Society*.
- [20] C. Diaz, F. Stajano, J. Anderson, J. Bonneau, "Privacy-Enabling Social Networking Over Untrusted Networks," in *Proceedings of the 2nd ACM Workshop on Social Network Systems (2009)*, pp. 1-6.
- [21] A. Adya, G. Cermak, J. Howell, J. R. Douceur, J. R. Lorch, M. Castro, M. Theimer, R. Chaiken, R. P. Wattenhofer, W. J. Bolosky, "Farsite: Federated, Available, and Reliable Storage for an Incompletely Trusted Environment," *SIGOPS Oper. Syst. Rev.* 36, SI (December 2002), 1-14.
- [22] L. A. Cutillo, R. Molva, T. Strufe, "Privacy Preserving Social Networking through Decentralization," in *WONS 2009, 6th International Conference on Wireless On-demand Network Systems and Services*, Snowbird, Utah, USA, February 2009.
- [23] E. Jaho, I. Stavrakakis, "Joint Interest- and Locality- Aware Content Dissemination in Social Networks," in *WONS 2009, 6th International Conference on Wireless On-demand Network Systems and Services*, Snowbird, Utah, USA, February 2-4, 2009.
- [24] E. Yoneki, J. Crowcroft, P. Hui, "Bubble Rap: Social-based Forwarding in Delay Tolerant Networks," in *Proceedings of the 9th ACM International Symposium on Mobile Ad-hoc Networking and Computing*, May 26-30, 2008, Hong Kong, Hong Kong, China.
- [25] A. D. Joseph, B. Y. Zhao, J. D. Kubiatowicz, J. Stribling, L. Huang, S. C. Rhea, "Tapestry: A Resilient Global-Scale Overlay for Service Deployment," *IEEE Journal on Selected Areas in Communications*, vol. 22, no. 1, pp. 41-53, January 2004.
- [26] A. I. T. Rowstron, P. Druschel, "Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems," *Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms Heidelberg*, p.329-350, November 12-16, 2001.
- [27] K. Aberer, "P-Grid: A Self-Organizing Access Structure for P2P Information Systems," *ACM Special Interest Group on Management Of Data (SIGMOD) Rec.* vol. 32, no. 3, pp. 29-33, September 2003.
- [28] A. Datta, K. Aberer, M. Hauswirth, "Efficient, Self-Contained Handling of Identity in Peer-to-Peer Systems," *IEEE Transactions on Knowledge and Data Engineering* vol. 16, no. 7, pp. 858-869 (2004).
- [29] M. Conti, M. Kumar, "Opportunities in Opportunistic Computing," *Computer*, vol. 43, no. 1, pp. 42-50, January 2010.
- [30] A. Chaintreau, A. Mtibaa, A. Pietilainen, C. Diot, E. Oliver, J. LeBrun, "Are You Moved by Your Social Network Application?," in *Proceedings of the 1st Workshop on Online Social Networks*, August 18-18, 2008, Seattle, WA, USA.
- [31] A. Pietiläinen, C. Diot, E. Oliver, G. Varghese, J. LeBrun, "Mobiclique: middleware for mobile social networking," *WOSN 2009: Proceedings of the 2nd ACM Workshop on Online Social Networks (New York, NY, USA)*, ACM, 2009, pp. 49-54.
- [32] G. Colombo, R. M. Whitaker, S. M. Allen, "Uttering: Social Micro-Blogging Without the Internet," in *Proceedings of the 2nd International Workshop on Mobile Opportunistic Networking (Pisa, Italy, February 22 - 23, 2010)*. *MobiOpp '10*. ACM, New York, NY, 58-64.
- [33] J. Komzak, M. Dzbor, M. Eisenstadt, Y. Vogiazou, "Buddyspace: Large-Scale Presence for Communities at Work and Play," *Appeared in: The Workshop 'The Role of Online Community Spaces in Shaping Virtual Community Interactions', in the International Conference on Communities and Technologies 2003, September 19th - 21st 2003, Amsterdam, The Netherlands*.