A Generic Agent-based Peer-to-Peer Infrastructure for Social-mobile Applications

Steffen Kern, Torsten Dettborn, Ronny Eckhaus, Yang Ji,
Christian Erfurth, Wilhelm Rossak
Friedrich Schiller University Jena
Computer Science Department
Ernst-Abbe-Platz 2, 07743 Jena, Germany
steffen.kern@informatik.uni-jena.de

Peter Braun
the agent factory GmbH
Wildenbruch Str. 15
07745 Jena
braun@the-agent-factory.de

Abstract:
This paper presents a generic agent-based framework for social-mobile applications, which has been developed as part of an ongoing linkage project. The *MobiSoft* project is driven by the vision of facilitating, augmenting, and promoting human social interaction by electronic personal assistants during face-to-face encounters. Possible areas of social mobile applications include the establishment of groups or communities based on shared interests or goals, the exchange of information such as personal profiles, news, private sales, or any kind of recommendations, and the preselection of possible communication partners in social networks. As part of this ongoing project, we have developed a new agent-based peer-to-peer software infrastructure for social-mobile applications. It focuses on strategies for information exchange in mobile ad-hoc networks and techniques for information representation using semantically rich languages based on existing standards. In contrast to existing approaches for the mobile Internet that are based on the client/server paradigm, the project proposes to use a completely decentralized approach (peer-to-peer) and to use mobile agents as intelligent information carrier.

1 Introduction

Mobile devices, in particular mobile phones, have become part of our daily life and it can be expected that they will become even more widespread in the near future. They have become our indispensable companions and are already powerful enough to serve as personal assistants to help us managing our appointments, contact lists, and personal tasks.

The next generation of mobile devices will support wireless network technologies that allow for the establishment of personal area networks (PAN) in order to exchange information with others in close proximity. The proposed work is driven by the vision of facilitating, augmenting, and promoting human social interaction by electronic personal assistants during face-to-face encounters. Areas of social mobile applications include:

- The establishment of groups or communities based on shared interests (work, hobbies) or activities and goals (such as to reduce travel costs by sharing a taxi).
- The exchange of information, such as personal profiles, news, private sales, or any kind of recommendations.
- The preselection of possible communication partners in social networks and the coordination of shared task lists and diaries by automated negotiations.

Our approach can be seen as complementary to more traditional techniques for information discovery on mobile devices.

- At first, instead of a client-server based communication approach known from the Internet, we propose a decentralized, peer-to-peer like technique, which is based on the notion of proximity.
- Second, the process of searching for and disseminating information is pro-actively initiated by the personal assistant rather than by the human user.
- Finally, the goal of our approach is mainly the establishment of social interactions rather than just information distribution.

In this project, we address several research problems that are at the intersection of distributed computing, mobile ad-hoc networks, and information representation using standardized languages and ontologies.

The rest of this paper is structured as follows: The next section outlines the goals of our research project and Sec. 3 discusses similarities and differences to other projects. The following two sections describe the architecture of our approach and the current state of our implementation. Finally, the last section gives an outlook to future development.

2 The MobiSoft Project

The MobiSoft project is an ongoing linkage project including Friedrich Schiller University Jena, the agent factory GmbH, and Godyo AG and is funded by the Thuringian Ministry of Economy, Technology and Labor.

In MobiSoft we aim at application scenarios, in which humans travel around and meet at specific places, for example shopping malls, sports stadiums, public transport, museums, libraries, conferences, lecture halls, etc. Although it might be helpful and interesting, people rarely start talking to complete strangers, because of inhibitions, social barriers or simply a lack of time. Otherwise, if people knew each other, they would more freely

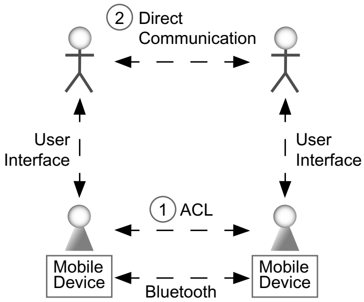

Figure 1: Interaction between agents running on mobile devices and their users in our approach. At first, agents start to communicate together before in a second step, agents delegate responsibility for further communication back to their users.

exchange information and, therefore, spread and receive pieces of useful information that could further be combined with already existing information and forwarded to others.

