Guest Editorial: Special Issue on Clouds for Social Computing

Surya Nepal, Athman Bouguettaya, *Fellow, IEEE*, and Cécile Paris

✦

IN recent time, two complimentary Internet based research areas are emerging: social computing and cloud computing. On the one hand, social computing empowers individual users with relatively low technological sophistication to use the web to engage in social interactions, contribute their expertise and share their content, experiences and opinions. On the other hand, cloud computing offers everyone sophisticated computing infrastructures and resources as utilities, so that individual users with relatively low computing knowledge can have at their disposal a high performing computing infrastructure (e.g., compute, storage, applications, etc.) with little investment. Together, these two complementary technological advances form the backbone of our digitized world, when coupled with the rise of sensors, mobile devices and the internet of things. Of course, they also face significant challenges. We briefly look at each area and describe how they complement each other.

The *Social Web* has become an important means of communication for everyone: individuals, organizations, and governments all use it to disseminate and share information, offer opinions and engage in discussions. This medium creates large social networks through which vast amount of information flows quickly and easily. Many events are now first reported on the social web (e.g., Twitter, FaceBook, etc.). This is true for emergencies such as fires or sudden riots, but also news items. For example, the news of the hit on Bin Laden first broke out on Twitter long before the US president officially announced it on the public media. Another instance is the recent *Occupy of Wall Street* protests which spread out quickly and widely to a larger population due to its Facebook page.

The growth and popularity of social networks pose unique challenges in terms of scalability, maintenance and management. Managing and processing a network with millions of edges (e.g., LinkedIn), distributing status updates to millions of users (e.g., Twitter, Facebook, WhatsUp, etc.), and distributing user generated content to millions of users spread across the globe are some of the practical challenges posed by social networks, or the social web in general, where cloud computing can contribute solutions. An important research area is thus to develop distributed cloud architectures that can handle sustained traffic generated by millions of users in social networks.

One of the interesting features of social networks is their growth characteristics: *sudden* and *unpredictable*. Social networks grow in members and contents following the principle of human dynamics—bursts followed by a long tail. For example, social media platforms such as Twitter go abuzz when a certain event occurs (e.g., Arab Spring [1], Haiti Earthquake [2], New Zealand Earthquake [3], etc.). To address these sudden needs, high computing resources are required during periods of burst characterized by uncertainty and unpredictability. As a direct result, scalability also becomes a problem, i.e., how resources are scaled to cope with the needs of a sudden burst. Cloud computing, with its inherent features of resources elasticity and having applications, software and hardware provided as a service over the Internet [4], becomes a natural choice to address this problem. This is thus the key point of intersection between the social web and cloud computing.

Mobile devices, smartphones in particular, have penetrated our social life, compounding this social communication revolution and its need for scalability of processing power. A recent survey report from TripAdvisor[1] suggests that 87 percent of global travellers use smartphones while on holiday, and 35 percent of US travellers use smartphones for accessing social media. One of the problems with smartphones is the limitation of their battery power. The increasing complexity of new mobile applications puts a heavy load on batteries. To overcome this problem, some techniques off-load mobile computations to software clones of real devices in the cloud. This process involves communication among many entities like mobile devices, their clones, the cloud providers, and mobile network operators. Establishing trust in the communication chain, more specifically with the cloud provider and the mobile network operators, is challenging. A lot of private information could be eavesdropped on and revealed. Yet, as people increasingly employ mobile devices to communicate and interact, both in their private and business lives, secure cloud computing infrastructures are crucial to capturing and processing the data they generate on-the-go. This is the second point of intersection between the social web and cloud computing.

More generally, researchers in social computing have started exploiting the availability of cloud computing to crunch the large amount of data resulting from the digitization of social behavior—i.e., the capture of all social behaviors in a digitized world. Data collected from our daily

- S. Nepal and C. Paris are with CSIRO Computational Informatics, Sydney, Australia. E-mail: {surya.nepal, Cecile.paris}@ndsu.edu.
- A. Bouguettaya is with the School of Computer Science and Information Technology, RMIT University, Melbourne, Australia. E-mail: athman.bouguettaya@rmit.edu.au.

1. http://ir.tripadvisor.com/releasedetail.cfm?releaseid=808058.

activities, such as groceries shopping and surfing the Internet, reveal a lot about our behavior. Based on past buying behavior, our local supermarket knows what we would likely be buying the following Christmas. The introduction of social networks adds yet a new dimension to this process of digitizing social behavior: it reveals our friends, affiliations, beliefs, thoughts, and opinions. It also results in a large data set (known as a *Big Data* problem) which needs to be analyzed to yield meaningful results. The challenge is therefore to be able to achieve the full potential of digitized social behavior. This requires powerful computing and storage infrastructures. Researchers in cloud computing have started developing new efficient and effective big data processing algorithms using clouds. This is the third point of intersection between the social web and cloud computing.

