

04441 Abstracts Collection
Mobile Information Management
— Dagstuhl Seminar —

Margaret Dunham¹, Birgitta König-Ries², Evaggelia Pitoura³, Peter Reiher⁴
and Can Türker⁵

¹ SMU - Dallas, US

² Univ. Karlsruhe, DE

³ Univ. of Ioannina, GR
`pitoura@cs.uoi.gr`

⁴ Univ. California - Los Angeles, US
`reiher@cs.ucla.edu`

⁵ ETH Zürich, CH

Abstract. From 24.10.04 to 29.10.04, the Dagstuhl Seminar 04441 “Mobile Information Management” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

04441 Executive Summary – Dagstuhl Seminar “Mobile Information Management”

From October, 25th until 29th, 2004, a Dagstuhl seminar on mobile information management took place. The seminar was attended by 34 researchers from Europe and North America with backgrounds ranging from database systems, mobile information systems, geographic information systems and business informatics to wireless networks and security.

The goal of the seminar was to bring together researchers, especially from the area of databases and information systems, to identify open problems and new challenges in data, service, and user management in mobile information processing environments. We are particularly interested to provide a forum for discussing the consequences of the mobility of users and devices on today’s and future data management systems. The first aim of this discussions was to clearly identify these consequences. The second and more challenging aim was to determine where existing solutions can be applied, where mobility raise truly new challenges, and which of these challenges are there to last.

Keywords: Moving objects and mobile users, mobile data dissemination and delivery, mobile data replication and synchronization, discovery and composition of mobile services, mobility awareness and adaptability, location-dependent, context-based querying services.

Joint work of: Dunham, Margaret; König-Ries, Birgitta; Pitoura, Evaggelia; Reiher, Reiher; Türker, Can

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2005/257>

04441 Working Group – Research Issues in Mobile Querying

This document reports on key aspects of the discussions conducted within the working group. In particular, the document aims to offer a structured and somewhat digested summary of the group’s discussions. The document first offers concepts that enable characterization of “mobile queries” as well as the types of systems that enable such queries. It explores the notion of context in mobile queries. The document ends with a few observations, mainly regarding challenges.

Keywords: Mobile Queries

Joint work of: Breunig, Martin; Jensen, Christian S.

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2005/165>

04441 Working Group – Towards a Handbook for User-Centred Mobile Application Design

Why do we have difficulties designing mobile apps? Is there a “Mobile RUP”?

Keywords: User-Centred Mobile Application Design

Joint work of: Boll, Susanne; Breunig, Martin; Davies, Nigel; Jensen, Christian S.; König-Ries, Birgitta; Malaka, Rainer; Matthes, Florian; Panayiotou, Christoforos; Saltenis, Simonas; Schwarz, Thomas

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2005/166>

04441 Working Group – Description and Matching of Services in Mobile Environments

Service oriented computing is a new paradigm that is especially interesting in mobile environments. As a characteristics, functionality is hidden behind an interface and described as a black box with the help of a service description language. This enables participants of the network to enlarge the limited capabilities of their devices by using services provided by others. As service requestors and providers are not fixedly tied together but are dynamically matched and bound, this architecture is especially advantageous in mobile environments and their constantly changing situation.

Keywords: Service Description, Mobile Environment

Joint work of: Grünbauer, Johannes; Klein, Michael; Koloniari, Georgia; Samaras, George; Türker, Can

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2005/167>

04441 Working Group – Research Issues in Mobile Transactions

This document discusses three scenarios for databases with mobile clients, summarizes typical applications and requirements for each of the three scenarios, and outlines the open research issues which should be solved within each of the three scenarios. While the first scenario consists of mobile clients that are connect to a wired network, the second scenario consists of a network of mobile clients with a single-hop distance to each other but without a wired network, and the third scenario considers a network of mobile clients some of which are in multi-hop distance.

Keywords: Transactions, mobile clients, multi-hop wireless networks

Joint work of: Böse, Joos-Hendrik; Böttcher, Stefan; Gruenwald, Le; Pitoura, Evaggelia; Reiher, Peter; Samaras, George; Schwarz, Thomas; Türker, Can

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2005/168>

04441 Working Group – Some Open Aspects of Mobile Ad-hoc NETWORK, Peer-to-Peer, and Self-organizing Systems

This document summarizes the results of the working group discussion on Mobile Ad-hoc NETWORKS (MANETs), peer-to-peer (p2p) and self-organizing systems held at the workshop on Mobile Information Management in October 2004 in Dagstuhl, Germany. The goals of the discussion group and of this document were to study the characteristics of these systems and to identify some of the open research issues that need to be solved in order for such systems to be implemented successfully.

