Rahul Chini Dwarakanath, Ralf Steinmetz: A Decentralized System for Privacy-Preserving Context Exchange: Facilitating a Better Work-Life Balance. In: Network Protocols (ICNP), 2014 IEEE 22nd International Conference on, PhD Forum, p. 489-491, October 2014. ISBN 978-1-4799-6203-7.

A Decentralized System for Privacy-Preserving Context Exchange: Facilitating a Better Work-Life Balance

Rahul Chini Dwarakanath, Ralf Steinmetz
Multimedia Communications Lab, TU Darmstadt
Darmstadt, Germany
Email: {rahul.chini.dwarakanath, ralf.steinmetz}@kom.tu-darmstadt.de

***Abstract*—In recent times, the working society has been plagued by a *work-life imbalance* as a result of the added flexibility introduced by advanced information and communication technology. A technically viable approach to improve one's work-life balance is to control the communication mediums depending on the respective user contexts, and consequently help the users maintain their concentration. To achieve this, an important prerequisite is the efficient exchange of context information among the users. Therefore, this paper motivates a novel decentralized approach which facilitates context data exchange as per prevailing conditions on privacy and confidentiality. We discuss the governing criteria surrounding user contexts and draw the main design challenges with respect to the proposed system, accordingly.**

I. Introduction

Information and communication technology (ICT) has seen radical changes over the past 10 years. The number of different communication mediums like push-email and messenger services like WhatsApp, Facebook, etc. has rapidly increased and their availability has improved drastically. This widespread influence of advanced ICT has revolutionised the working lifestyles of modern-day society, introducing flexibility in time and place into an otherwise monotonous 9-to-5 job. However, these very changes have led to the merging of one's work and private lives, creating a work-life imbalance with undesirable effects such as mood changes, stress, irritability, or even a complete burnout [1], [2].

Most current solutions to this problem involve legal intervention, where the time spent working outside stipulated working hours are accounted for monetary or other means of compensation. The car manufacturing company Volkswagen has introduced a rule that mandates an automatic shutdown of company email servers during non-working hours [3]. However, these solutions do not deal with the disturbances caused by private issues in one's work life. Furthermore, they are very scenario-specific, in that they entail fixed measures without considering each individual's needs and interests.

For a better understanding of the core issue, consider Fig. 1, which illustrates a typical scenario that can be detrimental to one's work-life balance. Mark is currently sitting in a meeting with a prospective client. As it turns out, the meeting extends beyond the scheduled end time, leaving Mark stressed since he is supposed to pick up his son, Andy, from school. At home, his wife, Sarah, expects both him and Andy to arrive any time soon. However, as time passes by, she starts getting anxious and tries contacting Mark on his phone. Out of professional courtesy, Mark does not attend to the calls but nevertheless, gets further stressed due to a series of missed calls.

/978-1-4799-6204-4/14/$31.00 © 2014 IEEE

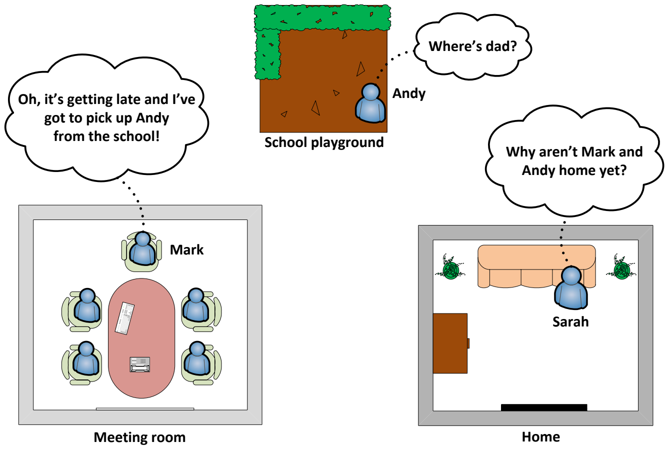
Oh, it's getting late and I've got to pick up Andy from the school!
School playground
Where's dad?
Andy
Why aren't Mark and Andy home yet?
Mark
Sarah
Meeting room
Home

Fig. 1. Illustration of a motivational scenario

From a technical perspective, a potent solution is the implementation of a context-aware system which regulates the *information flow* by identifying the users' context. This will relieve the user of certain relatively insignificant tasks and help him/her to maintain concentration. All communication mediums like phone calls, emails, SMS/WhatsApp, and push-based notifications attribute to the information flow towards a particular user.

