Marc-Olivier Killijian · Nicolas Rivière · Matthieu Roy

Experimental Evaluation of Resilience for Ubiquitous Mobile Systems

Received: date / Accepted: date

Abstract In this position paper, we address the needs and motivations for practical evaluation of ubiquitous mobile systems. We are particularly interested in the validation of resilience mechanisms for such systems. We advocate the use of real experiments that complement simulation for the evaluation of actual prototypes. Indeed, experimental validation is particularly appealing when resilience properties such as safety or availability are an issue. We then propose and discuss some technological trails for the implementation of such an evaluation platform and its validation.

Keywords Experimental evaluation · Validation · Mobile systems · Ubiquitous computing · Resilience

1 Introduction

Weiser's vision of ubiquitous computing [Wei99] is that of a society where technology has merged with the physical environment, and where systems have become ubiquitous, offering services to users while remaining invisible. This stealthiness implies two important aspects. First, it assumes excellent interfaces and a rigorous design. Second, it requires a level of trustworthiness such that users can forget the very presence of the technology, under any conditions and threats. To allow ubiquitous systems to achieve that level of trustworthiness, resilience mechanisms aimed at improving their security, reliability and availability must be devised.

Mobility has a significant impact on resilience of ubiquitous computing. Mobility of the physical devices implies open and highly dynamic environments, stressing usage conditions with multiple interactions and multiple sources of failures. Thus, new risks stem from mobility, along with new threats which can be both accidental and malicious. Indeed, risk of loss, theft or physical damage is exacerbated by mobility, but risks of energy (and other resources) exhaustion or dangerous encounters (such as malicious service discovery and usage) are new to mobile ubiquitous systems.

Nevertheless, positive impacts of mobility and ubiquity ought to be taken into account as well. For instance, heterogeneity of the services available in a ubiquitous system can be regarded as a source of diversity that resilience mechanisms can benefit from. Using specific fault-tolerance mechanisms such as triple modular replication with voting, a system can benefit from diversity to tolerate common mode failures. Without diversity, i.e. with equivalent replicas of a service, a single fault might lead the various replicas to fail in the same manner, in a common mode. Taking diversity into account, and using an appropriate mechanism, this risk is reduced.

Our main goal is the study of both positive and negative impacts of mobility and ubiquity on system resilience. To this respect, we believe that it is essential to practically validate and experimentally evaluate the proposed solutions for increasing the resilience of mobile ubiquitous systems.

2 Evaluation of Resilient Ubiquitous Mobile Systems

To the best of our knowledge, little research has been done on the evaluation of ubiquitous systems. Most of the literature in this domain concerns evaluation of users experience and human-computer interfaces [SK01]. However, some work is also looking at defining appropriate metrics for the evaluation of distributed applications running on ubiquitous systems [BKL01, CCKM01].

[BR01] is looking at a general approach to evaluating ubiquitous systems. In the paper, the authors ar-

---

This work was partially supported by the Hidenets project (EU-IST-FP6-26979), and the ReSIST network of excellence (EU-IST-FP6-26764).

M.-O. Killijian, N. Rivière, M. Roy
LAAS-CNRS
Université de Toulouse
7, avenue du Colonel Roche
31077 Toulouse Cedex 4, France
E-mail: marco.killijian@laas.fr

gue that quantitative measurements should be complemented with qualitative evaluation. The argument is that there is a number of problems for which evaluation cannot be easily quantified. Thus an evaluation should be conducted using an hybrid quantitative/qualitative strategy.

The area of resilient computing proposes a number of contributions concerning the evaluation of distributed systems. Analytical evaluation is probably the most popular technique, such as within the Assert project [ABC$^+$05] in the avionics application domain. More recently, experimental evaluation started to gain attention. The approach taken is often based on dependability benchmarking, for example the DBench IST project [KMM$^+$05] addresses dependability benchmarking of operating systems.