Keywords: MANET, P2P

Joint work of: Böse, Joos-Hendrik; Böttcher, Stefan; Gruenwald, Le; Marrón, Pedro José; Obreiter, Philipp; Pitoura, Evaggelia; Reiher, Peter; Sattler, Kai-Uwe; Seliger, Frank

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2005/217>

04441 Working Group – Business Models

This document presents the results from the working group “Business Models” (WG 1) discussing the seminar topic under market aspects. Main objectives were developing a shared view of the field (i.e. matching business, technology and research interests), identifying key players and relations between them (including barriers and lessons learned so far from the 2G/2.5G market), discussing the role of business models and value chains and applying them to P2P-networks. In the concluding round the difference between traditional DB/IS applications and mobile applications is explored.

Keywords: Business Models

Joint work of: König-Ries, Birgitta; Lehner, Franz; Malaka, Rainer; Matthes, Florian; Obreiter, Philipp; Pousttchi, Key; Seliger, Frank; Turowski, Klaus; Veijalainen, Jaro

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2005/169>

Context-aware multimodal mobile multimedia applications

Susanne Boll (Universität Oldenburg)

The short introduction presents our ongoing research in the field of user-centered mobile multimedia applications and systems. Starting point is our Niccimon platform for the flexible and rapid development of location-aware mobile applications. The platform features a modular system design which provides application-independent modules that provide functionality for tasks typically needed in mobile applications such as location-sensing, location-visualization, presentation of points of interest, and multimodality. A first sample application – a mobile multimedia tourist guide – illustrates how these modules are used to develop mobile applications. Introducing personalization, especially in the field of mobile multimedia content, we developed the MM4U framework for dynamic creation of personalized content in the heterogeneous mobile world. A personalized mobile tourist guide application illustrates the framework’s usage. Adding multimodality to the system, an auditory user interface supports acoustic navigation and orientation support. Our mobile location-aware mobile paper chase demonstrates the usage of an auditory user interface in a game scenario. The concept of context-aware Points of Interest are a further step towards the development of innovative location- and context-aware mobile applications.

Keywords: Mobile, multimedia, context, multimodal

Overview of research interests: Mobile access to geoservices and 3D/4D Information Systems

Martin Breunig (Hochschule Vechta)

In geoscientific modeling moving geo-referenced objects are involved such as moving rocks during landslides. The analysis of movements should be efficiently supported by mobile data acquisition in the field and by mobile database support. The challenge in the future is data access and analysis for 3D/4D geoinformation systems from mobile client applications.

Keywords: Mobile data access, geoservices, 3D/4D GIS

Transactions in mobile ad-hoc networks

Joos-Hendrik Böse (FU Berlin)

Business transactions between mobile clients e.g. the exchange or selling of ring tones or music files demands for guaranteed atomicity like money atomicity and goods atomicity. While in fixed networks atomicity guarantees can be achieved using a central coordinator, in a mobile network a single coordinator is not sufficient, because due to node movement every participant of a transaction can disappear at any time. To allow nodes to execute transactions in an atomic manner log data about the state of transactions must be available for all participants, even if they are not available at the same time.

This talk provides an introduction to our approach to support atomicity in mobile ad-hoc networks.

We propose to replicate log data over a set of nodes in a mobile ad-hoc network and thus provide a virtual shared log space (SLS). The SLS is used to make the state of a running transaction available for involved nodes even if one or more participants of the transaction are temporally not available.

Keywords: Transactions, atomicity, MANETs

Optimized XML Data Management for Mobile Transactions

Stefan Böttcher (Universität Paderborn, D)

This paper summarizes the research issues in the area of XML data management and mobile transactions performed at the University of Paderborn.

We shortly describe the problems worked at and the basic solution ideas, and we refer to the further literature that provides detailed descriptions of the elaborated solutions.

Keywords: XML, query optimization, mobile transactions

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2005/268>

Distributed systems support for mobile and ubiquitous computing

Adrian Friday (Lancaster University)

In this brief talk I present the work I'm doing in distributed systems for deployable Ubiquitous computing, within the context of the Equator project. Equator is building mixed reality, mobile & ubicomp experiences. Experience suggests that building such things is hard, requiring many man-months of work to create and orchestrate behind the scenes. To make Ubiomp truly ubiquitous, toolkits are required offering appropriate abstractions for authoring such environments. An early attempt in this direction is the Equip shared dataspace platform and Equator Component Toolkit (ECT).

Try ECT for yourself!

<http://www.crg.cs.nott.ac.uk/~cmg/Equator/Downloads/docs/ect/webstart/index.html>

Keywords: Distributed systems, middleware, ubicomp, equip, ect, equator

Managing Transactions in Mobile Ad-Hoc Network Databases

Le Gruenwald (University of Oklahoma - Norman)

In a mobile ad-hoc network (MANET), mobile hosts can move freely and communicate with each other directly through a wireless medium without the existence of a fixed wired infrastructure. MANET is typically used in battlefields and disaster recovery situations where it is not feasible to have a fixed network. Techniques that manage database transactions in MANET need to address additional issues such as host mobility, energy limitation and real-time constraints. This talk provides an overview of our solutions for transaction management, data caching and data replication that reduce the number of transactions missing deadlines while conserving and balancing the energy consumption by the mobile hosts in the system.