In this project we aim at supporting such a very human behavior of information exchange with software agents as user representatives that reside on mobile devices and to which users have delegated the task of finding proper human communication partners, compare Fig. 1. We conceive software agents to be small entities that are situated in a networked environment of mobile devices. Agents are able to react on percepts from the environment about other agents in proximity and then autonomously start information exchange with them. Communication between agents is based on messages, which are annotated with semantic information as defined in an *agent communication language* using high-level communication protocols such as negotiations. For more information about software agents in general, we refer to [Woo02].

This information exchange works transparently for the user only in the first step, in which the agents exchange information, such as user profiles, or negotiate best interaction time. Later, the agents inform their respective users about the potential communication partner and let them decide on further steps. By this, we overcome existing inhibitory behavior of humans by delegating this task to software agents, while the agents' goal is to find *proper* communication partners and *interesting* information. The project focuses on developing a new framework for social peer-to-peer information exchange in mobile ad-hoc networks. It has the following key aims:

- Develop hybrid information exchange techniques, in which agents pro actively distribute information to as well as reactively receive interesting information from other agents, taking into account the specific limitations of mobile ad-hoc networks.
- Develop methods to describe user profiles, interests and information using semantically rich languages, which are based on existing standards known from the Semantic Web. Develop techniques to match user profiles while taking into account the specific hardware limitations of mobile devices.

We are aware of several additional research issues, for example in the area of privacy protection and human-computer interaction to make this type of application both useful and acceptable by users. We see this project as a first step in which we aim at developing the framework and technical infrastructure that will also enable later studies of those issues in detail. The research proposed in this project will offer new technology solutions for social mobile applications and to gain access to a real test-bed and customers to demonstrate and evaluate research outcomes.

3 Related Work

Current approaches for mobile social applications [Smi05] are based on central servers. For example, www.dodgeball.com and www.playtxt.net are social mobile networks to locate friends, friends of mutual acquaintances or other people with matching profiles. In those applications a user has to provide his or her current location manually, whereas in the Reno system [Smi05] the current location is determined via GSM technology.

In this project we try to combine the advantages of both existing approaches, while avoiding a centralized architecture. On the one hand, we continue to use the concept of *places* rather than *locations*. A *place* is a logical description of an area such as *soccer stadium*, whereas a *location* is given by exact coordinates or cells. On the other hand, it is essential for our approach that proximity of humans can be determined transparently for users.

Therefore, we focus on applications based on mobile ad-hoc networks (MANETs). A MANET is a collection of mobile devices (nodes), which can communicate with each other over a wireless network, for example WiFi or Bluetooth. By definition, mobile ad-hoc networks have no fixed infrastructure, that is, all typical network functions need to be coordinated by the network nodes in a distributed manner. While WiFi is more suited for high-end mobile devices and for scenarios which demand long range ad-hoc networks, Bluetooth is a promising communication technology for short range ad-hoc networks like personal area networks (PAN) on which we focus in this project. A personal area network can be seen as a digital space around a person, whose size depends on the underlying wireless transmission technique. If two digital spaces overlap, people can virtually see each other, that is, their mobile devices are able to exchange information. The concept of MANET is very appealing both for research and industry, as it enables a new class of application that has not been possible so far. In a PAN we do not need an explicit notion of places and provision of proximity information is an inherent network function.

Most research into mobile ad-hoc networks has focused on the problem of multi-hop routing data packets to enable Internet-like applications in ad-hoc networks. In particular, they address the issues of how to enable peer-to-peer like applications on mobile devices [OAdM+05] to share files [CL02, GD04], MP3 play-lists [Wib04], or information dissemination of homogeneous data with one application such as traffic information [OW04]. Most of those approaches were mainly for enabling information exchange triggered manually by users [Wib04] or make information dissemination completely independent of the user [OW04]. In contrast, our project aims at the establishment of social interactions by use of MANET.