However, in the ubiquitous and mobile computing area, evaluation of resilience mechanisms, protocols, or prototypes remains an open problem. In most cases, the proposed algorithms are evaluated and validated using wireless network simulators such as NS-2 [NS2], Glomosim [GLO] or JiST/SWANS [BHvR05]. Since simulators use a model of physical components, such as network cards and location systems, this raises concerns as to the representativity of the assumptions that underlie the simulation [CSS02].

Little work concerning the evaluation of algorithms in a realistic environment is available. Of course, this is due to practical reasons: how can one realistically set up a platform allowing the execution of algorithms dealing with mobility and ubiquity?

In this paper, we elaborate on the development of a realistic platform, at the laboratory scale, to evaluate and validate fault-tolerance algorithms (e.g., group membership and replication protocols, backup mechanisms, etc.) targeting systems comprising a large number of communicating mobile devices equipped with various sensors and actuators. Our goal is to have an experimentation platform allowing for reproducible experiments (including mobility aspects) that will complement validation through simulation. As we will see, an important issue within this platform is related to changes of scale so as to emulate as many various systems as possible.

3 Building an Experimental Evaluation Platform

3.1 Target Applications and Systems

The mobile ubiquitous systems we seek to evaluate fall into several categories: automotive applications such as distributed black-boxes, platooning, etc. [BCF$^+$07]; cooperative services with both Internet-connected and disconnected periods [KPB$^+$04]; smart-spaces; e-health nano-devices. These very diverse systems involve different notions of mobility and represent different scales. In the remainder of this section we categorize more precisely the type of mobile systems we want to evaluate and elaborate on their scale.

3.2 What does Mobility Mean?

We think that when discussing about mobile systems, we should distinguish between the *goal* and the *reason* of mobility. As far as the goal of mobility is concerned, two notions are distinct: *physical mobility*, that deals with the movement of devices within the physical space, and *logical mobility* that relates to service roaming within the resource space, i.e. mobile services or software agents.

Issues related to the reason of mobility differ significantly depending on whether mobility arises as a consequence of user mobility (in the remainder, this will be referred to as *passive mobility*), as is the case with devices such as mobile phones, or whether it is controlled by the device or service itself (*active mobility*), as with autonomous robots for example. With this experimental platform, we are interested in the evaluation of the resilience of both mobile services and devices in the context of ubiquitous systems whose mobility can be either active or passive.

3.3 A Matter of Scales

The experimental evaluation platform is composed of both fixed and mobile devices. Technically speaking, the mobile devices are constituted by some programmable mobile hardware able to carry the device itself, a lightweight processing unit equipped with one or several wireless network interfaces, radio-frequency identifiers (RFIDs), a positioning device and a smart-card reader. The fixed counterpart of the platform contains the corresponding fixed infrastructure: an indoor positioning system, wireless communication support, an identification system for RFIDs (RFID readers), as well as some fixed servers.

We consider a platform set up in a room of approximately 100 square meters where mobile devices can move around. The room provides services, some of which are available in the whole room while others are limited to specific regions. Mobile devices can also offer services within a limited range and can access the services available in their environment.

By changing scales, we plan to emulate systems of different sizes. Hardware modeling of this type of system requires a reduction or increase of scale to be able to conduct experiments within the laboratory. To obtain a realistic environment, all devices must be modified according to the same scale factor. As mentioned above, we plan to conduct very diverse experiments, for example:

- A system of communicating vehicles (VANET) with intermittent connection to roadside infrastructure (e.g., there could be an access-point at each toll-gate on a motor-way).

- A system of nano-robots subject to mobility (as they are carried by the blood stream, for instance) offering services (e.g., delivery of a certain amount of medicine in a specific region of the human body), while interacting with their environment and communicating with each other (using chemicals [CHSL06], wired or wireless networks[Fre99]).

