Issues in Data Management for Pervasive Environments

Sasikanth Avancha, Dipanjan Chakraborty, Harry Chen, Lalana Kagal,
Filip Perich, Timothy Finin, Anupam Joshi
Department of Computer Science & Electrical Engineering
University of Maryland, Baltimore County
1000 Hilltop Circle, Baltimore, MD
{savanc1, dchakr1, hchen4, lkagal1, fperic1, finin, joshi}@cs.umbc.edu

Abstract

A new vista is open for research in the area of data management in pervasive environments where mobile devices gather and exchange information from other devices based on their current context. Each device is both a data source and a data consumer pursuing its individual and collective tasks. Our research group is investigating various ideas aimed at realizing pervasive computing systems based on the cooperation of autonomous, dynamic and adaptive components located in vicinity of one another. These systems will enable devices to discover information, compose it, and analyze it for consistency, all in a secure and trusted manner.

I. Introduction

In the past year or two, the research community has seen plenty of hype associated with wireless / pervasive / mobile / ubiquitous computing. The ongoing enhancements of handheld, embedded and wearable devices, together with the advent of pervasive connectivity and nanosensors, represent new paradigms for interaction among devices. Some of the pervasive connectivity is due to the infrastructure based on 2.5/3G cellular networks. In such environments, mobile devices are typically viewed as consumers of goods and information only. The information or goods come from servers on the wired side. Particularly in the m-commerce vision, cell phones or wirelessly connected PDAs became mobile storefronts for e-tailers – essentially an incremental change in the present e-tailing idea. We are all familiar with the ads of people buying goods via their cell-phones from the beach, or getting the latest traffic conditions from some server. In such systems the traditional client–proxy–server interaction is an appropriate model perhaps with the "client" database being extremely lightweight [4] or having a (partial) replicate of the main database on the wired side [27]. This approach has been developed by the academia over the last five–six years in contexts such as web access from mobile platforms (for instance[14], [17], [5], [18], [21], [3], [15]) or transaction support for database access(e.g. [11]). In some sense, one can think of this as a *supermarket approach*, where a few identified service providers exist and the service traffic is one–way.

An alternative approach will be necessary with the spread of short–range narrow-band systems such as Bluetooth [26], which allow for devices in a "vicinity" to spontaneously network with one another. Mobile devices will become autonomous, self-describing, highly interactive, and adaptive with respect to their environment. The hardware and software components will automatically become aware of each other, establish basic (wireless) communication, exchange information about their basic capabilities (e.g., services they can offer) and requirements (e.g., payments they need), discover and exchange "APIs", and learn to cooperate effectively to accomplish their individual and collective goals. The mobile devices will no longer be limited to consuming services and information only, but will also provide services/information themselves. Consequently, there will be no explicit clients and servers – but peers that can be both consumers and providers of different services and data. Mobile ad-hoc networks will thus enable the transformation of palmtops, embedded and wearable computers, into a truly pervasive environment that people use subconsciously.

This idea of "ad-hoc" teams of entities that are dynamically formed to pursue individual and collective goals can be used to create the software infrastructure needed by the next generation of mobile applications. These will use the emerging third and fourth generation broadband wireless systems, as well as short range narrowband systems such as Bluetooth. Heretofore, the software component of mobile computing has lagged behind its hardware (communication, computing and networking) aspects. Much of the research in the data management area is often limited to allowing applications built for the wired world (web, databases, etc.) to run in the wireless domain using proxy based approaches. Our research seeks to move beyond this and provide an architecture which uses the power of mobility and ad-hoc wireless connectivity to enable novel applications.