In our project, we identify and focus on two key problems, which are related to the two main goals of our project as mentioned previously:

At first, we consider different strategies for information exchange [FZ96]. *Flooding techniques* are based on the concept of broadcasting information units to all available nodes in the network. In general, flooding is a very simple technique that can be considered to be not appropriate in large networks, because of high resource usage in terms of bandwidth and energy consumption. The *publish/subscribe architectures*, well known from the Internet and also used for mobile applications connected to central servers on the Internet, are useful if the server (broker) can predict what type of information may be useful for the client. The more accurate this predication is the less data is sent superfluously over the network [EY04]. Profiles and application dependent requests are used to describe the clients needs and desires. Although, publish/subscribe architectures are very attractive, major requirements of this model, such as orderedness, consistency, and completeness make it difficult, if not impossible, to realise it in mobile networks. For example [CDKR02] [HGM01] propose a publish/subscribe model for peer-to-peer and mobile networks without addressing the three requirements mentioned previously. *Epidemic dissemination* sends each information unit to a randomly chosen group of nodes. This dissemination approach enables messages to propagate quickly in the network and it is very robust against the node and network link failures. This approach works completely decentralised and must be seen in contrast to IP multicast techniques in which a spanning tree has to be set up from the source to all receiver nodes. In those techniques the node or network link failures result in loss of messages, whereas in the epidemic-based approaches the messages can be delivered to a node via multiple (redundant) paths. For more information on those algorithms we refer to [KMG03, VvRB03, BHO+99, DHA03, KSSV00]. So far, the epidemic-based algorithms have only been studied as a general replacement for traditional routing and multicast algorithms in mobile ad-hoc networks. Finally, *proximity-based algorithms* send information units only to neighbours, that is, other nodes in close proximity. At the moment it is not clear which of these approaches works best under which circumstances. Early results are only based on simulations [NDK04] of small networks and focus on performance metrics rather than qualitative comparisons of the approaches.

The second aspect of our projects deals with the problem of information representation in open environments which are not specific to a single application domain as existing approaches [OW04]. We aim to apply information dissemination approaches to distribute information about users, user interests and similar information. To make this approach flexible, extendable, and to base the matchmaking process between user interests and roaming information units, it is necessary to use semantically rich languages. Although there has been a lot of research done in the area of matchmaking of user interests and profiles [KSvH04] and the creation of social groups, today's available techniques are still quite simple. For example, in the Internet we find www.tribe.net that is a Web site enabling people to find other people based on their interest. The users must create a profile, can publish recommendations for restaurants etc., and establish or join tribes (online communities) dedicated to a specific topic. The description of user interests is based on a list of keywords. For mobile devices, we find www.upoc.com that can be used to establish communities on the Web and then send short messages to all members of a group or a content channel. These approaches compare user interests by comparing keywords. The Web site www.jambo.org provides a software for WiFi-enabled mobile devices so that users can locate each other based on keyword-based profile matching. Neither approach uses semantic descriptions of user profiles and preferences, but only simple text-based approaches. The first two approaches mentioned are based on central servers in the Internet, whereas the last approach requires a WiFi managed network. Other projects such as Webhound/Webdoggie [SM95] and HOMR/Ringo/Firefly [LMM94] use similar approaches. Friend of a Friend Finder (FOAF) is a project that aims to share information about persons in the Internet. The language used in this project is RDF that provides a means to describe data and metadata. RDF defines a simple data model which consists of resources and statements that link two resources, comparable to a subject-verb-object relationship. A statement is called a triple, which consists of subject and object, and the predicate that plays a role of the verb mentioned previously. In a so-called FOAF file, a user describes his personal data and which other persons this user knows using RDF. With a help of these links to other persons, a search engine can now create a graph of who knows whom. However it is not possible to describe user interests and preferences with FOAF.

4 Architecture

This section covers some core issues of the architecture of our system. At first, we describe the network structure of the targeted scenarios, followed by a short summary of possible network overlays. Afterwards we present roles of agents found in our system. This sections ends with some remarks on communication and interaction protocols/techniques.

4.1 Network Layout and Overlay

We first describe two network layouts which are most likely to be found in our application scenarios, namely *Semi Ad-Hoc* and *Full Ad-Hoc*. Afterwards, we take a look at network overlay techniques to handle these two kinds of networks.