In the above example scenario, when Sarah tries to call Mark, the current contexts of both users, Sarah and Mark, must be compared in order to determine the next course of actions with respect to the phone call: to block the call by *automatically* recognizing that Mark is currently (more) busy, and inform Sarah about the same, so that she could go and pick up Andy by herself. This should not entail any loss of attention on Mark's part since one of the main goals of this work is to maintain a person's concentration and in turn, the work-life balance.

Given the recent advancements in sensor technology, user context recognition through the processing of relevant sensor data and appropriate reasoning has been explored quite extensively over the past few years [4]. In our work, we assume that the user contexts can be determined to a significant degree and made available in the form of user-specific context models.

Therefore, in order to realize one such system, an efficient exchange of these user context models has to be ensured in accordance with user privacy and data confidentiality. An *omniscient* central server would indeed be a simple solution to support storage and exchange of the context models. However, a central entity fails to assure user privacy, is inconvenient for quick access, and is vulnerable to the 'single-point-of-attack' problem. Hence, in this paper, we motivate the implementation of a viable decentralized system that allows for a privacy-preserving exchange of the context models.

This paper is structured as follows. Section II provides a brief overview of the assumptions and main challenges with respect to the context model exchange in such a system. Section III presents the proposed approach and associated design challenges. Section IV touches upon relevant related work and finally, a list of the main research questions concludes the paper.

II. Exchange of Context Models: Prerequisites and Challenges

Context is a multidimensional concept comprising physical, temporal, and user-related aspects, as recalled by Bellavista et al. [5]. A user's context depicts the environment or the situation which the user experiences at a particular point of time. A user's context model comprises the user's current activity, the current location, the available communication mediums, the current mood and stress level, and other relevant parameters.

The exchange of context models should not interfere with the user's daily activities. Furthermore, the context information of each user has to be consistent through the entire system and its availability must also be guaranteed. This introduces certain technical challenges with respect to the variability of the user environment, privacy concerns, and computational constraints. As a result, the designed system has to adapt to the changing environmental conditions in order to function as a robust platform for context exchange.

A. Dynamic User Contexts

Everyday life is substantially dynamic and susceptible to various changes. User contexts vary depending on the user location, the people with whom the users interact, and the time of the day or week. A typical employee has high-level contexts such as attending meetings, working on a computer, relaxing with family/friends, watching TV, reading a book, pursuing sports, etc. User mobility adds to the dynamic nature of the environment, posing certain challenges with respect to the consistency and accuracy of the exchanged context models.

B. Privacy and Trust

Context exchange has to take place in adherence to the relationship between the users and based on prevailing trust and confidentiality. An adaptive in-network processing of the user context models should be performed, such that the visibility scope of the user contexts can be varied based on the relationship level and the corresponding trust. For example, a close friend may receive detailed context information, whereas an old acquaintance may only receive a partial or restricted set. The visibility scope of a user's context model has to be adjusted to suit the user's preferences.

C. Computational Constraints

Context reasoning is a computationally intensive task requiring machine learning and filtering techniques to obtain more accurate and high-level information. As a result, the varying memory and computational limitations of the sensor devices as well as energy considerations introduce further variables towards the environmental dynamics. Depending on the performance and availability of these devices, it may be more pragmatic to delegate some of the context reasoning tasks to a different processing device.