Clouds can be classified into three categories, based on their deployment model: private clouds, public clouds, and community clouds. Of particular interest for this special issue is the community cloud, a collaborative effort where individuals or organizations from a specific community come together to build and share a common infrastructure. This is only possible when we can bring people with common objectives together. The purpose of social networks is specifically to bring people together and form communities. It is thus not surprising that research in social networks and cloud computing intersect in community clouds. Indeed, one could leverage social networks to build a community cloud. This is referred to as a *social cloud*. This is the fourth point of intersection between the two emerging areas of social web and cloud computing.

A fifth interesting area that intersects cloud computing and social computing is *trust*: understanding, modeling, building, and maintaining it. In the social web, "trust" means trusting three different entities: the content, the content provider (i.e., the user, who produces content on the social web), and the networking site provider [5]. Different techniques have been developed to establish trust for these entities [6]. In cloud computing, "trust" traditionally meant relying on the service providers to fulfill all the promises made in their service level agreements (SLAs). In recent times, however, trust in cloud computing has meant more than that. When dealing with confidential data, trust goes beyond the SLAs: cloud computing users want to ensure that their confidential and sensitive data will be handled in appropriate ways. Service providers must gain their users' trust in this respect. Different types of trusted storage services are being developed for the cloud environment to address this concern [7]. Trust is a complex concept, and specific solutions may not be applicable to all application domains; hence further research is needed in this area, and there may be productive cross fertilization between the research of trust on the social web and trust in cloud computing, two research areas which are very active.

The needs of the modern enterprises are driving the convergence of social computing and cloud computing with their enterprise IT systems. Enterprises are using social networks such as Facebook and Twitter to communicate with their customers in real time. Cloud computing vendors like Salesforce, Amazon and Microsoft are providing everything from customer relationship management (CRM) to enterprise resource planning (ERP) applications to platforms for deploying enterprise-specific customized applications, and compute and storage infrastructures. In order to reap the full potential of social computing and cloud computing, enterprises should be able to seamlessly integrate their IT systems with social web and cloud services. This raises an important question: How to integrate such systems without violating the privacy of individuals and increasing the attack space for cyber attacks? Answering this question requires further research. This is the sixth point of intersection between the two emerging areas of social web and cloud computing.

In this special issue, we present a number of selected papers that provide the state-of-the-art research, development and deployment efforts of social computing in clouds or using clouds for social computing bringing these two areas together, illustrating some of the points of intersections mentioned above.

The first three papers are about social clouds, where researchers exploit the implicit and explicit trust built in social networks to form trusted cloud environments, a specific form of community clouds. Mohaisen et al., in their paper entitled "Trustworthy Distributed Computing on Social Networks", propose a social cloud in which computing nodes are governed by social ties driven from a bootstrapping trust-processing social graph. This paper introduces the social cloud in some detail and thus provides a good background for the next two papers. We note that, while the authors have validated their proposed model using real-world social graphs, the deployment of such social clouds in real-life scenarios is still at least a few years away from deployment. In the next paper, "An Optimized Computational Model for Task-Oriented Multi-Community-Cloud Social Collaboration", Hao et al. propose a model for task-oriented multi-community-cloud collaboration, *TOMC*, to efficiently build a trustworthy community cloud using the inherent features of social networks. The model is optimized from four aspects: minimizing the sum of access cost and monetary cost, and maximizing the security-level agreement and trust among the community clouds. In addition, they address the issue of scalability of their proposed solution. Finally, the authors propose an efficient and comprehensive selection algorithm to extract the best group of community clouds in *TOMC*4. The final paper on social clouds is that of Caton et al., entitled "A Social Compute Cloud: Allocating and Sharing Infrastructure Resources via Social Networks". In this paper, the authors propose a social compute cloud where the provisioning of cloud infrastructure occurs through "friend" relationships. Their approach is to make use of the explicit trust exhibited through friend relationships to establish trust and accountability in social cloud. This approach takes users' preferences into account and deals with them in real time. The paper focuses on four specific challenges: supporting multiple operating systems and environments (a technical challenge); sharing individual profile information through a trusted third party (a challenge related to leveraging social structure); providing incentives to participate in such a social cloud environment, where there are no obvious economic benefits; and, finally, maintaining the platform running in a dynamic environment where participants may withdraw at any

time. The paper highlights these challenges and proposes some specific solutions.