What we call a *Semi Ad-Hoc* network is a mixture of a classic, hard-wired network and highly dynamic ad-hoc networks mentioned previously. The core network consists of a number of workstations surrounded by a high number of mobile devices connected via different network technologies like WiFi, Bluetooth or UMTS/GPRS. For example, this kind of network structure can be found in a shopping mall, in which shops or restaurants establish a network of stationary workstations providing Web services for product information or bookings. This network can be accessed by mobile devices through hotspots. Handling of Semi Ad-Hoc networks is rather easy as compared to Full Ad-Hoc networks because workstations can be used for routing purposes and service propagation. With a fixed core network the establishment of a stable overlay structure is possible, which can help to conquer the systems complexity. For example, an agent started from a mobile devices can find offered services rather easy as they are proposed by a number of fixed hosts which are known throughout the system. See Fig. 2 for an illustration of that case.

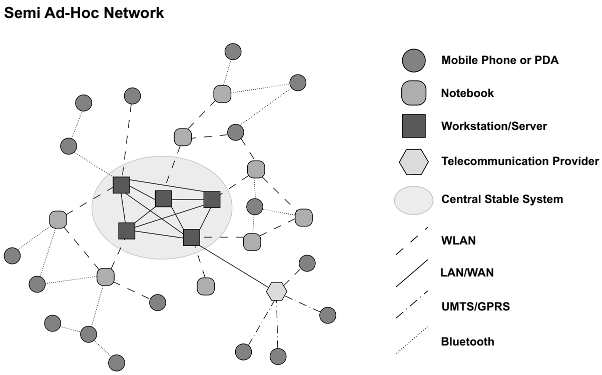

Figure 2: Structure of a Semi Ad-Hoc Network

The second case of *Full Ad-hoc* networks was outlined before in this paper and is the more interesting and difficult one. Not only must we deal with small and instable network connections with a very short range, the network is formed of mobile devices with limited hard- and software resources. Service propagation and discovery must be performed in a different way as in the first case. No central server is present for hosting a service directory. Interaction between two peers is mainly performed by the interaction of two or more agents, that is, service propagation is an agent task. Our system will allow a completely transparent, context aware connection to one of the available networks considering user preferences, network quality, etc. This type of network is typical for any kind of places without any communication infrastructure.

Network Overlay

For a Semi Ad-Hoc Network structure, we plan to apply a peer-to-peer overlay using the TLS routing protocol [BL04]. This protocol was developed to allow P2P networking in highly dynamic networks and guarantees time complexity of $O(log\ n)$ for joining and leaving of peers, when n is the number of existing peers. The main idea behind TLS is to structure peers in a forest of binary trees. New peers are inserted as leafs and the longer a peer resides inside the network the closer it will move towards the root node. Thus, stable machines such as servers are at the root of such a tree and will carry most of the message routing load.

Due to the fact that the TLS system is still under development, we have considered other P2P systems like Chord [SMK$^+$01], CAN [RFH$^+$01] and Pastry [RD01] even if they do not deliver the same runtime bounds as TLS does. Additionally, the routing protocols of all those systems should be adapted to consider the heterogeneous nature of our networks, for example lower routing load on mobile devices and increase load on central and stable parts of the network.

For Full Ad-Hoc networks, a complex overlay like P2P is surely not applicable, but we believe that overlay structures might nevertheless proved to be useful in such networks. For example, routing optimization to save battery power of mobile devices may rely on overlay structures. We are inspired by overlay techniques for wired networks and try to derive more lightweight concepts for our purpose. Beside that, we consider more straightforward techniques like broadcasts or multicast.

4.2 Agent Roles

The concept of the framework proposed in this project for realising wireless information dissemination systems is based on cooperative, autonomous, dynamic, and adaptive software agents that are located on mobile devices. We mainly distinguish two roles:

At first, mobile agents [BR05] are used as information carriers. They are injected into the system and then roam autonomously from peer to peer to distribute the information they carry as their data. Mobile agents are aware of their environment, that is, devices and other agents in their vicinity, and communicate spontaneously to other agents. Information is represented as RDF statements, using only standardised ontologies described in OWL. Besides pure information dissemination, these mobile agents are also used to distribute user requests to other devices and then bring results back to the user. The mobile agents code can be considered to define the dissemination strategy, for example the scope of the information unit, that is, the distance from the information source, the spatial direction and temporal freshness [MM04].

