

Context-aware Business Application Service Co-ordination in Mobile Computing Environments^{*}

Heikki Helin¹, Matthias Klusch², António Lopes³, Alberto Fernández⁴, Michael Schumacher⁵, Heiko Schuldt⁶, Federico Bergenti⁷, and Ari Kinnunen⁸

¹ TeliaSonera Finland Oyj, Finland
Heikki.j.Helin@teliasonera.com

² Deutsches Forschungszentrum für Künstliche Intelligenz (DFKI), Germany
klusch@dfki.de

³ Associação para o Desenvolvimento das Telecomunicações e Técnicas de
Informática (ADETTI), Portugal
antonio.lopes@we-b-mind.org

⁴ Universidad Rey Juan Carlos (URJC), Spain
alberto.fernandez@urjc.es

⁵ Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
michael.schumacher@epfl.ch

⁶ University for Health Sciences, Medical Informatics and Technology (UMIT),
Austria
heiko.schuldt@umit.at

⁷ FRAMETech S.R.L., Italy
bergenti@cs.unipr.it

⁸ EMA Group, Ltd., Finland
ari@ema.fi

Abstract. The research project CASCOS will implement, validate, and trial value-added support for business services for mobile workers and users across mobile and fixed networks. The vision of the CASCOS approach is that ubiquitous application services are flexibly co-ordinated and pervasively provided to the mobile users by intelligent agents in dynamically changing contexts of open, large-scale, pervasive environments. The essential approach of CASCOS is the innovative combination of intelligent agent technology, semantic Web services, peer-to-peer, and mobile computing for intelligent peer-to-peer (IP2P) service environments. The services are provided by software agents exploiting the co-ordination infrastructure to efficiently operate in highly dynamic environments.

1 Introduction

In this paper we describe work done and research challenges towards supportive infrastructure for business application services for mobile workers and users

^{*} This work has been supported in part by the European Commission under the project grant FP6-IST-511632-CASCOS.

across mobile and fixed networks in a research project CASCOM (Context-Aware Business Application Service Co-ordination in Mobile Computing Environments) [3]. Our driving vision is that ubiquitous business application services are flexibly coordinated and pervasively provided to mobile workers and users by intelligent agents in dynamically changing contexts of open, large-scale, and pervasive environments.

Our approach is a combination of agent technology, Semantic Web service coordination, P2P, and mobile computing for intelligent peer-to-peer (IP2P) mobile service environments. IP2P environments are extensions to conventional P2P architectures with components for mobile and ad hoc computing, wireless communications, and a broad range of pervasive devices. Basic IP2P facilities come as web services, while their reliable, task-oriented, resource-bounded, and adaptive co-ordination-on-the-fly characteristics call for agent-based software technology. A major challenge in IP2P environments is to guarantee a secure spread of service requests across multiple transmission infrastructures and ensure the trustworthiness of services that may involve a variety of providers. The services of our infrastructure are provided by peer software agents exploiting the co-ordination infrastructure to efficiently operate in dynamic environments. The IP2P infrastructure includes efficient communication means, support for context-aware adaptation techniques, as well as dynamic and secure service discovery and composition planning.

Given that our architecture builds on an assumption that users are providing services to other users, it is essential that these services will work on a broad range of pervasive devices. Therefore, we will concentrate on solutions that can be applied to mobile devices lacking processing capabilities of their office counterparts. Another point that is frequently ignored when adapting software to mobile devices is the limitation of wireless communication paths that these devices typically employ. In our architecture, the services will be adapted not only to the constraints of mobile devices but also to the constraints of wireless communication paths (e.g., by optimizing the communication over wireless connection). In the latter case, the concept of seamless service experience is essential. With seamless service experience we refer to an environment in which users have an easy and seamless access to electronic services, applications, and information anywhere and anytime.

The rest of this paper is organized as follows. In Section 2 we describe motivating use case scenarios. Section 3 gives an overview of technologies on which our architecture builds on. Then, in Section 4 we introduce our architecture. Finally, Section 5 concludes the paper.

2 Use Case Scenarios

We have developed several use case scenarios in order to gain more deep understanding of requirements, behavior, and needs of mobile users and workers. Further, these use case scenarios are important to investigate detailed require-

ments for the implementation and execution of agents and multi-agent systems in next-generation IP2P environments.

