

Scientific Report of the 4th MiNEMA Workshop

July 2-3, 2006. Sintra, Portugal

Boris Koldehofe Luís Rodrigues
University of Stuttgart University of Lisbon

Hugo Miranda
University of Lisbon

July 20, 2007

Abstract

The MiNEMA Workshops intend to bring together a mixture of young and senior European researchers working on middleware for network centric and mobile applications. They help to foster further collaboration between existing MiNEMA members, with special emphasis on PhD students, to advertise and widen participation in the MiNEMA network and to establish links with the industry. This document reports the organisation of the the 4th MiNEMA workshop by the University of Lisbon.

1 Summary

The 4th MiNEMA workshop was held in Sintra, Portugal, from the afternoon of July 1st to the the afternoon of July 3rd, 2006. Accommodations and meeting rooms have been arranged at Hotel Tivoli Sintra.¹ The workshop had four half day working sessions. Each morning session was opened by an invited talk from a relevant R&D company operating in Portugal. The detailed programme of the workshop is presented in Appendix A.

The scientific organisation of the workshop was separated from the local organisation. The local arrangements were prepared by Hugo Miranda and Luís Rodrigues from the University of Lisbon. Boris Koldehofe, from EPFL, was responsible for the scientific program.

The workshop had 41 participants from 14 countries and from 19 distinct institutions. It is interesting to note that two participants came from *SAP Research*, *CEC Karlsruhe*, a research unit of a private company; this reflects the fact that the network has gathered interest outside academia.

2 Scientific Content

The workshop followed the same model of the 3rd MiNEMA workshop in Leuven, which is briefly outlined here. The scientific program was formed by research papers which are presented and discussed during the workshop. Submission of papers presenting early results, namely from PhD students, was encour-

¹<http://www.tivolisintra.com>

aged in order to foster discussion and provide feedback to the authors. Each paper, in order to be accepted, was reviewed by the workshop program committee.

Categories for Contributions Papers could be submitted to one of the following categories:

Research Research papers presenting new, still preliminary research on Middleware for mobile environment. The recommended length of the submission was between 3 and 5 pages.

Study Topic A study topic identifying potentially open questions in the domain, and setting out a plan for how to answer those questions during the workshop. These submissions could not exceed 3 pages.

Lessons Learned A summary of complete research tracks and/or projects. This work could be submitted by sending a position paper, optionally complemented with a report or reference to an existing publication describing the work. The recommended length of the submission was between 3 and 5 pages.

MiNEMA PhD workshop report Presentations reporting on the results of workshops among PhD students sponsored by the MiNEMA programme.

Presentation and Advertisement For each accepted submission, one author was invited for giving a 30m presentation at the workshop.

The authors of accepted submissions were requested to revise their contribution by considering the feedback of reviewers. The proceedings of the workshop, composed of all submissions accepted, were published as a technical report of the Departamento de Informática - Universidade de Lisboa [1]. The abstracts of the papers accepted are presented in Appendix B.

2.1 Programme Committee

The program committee was responsible for the evaluation of submitted contributions. It was chaired by Boris Koldehofe and composed by:

- Filipe Araújo, University of Coimbra, Portugal
- Paolo Costa, Politecnico di Milano, Italy
- Paul Grace, University of Lancaster, UK
- Sidath Handurukande, NEC Research Labs, USA
- Boris Koldehofe, EPFL, Switzerland (Chair)
- Sam Michiels, KU Leuven, Belgium
- Sasu Tarkoma, University of Helsinki, Finland
- Stefan Weber, Trinity College, Ireland

Besides the members of the programme committee, the list of reviewers also included:

- Arnas Kupsys, Eurospot, Switzerland
- Iqbal Mohomed, University of Toronto, Canada

2.2 Analysis of the Submissions

This workshop received a total of 18 submissions, of which 8 were research papers, 9 study topic papers and 1 lessons learned papers. The program committee decided to accept 15 of the 18 submissions. One paper was not presented because no author attended the workshop.

3 Conclusions

The workshop had a significant number of submissions and participants from many different institutions. Furthermore, presentations raised fruitful discussion that are likely to contribute in the future for joint projects between different members. Therefore, we can state that the workshop has achieved its goals entirely.