To better illustrate our research goals, let us consider the following scenario that the architecture will enable. It is 5:40 in the afternoon, and Bob's working day at his new job is just ending. As he is getting ready to leave the office, his phone rings. It is his new friend Jane asking him to meet her at the local shopping mall. Bob agrees to meet her, and notifies his palmtop about the decision. In addition, Bob asks the palmtop to find directions to the mall, as he has never been there before. While he is walking through the building toward the parking lot and ultimately toward his car, the palmtop is able to connect to the office network infrastructure and fetch the appropriate directions through a service/information broker. Once in the car, Bob reads and follows the instructions. However, he feels that the traffic is not moving fast enough; he would like to get to his destination quicker. He instructs his palmtop, which can now connect to the cars around him or passing him, to ask them whether they know about a faster route to the mall. The palmtop contacts its neighbors and returns with an alternative map, which is longer but which avoids the afternoon traffic jam that is building up on the current route. Bob therefore takes the different roads and arrives at the mall's entrance twenty minutes before the expected time. He decides to use the extra time by checking out the local stores to see if he can get a good deal on some small gift for Jane. It is forty minutes later and Jane finally arrives. After exchanging greetings, they decide to walk to a quiet place for dinner. Bob asks his palmtop to suggest available restaurants and lets Jane pick one. She chooses the closest Italian restaurant, which indicates it has an available table with no waiting period. Thus, they get seated immediately and spend several hours while eating and chatting. In the meantime, Bob's palmtop learns that it will stay in the given location for a while and it starts to autonomously interact with other devices in its vicinity. For instance, it could share traffic condition related information it might have cached while Bob was driving over to the mall. It could also obtain and cache business cards of other people in the mall by matching Bob's profile with theirs. As an other example, Bob's palmtop could receive an e-mail with Word attachment that it would like to print on the restaurant's printer, which accepts PS files only. Therefore, it needs to first discover one or more converter services before it can print it. After the meeting with Jane, Bob walks down to the garage and drives back to his apartment.

II. Architecture Overview

In the example above, we can see that both Bob's and Jane's palmtops heavily utilize the locally available resources as well as change their interaction modes based on the particular context to satisfy any implicit or explicit query that may be posed. Each palmtop utilizes the knowledge of its owner's interests and preferences (i.e., a profile containing facts and rules) that were either explicitly provided or that it has learned. It uses this knowledge to determine the appropriate actions it requires to satisfy a given request. Moreover, it also plans and executes these actions to obtain a maximum utilization.

Mobile devices, therefore, have to be able to discover other devices, their capabilities, and information that they provide based on the current context. For example, when Bob is driving on the highway, his palmtop attempts to discover traffic condition services provided by cars driving along. This requires service discovery based on rich description of requests and services. We have designed DReggie[7], [1] and enhanced the Bluetooth Service Discovery Protocol (SDP) [2] to enable service discovery using semantic-based matching. The available services may need to be composed into a larger service in order to satisfy a given request. As illustrated above, in order to print the Word attachment, the converter and printer services need to be composed. We are developing a system called Anamika that will enable the composition and execution of services in ad-hoc environment. The device's neighborhood changes dynamically. Thus, the services and information currently available may not be accessible when actually needed. In order to insure the required information is available to the user, the device must be able to plan and execute queries accordingly. In addition, it must be able to apply caching heuristics to answer future user queries even in the absence of information source in its vicinity. The MoGATU project [22] presents our preliminary solutions to these problems. Moreover, the mobile device must also be able deduce various context information to operate efficiently. Inferring context from numerous sensing devices may present challenges when sensors provide potentially contradictory evidence. We developing a mechanism that can fuse diverse evidence and reason over it to obtain context as a part of our architecture. Finally, every interaction among devices must be secure and all actions must only be performed by "trusted" entities. Our work on Vigil[16], an ongoing research project on distributed trust management, will provide a solution to security and trust problems in mobile pervasive environments.

Enhanced Bluetooth Service Discovery Protocol

The existing service discovery protocol (SDP) in Bluetooth is based on simple integer matching. Each service and its corresponding attributes are represented by Universally Unique Identifiers (UUID). This approach does not allow for rich description of services leading to an inefficient service discovery process for most future ad-hoc networks. This is because UUID based matching provides only *yes/no* type responses. Our enhanced SDP attempts to remedy these limitations by allowing inexact matching and responses. The principal components of the enhanced SDP system are service ontologies described in DARPA Agent Markup Language (DAML) [12] and a Prolog-based reasoning engine. Services and service requests are also described in DAML and they must follow the service ontology. *Semantic information* from the service ontology and various service instances are loaded into the Prolog knowledge base. Intensional knowledge in the form of Horn clauses is used along with extensional knowledge to determine closest possible matches to a given service request.

Anamika: A simple Reactive Service composition system over Bluetooth

Project Anamika addresses the problem of service composition in a mobile ad-hoc environment. The traditional service composition architectures [6], [10], [23] are built over the fixed network infrastructures. They assume the presence of a centralized, highly available manager that handles complex queries requiring multiple services to be executed in a certain manner. This assumption does not hold in an ad-hoc environment. In addition, the central broker-based model cannot take advantage of the changing network topology and resource variations of currently available mobile devices efficiently. In this project, we are designing a distributed architecture to facilitate service composition in an ad-hoc environment. We are using DAML-S[9] to describe composite services and process models in a semantically rich manner. Our architecture takes advantage of the varying resources in an ad-hoc environment and topology to determine a platform that is most apt to carry out the integration and execution of composite services. The components or services that participate to build a composite service are assumed to be distributed across the mobile nodes. Thus multiple nodes may carry out service compositions in such an environment, making the architecture immune to central point of failure. Current implementation of our system is based on Bluetooth using IBM's Bluedrekar driver.