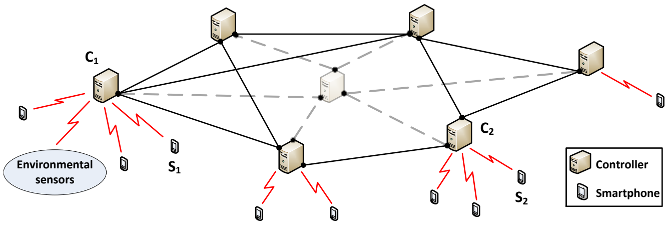

Fig. 2. The proposed system with a distributed peer-to-peer based architecture

III. Proposed Approach

In order to address the aforementioned challenges and to circumvent the problems experienced with a central entity, the system envisioned in this paper takes on a distributed peer-to-peer based approach as shown in Fig. 2. Basically, a set of so-called *controllers* assumes the functionality of a central server, where the controllers assist in the processing of sensor data for context reasoning, and the storage of the user context models. These controllers can be local servers, desktop machines, or any high-performance devices which do not have any power constraints. As the initial approach for the proposed system, we assume that the controllers form the overlay network and are exclusively responsible for context model distribution between users depending on the current smartphone-controller associations and, as repeatedly emphasized thus far, in compliance with the user- (and/or company-) specified policies for privacy and confidentiality.

For a general overview of the functionality of the proposed system, let us consider the following example using Fig. 2. The smartphones S_1 and S_2 are connected to the (trusted) controllers C_1 and C_2, respectively. The controllers (process and) store the user context models, M_1 and M_2, corresponding to the respective smartphones. If S_2 calls S_1, M_2 (or a partial set thereof) has to be made available at C_1 for comparison. In case of a reactive set-up, C_1 will search for the controller responsible for M_2 (i.e. C_2) using the overlay network's lookup functionality. In a proactive set-up, M_2 will be available at C_1 beforehand, updated at regular intervals of time. Given such a system, the following design challenges come into picture, complementary to the ones mentioned in the previous section.

Data availability. Whether the context model distribution is done in a reactive pull-based manner, relying solely on a data lookup and retrieval functionality, or a proactive push-based manner, using a publish/subscribe mechanism, is a decisive design question. For instance, the feasibility of delaying a call by a few milliseconds, so that the requisite context models can be queried, has to be compared against the overhead of having a constant exchange of context models between the controllers. The *up-to-dateness* of the context models primarily governs such a proactive set-up, whereas a reactive set-up must be highly responsive so that the requisite context model is available in tenable time. The overlay network should also support a range-search functionality which will allow a controller to search through multiple context models which satisfy one or more criteria.

Device associativity. The association of a particular smartphone with a controller primarily depends on the location of the smartphone and the controller, and the trust factor between the two. For example, a user would completely trust a controller located at the workplace; however, an unknown controller located at a client's main office would probably not be trusted to the same extent. In this work, we assume that the requisite sensor data for context reasoning are available from environmental sensors as well as sensors present in the smartphones. The processing of the sensor data and storage of the generated user context models may then take place on the smartphones themselves or can be delegated to a controller, depending on the battery power levels of the smartphones (as well as the sensors), the availability of controllers, and the prevailing trust levels.

Controller accessibility. The positioning of the controllers as well as their discovery are crucial design challenges for the proposed system. The controllers can also be envisaged as so-called *cloudlets* [6], which are spawned as per user requirements. In such a case, a central controller could be considered which would facilitate bootstrapping, coordinate the spawning of these cloudlets as well as store global user context information (e.g. coarse user location) for further assistance in spawning. Of course, a key challenge is it to achieve the same in a decentralized manner.

IV. Related Work

To the best of our knowledge, one such system has not been developed so far. Indeed, as pointed out by Bellavista et al. [5], a considerable number of decentralized architectures have been proposed in the literature for context data distribution. However, our system particularly differs from these, in that it addresses an adaptive exchange of user contexts based on prevailing conditions on privacy and confidentiality. Furthermore, the existing context distribution architectures do not adequately address the other design challenges discussed above.

With respect to privacy, the field of online social networks has relevant research work closely related to our work. Through *Safebook* [7], the authors present a decentralized peer-to-peer based online social networking system which employs the existing trust between users. Concentric rings of nodes called *matryoshkas* are used to enable trust-based lookup and retrieval of data. *DECENT* [8] differs from *Safebook* and other pure peer-to-peer based architectures, where it combines cryptographic data protection with flexible attribute policies for confidentiality, integrity, and availability. Our system resembles the one proposed by Shakimov et al. [9], where the users store their personal data in local machines called the *Virtual Individual Servers*, which effectively self-organize to form peer-to-peer overlay networks.