The specified use case scenarios are (1) an emergency healthcare scenario, (2) a telemonitoring and e-inclusion scenario, and (3) a shopping mall scenario. These use case scenarios show a large potential for commercial exploitation of our architecture. Scenarios are outlined below. Detailed information about the scenarios can be found from [4].

2.1 Emergency Healthcare Scenario

The emergency healthcare scenario is based on the fact that people on the move, for example, traveling in foreign countries for business or holidays, may get into situations where they need medical assistance because of a sudden disease or emergency. Currently, these sorts of episodes are neither tackled nor realized in this form in practice and no software system is presently widespread in use to address them.

Our emergency healthcare scenario consists of two stories. In the first story, a man on a business trip suddenly observes serious ailments with his stomach unknown to him. Fortunately he has installed our system to his smart phone and he is directed to a sophisticated healthcare institution where our infrastructure is also available. In addition, our infrastructure allows for accessing information from the record of this person stored in her/his home country in order to avoid redundant and unnecessary examinations and even to negotiate on the coverage of costs with his/her health insurance. In the second story, tourists from Finland are having summer vacation abroad. During the trip, one of them seriously suffers from pain in the upper part of her body. Although she has installed our system to her PDA, our infrastructure is not fully supported at the present location, which changes the story significantly from the first one. However, even in this case, our system proves to be helpful. Essentially, although repeated examinations cannot be avoided, the availability of our system allows the patient for instance to contact another physician for a second opinion. In addition, a service provider in the home country can be contacted for organizing the transfer back home.

As results of a medical emergency, persons (patients) not only need medical treatment, they also need informational as well as sometimes transportation assistance either directly or indirectly. In the latter case, assistance in the form of information is usually required by the physicians, hospitals, or healthcare professionals involved. One implication of these complex requirements is the need for on-demand initiation, co-ordination, and supervision of various activities represented either through persons, or non-human actors (i.e., agents and services).

2.2 Telemonitoring and e-Inclusion Scenario

The telemonitoring and e-inclusion scenario relates to improving people's life conditions and to reduce the costs of long hospitalizations. In this scenario, the underlying idea is to allow elderly people, chronic patients or high risk patients to stay at home and to benefit from a remotely and automated medical supervision.

The scenario assumes the existence of universal communication infrastructures mainly used by mobile devices. The infrastructure in use has to provide a flexible platform for different monitoring applications and must guarantee a high level of reliability. A core task of monitoring systems in healthcare is to co-ordinate and process various data streams. These data streams are either automatically generated by different sensors or manually by user (patient) input. Data has to be collected, propagated, processed, and stored in a distributed manner. Intelligent personal and infrastructure agents are candidates to fulfill data processing. In addition, these agents not only have to detect critical situations from the sensor data streams but also to determine how to handle these situations (e.g., proposal of nearest healthcare institution based on the current location of a patient, etc.). However, an important fact is that the system has to support intermitted connectivity, which means that disconnected parts of the system are still able to operate locally. This applies especially to all devices or agents in the patient's immediate vicinity.

Apart from the infrastructure for transferring different medical measurements to the place where decisions about medical treatments are made, the system should also support health care service providers to attend large numbers of patients efficiently. Throughout this scenario, connections between end users are normally characterized by a one-to-one communication. Thus, with every new patient, the efforts for the care-giver to provide thoroughly attendance increase and eventually become very extensive. Extending the communication pattern to a one-to-many communication allows the caregiver to build virtual groups of patients with similar characteristics, thus opening up new possibilities for medical attendance. It is important that this extension is compatible with the highly individual and personal physician-patient relationship, that is, this relationship should not be impaired. Instead, it should be an orthogonal add-on.

2.3 Shopping Mall Scenario

The shopping mall scenario is based on the existence of commercial areas that can provide several different services to the client, such as shopping malls, airports, and city centers. These are usually large areas with a lot of different services offered to their clients. In this scenario, user-centric information combined with context-aware information will allow the personal agent to easily deliver the most adequate information and services to the end user.

The scenario assumes the existence of a software system on these large commercial areas that would help clients with small devices to access context-aware networked agent applications providing a range of services that are specifically related with the place they are in. For example, a client entering a shopping mall could look for restaurants, shops, movies of his/her preference, and possibly interact with them. Alternatively, it is desirable that clients receive information about relevant events taking place within the commercial area—not just any event, but events related to the clients' profiles. The scenario is an open system of intelligent agents, Web services, end user devices and external systems.