Another aspect modelled within agents code is their behaviour when they meet other information units. For example, it might be reasonable for agents to modify their content while roaming. The reason to employ mobile agents in this project is simply its beneficial design paradigm compared with other more traditional paradigms such as client/server. It is not necessary to define complex network protocols for the transmission of information units because mobile agents carry the protocol. Therefore, with this approach we are able to implement existing strategies for information dissemination and are open for future enhancements. In addition, mobile agents have been proven theoretically to work very well in mobile environments and to be in particular robust against network failures.

Thus, our approach enables information units to be spread in two ways. First, they virtually hop in a form of mobile agents from device to device. Second, as the users move physically through an environment, they distribute information to new areas. The latter is of particular importance in the targeted network environment because we have to assume unconnected piconets. Mobile devices are therefore used to distribute information from one piconet to another. This is comparable to Smart-Tag [BLB02] and RuralWiFi (www.medialabasia.org), where buses are used to carry emails from rural areas to cities.

Second, stationary agents are used on mobile devices as user representatives. They act as data consumers as well as data providers, to protect users private profile and to protect users against information spamming. These agents know user interests as described in profiles, which are grounded on a semantically rich language in terms of beliefs, desires, and intentions a model that has been used successfully in multi-agent systems for years. Whenever they receive new information from roaming mobile agents, they have to match them to user interests on a semantic level. New information will be presented to the user only if it is important and related to the current context. Otherwise, it will be inserted to a pool of pending information units. The agents will improve their internal user model by learning from user feedbacks about the relative importance of information. Other information units, not important to the user at the moment, will be tolerated to some extent to support the general dissemination strategy of our approach. The user agents also sort out out-dated or falsified information from time to time, if necessary. Access to the user profiles is protected by the user agents. Whenever a mobile agent gains access to sensitive profile information, the user is requested for acknowledgement. Again, these user agents are able to learn users intentions and in the longer term act autonomously on behalf of their owner and provide access to the user profile in a very fine-grained way.

In general, this project adopts a multi-disciplinary approach involving new technologies related to software agents and semantic information representation. The research will be carried out at the intersection of those areas with an input from our previous research on architectures for mobile agent systems, and state-of-the-art technology solutions of industry partners. The project not only aims at the development of the core infrastructure and framework as described previously, but also to evaluate its applicability in a real world scenario. The research will be carried out in close collaboration with industry partners.

4.3 Communication and Interaction

Our architecture supports several kinds of communication techniques for software agents, such as *message passing* or *information space*, because the choice for one model depends on the application and the underlying network environment.

In a semi ad-hoc network, communication can be supported by workstations and servers that can cache and forward messages or provide a place where agents can meet with each other. This art of communication is related with the communication model of an *information space*. All agents share a single space where they can exchange information and knowledge with each other. In this way, every mobile agent posts into the common information space the information about its profile, for example information about its work or the preferences of its owner. This information will be described with semantically rich languages as RDF or OWL. If an agent knows the name or address of a potential communication partner, communication can change to *message passing*. This form of direct, asynchronous communication between two entities is based on sophisticated interaction protocols and agent communication languages (ACLs) defined by FIPA.

The next step in our project is to find solutions to solve the problem of location-transparent communication between entities in mobile ad-hoc networks. As mentioned above, in a *semi ad-hoc network* this problem is not extremely serious as stationary workstations/ can always support reliable communication.

5 State of the Implementation

Due to the fact that we are at the beginning of our project, implementation is just about to start. We have conducted several tests to analyse some aspects of the architecture under consideration. For example, we run some evaluations to compare existing P2P protocols and we performed early experiments with mobile devices focusing on Bluetooth communication in order to find and interact with each other in an Ad-Hoc network. Other project partners have designed and tested ontologies for user preferences, service description and comparison and analysed techniques for service propagation, lookup and usage, which are outside the scope of this paper.