These two experiments are good examples of the scale diversity the experimental platform should cope with: the VANET application needs to be scaled down while the e-health experiment has to be scaled up. For instance, if we consider the VANET experiment within a 16x6 meters room (a typical 100 square-meter lab experimental room), we can estimate the scale factor to be around 50 (for a 800x300 meter simulated environment). Conversely, the nano-robot experiment would typically involve ⌀1 µm nodes able to communicate with each other in a range of tens µm and thus an increase factor in the order of 10^6 (for a simulated environment composed of few ⌀20 µm vessels). In the remainder of this article, we discuss the technical implications of this scale challenge.

3.3.1 Precise Indoor Location

A first technical issue concerns indoor location. Indeed, most of the applications and ubiquitous systems we plan to experiment on this platform benefit from some kind of geo-positioning, usually using a GPS device. As the platform will be built indoor, within our laboratory, GPS devices are not able to receive the GPS satellites signal and thus cannot be used. An indoor location solution is thus necessary. Again the question of scaling can be translated into an precision problem: how precise do we need the indoor location system to be? For example, if we come back to the VANET experiment, a typical GPS in a moving car is 5-20m precise. Then, indoor we need to have 10-40cm precision.

Several technologies are currently available for indoor location [HB01], mostly based either on scene analysis (e.g. using motion capture systems) or on triangularization (of RF and ultrasound [SBGP04] or wireless communication interfaces [CKSW03]). For us, the main metrics are obviously precision, accuracy, but also size and weight of the devices.

3.3.2 Realizing Physical Mobility

One of the most important question we have to answer when designing this platform is how to realize mobility, i.e. how to make the devices actually mobile. Obviously, when conducting experiments, we cannot have a human operator behind each device, mobility has to be automated. This is why we are considering the use of simple small robots platforms in order to carry around the platform devices. These robots task is to “implement” the mobility of the nodes. If, for a given experiment, the emulated nodes are actively mobile, the devices need to have control over the mobility and thus needs to communicate with the robot controller. Conversely, if the nodes are passively mobile, there is no need for any communication between the robots and the devices they carry and the robot controller can be much simpler as it doesn't need any communication means.

We are currently evaluating several robot platforms. The questions we have to answer during this selection are:

- Load: is the robot able to carry the device which is typically composed of a PDA or a laptop, a few actuators/sensors, a few RFIDs, etc.
- Accuracy: as we want the experiments to be reproducible, how accurate are the movements of the robot under the specified load?
- Speed: under the load and the accuracy specified above, what is the maximum speed the robot can attain?
- Energy: with fully-charged batteries, how long can last an experiment or how many experiments can be conducted?
- Scope: is the robot controller communication-enabled (in order to implement active mobility) ?
- Ease-of-use: how easy is the robot platform to program for a given experiment?
- Size: how big/heavy is the robot? This is important as it can physically limit the number of devices participating in an experiment.

3.3.3 How to Scale Down Wireless Communication?

The most important design issue for the platform concerns wireless communications. First, the communication range of the participants (mobile nodes and infrastructure access-points) has to be scaled according to the experiment being conducted. For example, with the VANET experiment, a typical automobile have a wireless communication range of a few hundred meters, say 200m[1]. With a scale reduction factor fixed at 50, the mobile devices communication range has to be limited to 4m. However, to cope with other experiments and other scale reduction factors, this communication range should ideally be variable. A second problem concerns communication obstacles. For certain experiments, it might be necessary to simulate obstacles (like tunnels, mountains, buildings, etc.) and their impact on wireless communications. For example, when entering a tunnel, a car can suddenly be disconnected from its neighbors but also become connected to another network partition. The platform must simulate those communication obstacles. To summarize, there are two issues to be solved: communication range reduction and obstacle simulation.

---

[1] The 802.11p Wireless Access in the Vehicular Environment (WAVE) standard is still under discussion but it will probably define a communication range of several hundred meters. We consider here one of the most restrictive options.

Concerning communication range reduction, some WiFi network interface drivers implement an API for reducing their transmission power. However, the implementation of this feature is often limited and many interface drivers only propose to choose transmission power between a few pre-selected values. Theoretically, a limited transmission power should accordingly reduce transmission range. Yet, this assumption has to be confirmed practically. This approach could also be applied using short range wireless technologies, such as ZigBee, Wibree[2] or Bluetooth.