Context-awareness plays an important role in this scenario. Available context information that is provided by a wide range of sensors is collected and is then acquired by agents. Agents use this information to discover, compose and execute services in the most adequate way for the specific end user. Agents use this collected information to create their “perceptions” of the surrounding environment. Ultimately, the assembling and analysis of the collected information allows the agent to perceive such situations as the end user affective state, end user activity, end user interest/boredom levels over a certain object/subject and end user focus of attention. Equipped with such perceptions, the agent is able to adapt to the new reality and guide its activity accordingly, thus providing better services to the end user and other agents.

Apart from being focused in context information, the shopping mall scenario is focused on end users characteristics and preferences. Clients have a profile, which has a set of properties related to their personal information and to their preferences. When entering a shopping mall, they would usually like to receive information that might be relevant to their profiles. Further, the scenario provides simple methods that users may use to provide feedback to their personal agents. It is also important to provide means for personal agents of different clients to exchange information about their users, preserving the privacy and security of sensible information.

3 Technical Approach

Software agents will be a key technology to address the challenges of our architecture as they offer an adequate abstraction for dealing with services from pervasive devices in IP2P environments. On the other hand, IP2P networks provide a suitable environment for agents to collaborate as peers sharing information, tasks, and responsibilities with each other. Agents can help manage the complexity of P2P networks, and they can be used to improve the functionality of conventional P2P systems and protocols. The inherently autonomous nature of intelligent agents helps achieving peer node autonomy, which is a requirement to operate efficiently in highly dynamic environments. Our innovations in the agent domain will concern the development of context-aware agent-based Semantic Web services, and flexible resource-efficient co-ordination of such services in the nomadic computing field.

Using agents in wireless environments has been an active research area in the past few years [17]. Several researchers have addressed these issues by developing agent platforms for resource-poor devices enabling them to run agent-based software (e.g., [2, 16]). We will build on the previous work by using the developed agent platforms as a basis of our architecture. However, the IP2P aspects are typically insufficiently taken into account and thus our research represents a relevant advancement in this direction. Wireless communication in agent systems has been addressed in many levels [8]. However, agent communication methods for wireless environments typically assume proxies in the fixed network. We,

on the other hand, will provide solutions for agent communication in wireless environments with minimal assumption of fixed infrastructure.

Service co-ordination mechanisms of P2P systems can be applied to multi-agent systems to improve their efficiency. Although this may be accepted on a conceptual level, the combination of agents and P2P environments certainly deserves more innovative research and development, especially regarding nomadic environments. However, many modern P2P overlay network algorithms (e.g., [11, 13, 15, 19]) lack support for rapid node movements and expect that the network topology remains relatively static. This assumption no longer holds in highly dynamic environments where a node providing a service may be mobile and connected to the overlay network using a wireless connection. However, some distributed hash table algorithms taking mobility and wireless environments into account has been developed (e.g., M-CAN [10] and Warp [18]). But, the dynamic topology of IP2P networks, the fluctuating QoS of wireless network connections, and the limited capacity of mobile devices connected to such networks pose several challenges that typically have been addressed inadequately in service discovery architectures.

The problem of service co-ordination can be split into several subproblems: service discovery, service composition planning, execution of composite services, service execution monitoring, and failure recovery. We will advance the state of the art in these research areas by carrying out innovative research on how these problems can be solved in open, secure IP2P environments taking into account resource-poor devices. Most existing service discovery technologies focus on matchmaking algorithms [9]. Despite the large efforts made by the research community so far to semantically describe and reason on Web services, efficient methods for reasoning on such descriptions still remain to be invented, and to be widely adopted by industry. That particularly holds for flexible and efficient matching algorithms to be performed in large scale and resource limited IP2P environments.

Service composition and planning could be addressed using existing artificial intelligence planning formalisms. However, these methods were developed for problems where the number of operators is relatively small but where plans can be quite complex. In contrast, in Web service composition for open, large-scale IP2P service environments planning methods that can deal with huge number of possible service are required. However, plans are not necessarily very complex. This means that planning methods must follow more closely the structure of the service directories rather than be geared to generating highly complex plans. We will develop planning mechanisms that establish plan fragments directly on top of the service directory to solve this problem.