A common theme through the above three papers is that the social cloud is gaining increasing attention from the research community, as it provides a solid framework for building trusted community clouds using the explicitly exhibited trust in social networks.

As already mentioned, people increasingly employ mobile devices over which data is continuously transferred, typically to and from the cloud. This process is prone to security and privacy problems, as one can eavesdrop on the communication, and private or sensitive information can be revealed. Ardagna et al. address this problem in their paper: "An Anonymous End-to-End Communication Protocol for Mobile Cloud Environments". They leverage the social network properties, more specifically trust among friend nodes, to develop anonymous communications among all entities in the communication chain, including the cloud provider and mobile network operator.

In their paper entitled "Probabilistic Diffusion of Social Influence with Incentives", Doo and Liu also look at social networks' properties, this time to study how to best disseminate information on the social web. Whether it is for a social or political event, a marketing campaign or an emergency scenario, it is important to know how to disseminate information on the social web so that it has the desired impact. This requires identifying the influential nodes in the network and providing the right incentives for them to propagate the desired message. Different types of model have been proposed in the literature, such as Heat Diffusion Model and Epidemic Model. Heat diffusion models in particular have been widely used in recent time. However, they have limitations because they assume (1) that there is a uniform distribution for all nodes; (2) that all high degree nodes are active; and (3), that there is no need to consider incentives. Doo and Liu address these limitations in their work, introducing multi-scale rewards as incentives. Their work might also be exploited in social clouds.

Finally, Khalid et al. propose a cloud based system, *Omni-Suggest*, to compute venue recommendations for users on-the-go, taking into account real-time socio-physical factors such as one's friends network, the friend types, weather, and traffic. Recommendation systems have become an essential part of information systems in today's interconnected world, where the quantity of information is soaring every day. The information deluge makes it challenging for users to sift through the vast amount of available data and find what they want. Recommender systems have been successfully used in different application domains (e.g., social, health, entertainment, commercial, etc.) for a variety of purposes, ranging from finding the right partners on dating websites to finding good places for holidays. In recent times, venue recommendation has gained attraction due to the popularity of smartphones and the emergence of new social networks such as *Foursquare* and *Gowalla*. However, many of these recommendation systems suffer from data sparseness, cold start, and scalability problems. OmniSuggest employs ant colony algorithms, social filtering, and hub and authority scores to generate optimal venue recommendations.

The papers we present in this special issue illustrate the convergence and complementarity of the social web and cloud computing, while also demonstrating the variety of issues that arise. This represents the beginning of this convergence, though. Further research is needed to realize the vision of the social cloud, the convergence of social computing and cloud computing.

ACKNOWLEDGMENTS

REFERENCES

[1] G. Lotan, E. Graeff, M. Ananny, D. Gaffney, and I. Pearce, "The Arab Spring | the revolutions were tweeted: Information flows during the 2011 Tunisian and Egyptian revolutions," *Int. J. Commun.*, vol. 5, no. 31, pp. 1375–1405, 2011.

[2] B. G. Smith, "Socially distributing public relations: Twitter, Haiti, and interactivity in social media," *Public Relations Rev.*, vol. 36, no. 4, pp. 329–335, 2010.

[3] J. Yin, A. Lampert, M. Cameron, B. Robinson, and R. Power, "Using social media to enhance emergency situation awareness," *IEEE Intell. Syst*, vol. 27, no. 6, pp. 52–59, 2012.

[4] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," *Commun. ACM*, vol. 53, no. 4, pp. 50–58, 2010.

[5] S. Nepal, C. Paris, and A. Bouguettaya, "Trusting the social web: Issues and challenges," *World Wide Web*, pp. 1–7, 2013.

[6] W. Sherchan, S. Nepal, and C. Paris, "A survey of trust in social networks," *ACM Comput. Survey*, vol. 45, no. 4, p. article 47, 2013.

[7] J. Yao, S. Chen, S. Nepal, D. Levy, and J. Zic, "Truststore: Making amazon s3 trustworthy with services composition," *Proc. 10th IEEE/ACM Int. Conf. Cluster, Cloud Grid Comput.*, pp. 600–605, 2010.

