

Context-Aware Profiles

A. Ferscha, M. Hechinger, A. Riener, H. Schmitzberger
Institute for Pervasive Computing
Johannes Kepler University, A-4040 Linz

M. Franz, M. dos Santos Rocha, A. Zeidler
Corporate Technology, Software & Engineering,
Architecture, Siemens AG, D-81739 Munich

Abstract

A spatial proximity based spontaneous interaction approach for "shared nothing" peer-to-peer systems is proposed. Mobile peers that come spatially close to each other wirelessly exchange self describing profiles encoded in an XML dialect. Each peer individually conducts a similarity analysis of the semistructured data contained in the received profile, and, in case of the similarity exceeding a predefined threshold, notifies an application to start executing. Profiles are organized hierarchically, allowing for an efficient interaction check based on peer roles in the first step, and a more detailed analysis of peer preferences in the second step. A novel context-aware profile description language (PPDL) has been developed, expressing peer preferences with respect to the particular situation (like time, geolocation, owner, environmental conditions, etc.) in which the interaction among peers is attempted. The context of a peer is modelled symbolically within its profile, so that sensors can verify the actual context and use the context specific interpretation of the profile prior to similarity analysis. A fully functional software framework has been developed implementing context-aware profiles. Similarity metrics and profile matching strategies are presented and assessed. In a flexible manufacturing systems applications scenario we demonstrate its potentials in the context of wireless peer ensembles.

Context-Aware Profiles, Semistructured Data, Peer-to-Peer, Profile Matching, Proximity, Self Description

1 Motivation

Since the number of technological artefacts in pervasive computing environments continues to increase and the appliances work together in a steady higher amount, the necessity for automatically adapting to changed/modified surroundings is a major research field with high capabilities. Second, the today's estimation how the technologies, interfaces and

communication standards in the future will evolve, becomes also more difficult. Solving these problems requires the device's ability of self description and in succession the introduction of self description language constructs.

As we propose a profile-based configuration of devices it would be an interesting feature, that these XML-based configuration files are context-aware, meaning that the content of a profile adapts dynamically and at runtime to changed environment-conditions (as measured by any number of autonomous sensors) and is also receptive against users's needs.

This paper is organized as follows: Section 2 gives an account of related work done in the topic *context-awareness*, *semantics* and *semistructured data*. Section 3 presents the *Pervasive Profile Description Language* (PPDL). Furthermore the process of specifying profiles and domains in the PPDL as well as expressing interests is illustrated; another subsection accounts the profile matching and similarity analysis process. Chapter 4 validates the approach in an empirical study involving context aware profiles for peers setting up a flexible manufacturing system. Conclusions are drawn in section 5.

2 Related Work

With the understanding of context-awareness (Dey, [5]) as systems that gather relevant informations from the environment (place, time, temperature, humidity, etc.) and performs according reactions, context-aware profiles need to be developed.

Some of the profile-exchange and comparison-based spontaneous interaction issues have been studied in the literature. In [13] and [14] a relatedness analysis tool has been introduced and applied to compare drafted regulations with the associated public comments. An example using a drafted regulation on rights-of-way access and the comments received by the Access Board

has been employed for an illustration of the prototypical analysis tool. Comparison is done not only by using traditional term matching algorithms, but also with a combination of feature matches, and not only content comparison but also structural analysis has been implemented. The comparison framework should help the reviewing of comments with respect to provisions in analysed drafts.

Chen describes an approach for structuring and visualising a hypertext-based information space [4]. In the work a unifying framework for extracting structural patterns from a range of proximity data concerning the three fundamental relationships a) hypertext linkage, b) content similarity and c) browsing patterns has been presented. The framework emphasizes the role of users interest in dynamically structuring the underlying information space.

In [20] A. Toninelli et. al. (2005) proposed a middleware that supports the intelligent discovery of context-aware services for mobile users by providing an automatic tool that enable users to specify service queries at a high level of abstraction and a semantics-based matching functionality.

