

FADA: Formalisms and Algorithms for Resilient Services Design in Ambient Systems*

Matthieu Roy
LAAS-CNRS
Université de Toulouse
Toulouse, France
roy@laas.fr

François Bonnet
IRISA
Campus de Beaulieu
Rennes, France
fbonnet@irisa.fr

Marc-Olivier Killijian
LAAS-CNRS
Université de Toulouse
Toulouse, France
mkilliji@laas.fr

Sara Tucci Piergiovanni
University of Roma
La Sapienza
Roma, Italy
sara.tucci@dis.uniroma1.it

Leonardo Querzoni
University of Roma
La Sapienza
Roma, Italy
querzoni@dis.uniroma.it

Silvia Bonomi
University of Roma
La Sapienza
Roma, Italy
silvia.bonomi@dis.uniroma1.it

1. INTRODUCTION AND OBJECTIVES

The advent of massively distributed systems in which every user carries an entity with full computing power, communication, storage and positioning capabilities, allows us to envision many new applications and services, truly decentralized by nature and tightly coupled to the position of entities. Nevertheless, the deployment of such systems imposes the definition of formalisms that capture the new features of these systems and allow algorithms, mechanisms and architectures to be developed and reasoned about.

Indeed, for the first time ever, there exists a strong coupling between the graph of possible interactions between entities and the geographical distribution of nodes. The strength of the classical Internet model was to abstract all communication links into a complete graph, where any two nodes willing to cooperate could open a connection. This clique-based system modeling can no longer be applied to distributed dynamic systems where entities can communicate using short-range wireless technologies.

As pointed out by the ReSIST NoE research agenda¹, two main issues arise when trying to deal with such large-scale and ambient systems:

- The evolvable nature of mobile systems imposes any mechanism built for them to be resilient to mobility- and failure-induced changes to the composition and/or topology of the system.
- New features of such systems are not represented in traditional distributed models, namely their dynamicity and locality, and more generally the geographical distribution of entities.

The overall objective of FADA is to propose formalisms, algorithms and architectural patterns towards solving those issues.

*This work is funded by the ReSIST Network of Excellence (FP6 IST contract 026764)

¹ReSIST (Resilience for Survivability in IST) website:
www.resist-noe.org

2. SCIENTIFIC APPROACH

In this work, we address the problem of building resilient services on a distributed, mobile system in a cooperative way. In order to have a complete and sound solution, our approach focuses on the following aspects of the problem:

- Formalization of the system into a model that goes beyond traditional distributed systems models, by including mobility and geographical information, and by adapting failure models,
- Architectural design definition of the basic building blocks (i.e., abstractions) for cooperative services implementation, based on the above model,
- Development of algorithms to implement cooperative resilient services, despite limited knowledge, arbitrary size, mobility patterns and an unbounded number of failures,
- Assessment of the proposed algorithms using simulators, and then on experimental platforms [4].

In this work, we follow a horizontal and a vertical approach. Horizontally, we started to design an architectural landscape for cooperative services on top of a system composed of mobile nodes, by identifying basic building blocks, providing relative abstractions and implementation algorithms. As presented below, we focused on the traditional abstractions of distributed systems, namely coordination, synchronization, agreement and reliable communication. Vertically we will focus on two main applications:

- a reliable, geographically *localized blackboard* system, modeled as a reliable storage service,
- a *traffic jam avoidance* application, where the participants locally evaluate congestion of roads. Concerning the distribution of events associated with this application, we are investigating whether events dissemination should be implemented using a push or a pull scheme.

3. RESEARCH GOALS

To address the above problems, the first step consists in identifying a suitable formalization of the system that takes mobility into account. In this work, we are interested in *passive* mobility: entities move independently of the problem to be solved. This mobility model represents a system where users carry a device (phone, computer); the path followed by a device is bound to the user’s movements and the algorithms implemented on the device must provide the service despite various movement patterns.

From the theoretical point of view, the aim of this work is to show that using a formal model of the system allows the definition of conditions on mobility patterns and node density to be able to implement basic resilient services that can be formally proven correct.