An alternative solution, based on topology emulators [JN07], could also be envisaged. An emulator is connected through a switch to nodes (one access point per mobile device). This emulator simulates nodes connectivity according to the topology. Notice that this hybrid solution, between simulation and a real experiment, is not really adequate in the plain experimental platform we are developing.

The last and most appealing solution is the one we are currently investigating. It consists of using signal attenuators between the WiFi network interfaces and their antennas. An attenuator is an electronic device that reduces the amplitude or power of a signal without appreciably distorting its waveform. Attenuators are usually passive devices made from resistors. The degree of attenuation may be fixed, continuously adjustable, or incrementally adjustable. In our case, the attenuators are used to reduce the signal received by the network interface and are actually much more efficient than just disconnecting the antenna. Then, we plan to use spectrum analyzers in order to characterize the signal reduction with respect to the degree of attenuation.

The second technical problem can be solved physically using Faraday's cages (or tents[3]). An alternative solution based on topology emulation could also be envisaged but again we would prefer not to rely on this type of technical solution. In any case, it will be important to take into account room reflection, as it could affect the awaited results.

Therefore, these different technical solutions need to be evaluated according to several metrics:

- Range accuracy: how accurate is the actual transmission range compared to the specified limit?
- Setting accuracy: how accurate is the lab setup compared to a real experiment? Indeed, lab walls, floor and ceiling will have an impact on the wireless communication, so these effects needs to be evaluated as well.
- Obstacles accuracy: how good are the simulated (physical or logical) communication obstacles?

3.3.4 Validation

Validation of the platform is an important issue to gain confidence in the platform and its results. We will thus need to evaluate both the representativity, validity and reproducibility of the experiments. Validity compares the results obtained with the platform against "reality" and reproducibility measures the stability of the results in a sequence of repeated experiments. The approach we plan to use is inspired by benchmarking. A benchmark application for which data is available (on Crawdad[**?**] for example) is experimented on the platform. The results can thus be compared with the real data to evaluate their consistency. Concerning reproducibility, one must remember that the applications we are interested in involve many nodes. Henceforth concurrency should probably introduce some variability in the obtained results during a series of experiments. However, the benchmarking technique mentioned above could be repeated in order to measure the variability of the results and compare it with results obtained from simulators and/or real experiments.

4 Conclusion and Future Work

We believe that experimental evaluation of resilient ubiquitous systems is necessary to complement the usual analysis based on simulation. In this article we elaborate on the technical implementation details for an experimental platform for mobile ubiquitous systems. With such a platform we plan to evaluate mobile systems coming from diverse application areas: automotive, health, smart spaces, etc. This variety introduces a big scaling challenge. We discuss three different scaling problems: indoor location, actual mobility and wireless communication. We present a few solutions to these challenging tasks, and propose some evaluation metrics. We have also to design a validation technique for this platform. To conclude, although the problem of indoor location seems easier to solve as much work is available on the subject, the two last problems are still open. We are currently evaluating several technical solutions to these challenges.

References

[ABC$^+$05] J. Arlat, M. R. Barone, Y. Crouzet, J-C. Fabre, M. Kaaniche, K. Kanoun, S. Mazzini, M. R. Nazzarelli, D. Powell, M. Roy, A. E. Rugina, and H. Waeselynck. Dependability needs and preliminary solutions concerning evaluation, testing and wrapping. Technical Report 05424, LAAS, Toulouse, 2005.

[BCF$^+$07] I. De Bruin, A. Casimiro, L. Falai, P. Frejek, S.H. De Groot, G. Huszerl, M. Kaaniche, M.-O. Killijian, R. Lippman, E.V. Matthiesen, M. Roy, H. P. Schwefel, and I. E. Scinnset. Specification hidenets laboratory set-up scenario and components. Technical Report 07187, LAAS, Toulouse, 2007.