However, none of the above systems account for high responsiveness in data retrieval, or the partial or range-specific search for data. Furthermore, these works are all based on the participatory nature of user interaction, whereas our system works in an opportunistic manner, which introduces additional design challenges, as discussed in the previous sections. Previous efforts concerning fast lookup and range-search in overlay networks such as *HPM* [10] and *LIGHT* [11], as well as works in the field of the Semantic Web like *3rdf* [12] provide some preliminary groundwork for our system.

V. Main Research Questions

The prime objective of this work is to devise a feasible decentralized system that supports a privacy-preserving exchange of context data. The main research questions, which must be answered henceforth, are:

1) What kind of approach is most viable for content data exchange: a pull-based lookup functionality, a push-based publish/subscribe mechanism, or a hybrid approach?
2) What kind of indexing schemes should be employed in the overlay network so as to support an adaptive privacy- and confidentiality-preserving exchange of context data between the controllers, without adversely affecting the system performance?
3) How should the controllers be positioned and how can the associativity between the smartphones and the controllers be resolved? How can this be achieved in a decentralized manner?

Acknowledgment

This work is co-funded by the Social Link Project within the Loewe Program of Excellence in Research, Hessen, Germany.

References

[1] W. B. Schaufeli, M. P. Leiter, and C. Maslach, "Burnout: 35 Years of Research and Practice," *Career Development International*, vol. 14, no. 3, pp. 204–220, 2009.

[2] G. M. Alarcon, "A Meta-Analysis of Burnout with Job Demands, Resources, and Attitudes," *Journal of Vocational Behavior*, vol. 79, no. 2, pp. 549 – 562, 2011.

[3] "Volkswagen turns off Blackberry email after work hours," www.bbc.com/news/technology-16314901, 2011, [Online; last accessed 05-July-2014].

[4] C. Bettini *et al.*, "A Survey of Context Modelling and Reasoning Techniques," *Pervasive and Mobile Computing*, vol. 6, no. 2, pp. 161 – 180, April 2010.

[5] P. Bellavista *et al.*, "A Survey of Context Data Distribution for Mobile Ubiquitous Systems," *ACM Comput. Surv.*, vol. 44, no. 4, pp. 1–45, Sep. 2012.

[6] T. Verbelen *et al.*, "Cloudlets: Bringing the Cloud to the Mobile User," in *Proceedings of the third ACM workshop on Mobile cloud computing and services*, June 2012, pp. 29–36.

[7] L. A. Cutillo, R. Molva, and T. Strufe, "Safebook : A Privacy Preserving Online Social Network Leveraging on Real-Life Trust," *IEEE Communications Magazine Consumer Communications and Networking Series*, pp. 94–101, December 2009.

[8] S. Jahid *et al.*, "DECENT: A Decentralized Architecture for Enforcing Privacy in Online Social Networks," in *IEEE International Conference on Pervasive Computing and Communications Workshops*, March 2012, pp. 326–332.

[9] A. Shakimov *et al.*, "Privacy, Cost, and Availability Tradeoffs in Decentralized OSNs," in *Proceedings of the 2nd ACM workshop on Online social networks*, August 2009, pp. 13–18.

[10] M. Amad *et al.*, "HPM: A Novel Hierarchical Peer-to-Peer Model for Lookup Acceleration with Provision of Physical Proximity," *Journal of Network and Computer Applications*, vol. 35, no. 6, pp. 1818 – 1830, 2012.

[11] Y. Tang, S. Zhou, and J. Xu, "LIGHT: A Query-Efficient yet Low-Maintenance Indexing Scheme over DHTs," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 1, pp. 59–75, Jan 2010.

[12] L. Ali, T. Janson, and G. Lausen, "3rdf: Storing and Querying RDF Data on Top of the 3nuts Overlay Network," in *22nd International Workshop on Database and Expert Systems Applications*, Aug 2011, pp. 257–261.