Surya Nepal received the BE degree from the National Institute of Technology, Surat, India, ME degree from the Asian Institute of Technology, Bangkok, Thailand and the PhD degree from RMIT University, Australia. He is a principal research scientist at CSIRO Computational Informatics. His main research interest is in the development and implementation of technologies in the area of distributed systems and social networks, with a specific focus on security, privacy and trust. At CSIRO, he undertook research in the area of multimedia databases, web services and service oriented architectures, social networks, security, privacy and trust in collaborative environment and cloud systems and big data. He has more than 100 publications to his credit. Many of his works are published in top international journals and conferences such as *Very Large Data Base Endowment*, ICDE, ICWS, SCC, CoopIS, ICSOC, *International Journals of Web Services Research*, *IEEE Transactions on Service Computing*, *ACM Computing Survey* and *ACM Transaction on Internet Technology*. He leads a distributed system team at CSIRO comprising research staff and PhD students. He is a co-inventor of two patents including Trust Extension Device (TED) that provides portable and mobile trusted computing environment in an un-trusted host. He has edited two books: "*Security, Privacy and Trust in Cloud Systems*" (Springer) and "*Managing Multimedia Semantics*." He has edited many international journal special issues such as "Trusting Social Web" for Springer *WWW Journal* and "Security, Privacy and Trust in Cloud Systems" for *International Journal of Cloud Computing*. He serves as program chairs and program committee members in many international conferences and workshops. He has also delivered talks/tutorials/keynotes on trusted systems in national and international avenues.

Athman Bouguettaya received the PhD degree in computer science from the University of Colorado at Boulder in 1992. He is a professor and the head of School of Computer Science and Information Technology, RMIT University, Melbourne, Australia. He was previously a science leader in Service Computing at CSIRO ICT Centre, Canberra. Australia. Before that, he was a tenured faculty member and Program director in the Computer Science Department at Virginia Polytechnic Institute and State University (commonly known as Virginia Tech). He is a founding member and past the president of the Service Science Society, a non-profit organization that aims at forming a community of service scientists for the advancement of service science. He is or has been on the editorial boards of several journals including, the *IEEE Transactions on Services Computing, ACM Transactions on Internet Technology*, the *International Journal on Next Generation Computing, VLDB Journal, Distributed and Parallel Databases Journal*, and the *International Journal of Cooperative Information Systems.* He is also on the editorial board of the Springer-Verlag book series on Services Science. He was a guest editor of a number of special issues including the special issue of the *ACM Transactions on Internet Technology* on Semantic Web services, a special issue of the *IEEE Transactions on Services Computing* on Service Query Models, and a special issue of the *IEEE Internet Computing on Database Technology* on the web. He was a program chair of the 2012 International Conference on Web and Information System Engineering, the 2009 and 2010 Australasian Database Conference, 2008 International Conference on Service Oriented Computing (ICSOC) and the IEEE RIDE Workshop on Web Services for E-Commerce and E-Government (RIDE-WS-ECEG'04). He has published more than 180 books, book chapters, and articles in journals and conferences in the area of databases and service computing (e.g., the *IEEE Transactions on Knowledge and Data Engineering*, the *ACM Transactions on the Web, WWW Journal, VLDB Journal*, SIGMOD, ICDE, VLDB, and EDBT). He received several federally competitive grants in Australia (e.g., ARC) and the US (e.g., NSF, NIH). He is a fellow of the IEEE and a distinguished scientist of the ACM.

Cécile Paris received the bachelor's of arts degree in computer science from the University of California, Berkeley, California, and the master's and PhD degrees in artificial intelligence (natural language processing) from Columbia University, New York. She is a principal research scientist and a research team leader at the CSIRO Computational Informatics, where she leads the research on language and social computing. Her main research interests lie in the areas of social computing, language technology, trust, user modeling and human-computer interaction. She studies language and communication issues. She worked in research laboratories in both the US and the UK before joining CSIRO. Besides her appointment at CSIRO, she is an adjunct professor at ANU and Macquarie University. She has led a number of projects at CSIRO, working with a variety of organizations, including the Australian Department of Human Services, the Australian National Mental Health Commission, the New South Wales State Library, the Office of Naval Research (ONR), and Boeing. She has published widely (over 200 refereed articles), and she is very active in the research community in Australia and abroad. She is the current chair of the Australian professional organization in human computer interaction, CHISIG, and received the CHISIG medal in 2011. She is very active in the research community at large, both in Australia and internationally, serving on numerous conference and workshop committees, on reviewing boards of grant-giving bodies and journals. Together with Surya Nepal and Athman Bouguettaya, she recently edited a special issue of the Springer *WWW Journal* on "Trusting Social Web."