Appendixes

A Final Programme

Saturday, 1st July

19:00 Welcome reception

Sunday, 2nd July

09:00-10:15 Opening and Keynote
Orlando Remdios (Siemens Portugal)
Coffee Break

10:45-12:15 Session 1
A Service-Oriented Middleware for Integration and Management of Heterogeneous Smart Items Environments. Jürgen Anke (SAP Research, CEC Dresden / Dresden University of Technology), Jens Müller, Patrik Spieß and Leonardo Weiss Ferreira Chaves (SAP Research, CEC Karlsruhe)
Service Discovery in Highly Heterogeneous Mobile Environments. Carlos Flores, Gordon Blair and Paul Grace (Lancaster University)
Deploying Services over Ad Hoc Networks through Context-based Migration. Oriana Riva (University of Helsinki, Department of Computer Science), Tamer Nadeem (University of Maryland, Department of Computer Science), Cristian Borcea (New Jersey Institute of Technology, Department of Computer Science) and Liviu Iftode (Rutgers University, Department of Computer Science)

Lunch

14:00-14:30 Ph.D Student Workshop discussion

14:30-15:30 Session 2
Pervaho: A Development & Test Platform for Mobile Ad hoc Applications. Patrick Eugster (Purdue University), Benoît Garbinato and Adrian Holzer (University of Lausanne)
Reliable Probabilistic Broadcast in Wireless Ad-Hoc Networks. Vadim Drabkin, Roy Friedman and Marc Segal (Technion I.I.T)

Coffee Break

16:00-17:00 Session 3
Reconfigurable Architecture for Group Communication Support in Hybrid Networks. José Mocito and Luís Rodrigues (University of Lisbon)
Multi-homed Internet access in ad hoc networks using host identity protocol. Mikko Särelä (Helsinki University of Technology)

19:00 Workshop Dinner

Monday, 3rd July

- 09:00-10:00 Keynote
An approach for creating and deploying services for mobile devices in a digital city. Alexandre Sousa (Associação Porto Digital, ISMAI, ParadigmaXis)
- Coffee Break
- 10:20-11:20 Session 4
A Modular Solution to Fair Exchange for Peer-to-Peer Middleware. Benoît Garbinato and Ian Rickebusch (University of Lausanne)
Coordination of Autonomous Mobile Entities. Mélanie Bouroche and Vinny Cahill (DSG, Trinity College Dublin)
- 11:40-12:40 Session 5
Practical DHT lookup in Mobile Ad-hoc Networks. Raphaël Kummer, Peter Kropf and Pascal Felber (University of Neuchâtel)
TeenyLime: Transiently Shared Tuple Space Middleware for Wireless Sensor Networks. Paolo Costa (Politecnico di Milano), Amy L. Murphy (University of Lugano) and Gian Pietro Picco (Politecnico di Milano)
- Lunch
- 14:15-15:45 Session 6
Adaptation to Connectivity Loss in Pervasive Computing Environments. Manel Fredj (INRIA-Rocquencourt, France), Apostolos Zarras (Dept. of Computer Science, University of Ioannina, Greece), Nikolaos Georgantas and Valérie Issarny (INRIA-Rocquencourt, France)
Geotumba: Geographic Information Retrieval System for Mobile Devices. Sérgio Freitas, Ana Afonso and Mário Silva (Universidade de Lisboa, Faculdade de Ciências)
Replication Management in Low Availability, High Latency Wireless Networks. Kulpreet Singh and Siobhan Clarke (DSG, Trinity College Dublin)

B Abstracts of the Accepted Papers

A Service-Oriented Middleware for Integration and Management of Heterogeneous Smart Items Environments

Jürgen Anke, Jens Müller, Patrik Spieß and Leonardo Weiss Ferreira Chaves

In this paper, we describe concepts for a middleware that enables communication between smart items networks and business applications. Smart items networks consist of RFID systems, wireless sensor networks, and embedded systems. The integration of such devices into business applications is challenging, since each device type has its own protocols and programming interfaces. Ex-

isting middleware simplifies development of software which runs on smart items but not integration with back-end systems. We propose middleware concepts that abstract from the smart items network on a serviceoriented level to ease integration into business applications. Our middleware allows the development of business applications which are able to deploy, run and query services from the network without any knowledge of the underlying smart items network.

Service Discovery in Highly Heterogeneous Mobile Environments

Carlos A. Flores Cortés, Gordon S. Blair and Paul Grace

Service discovery protocols make it possible for participating nodes in a network to locate and advertise services with minimum user intervention. However, because it is not possible to predict at design time which protocols will be used to advertise services in a given context/ environment, it is now becoming clear that dynamic discovery mechanisms are required by mobile nodes to cope with the heterogeneity of discovery platforms. Furthermore, mobile nodes are characterised by limited resources, so this problem must be tackled in an efficient manner. Existing adaptive mobile middleware solutions such as ReMMoC and INDISS have investigated this style of dynamic discovery. However, significant development effort is required for both systems i.e. components or translator development in order to integrate new protocol implementations. In this paper we present a component-based service discovery framework for the development of an adaptive multi-personality service discovery middleware. This supports a common architecture for individual discovery protocols to enhance configurability and reconfigurability of the framework, and minimise resource usage through component re-use. Finally, we evaluate this framework by investigating the development of four existing discovery protocols using our approach.