Context-Aware Knowledge Integration and Inconsistency Management

The ability to use implicit situational information, or context, is crucial to the development of "intelligent" mobile applications. Instead of relying on the explicit user inputs, context-aware applications are capable of providing tailored services and information by exploiting user context (e.g., location context, social context and computing context). For applications to perceive context, sensors are used to acquire various contextual information in the environment [24], and inference procedures are employed to aggregate contextual knowledge from raw data [20], [8]. Contextual information is inherently distributed and heterogeneous. Individual context-aware applications, in particular ones that operate under resource-limited mobile devices, have limited capability to acquire contextual information and to reason about context. To overcome these problems, we are prototyping an agent-based knowledge management system as a middleware infrastructure to enable context-awareness in mobile applications. This will provide support for a knowledge integration, knowledge consistency, and spatial and temporal reasoning.

MoGATU: Serendipitous Query Routing and Processing

Existing mobile information access systems require the support of wired infrastructure, thus restricting the flexibility of information exchange among peer mobile devices. In *ad-hoc* environments, additional challenges – described in [22] – to the distributed database framework arise. We have designed and implemented a framework prototype, MoGATU, that addresses these challenges and issues through the use of DAML-S. The primary component of our framework is *InforMa* – a powerful information manager that allows applications to query and obtain responses from their dynamically changing vicinity. In addition, every entity in the framework may implement instances of information providers that act as sources of domain-based information.

Vigil

The current mobile ad-hoc environments provide a limited amount of security in terms of encryption and authentication, usually as part of the particular wireless technology. These solutions do not address the important issue of trust, which is required to enable secure interaction among various devices in ad-hoc environments. We have designed Vigil, which uses existing security methods along with distributed trust, to enable authentication and access control in pervasive environments. Vigil complements SPKI [13] and Role Based Access Control [19] [25] with *trust management*. Every space, whether it is a room, a part of a room, building etc. is managed by an *environment controller*. The controller enforces a *security policy* which consists of rules for authenticating users and assigning domain-specific roles. Every resource in the space registers with the controller and trusts the controller to manage its access control. When a user enters the space, she/he registers with the controller using her/his certificates. These certificates are validated by the controller using certain rules. Based on her/his credentials, the user is assigned a set of roles and is given access rights associated with those roles. These are basic role based access rights. Access rights of users are dynamic and can be changed by delegations and revocations without affecting the roles of the user.

III. Conclusions

In this paper we have described several research challenges that data management systems in the future mobile environments will face. We have outlined the work that our research group (http://research.ebiquity.org/) is doing in these areas. Our initial work has started to address many of the issues that have been raised in this paper. We continue to enhance our systems and build their underlying theoretical foundations.

References

[1] K. Arnold, A. Wollrath, B. O'Sullivan, R. Scheifler, and J. Waldo. *The Jini specification*. Addison-Wesley, Reading, MA, USA, 1999.

[2] S. Avancha, A. Joshi, and T. Finin. Enhancing the Bluetooth Service Discovery Protocol. Technical report, August 2001. TR-CS-01-08.

[3] H. Bharadvaj, A. Joshi, and S. Auephanwiriyakyl. An Active Transcoding Proxy to Support Mobile Web Access. In *Proc. of the IEEE Symposium on Reliable Distributed Systems*, October 1998.

[4] C. Bobineau, L. Bouganim, P. Pucheral, and P. Valduriez. PicoDBMS: Scaling down Database Techniques for the Smartcard. In *Proc. of the 26th International Conference on Very Large Databases*, 2000.

[5] E.A. Brewer, R.H. Katz, Y. Chawathe, A. Fox, S.D. Gribble, T. Hodes, G. Nguyen, T. Henderson, E. Amir, H. Balakrishnan, A. Fox, V. Padmanabhan, and S. Seshan. A network architecture for heterogeneous mobile computing. *IEEE Personal Communications Magazine*, 5(5):8 – 24, 1998.