Agent-based IP2P applications may be largely pervasive thus inherit the main characteristic of minimally intrusive pervasive applications: Context-awareness (e.g., [6]). These concepts have been intensively investigated in many contexts. However, there are neither well defined, nor commonly agreed concepts of context, situation, and context-awareness in P2P environments, not to speak about the notion of situation-aware agents and multi-agent systems in IP2P environ-

ments. We will investigate these issues in the context of IP2P environment and will develop context-aware agents providing various business application services.

Several other similar architectures than ours have been developed in the past. The Agents2Go [12] is an agent-based architecture which uses location as a context-information in order to provide services to mobile users. The Agents2go architecture is based on (centralized) server, whereas we employ more distributed P2P paradigm. Further, Agents2go architecture does not support service coordination (i.e., discovery, composition planning, execution of composite services, execution monitoring, and failure recovery), but only discovery and execution. The Ronin and DReggie system [5] is an agent discovery architecture which builds on Jini and provides semantic service matching functionality. Context-information is not taken into account in this architecture, and as in Agent2go, the architecture does not support service coordination. Yet another similar architecture is MyCampus [14]. MyCampus is an agent-based environment for context-aware mobile services. The main difference between MyCampus and our architecture is that MyCampus is based on centralized server. Further, MyCampus supports only semantic service discovery and not other parts of service coordination.

4 Conceptual Architecture

In this section we introduce our conceptual architecture of the agent-based IP2P service co-ordination infrastructure. This infrastructure will be the basis of our architecture providing functionality such as efficient, secure, and reliable communication independent of the access technology, efficient and secure service discovery, and context monitoring.

Our architecture is an extension to conventional P2P architectures with a support for mobile computing and a broad range of mobile and pervasive devices. Users are logically connected to the system either by using a wireless access network or by using a wireline access network. In the former case, the user most likely has a mobile device with a necessary equipment to employ the selected access technology, whereas in the latter case, a desktop computer may be used.

4.1 IP2P Network Infrastructure

The IP2P network infrastructure is logically situated on top of the combination of various wireless and wireline access networks providing seamless mobility between access technologies (see Figure 1). In general, seamless mobility means that roaming from one location to another possibly switching the underlying access technology occurs without inconvenience to the user. Such a feature will be an important feature of the future nomadic applications and therefore an essential part of our architecture. Figure 1 depicts a situation where the user drives through an urban area having available several different access technologies.

On top of seamless mobility environment, we will build a P2P overlay network architecture (see Figure 2), which takes into account characteristics of wireless networks and resource-poor mobile devices. The fact that many P2P systems




Fig. 1. Seamless mobility




Fig. 2. P2P Overlay network

and algorithms are designed for fixed network environments implies that these solutions are not directly applicable in environments where wireless communication paths are involved. The dynamic topology of P2P networks, the fluctuating QoS of wireless connections, and the limited processing capacity of many of the devices connected to such networks, pose several challenges that have been addressed inadequately so far. Some P2P platforms consider that a client may be situated in a resource-poor device (e.g, JXTA [7]), but the communication over slow wireless connections is typically insufficiently taken into account. However, agents should be able to communicate efficiently with one another also in wireless environments. Sometimes, efficiency is not as important an aspect as, for example, reliability. Our architecture will provide necessary, efficient, and reliable agent communication means for IP2P environments.

Agent-based service discovery in IP2P environments is difficult due to the lack of a fixed infrastructure configuration and support of dynamic topologies with a changing set of members. As nodes of a P2P network move, new agents, devices, and hosts may enter, or leave the network, the interconnection patterns




Fig. 3. Semantic overlay network

among them change, so new routes must be dynamically discovered and maintained with minimal routing overhead and bandwidth cost. Another challenge here is that distributed hash table (DHT) algorithms provide only keyword-based searches. However, such search method is very limited and does not fit well together with searches based on semantics. By our knowledge, for DHTs, even for multiple keyword queries, there are not known efficient approaches (P2P architectures supporting semantic service discovery are typically based on flooding (e.g., [1])). Therefore, we may need to consider an architecture, where we have a P2P infrastructure, where the directory service plays a central role. The P2P system would be used mainly for publishing and searching meta-services, which are needed for the coordination architecture to be functional for all coordination tasks. The (functional) services would then be discovered through the directory service, simply because it is only possible to expose a key on a P2P infrastructure (i.e., no semantic information can be taken into account).