Deploying Services over Ad Hoc Networks through Context-based Migration (Extended Abstract)

Oriana Riva, Tamer Nadeem, Cristian Borcea and Liviu Iftode

Ad hoc networks are formed either spontaneously (e.g., vehicles on the road) or deployed for specific tasks (e.g., monitoring a certain region during an emergency situation). Unlike the traditional view of ad hoc networks as data carriers between a mobile device and an Internet server or between two mobile devices, we believe that ad hoc networks can be effectively employed to provide a new class of services specific to ubiquitous environments. Besides providing data already stored in the network, these services can acquire real-time information from nodes located in the immediate proximity of geographical regions, objects, or activities of interest. For instance, a mobile ad hoc network of vehicles can provide traffic information on highways, while an ad hoc network of intelligent video cameras can transmit images from the proximity of a disaster area.

Pervaho: A Development & Test Platform for Mobile Ad hoc Applications
Patrick Eugster, Benoît Garbinato and Adrian Holzer

This paper introduces Pervaho, a platform for developing and testing mobile ad hoc applications. The Pervaho platform is founded on a location-based publish/subscribe service, which allows mobile peers to interact in a flexible and anonymous manner, based on their collocation at the time of the interaction. To validate the semantics of location-based criteria - a general issue in mobile ad hoc applications - we also propose a phone motion simulator as second cornerstone of our platform. Finally, we evaluate the usability of Pervaho by developing a concrete application with and without Pervaho.

Reliable Probabilistic Broadcast in Wireless Ad-Hoc Networks
Vadim Drabkin, Roy Friedman and Marc Segal

This paper presents a probabilistic broadcast protocol for wireless ad-hoc networks. It is designed to take advantage of the network's connectivity information. One of the factors that this protocol uses in order to set the broadcast probability of the node is the number of observed neighbors that the node detects at a given time interval. This protocol does not rely on having an overlay and every node broadcasts a message according to its probability. The paper also includes a detailed performance evaluation by simulation.

Reconfigurable Architecture for Group Communication Support in Hybrid Networks
José Mocito and Luís Rodrigues

Today's applications are developed to support all sorts of devices and networking solutions. For instance, in distributed applications each participant may be executing in a different type of device, with its specific resource potentialities and constraints. With this diversity in mind it is increasingly more important to design and implement adaptive communication protocols, that can be reconfigured in regard to, not only the local context, but also as a function of the context of the remaining participants. This paper presents an architecture that makes use of a reconfigurable communication stack to offer support for group communication in hybrid networks, where participants may execute in wired networks, MANETs, or both simultaneously.

Multi-homed Internet Access in Ad Hoc Networks using Host Identity Protocol
Mikko Särelä

An ad hoc access network may have several Internet Gateways that are located in different parts of the Internet topology. In this paper, we present a multi-homed solution to the global connectivity problem using the Host Identity Protocol (HIP). The solution, taking advantage of the multihoming capabilities of HIP, allows a MANET node to maintain several locators e.g. IP addresses to the corresponding node and to automatically switch between them, when

any of them becomes unusable. The main benefit comes from the fact that the return routability signaling about the locators can be done well in advance. In addition, the Internet Gateways can act as NAT devices between the Internet and the ad hoc network and allow the ad hoc network routing to utilize the cryptographically secure host identifiers HIP provides.

An RFID Enabled Wide Area Gaming Architecture

Kevin Curran, Derek Harkin, Peter Mee and Ian Tierney

Wide Area Gaming is a branch of computer games where the real world location of the players and actual geography of an area becomes elements in the game. Wide area gamers have harnessed the power of GPS, high-resolution screens of handheld computers, cameras and the latest mobile phones to play games across towns and cities, where they become spies, vampire slayers, celebrities and even Pac-Man. The concept behind this project was to address the deficiencies in the technologies used in Wide Area Gaming and to determine if RFID could be utilized as a location determining technology. The application developed demonstrates that with the RFID tags having a known fixed location, the RFID system can accurately provide the actual or relative location of the tag and therefore the person scanning the tag.

A Modular Solution to Fair Exchange for Peer-to-Peer Middleware

Benoît Garbinato and Ian Rickebusch

In this paper, we propose a modular and fully decentralized solution to solve the fair exchange problem in the context of a middleware dedicated to peer-to-peer applications. Our motivation roots in the observation that fair exchange is a key problem in settings where mutually untrusted peers are willing to exchange critical digital items in the absence of a trusted third party, which is typically the case in peer-to-peer applications. Our solution is based on three key building blocks, namely a perfect link module, a tamperproof secure box and a module solving the well-known Byzantine Agreement problem. The tamperproof secure boxes need not communicate directly with each other and are only required in a limited number of key steps of our algorithm. Our approach has the advantage to allow fair exchanges to complete even though truly malicious processes have aborted.