Our industry partner made first steps in migrating the mobile agent toolkit Tracy2 [BR05] to the Java2 Micro Edition (J2ME) environment. Due to the highly modular architecture of Tracy2, which is based on a micro kernel for agent execution accompanied by several plugins for advanced functionality, they can easily migrate component after component. So far, a core agency executing several agents in parallel is already running on Nokia and Sony-Ericsson mobile phones. Most of these experiments aim at understanding the techniques under consideration and to direct our decisions for further development.

6 Conclusions and Outlook

Social-mobile applications are capable to change the way people perceive their environment and interact with each other. Using a decentralized approach, mobile agents as information carriers and semantically rich profile and service descriptions, we are optimistic to handle the complexity of such systems.

Our ongoing work is focused on implementing and testing a prototype for the network overlay. Further, we develop ontologies for user profiles and agent interaction and our industry partner continues to migrate more Tracy2 functionality to J2ME. Another important aspect of our future research will be the analysis and test of interaction techniques between mobile devices to find a suiting one for our Ad-Hoc networks scenarios.

As part of the project, we aim at the establishment of a test-bed at the university campus and a nearby shopping center. This prototype will provide a wide variety of services (student, shopping and event information, community functions, information exchange, etc.) to demonstrate the capabilities of the system. Beside that, this test-bed should allow us a comprehensive performance evaluation.

Acknowledgements

The work presented in this paper is partially funded by the Thuringian Ministery of Economy, Technology and Labor under grant FKZ B 509-04005.

References

[BHO+99] Kenneth P. Birman, Mark Hayden, Oznur Ozkasap, Zhen Xian, Mihai Budiu, and Yaron Minsky. Bimodal multicast. *ACM Transaction on Computer Systems*, 17(2):41–88, 1999.

[BL04] Francesco Buccafurri and Gianluca Lax. TLS: A Tree-Based DHT Lookup Service for Highly Dynamic Networks. In *CoopIS/DOA/ODBASE 2004, LNCS 3290*, pages 563–580. Springer Verlag, October 2004.

[BLB02] Allan Beaufour, Martin Leopold, and Philippe Bonnet. Smart-Tag based data dissemination. In *Proceedings of the First ACM international workshop on Wireless sensor networks and applications, Atlanta (USA), 2002*, pages 68–77. ACM Press, 2002.

[BR05] Peter Braun and Wilhelm R. Rossak. *Mobile Agents–Basic Concept, Mobility Models, and the Tracy Toolkit.* Morgan Kaufmann Publishers, 2005.

[CDKR02] Miguel Castro, Peter Druschel, Anne-Marie Kermarrec, and Antony Rowstron. SCRIBE: A large-scale and decentralized application-level multicast infrastructure. *IEEE Journal on Selected Areas in Communications*, 20(8), 2002.

[CL02] Oliver P. Waldhorst Christoph Lindemann. A Distributed Search Service for Peer-to-Peer File Sharing in Mobile Applications. In *Proceedings fo the Second International Conference on Peer-to-Peer Computing (P2P02)*. IEEE Computer Society Press, 2002.

[DHA03] Anwitaman Datta, Manfred Hauswirth, and Karl Aberer. Updates in highly unreliable, replicated peer-to-peer systems. In *In Proceedings of the 23rd International Conference on Distributed Systems (ICDCS 2003), Providence (USA), May 2003*, pages 76–87. IEEE Computer Society Press, 2003.

[EY04] Jean Bacon Eiko Yoneki. An Adaptive Approach to Content-Based Subscription in Mobile Ad-Hoc Networks. In *Proceedings of the 2nd IEEE Annual Conference on Pervasive Computing and Communication Workshops (PERCOMM04)*. IEEE Computer Society Press, 2004.

[FZ96] Michael J. Franklin and Stanley B. Zdonik. Dissemination-based information systems. *Data Engineering Bulletin*, 19(3):20–30, 1996.

[GD04] Bharat Bhargava Gang Ding. Peer-to-Peer File-sharing over Mobile Ad-hoc Networks. In *Proceedings fo the 2nd IEEE Annual Conference on Pervasive Computing and Communications Workshops (PERCOMM04)*. IEEE Computer Society Press, 2004.