4.2 Agent Architecture and Mechanisms for Service Co-ordination

The IP2P service agent and multi-agent system architecture is logically situated on top of the IP2P network infrastructure. This can be called semantic overlay network (see Figure 3). The available agents, and thus available services, depend on clients connected to the system. That is, only those agents or services are available that the (connected) clients are willing to provide. However, our architecture provides support for innovative service description, agent-based service discovery and mediation, service composition planning, service execution, and generic mechanism for situation-aware agents.

Agent-based IP2P applications in our architecture are context-sensitive and context-aware. With the context-sensitivity we refer to the ability of a device to detect its current context and changes in contextual data, such as data on

the characteristics and state of used devices, the network, the agent system, and the user. With the context-awareness we refer to the ability of applications to capture and analyze the context and interrelationship between users' actions and context-sensitive devices over time. One main research challenge here concerns solutions to the problem of how to develop and incorporate the concepts of context-sensitivity and context-awareness into the adaptive behavior of an individual agent and a multi-agent system in a given IP2P environment and application. Context-aware agents have to flexibly and timely react on changes that are prompted, for example, by variations in resource availability as a result of different types of failures, the deployment of new services, devices, or agents, as well as by variations in patterns of usage or mobility. Context information will be taken into account in service discovery, composition and execution mechanisms.

Initiatives such as Web services and the Semantic Web aim at shifting the burden of doing certain mechanical and monotonic things from humans to computers. Web services operate over the ubiquitous infrastructure of the Internet. Web service may pragmatically be defined as a task-oriented, coarse-grained, XML-based business application that is network-accessible via an API from anywhere with one or multiple protocols of the IP suite, and that may be composite. A Web service provides a fixed function interface that may be invoked by any client; it is usually enlarged with descriptive metadata for such service consumers. We are looking for a decentralized P2P approach that would significantly improve the scalability of co-ordination and availability of services.

On the semantic overlay network layer we will investigate on efficient service composition planning in IP2P environments. When several services are attempting to access and use one and the same resource, it is necessary to introduce co-ordination mechanisms to coordinate this access. We will develop Semantic Web service co-ordination mechanisms specifically for the IP2P environment and the co-ordination problems that occur in the selected application scenario, in particular optimization of competing service configurations.

Service execution monitoring addresses activities that seek to establish the extent to which progress is being made according to some reference. This enables that a timely action can be taken in view of detected changes or deficiencies. As especially in open large-scale IP2P settings it may not always be known what specific information is required to this purpose, monitoring systems should allow for run-time configurability and provide collected data for further information extraction when the infrastructure provides appropriate instrumentation. First research results indicate how domain-independent monitoring services can improve failure tolerance of open environments and effectively complement integrated exception handling within agents. Furthermore, for important classes of systemic dysfunctions monitoring services can provide an efficient alternative to complex protocols aimed at preventing such pathological states altogether. Useful monitoring of "implicit" service level agreements can be performed even when no sufficient details for monitoring were been specified or no explicit service level agreements were made. Research has also begun to address monitoring support for run-time scheduling and re-scheduling tasks in dynamic IP2P sce-

narios. As an example characterizing the current state of the boundary between research and deployment, the OWL-S white paper accompanying the current specification maintains automatic Web service execution monitoring as one of the four prominent goals in terms of the kinds of tasks OWL-S is expected to enable. This topic, however, still has not been tackled to date. By providing domain-independent monitoring services, we will thus contribute both to the consolidation and the furthering of the state of the art in this key area.

5 Conclusions

CASCOM is an IST-FP6 Specific Targeted Research Project with an objective to develop, implement, validate, and trial of agent-based service co-ordination infrastructure for innovative Semantic Web service discovery, composition, and execution across mobile and fixed P2P service networks.

The project will carry out highly innovative research aimed at providing a coherent framework for agent-based data and service co-ordination in open IP2P environments. For this purpose, we will integrate and extend existing technologies in areas such as agent-based mobile computing, service co-ordination (description, discovery, composition), and P2P computing in nomadic environments. A generic, open IP2P service environment with its agents and co-ordination mechanisms will be prototypically implemented. Further, we will deliver a fully functional proof-of-concept system showing how services for mobile users and workers can be developed using the co-ordination framework. The implementation will be developed using best practices, including “open source for open standards” development, developing generic re-usable components of an open distributed architecture.