This work is partly supported by National Science Foundation Grants EIA-9973465 and IIS-0312746

Keywords: Transaction management, data caching, data replication, mobile ad-hoc networks, mobile databases

Adaptive Services in the Automotive Environment

Johannes Grünbauer (TU München)

In this short presentation a brief overview on some challenges in the automotive engineering domain is given. Since a car can be seen as a “big mobile device”, most of the problems of the common mobile devices exists here, too. One special challenge is that a car is a distributed systems for itself. A lot of functionality is distributed over up to 80 different electronic control units connected via several bus systems.

Together with the possibility to communicate with external systems using for example a UMTS connection, the mobile device “car” is getting rather complex. One has to deal with a multi-modal human machine communication, context awareness and security issues on different levels. So, there arises a strong need to use a service oriented development of software running in a car. This and the challenges mentioned above are addressed within the project MEwaDis at the Technische Universität München (<http://www4.in.tum.de/~mewadis>), in which research on the modelling of adaptive services is done.

Keywords: Service-based Software Engineering, Context, Adaptivity, Automotive

Efficient, self-contained handling of identity in peer-to-peer systems

Manfred Hauswirth (EPFL – Lausanne)

Identification is an essential building block for many services in distributed information systems. The quality and purpose of identification may differ, but the basic underlying problem is always to bind a set of attributes to an identifier in a unique and deterministic way. Name/directory services such as DNS, X.500, or UDDI are a well-established concept to address this problem in distributed information systems. However, none of these services addresses the specific requirements of peer-to-peer systems with respect to dynamism, decentralization and maintenance. We propose the implementation of directories using a structured peer-to-peer overlay network and apply this approach to support self-contained maintenance of routing tables with dynamic IP addresses in structured P2P systems. Thus we can keep routing tables intact without affecting the organization of the overlay networks, making it logically independent of the underlying network infrastructure. Even though the directory is self-referential, since it uses its own service to maintain itself, we show that it is robust due to a self-healing capability. For security we apply a combination of PGP-like public key distribution and a quorum-based query scheme. We describe the algorithm as implemented in the P-Grid P2P lookup system (<http://www.p-grid.org/>) and give a detailed analysis and simulation results demonstrating the efficiency and robustness of our approach.

Keywords: Peer-to-peer systems; Identity Handling; Self-Contained, Self-Maintaining, Decentralized Directory Service; Distributed Hash Tables, Dynamic Resilience

Joint work of: Aberer, Karl; Datta, Anwitaman; Hauswirth, Manfred

Decentralized Information Management

Manfred Hauswirth (EPFL - Lausanne)

Management of information in decentralized information systems has many connections to mobile information management: The participants are autonomous and unreliable, the communication network is unreliable and insecure, global knowledge and coordination cannot be assumed and scalability and robustness are key design issues. At EPFL we are trying to address these problems from the angle of peer-to-peer (P2P) systems. We think that many of the solutions coming from this domain can also be applied to mobile information management. In particular we are interested in building P2P infrastructures (our DHT-based P2P system P-Grid, <http://www.p-grid.org/>), exploring self-organization as an organizational principle to achieve scalability and robustness, incrementally achieving global semantic agreements through local interactions (semantic gossiping), semantic routing, large-scale information integration, semantic web services, distributed information retrieval and ranking (“decentralized Google”), and decentralized trust and reputation management.

Keywords: Decentralization, peer-to-peer, self-organization, trust and reputation management

Joint work of: Hauswirth, Manfred; Aberer Karl

Mobile services - moving objects

Christian S. Jensen (Aalborg University)

In this presentation, I gave only a very brief overview of most of the topics that I have been working on that relate to the topic of the seminar. These include: acquisition and subsequent provisioning of geo-related context data, e.g., destinations and routes of mobile users along with usage meta-data; query and update efficient indexing of the past positions and the current and anticipated future positions of moving objects; query processing for moving objects, including the processing of nearest and reverse nearest neighbor queries, skyline queries, and so-called in-route queries where the objects are embedded in Euclidean space or transportation networks.

In order for the presentation to not be entirely abstract, I covered a single topic, tracking of moving objects, in a bit of detail.

Specifically, in step with the continued advances in wireless communications, geo-positioning, and consumer electronics in general, an infrastructure is emerging that enables services and applications that rely on the tracking of the continuously changing positions of entire populations of mobile-service users. This scenario is characterized by large volumes of updates, for which reason location update technologies become important.

I introduced an approach to tracking that offers precision guarantees for the positions of the moving objects that are stored on the server side. In this approach, the server and each mobile client share a representation of the client's changing current position, and each client issues an update to the server when its position as given by this shared representation deviates from its GPS position by more than a certain threshold.