---

[2] `www.wibree.com`

[3] see, e.g., `www.faradaycages.com`

[BHvR05] R. Barr, Z. J. Haas, and R. van Renesse. Scalable wireless ad hoc network simulation. In *Handbook on Theoretical and Algorithmic Aspects of Sensor, Ad hoc Wireless, and Peer-to-Peer Networks*, chapter 19, pages 297–311. CRC Press, 2005.

[BKL01] Prithwish Basu, Wang Ke, and Thomas D. C. Little. Metrics for performance evaluation of distributed application execution in ubiquitous computing environments. *Workshop on Evaluation Methodologies for Ubiquitous Computing at Ubicomp'01*, 2001.

[BR01] Mark Burnett and Chris P. Rainsford. A hybrid evaluation approach for ubiquitous computing environments. *Workshop on Evaluation Methodologies for Ubiquitous Computing at Ubicomp'01*, 2001.

[CCKM01] Paul Castro, Alvin Chen, Ted Kremenek, and Richard Muntz. Evaluating distibuted query processing systems for ubiquitous computing. *Workshop on Evaluation Methodologies for Ubiquitous Computing at Ubicomp'01*, 2001.

[CHSL06] Adriano Cavalcanti, Tad Hogg, Bijan Shirinzadeh, and Hwee C. Liaw. Nanorobot communication techniques: A comprehensive tutorial. In *IEEE ICARCV 2006 International Conference on Control, Automation, Robotics and Vision*, page 6, 2006.

[CKSW03] N. S. Correal, S. Kyperountas, Q. Shi, and M. Welborn. An uwb relative location system. In *Proc. of IEEE Conference on Ultra Wideband Systems and Technologies*, November 2003.

[CSS02] David Cavin, Yoav Sasson, and André; Schiper. On the accuracy of manet simulators. In *POMC '02: Proceedings of the second ACM international workshop on Principles of mobile computing*, pages 38–43, New York, NY, USA, 2002. ACM Press.

[Fre99] Robert A. Freitas. *Nanomedicine, Volume 1: Basic Capabilities.* Landes Bioscience, 1999.

[GLO] Global mobile information systems simulation library
`http://pcl.cs.ucla.edu/projects/glomosim/`.

[HB01] J. Hightower and G. Borriello. A survey and taxonomy of location systems for ubiquitous computing, 2001.

[JN07] Morten N. Jensen and Anders Nickelsen. Evaluation of routing dependability in manets using a topology emulator. Master's thesis, Aalborg University, june 2007.

[KMM$^+$05] K. Kanoun, H. Madeira, F. Moreira, M. Cin, and J. Garcia. Dbench - dependability benchmarking. In *Proc. of the Lecture Notes in Computer Science (LNCS), Springer-Verlag, Fifth European Dependable Computing Conference (EDCC-5)*, April 2005.

[KPB$^+$04] Marc-Olivier Killijian, David Powell, Michel Banâtre, Paul Couderc, and Yves Roudier. Collaborative backup for dependable mobile applications. In *Proceedings of 2nd International Workshop on Middleware for Pervasive and Ad-Hoc Computing (Middleware 2004)*, pages 146–149, Toronto, Ontario, Canada, October 2004. ACM.

[NS2] The network simulator
`http://www.isi.edu/nsnam/ns/`.

[SBGP04] Adam Smith, Hari Balakrishnan, Michel Goraczko, and Nissanka Bodhi Priyantha. Tracking Moving Devices with the Cricket Location System. In *2nd International Conference on Mobile Systems, Applications and Services (Mobisys 2004)*, Boston, MA, June 2004.

[SK01] Mirjana Spasojevic and Tim Kindberg. Evaluating the cooltown user experience. *Workshop on Evaluation Methodologies for Ubiquitous Computing at Ubicomp'01*, 2001.

[Wei99] Mark Weiser. The computer for the 21st century. *SIGMOBILE Mob. Comput. Commun. Rev.*, 3(3):3–11, 1999.