From the practical point of view, we are providing an architecture for resilient services in mobile systems. Moreover, the basic building blocks provided within this architecture will be evaluated both on simulators and on reduced-size real systems, thus improving confidence in the pertinence of the formal system models and algorithms.

4. ARCHITECTURAL DESCRIPTION

Reasoning on constantly evolving systems imposes the need for suitable abstractions that capture the very nature of interactions between entities. Following the seminal work of Chockler *et al.* [3], we architected our model using three layers, as depicted in Figure 1 :

- The lower layer is composed of abstractions that are close to the hardware. [3] identified three abstractions for mobile systems:
 - A *timed local broadcast* service that sends messages in at most δ time units but can lose messages. This service refers to hardware properties of the communication medium.
 - A *collision detector*, i.e. an oracle which definition is close to classical failure detectors [2]. It encapsulates the additional formal assumptions needed to provide reliable algorithms such as consensus or reliable broadcasts.
 - A *wake-up service* that determines which node can send message to prevent collisions between messages.
- Additionally, since we are interested in the interactions of users with the physical world, we propose a fourth abstraction, the *Localization and Clock service*, that provides information available from a GPS-like device, i.e. localization information, and a global clock service.
- The intermediate level provides higher-level services that integrate physical (localization-based) and logical (network-based) information:
 - The *proximity map* exports a hybrid map of an entity’s neighborhood that indicates the position of nodes that are within communication range.
 - The *node failure/leave/join detection* service is meant to detect changes in the configuration of the distributed system.

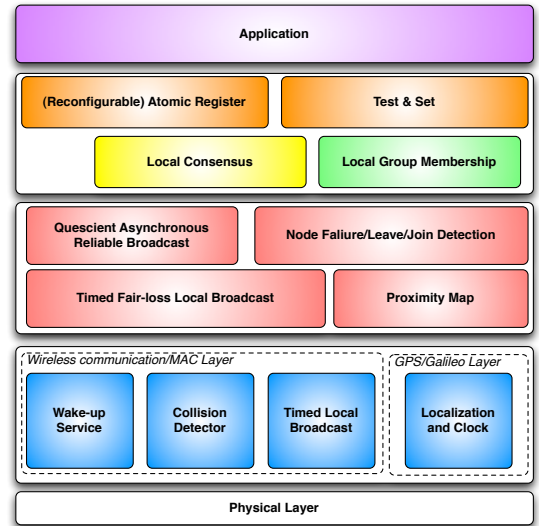


Figure 1: FADA architectural description

- The *time fair-loss* local broadcast is built using timed local broadcast and wake-up service: each message suffers from a maximal delay δ , but messages are not systematically lost.
- The *quiescent asynchronous reliable broadcast* service[1] is a broadcast service that ensures that messages are not sent forever, i.e. that retransmissions occur only a finite number of times.
- At the higher level, one can find simple reliable abstractions, such as geo-localized atomic registers, test&set operations, or localized consensus.

Current and further work includes formalizing the assumptions of the above building blocks, specifying and developing the middleware associated to them. Evaluation of the algorithmic solutions will be based on our proof-of-concept applications, namely a localized blackboard system and a traffic jam avoidance system.

5. REFERENCES

- [1] F. Bonnet, P. Ezhilchelvan, and E. Vollset. Quiescent consensus in mobile ad-hoc networks using eventually storage-free broadcasts. In *Dependable and adaptive distributed systems (DADS)*, pages 670–674, 2006.
- [2] T. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, 43(2):225–267, 1996.
- [3] G. Chockler, M. Demirbas, S. Gilbert, C. Newport, and T. Nolte. Consensus and collision detectors in wireless ad hoc networks. In *PODC 2005*, Las Vegas, Nevada, USA, 2005.
- [4] M.-O. Killijian, N. Rivi re, and M. Roy. Experimental evaluation of resilience for ubiquitous mobile systems. In *Proc. of UbiComp 2007 Workshops, Workshop on Ubiquitous Systems Evaluation (USE)*, pages 283–287, Innsbruck, Austria, 2007.