A key issue is how to represent the changing position of a moving object so that tracking can be done with as few updates as possible.

A range of tracking techniques were described that utilize progressively more information about an object's movement, ranging from assuming that an object remains in its most recently reported position to assuming that past movement trajectories of an object through a road network are known. I briefly covered experimental evaluations of the tracking techniques based on real GPS data and a corresponding real road network.

Keywords: Mobile services, moving objects, tracking, location-related context, indexing, query processing, update processing

Semantic Service Descriptions to Automize Service Oriented Computing

Michael Klein (Universität Karlsruhe)

One major challenge in service oriented computing is an automated service binding. This goal can only be achieved with the help of an appropriate service description language that goes beyond a simple interface description. In our work, we present such a language, the DIANE service description (DSD). It integrates additional information into the description that is necessary to automate the service usage process. First, we integrate unique and structured real world semantics into the description by using layered ontologies. Second, we integrate the configuration semantics into the service description, i.e. we show how the inputs and preconditions influence the result. This is done by using a purely state oriented service description where only state changes are described and an explicit description of the exchanged messages is omitted. We use variables to provide information how the state can be configured. Third, we integrate preference semantics into request descriptions, i.e. we enable the requestor to integrate all of his preferences of the desired service into the description. This allows us to generate a personal matcher for each request.

Keywords: Semantic Service Description

Languages for (Semantic) Service Description

Michael Klein (Universität Karlsruhe)

Service oriented computing (SOC) is a promising programming paradigm to create robust, context-aware applications in mobile environments independent from the underlying network. A main task here is to automate the complete process from service search over matching, selecting and invocation. The most important technique to achieve this is an appropriate description of the services. Therefore, in this talk, an overview over existing syntactic (like WSDL) and semantic (like OWL-S, WSMO, and DSD) service descriptions is given.

Keywords: Semantic Service Description, Semantic Web, Web Service, WSDL, OWL-S, WSMO, DSD

Joint work of: Klein, Michael; König-Ries, Birgitta

Mobile Peer-to-Peer

Georgia Koloniari (University of Ioannina)

Peer-to-peer systems are gaining increasing popularity as a scalable means to share data among a large number of autonomous nodes. Since the shared data are unstructured and they do not follow a global schema, XML-based descriptions of the data can be used to provide a uniform way to query the heterogeneous data. In our research, we are interested in designing a fully decentralized approach for the problem of efficiently routing path queries among the nodes of a peer-to-peer system. Our approach is based on (a) selecting and maintaining specialized data structures, called filters that efficiently summarize the content, i.e., the documents, of one or more node and (b) using these filters to build an overlay network that groups together nodes with similar content.

Keywords: Mobile Peer-to-Peer

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2005/222>

Research Overview: Resource Sharing in Dynamic Environments

Birgitta König-Ries (Universität Karlsruhe)

My work deals with resource sharing in dynamic environments. Within the context of this seminar, one specific instantiation of this research area is of interest: The sharing of data (or information) in mobile environments.

We have adopted the service oriented computing paradigm as the basis of our work. That is, we view resources as services that are offered and requested. Within this framework, we are addressing the following questions: 1. How can resource sharing be automated, i.e., how can service offers and requests be described in such a way, that a machine is able to determine whether a given offer matches a request? We develop DSD, the DIANE Service Description language along with appropriate matching methods, to meet this goal. 2. In the absence of infrastructural support, how can service offers and requests be brought together efficiently? Here, we look at different types of (semantic) overlays that are tailored to service matching and take the characteristics of the underlying networks, i.e. high dynamics, movement of nodes, etc., into consideration. 3. How can autonomous entities be stimulated to cooperate? Clearly, some kind of incentive scheme is needed here. We are particularly interested in evidence-aware reputation systems, an approach that works even in dynamic environments with no trusted third party.

More information about my work can be found here:
<http://www.informatik.uni-jena.de/~koenig>

Keywords: Resource Sharing, Dynamic Environments

Mobile Business and Services - A systematic, business oriented approach

Franz Lehner (Universität Passau)

In terms of business models, mobile commerce (MC) includes relations between customers, operators, ecommerce providers, payment providers, and other parties. There is a growing number of publications on MC sometimes labeled as research. But most of this work should be better called market studies or benchmarks. They highlight different aspects such as the number of prospective consumers, the estimated market volume, etc which do not really provide deeper insights. Well known sources for this kind of information are Arthur D. Little, Ovum, Durlacher, Merrill Lynch, Gartner, Forrester and others. They all try to estimate the perspectives of the mobile market(s) by using figures on mobile phone penetration, SMS usage, mobile internet access etc. Many of the remaining publications are technically oriented (for an overview see e.g. Prasad et al. 2000, Muller-Veerse 2000, Webb 1999). So what we are generally missing are

concepts and guidelines for developing the mobile business, which something the MIS field can undertake. It is the purpose of this paper to contribute to a solution by providing a structured view of the major areas from a business point of view.