André Bergholz presents in his PhD-thesis [2] a framework for quering semistructured data with the background of representing them by a set of partial schemas rather than by one complete schema. He conceptually modeled and implemented his approach as a prototypic software-framework on the basis of Constraint Satisfaction Problems (CSPs).

Declan O’Sullivan et. al. [17] (2004) examines the potential for using ontology-based semantics for supporting run-time semantic interoperability between pervasive computing applications by presenting a software framework that a) guides the design time identification of candidate mapping information for a set of ontologies and b) enables runtime determination of the actual mappings based on the context of the interoperability required.

Nanyan Jiang et. al. [11] (2004) presents a programming model that enables opportunistic application flows in pervasive environments (programming and management paradigm where the behaviours as well as the interactions of application elements - sensors, actuators, services, etc. - are dynamic, opportunistic, and context, content and capability aware). The presented model builds on content-based discovery and routing services and defines associative rendezvous (AR) as an abstraction for content-based decoupled interactions.

In [3] a digital image library is used for illustrating a number of methods of data manipulation on the basis of XML. Extensible Markup Language (XML)

has had the effect of increasing awareness of data design in the commercial world. On the other hand, in the relational database world, a structured XML document is regarded as semi-structured data because the structure is not as rigid, regular, or complete as the table structure requires.

In summary, the gathering and examination of context-aware attributes/profiles still remains a research challenge, while basic solutions have been proposed [14], [20], [2], [17].

3 The PDDL Approach

With a peer’s ability of self description and the introduction of the PDDL- (pervasive profile description language-) framework, containing mechanisms for context-evaluation, we create the possibility to manage applications not only by context-independent properties (like their spatial proximity), but also by the examination of context-aware attributes with according reactions, eg. by execution different applications regarding the actual environment properties (time, location, temperature, etc.).

With *compositional context* we follow an approach of managing situational information by mobile peers in an autonomous way: context constraints can be created by all possible set operations on geometrical objects, for example intersection or union. If mobile peers utilize appropriate sensor technology, these context constraints can be evaluated independently at runtime, with the result of executing context-aware software (see Figures 1, 2). Peer interaction in a layered architecture

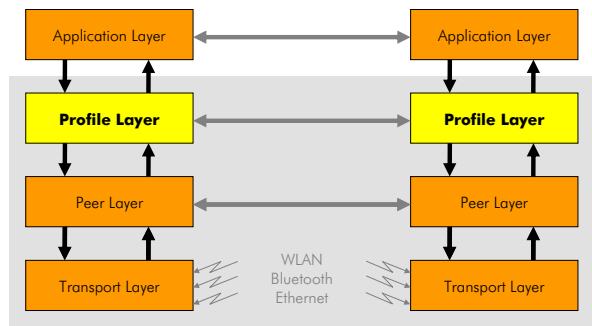


Figure 1. Framework layer structure

can be described step-by-step as follows:

- i) **sendAdvertising()** The PeerLayer periodically initiates sending an advertisement using the underlying TransportLayer.