Coordination of Autonomous Mobile Entities

Mélanie Bourroche and Vinny Cahill

Autonomous mobile entities, for example automated guided vehicles, are playing an increasingly important role in our everyday lives. Since these entities share their environment with each other and with humans, they need to coordinate their behaviour to ensure that strong safety constraints are respected. However, as sensing range and accuracy are inherently limited, and communication in wireless networks is unreliable, autonomous entities have access to only limited information about their environment and the current behaviour of

other entities. This makes ensuring system-wide safety constraints particularly challenging.

In this paper, we show how system-wide safety constraints can be translated into requirements on the behaviour of individual entities depending on the information available. We first present a formalism to express high-level system-wide safety constraints. We then introduce a notion of distributed responsibility allowing requirements on the behaviour of individual entities to be deduced. We show how this process can be applied to an example from the Intelligent Transportation System (ITS) domain.

Practical DHT lookup in Mobile Ad-hoc Networks
Raphaël Kummer, Peter Kropf and Pascal Felber

Various peer-to-peer (P2P) architectures for ad-hoc networks have been proposed over the last few years. Most of them are unstructured and use some form of flooding to locate content, because the physical constraints of the underlying network make the construction of arbitrary application-layer overlays impractical. In this paper, we study the problem of applying distributed hash tables (DHT) for mobile ad-hoc networks. Our approach to efficiently lookup content in such networks exploits physical proximity of peers when establishing and maintaining the DHT based routing tables. The efficiency of our method is demonstrated by simulation of large networks.

TeenyLIME: Transiently Shared Tuple Space Middleware for Wireless Sensor Networks
Paolo Costa, Amy L. Murphy and Gian Pietro Picco

Recent developments in wireless sensor networks (WSNs) are pushing scenarios where the application intelligence is no longer relegated to the fringes of the system (i.e., on a data sink running on a powerful node) rather it is distributed within the WSN itself. To support this scenario, we devised a novel tuple space model, called TeenyLIME, where we assume that devices are capable of independent computation, i.e., the application code is not confined to the powerful sinks, rather it is deployed on the devices. Further, tuple space operations are no longer used only for data collection, rather they are exploited for coordination of the devices themselves. Here, we describe the TeenyLIME model and its implementation in TinyOS for the MICA2 platform, and elaborate on alternate designs and research opportunities.

Adaptation to Connectivity Loss in Pervasive Computing Environments
Manel Fredj, Nikolaos Georgantas, Valérie Issarny and Apostolos Zarras

Pervasive computing environments aim at providing users with advanced services, dynamically composed out of networked services. In these open environments, availability of specific networked service instances cannot be guaranteed over time as users move and services leave and join the network accordingly. A major challenge in pervasive environments is thus to maintain services func-

tionalities despite the dynamics of the environment, which induces connectivity loss with service instances. In this paper, we analyse the requirements to make distributed composite services able to face connectivity loss, i.e., able to dynamically adapt their configuration according to the networking environment. We then discuss the adaptation of relevant techniques that originate in the fault tolerance domain to the specifics of pervasive computing.

Mobile Geotumba: Geographic information retrieval system for mobile devices
Sérgio Freitas, Ana Paula Afonso and Mário Silva

Current information retrieval (IR) systems are not designed to be used in mobile devices. This occurs, because existing IR systems only recently began to offer local services and largely ignore geo-spatial information that can be inferred from the analysis of Web pages and do not recognize the geographic context that can be extracted from the mobile devices. In order to integrate the geographic context in the IR systems, new methods, algorithms and software architectures have to be proposed. This article presents mobile Geotumba, a geographic information retrieval system, describing the main challenges of designing user interfaces for mobile devices, new methods for dening the user geographical context and visualizing the retrieved information.

Replication Management in Low Availability, High Latency Wireless Networks
Kulpreet Singh and Siobhán Clarke

Wireless networks with long distance links and intermittent connectivity have recently attracted some attention, especially in providing network connectivity in remote areas. Such networks are formed between access points which are connected via rarely available, very high latency links. In rural India, such networks are being used to provide connectivity between villages where the links are formed by buses commuting between villages. Such networks present new challenges in implementing distributed applications which require consistent replication of information. For example, a wiki replicated across villages is a useful application for building trans-village communities, but requires solutions for consistent replication of the wiki. In this paper we identify the challenges in supporting replication in low availability, high latency wireless networks, analyse why traditional solutions can not be used, and motivate new approaches to address the problem.

References

- [1] Boris Koldehofe. Proceedings of the 4th MiNEMA Workshop in Sintra. Technical Report DI/FCUL TR 06–10, Department of Informatics, University of Lisbon, June 2006. <http://www.di.fc.ul.pt/tech-reports/06-10.pdf>.