Keywords: Mobile Business

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2005/163>

Mobile Business and Services

Franz Lehner (Universität Passau)

Mobile Business, which represents the evolution of electronic business, encompasses the use of a variety of devices, particularly PDAs and mobile phones, to further business functions. The dynamic development in this field and the increasing integration in the internet have opened up an entirely new research and development field whose market potential makes it a relevant field. Our research activities are focused on business aspects and the development of business models as well as exploring the use of mobile technologies in new application areas (among them e-learning and knowledge management). The motivation for this effort is the rapidly increasing importance of mobile technologies in general and their relevance for other applications.

Within this field we are especially interested in P2P-technologies which can be described as a new paradigm within electronic communication. As the economic aspects of peer-to-peer systems have received little attention so far we try to investigate the business context of P2P application. Due to the success and popularity of applications like Napster, KaZaa and others a lot of ideas were created how to use P2P networks within an organisation as well as an instrument for electronic business deployment. Both areas can lead to new market opportunities and also to the foundation of new firms. As a mean of investigation we are about to develop a P2P application supporting eLearning activities; this application will hopefully bring first hand insights regarding communication mechanisms within P2P networks as well as technical and commercial aspects. The findings are intended to serve as basis for creating business models, decision making (eg. substitution of existing applications) or just for improving current work processes.

Robust Data Management

Wolfgang Lehner (TU Dresden)

Mobility gains more and more importance from a technological as well as social perspective. On the one hand, mobility is required from the personal and professional environment in order to keep pace with the developments in a global world. On the other hand, the existence of wireless networks and the success of cell-phones enable a wide usage of mobile communication infrastructure. While mobile devices (especially cell-phone) are becoming more and more a general vehicle to perform a wide spectrum of applications like internet browsing, etc. many, many issues are still unsolved in order to provide a technologically solid and well accepted mobile infrastructure. In the following, we focus on the term of robustness as a mean to achieve this goal. No only the general possibility to communicate via mobile devices using wireless networks is the question, but the reliable, secure, and finally simple way of doing it must be the core research in the context of mobile environments.

Keywords: Robustness, Reliability, Partial System Failure, Linux, trusted platform

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2005/164>

Mobile and Pervasive Information for All

Rainer Malaka (Europ. Media Lab. - Heidelberg)

Mobile data management plays a role in multiple projects at the European Media Lab. In a number of research projects we focus on mobile users and ubiquitous and seamless information access. In some application scenarios, tourists or visitors of a city are to be supplied with location-based services. Agent-based infrastructures orchestrate services in CRUMPET. Ontologies describe the services and can also facilitate intuitive user interaction. SmartKom is an example of such a system where multimodal interaction and semantic service integration are combined. In the home environment where ensembles of devices such as VCRs and TV sets should interact, we work on self organizing strategies that help to coordinate functions and presentations through multiple channels, modalities, and devices.

Ubiquitous Computing

Rainer Malaka (Europ. Media Lab. - Heidelberg)

Ubiquitous computing is an important area in computer science and engineering today. It is based on a number of technological trends and there are a number of related and similar definitions. In the recent years a variety of applications have been built and some of them are already leaving the research labs, entering the mass market. This talk discusses some of the trends in ubiquitous computing, presents some sample applications and projects, and will introduce issues for research and further discussions.

An Adaptive Cross-Layer Framework for Sensor Networks

Pedro José Marrón (Universität Stuttgart)

An intrinsic characteristic of current projects in the area of sensor networks is the heterogeneity of hardware and application requirements. In addition, the requirements of current applications are expected to change over time and become more and more heterogeneous. This makes developing, deploying, and optimizing sensor network applications an extremely difficult task, especially if they are to be deployed in hybrid environments where the capabilities of the available network infrastructure needs to be taken into account. In this talk, we present the architecture of TinyCubus, a flexible and adaptive cross-layer framework for TinyOS-based sensor networks that aims at providing the necessary system support to cope with the complexity of such systems. TinyCubus consists of three parts: a data management framework that selects and adapts both, system and data management components, a cross-layer framework that enables optimizations through cross-layer interactions, and a configuration engine that installs components dynamically. At the end, we demonstrate the feasibility of our architecture by describing and showing some evaluation results from a code distribution algorithm that optimizes its behavior by using application knowledge about the sensor topology.

Keywords: Sensor networks, adaptive systems, system support, generic framework

Mobile Information Management: Research Background & Research Interests

Florian Matthes (TU München)

In this 5-minute presentation I describe the research background and research interests of our group at TU Munich related to mobile information management.

We are particularly interested in software architectures for personal information management.