For end users, CASCOM architecture provides easy and seamless access to Semantic Web services. Further, CASCOM architecture responds to emerging needs of mobile users, by providing an environment in which a user may easily access the same services using her desktop computer as she would be able to do in an IP2P environment anywhere, at any time by using a variety of different kinds of mobile computing devices. This enables mobile workers to do their business wherever and whenever using a common collaboration environment. For network operators, CASCOM architecture aims towards vision of seamless service experience providing better customer satisfaction, which in turn helps to retain current customer relations as well as attract new customers. For service providers, CASCOM architecture provides an innovative platform for various mobile business application services.

References

1. F. Banaei-Kashani, C.-C. Chen, and C. Shahabi. WSPDS: Web services peer-to-peer discovery service. In *International Symposium on Web Services and Applications (ISWS'04)*, pages 733–743, Las Vegas, Nevada, USA, June 2004.

2. F. Bergenti, A. Poggi, B. Burg, and G. Caire. Deploying FIPA-compliant systems on handheld devices. *IEEE Internet Computing*, 5(4):20–25, 2001.
3. CASCOM Consortium. Home page. Available at <http://www.ist-cascom.org>.
4. CASCOM Consortium. Deliverable 3.1: Use case scenarios, 2005. Available at <http://www.ist-cascom.org>.
5. D. Chakraborty, F. Perich, S. Avancha, and A. Joshi. An agent discovery architecture using ronin and dreggie. In *First GSFC/JPL Workshop on Radical Agent Concepts (WRAC)*, NASA Goddard Space Flight Center, MD, Jan. 2002.
6. A. K. Dey. Understanding and using context. *Personal and Ubiquitous Computing*, 5(1):4–7, 2001.
7. L. Gong. JXTA: A network programming environment. *IEEE Internet Computing*, 5(3):88–95, 2001.
8. H. Helin and M. Laukkanen. Performance analysis of software agent communication in slow wireless networks. In R. Luijten, E. Wong, K. Makki, and E. K. Park, editors, *Proceedings of the Eleventh International Conference on Computer Communications and Networks (ICCCN'02)*, pages 354–361. IEEE, Oct. 2002.
9. M. Klusch and K. P. Sycara. Brokering and matchmaking for coordination of agent societies: A survey. In *Coordination of Internet Agents: Models, Technologies, and Applications*, pages 197–224, 2001.
10. G. Peng, S. Li, H. Jin, and T. Ma. M-CAN: A lookup protocol for mobile peer-to-peer environment. In *Proceedings of the 7th International Symposium on Parallel Architectures, Algorithms and Networks (ISPAN'04)*, pages 544–549, 2004.
11. S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *Proceedings of the ACM SIGCOMM 01*, Aug. 2001.
12. O. Ratsimor, V. Korolev, and T. F. Anupam Joshi. Agents2Go: An infrastructure for location-dependent service discovery in the mobile electronic commerce environment. In *Proceedings of the First International Workshop on Mobile Commerce*, pages 31–37, July 2001.
13. A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location and routing for large-scale peer-to-peer systems. In *Proceedings of the ACM/IFIP Middleware*, 2001.
14. N. M. Sadeh, T.-C. Chan, L. Van, O. Kwon, and K. Takizawa. Creating an open agent environment for context-aware m-commerce. In B. Burg, J. Dale, T. Finin, H. Nakashima, L. Padgham, C. Sierra, and S. Willmott, editors, *Agentcities: Challenges in Open Agent Environments*, pages 152–158. 2003.
15. I. Stoica, R. Morris, D. R. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the ACM SIGCOMM 01*, San Diego, California, Aug. 2001.
16. S. Tarkoma and M. Laukkanen. Facilitating agent messaging on PDAs. In *Fourth International Workshop on Mobile Agents for Telecommunication Applications (MATA-2002)*, pages 259–268, Barcelona, Spain, 2002. Springer.
17. S. Tarkoma, M. Laukkanen, and K. Raatikainen. Software agents for ubiquitous computing. In R. Khosla, N. Ichalkaranje, and L. Jain, editors, *Design of Intelligent Multi-Agent Systems: Human-Centredness, Architectures, Learning and Adaptation Series: Studies in Fuzziness and Soft Computing*, volume 162, pages 31–60. 2004.
18. B. Y. Zhao, L. Huang, A. D. Joseph, and J. D. Kubiatowicz. Rapid mobility via type indirection. In *Proceedings of Third International Workshop on Peer-to-Peer Systems (IPTPS)*, San Diego, CA, USA, Feb. 2004.
19. B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Technical Report CSD-01-1141, University of California at Berkeley, 2001.