[6] F. Casati, S. Ilnicki, L. Jin, V. Krishnamoorthy, and M. Shan. Adaptive and dynamic service composition in eflow. Technical Report, HPL-200039, Software Technology Laboratory, Palo Alto, CA, march 2000.

[7] Dipanjan Chakraborty, Filip Perich, Sasikanth Avancha, and A. Joshi. DReggie: Semantic Service Discovery for M-Commerce Applications. In *Workshop on Reliable and Secure Applications in Mobile Environment*, 20^{th} *Symposium on Reliable Distributed Systems*, October 2001.

[8] H. Chen and S. Tolia. Steps towards creating a context-aware agent system. Technical Report HPL-2001-194, Hewlett-Packard Labs, Palo Alto, CA, U.S.A., 2001.

[9] The DAML Services Coalition. DAML-S: Semantic markup for web services. http://www.daml.org/services/.

[10] S. Czerwinski, B. Zhao, T. Hodes, A. Joseph, and R. Katz. An Architecture for a Secure Service Discovery Service. In *Fifth Annual International Conference on Mobile Computing and Networks (MobiCom '99)*, pages 24 – 35, Seattle, 1999.

[11] M. Dunham, A. Helal, and S. Balakrishnan. A mobile transaction model that captures both the data and movement behavior. *ACM/Baltzer Journal of Mobile Networks and Applications*, 2(2):149–162, 1997.

[12] J. Hendler. DARPA Agent Markup Language. `http://www.daml.org`, 2000.

[13] IETF. Simple public key infrastructure (spki) charter: http://www.ietf.org/html.charters/spkicharter.html.

[14] A. Joshi. On proxy agents, mobility and web access. *ACM/Baltzer Journal of Mobile Networks and Applications*, 2000.

[15] A. Joshi, S. Weerawarana, and E. N. Houstis. On Disconnected Browsing of Distributed Information. In *Proc. of the Seventh IEEE Intl. Workshop on Research Issues in Data Engineering*, pages 101–108. IEEE, April 1997.

[16] L. Kagal, J. Undercoffer, A. Joshi, and Tim Finin. Vigil : Providing trust for enhanced security in pervasive systems, October 2001.

[17] R. H. Katz, E. A. Brewer, E. Amir, H. Balakrishnan, A. Fox, S. Gribble, T. Hodes, D. Jiang, G. T. Nguyen, V. Padmanabhan, and M. Stemm. The bay area research wireless access network (barwan). In *Proceedings Spring COMPCON Conference*, 1996.

[18] M. Liljeberg, M. Kojo, and K. Raatikainen. Enhanced services for world-wide web in mobile wan environment. http://www.cs.Helsinki.FI/research/mowgli/mowgli-papers.html, 1996.

[19] E. Lupu and M. Sloman. A policy based role object model. In *1st IEEE International Enterprise Distributed Object Computing Workshop (EDOC'97), Gold Coast, Queensland, Australia*, 1997.

[20] J. McCarthy. Notes on formalizing contexts. In Tom Kehler and Stan Rosenschein, editors, *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 555–560, Los Altos, California, 1986. Morgan Kaufmann.

[21] B. D. Noble, M. Satyanarayanan, D. Narayanan, J. E. Tilton, J. Flinn, and K. R.Walker. Agile application-aware adaptation for mobility. In *Proceedings of the 16th ACM Symposium on Operating System Principles*, 1996.

[22] F. Perich, S. Avancha, A. Joshi, Ye. Yesha, and K. Joshi. Query routing and processing in mobile ad-hoc environments. Technical Report TR-CS-01-18, CSEE, UMBC, November 2001. submitted for conference.

[23] C. Pullela, L. Xu, D. Chakraborty, and A. Joshi. Component based architecture for mobile information access. In *Workshop in conjunction with International Conference on Parallel Processing (ICPP).*, August 2000.

[24] D. Salber, A. Dey, and G. Abowd. The context toolkit: Aiding the development of context-enabled applications. In *CHI*, pages 434–441, 1999.

[25] R. Sandhu, E. Coyne, Hal L. Feinstein, and C. Youman. Role-based access control models. *IEEE Computer*, 20(2):38–47, 1996.

[26] The Bluetooth SIG. The Bluetooth Specifications. http://www.bluetooth.com/.

[27] Carl Tait, Hui Lei, Swarup Acharya, and Henry Chang. Intelligent file hoarding for mobile computers. In *Proc. of the First ACM International Conference on Mobile Computing and Networking - MobiCom'95*, 1995.