Keywords: Personal Information Management, Content Brokerage Solutions, Software Architectures

Time & Experience Based Personalization for the Moving User

Christoforos Panayiotou (University of Cyprus)

The needs of the wireless and mobile user regarding information access and services are quite different than those of the desktop user. This need is not about browsing the Web but about receiving personalized services that are highly sensitive to the immediate environment and requirements of the user. Personalization appears to be the most appropriate solution to this need. It comes into aid by creating personalized portals that are specific for the wireless user, which (a) are focus on the local content and (b) directly tones down factors that break up the functionality of the Internet/wireless services when viewed through wireless devices; factors like the “click count”, user response time (the “choice” factor) and the size of the wireless network traffic.

Furthermore, timing can be crucial for the moving user. E.g. it is lunch time and our user is hungry. He should first navigate through all the available services, find the restaurant services and then get the desired information. If the system took the timing of the user into consideration it could alter the order of the provided results in order to display first the restaurant services, and in doing so, it upgrades the user perceived quality of service. However timing isn't always as easily exploitable and so the question is how exactly we may exploit this information. Extending this idea, we have the notion of user experience where his interests change according his current activity (e.g. on vacation or at work?)

Timing and user experience are two factors which are quite significant. These are new factors to the personalization problem and introduced mostly because the needs of the moving users are not limited during the time he is in front of his office PC but around the clock all the year long, including weekends and vacations.

Joint work of: Panayiotou, Christoforos; Samaras George

Data Management for Mobile and Peer-to-Peer Systems

Evaggelia Pitoura (University of Ioannina)

This 5-min talk provides a short overview of my research work and interests in (a) consistency management and (b) resource location in mobile and (more recently) peer-to-peer systems.

In terms of consistency management, our work focuses on replication control. In mobile computing, connectivity limitations and the need to support

disconnected operations motivated the development of flexible replication control schemes that allow weak data operations to run at mobile handheld devices [4]. Weak reads access locally-available potentially out-of-date copies, while weak writes are only tentatively committed. To take advantage of the physical support for broadcast schemes [5]. Our more recent work considers consistency issues in p2p systems [2].

In terms of resource location, our research interests are in developing items in mobile computing [6] and in p2p systems [1,3].

References:

1. G. Koloniari, E. Pitoura: Content-Based Routing of Path Queries in Peer-to-Peer Systems. In EDBT 2004
2. E. Leontiadis, V. V. Dimakopoulos, E. Pitoura. Cache Updates in a Peer-to-Peer Network of Mobile Agents. In Peer-to-Peer Computing 2004
3. E. Pitoura, S. Abiteboul, D. Pfoser, G. Samaras, M. Vazirgiannis. DBGlobe: a Service-oriented P2P System for Global Computing. SIGMOD Record 32(3), 2003
4. E. Pitoura, B. K. Bhargava. Data Consistency in Intermittently Connected Distributed Systems. IEEE Trans. Knowl. Data Eng. 11(6), 1999
5. E. Pitoura, P. K. Chrysanthis. Exploting Versions for Handling Updates in Broadcast Disks. In VLDB 1999.
6. E. Pitoura, G. Samaras. Locating Objects in Mobile Coputing. IEEE Trans. Knowl. Data Eng. 13(4), 2001

Keywords: Mobile computing, peer-to-peer computing, consistency, replication, resource management

Mobile, Ad-hoc, Peer-to-Peer Self Organizing Systems: Data Engineering Issues

Evaggelia Pitoura (University of Ioannina)

This talk presents an overview of data management research in the area of mobile and peer-to-peer (p2p) computing as the underlying distributed computing models for ubiquitous systems. The first part of the talk surveys results in mobile information systems related to handling challenges related to mobile computing, in particular (a) physical mobility, (b) network connectivity, and (c) resource constraints. In terms of mobility, various distributed database architectures have been proposed for locating mobile objects and novel data models and index structures have been advanced in the area of moving object databases. In terms of network connectivity, much research has been conducted in the area of caching, replication and transaction management for supporting disconnected operation and intermittent connectivity. In terms of resource constraints, multi-tier models have been proposed for offloading functionality from the mobile devices to the stationary network and for building adaptable and resource-aware information

systems. The second part of the talk is a short introduction to p2p computing with emphasis on overlay networks and resource discovery. Some conclusions are drawn relating mobile and p2p research. The talk concludes with a short presentation of DBGlobe, our research project that builds on the peer-to-peer paradigm to create a data management system for querying data located on small mobile devices.

References:

1. E. Pitoura, S. Abiteboul, D. Pfoser, G. Samaras, M. Vazirgianis. DBGlobe: a Service-oriented P2P System for Global Computing. SIGMOD Record 32(3), 2003

Keywords: Mobile computing, peer-to-peer computing, data management

Security for Wireless Networks and Mobile Environments

Peter Reiher (Univ. California - Los Angeles)

This presentation describes general network security problems that beset mobile and wireless environments, as well as the particular problems introduced by the special characteristics of these environments.