[HGM01] Yongqiang Huang and Hector Garcia-Molina. Publish/subscribe in a mobile environment. In Sujata Banerjee, editor, *Proceedings of the 2nd ACM International Workshop on Data Engineering for wireless and mobile access, Santa Barbara CA (USA), 2001*, pages 27–34. ACM Press, 2001.

[KMG03] Anne-Marie Kermarrec, Laurent Massoulie, and Ayalvadi J. Ganesh. Probabilistic reliable dissemination in large-scale systems. *IEEE Transaction on Parallel and Distributed Systems*, 14(3):248–258, 2003.

[KSSV00] Richard M. Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vcking. Randomized rumor spreading. In *In Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS 2000), Redondo Beach (USA), November 2000*, pages 565–574. IEEE Computer Society Press, 2000.

[KSvH04] Thomas Kleemann, Alex Sinner, and Andreas von Hessling. Semantic User Profiles and their Applications in a mobile environment. In *Workshop on Artificial Intelligence in Mobile Systems at UbiComp 2004, Nottingham (UK), September 2004*, 2004.

[LMM94] Yezdi Lashkari, Max Metral, and Pattie Meas. Collaborative interface agents. In *Proceedings of the 12th National Conference on Artificial Intelligence, Seattle (USA), 1994*, volume 1, pages 444–449. MIT-Press, 1994.

[MM04] Franco Zambonelli Marco Mamei. Programming Pervasive and Mobile Computing Applications with the TOTA Middleware. In *Proceedings of the 2nd IEEE International Conference on Pervasive Computing and Communications. Orlando, FL (USA), March 2004*. IEEE Computer Society Press, 2004.

[NDK04] Silvia Nittel, Matt Duckham, and Lars Kulik. Information Dissemination in Mobile Ad-Hoc Geosensor Networks. In *Third International Conference on Geographic Information Science (GIScience 2004), College Park, ML (USA), October 2004*, Lecture Notes in Computer Science. Springer Verlag, 2004.

[OAdM+05] Jens Oberender, Frank Uwe Andersen, Herman de Meer, Ivan Dedinski, Tobias Hossfeld, Cornelia Kappler, Andreas Maeder, and Kurt Tutschku. Enabling Mobile Peer-to-Peer Networking. In G. Kotsis and O. Spaniol, editors, *Mobile and Wireless Systems*, volume 3427 of *Lecture Notes in Computer Science*, pages 219–234. Springer Verlag, 2005.

[OW04] A. Prasad Sistla Ouri Wolfson, Bo Xu. An Economic Model for Resource Exchange in Mobile Peer-to-Peer Networks. In *Proceedings of the 16th International Conference on Scientific and Statistical Database Management (SSDBM04)*. IEEE Computer Society Press, 2004.

[RD01] Antony Rowstron and Peter Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, pages 329–350, November 2001.

[RFH+01] Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard M. Karp, and Scott Shenker. A scalable content-addressable network. In *SIGCOMM*, pages 161–172, 2001.

[SM95] Upendra Shardanand and Pattie Maes. Social Information Filtering: Algorithms for Automating Words of Mouth. In *Proceedings of the Conference on Human Factors in Computing Systems*. ACM Press, 1995.

[Smi05] Ian Smith. Social-Mobile Applications. *Computer*, 38(4):84–85, 2005.

[SMK+01] Ion Stoica, Robert Morris, David Karger, Frans Kaashoek, and Hari Balakrishnan. Chord: A Scalable Peer-To-Peer Lookup Service for Internet Applications. In *Proceedings of the 2001 ACM SIGCOMM Conference*, pages 149–160, 2001.

[VvRB03] Werner Vogels, Robbert van Renesse, and Ken Birman. The power of epidemics: robust communication for large-scale distributed systems. *ACM SIGCOMM Computer Communication Review*, 33(1):131–135, 2003.

[Wib04] Mikael Wiberg. FolkMusic A Mobile Peer-to-Peer Entertainment System. In *Proceedings of the 37th Hawaii International Conference on System Sciences*. IEEE Computer Society Press, 2004.

[Woo02] Michael Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons, 2002.