Keywords: Security mobile computing wireless networks

Security Issues for Mobile Ubiquitous Environments

Peter Reiher (Univ. California - Los Angeles)

A short overview presentation on some issues of security in mobile ubiquitous environments that I am researching.

Short presentation of research interests

Simonas Saltenis (Aalborg University)

Advances in mobile communication technologies and positioning technologies enable the so-called location based services (LBS)—mobile internet services, where the geographical location of a user (or other objects) plays an essential part. Efficiently tracking, storing, updating, and querying continuously changing positions in a database system are challenges addressed by our research. Specifically data structures (indices) and query processing algorithms are explored in this novel setting.

While in traditional database systems data is assumed to be constant until it is explicitly updated, tracking the continuously changing positions of mobile objects with this assumption would lead to unacceptably high rates of updates if

the positional data were kept up to date. For this reason, positions of objects are represented as functions of time. This reduces the rate of updates and enables predictive near-future queries. Traditional spatial indices are adapted to support storing positions of objects as functions of time.

Remaining challenges involve making indices more update efficient, enabling efficient accumulating and querying of movement history, enabling efficient indexing of the positions of objects which move in road networks.

Personalization Portals for Wireless Users

George Samaras (University of Cyprus)

One of today's trends in modern computing is the enablement and support of moving users. Unlike desktop users, who have been around for a while, mobile users are a new and more demanding breed.

Technology provided for the first group are often found lacking for the later one. Personalization is such an example. The moving user differs from the desktop user in that his handheld device is truly personal. It roams with the user and allows him access to info and services at any given time from anywhere. As the moving user is not limited to a fix place and to a given time period factors such as time and current experience becomes increasingly important for him.

His context is now a function of time and experience and the goal of personalization is to match the local services to this context. In this work we exploit the importance of time and experience in personalization for the moving user and present a system that anticipates and compensates the time-dependant shifting of user interests. We clearly show that exploiting time and experience results in an enhanced and more effective personalization system.

Joint work of: Samaras, George; Christoforos Panayiotou

Scalable Infrastructures for Mobile Data

Kai-Uwe Sattler (TU Ilmenau)

Infrastructures for mobile data require scalable, distributed, time-shared storages. Structured P2P systems in the form of distributed hash tables (DHT) are a promising approach for building such platforms. However, for many applications the supported key lookup queries are not sufficient. Instead, techniques for managing and querying (relational) structured data are required. In this talk, we give a brief overview on data fragmentation and query processing techniques for DHTs.

Keywords: Mobile data, P2P systems, DHT

Digital World Models for Mobile Context-based Systems

Thomas Schwarz (Universität Stuttgart)

The Nexus project investigates methods to manage and use federated spatio-temporal models comprising of many local context models offered by independent data providers. The federation component provides a layer of abstraction and transparency, so that many different applications can benefit from a combined and consolidated data pool.

Keywords: Nexus, federated context models

Pervasive Computing at IBM

Frank Seliger (IBM - Böblingen)

Working in IBM Development as the Security Architect of "Pervasive Computing" Development team, I am a consumer of research rather than a producer. For a large part I can draw for my work on IBM's own Security Research that happens in the Watson, Zurich, Almaden, Tokyo, China and Haifa Research lab. IBM provides server and client side products explicitly supporting Pervasive Computing since 1998. These products address the Business to Consumer (telecommunication companies) as well as the Business to Employee (corporate) use cases. Consequently, the security challenges range from securing J2EE resources to protecting the owner's information on a client to protecting an application provider's services executing on a client.

Processing of Ontologies in Mobile Environments

Günther Specht (Universität Ulm)

Mobile information systems will considerably profit from the use ontologies as previously incoherent data is put into relationship for the users benefit. In this presentations we discuss three possible architectures for the mobile processing of ontologies. We show that despite todays weaknesses an approach based on mobile database systems will be preferable in the future. As a foundation we present and compare ways to convert ontologies into logic and relational database programs. Additionally when using ontologies on mobile clients interesting new requirements for replication and synchronisation have to be explored. We show that the synchronisation mechanism has to be extended with unification.

in german:

Mobile Informationssysteme können erheblich von Ontologien profitieren, indem zuvor zusammenhanglose Fakten für den Anwender in Beziehung gebracht werden. In dieser Arbeit untersuchen wir drei mögliche Architekturvarianten zur mobilen Verarbeitung von Ontologien. Wir zeigen, dass ein Ansatz,

beruhend auf einem mobilen Datenbanksystem, trotz seiner heutigen Limitationen in Zukunft zu bevorzugen sein wird. Bei der Verarbeitung von Ontologien auf mobilen Clients treten darüber hinaus interessante neue Anforderungen an Replikations- und Synchronisationsmechanismen auf, die um Unifikation erweitert werden müssen.

Mobile Transactions and Synchronization in Mobile Environments

and

Completely Distributed Transaction Processing

Can Türker (ETH Zürich)

Information access and processing in mobile environments requires a flexible and powerful transaction processing infrastructure that takes into account the peculiarities of mobile computing. Mobile transaction processing shall be able to transparently deal with frequent disconnections and occasional movements of mobile devices. This talk outlines state-of-the-art approaches to transaction processing in mobile databases.

Energy Consumption Tradeoffs for Compressed Wireless Data at a Mobile Terminal

Jari Veijalainen (University of Jyväskylä)

The high-end telecom terminal and PDAs, sometimes called Personal Trusted Devices (PTDs) are programmable, have tens of megabytes memory, and rather fast processors. In this paper we analyze, when it is energy-efficient to transfer application data compressed over the downlink and then decompress it at the terminal or compress it first at the terminal and then send it compressed over up-link. These questions are meaningful in the context of usual application code or data and streams that are stored before presentation and require lossless compression methods to be used. We deduce an analytical model and assess the model parameters based on experiments in 2G (GSM) and 3G (FOMA) network. The results indicate that if the reduction through compression in size of the file to be downloaded is higher than four per cent, energy is saved as compared to receiving the file uncompressed. For the upload case even two percent reduction in size is enough for energy savings at the terminal with the current transmission speeds and observed energy parameters. If time is saved using compressed files during transmission, then energy is certainly saved. From energy savings at the terminal we cannot deduce time savings, however. Energy and time consumed

at the server for compression/decompression is considered negligible in this context and ignored. The same holds for the base stations and other fixed telecom infrastructure components.

Keywords: Personal trusted device, energy consumption, data compression, wireless data transmission

On the development of global mobile environment

Jari Veijalainen (University of Jyväskylä)

The global mobile environment consists of hundreds of mobile access networks based on telecom air interfaces (GSM(+), PDC, W-CDMA, CDMA(2000), etc.) that facilitate well over one billion subscribers throughout the world. These digital communication systems and other global networks, like Internet, are designed in a cooperative manner by a Global Designer, an organization (such as ISO, OMA, ITU, ANSI, etc) that is responsible for certain technical area or geographic region. The companies and other actors participating in the work are often competing and cooperating simultaneously. The technical solutions agreed upon by a GD are results of a political process where the power constellation is more important than technically optimal solutions. The component systems of a cooperative global system exhibit usually various aspects of autonomy, such as design, execution, and communication autonomy, the latter being especially typical of mobile terminals.

In the global environment there seem to be four relatively independent spheres of concerns that have their own development laws, Regulatory Frameworks (laws, standards, recommendations+ the organizations issuing them), Business Models (strategies, BMs of various actors), Global Infrastructure (the wireless and wireline networks, terminals and mobile applications deployed at a certain point of time in the world), and Enabling Technologies (those being developed in the laboratories, but not yet deployed). The dynamics of the spheres is of utmost importance for understanding the future trends. Perhaps the most startling effects will come from the convergence of already existing networking technologies. What are the barriers for the advances in mobile applications and what are the disruptive technologies that might challenge the telecom technology and operators?

Keywords: Global network infrastructure, autonomy, disruptive technologies

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2005/221>

Mobile Peer-to-Peer Networks

Ouri Wolfson (Univ. of Illinois - Chicago)

In this talk we examine the dissemination of reports about resources in mobile peer-to-peer networks, where moving objects and sensors communicate with each other via short-range wireless transmission. Each disseminated report represents information about a spatial-temporal resource. Such a resource is available at a particular time and location, and is relevant to clients in the geographic vicinity, possibly for a limited period of time. Applications involving these types of resources occur in transportation, battlefield and emergency response situations, and social networks.

We propose an opportunistic dissemination paradigm, in which a moving object exchanges resources with encountered moving objects. We analyze several variants of this paradigm, and incentive economic models for collaboration.

The subject is related to peer-to-peer networks, resource discovery in mobile ad hoc networks, epidemic replication and routing, and publish/subscribe literature.

Keywords: Peer-to-peer, mobile ad hoc networks, publish/subscribe

Query Processing in (Mobile) P2P Networks

André Zeitz (Universität Rostock)

Researching in the field of Peer-to-Peer systems has become very popular in the last time. From the point of view of database researchers there is an important issue in this field that is missing at the moment: the query processing as it is known from database systems. Therefore we are interested in finding ways to realize basic concepts like managing and processing complex queries, a cost model for basic search operations, the facility to optimize queries or aggregate data. Especially in mobile environments data can be incomplete. In order to answer queries for such incomplete data we are interested in finding ways to answer those query in a reasonable manner. Another aspect we are interested in is to find algorithms in other research areas, e. g. psychology, demography or network theory, that can either be transformed for usage in P2P networks or that help us to understand different processes and finally to find new approaches for handling queries.

Keywords: Query Processing, Peer-to-Peer networks, Mobility