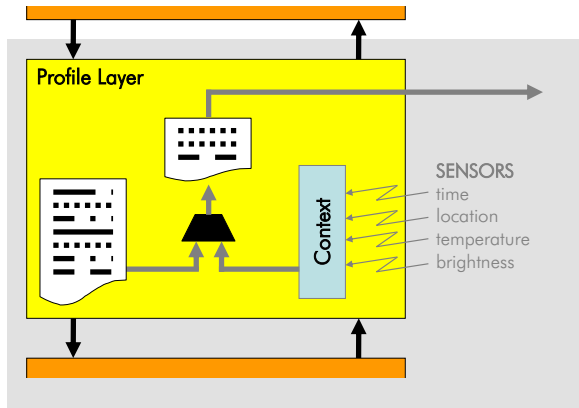


Figure 2. Context transcoding of profiles

- ii) **advReceived()** An advertisement of another peer is received by the TransportLayer and handed over to the PeerLayer.
- iii) **peerFound()** The PeerLayer establishes a connection to the corresponding peer and notifies the layers above.
- iv) **exchangeRoles()** In a first step the roles (eg. employee, student, etc.) are exchanged between the peers (each peer can be involved in one or more of these roles/domains).
- v) **exchangeProfiles()** A *personalized* profile, containing only the data corresponding to the received roles, is exchanged between the peers. Furthermore, the current context is used to modify the exchanged profile according to the specified context definitions of the profile.
- vi) **matchProfiles()** The exchanged (*personalized* and *contextualized*) profiles are matched against the interest specifications in the individual local profile (using different algorithms and/or context-dependent interest specifications for the matching process).
- vii) **notifyApp()** The matching process returns a similarity value in the range $[0, 1]$; if this value is higher than an adjustable threshold, applications in the upper layer are notified and the received profile is handed over for further application-specific processing.

The exchange of roles and profiles is well-established as a one-stage approach: Independent from context conditions, XML-data containing the entities' attributes are shared, matched and according to the result, actions are invoked.

As a novelty, we propose a two-stage profile exchange mechanism in the following form: In the first phase, the coarse grained parameter values of roles are exchanged, which allows a fast identification whether the peers belong to a collective or not. If this *fast matching analysis* is fulfilled (and only in this case!) a fine grained, elaborate profile matching algorithm will be applied on the peers profiles. As shown in Figure 3, this approach can be extended in a second phase to a multi-level profile exchange.

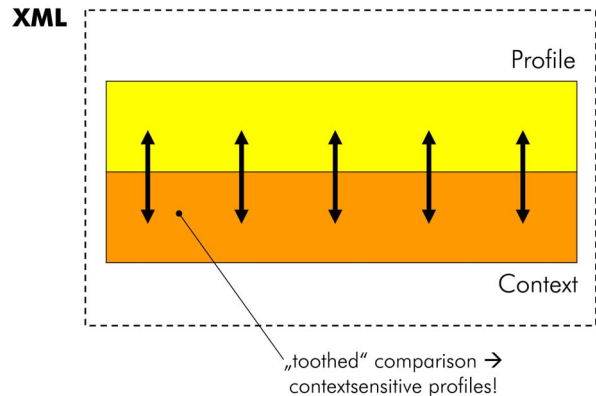


Figure 3. The multi-level profile exchange

3.1 Profile Layer

The profile approach suggests that coordination is highly dependent on profiles and hence on the peer's interests and preferences. One further step in the direction of dynamically coordinated mobile peers is to make coordination context-aware. Context-aware entities will be able to dynamically define their behaviour for different context situations. Such entities can therefore react depending on context properties like time, space, speed, etc. Thereby all possible types of context, without any restrictions, can be thought of as context dimensions.

We utilize profiles to describe features and attributes of persons, objects or devices in many cases. Usually, such profiles are stored locally on an entity. Moreover, profiles are used to provide ubiquitous entities with the ability to communicate with other entities in spontaneous built ad hoc networks and in durable infrastructure networks in an autonomous way. They are supposed to be the trigger for the exchange of interests, the transfer of customized information and the execution of actions depending on preferences and interests of the particular entities.

The PDDL describes profiles in the form of XML-documents because of the feature to be readable and

comprehensive for both users and devices, as well as the feature of being a standard in network communications. Domains are identified by domain specific namespaces (to avoid name clashes) that define an ontology for these domains.

Both, the integration of the profile-layer into the framework and especially the cooperation with the peer-layer, are illustrated in Figure 4. The main tasks of the profile layer are the embedding of the Pervasive Profile Description Language (PPDL), the embedding of an appropriate profile matching algorithm, the implementation of the profile-oriented peer-to-peer coordination process as basis for the P2P communication, and the management of applications setting up on top of the framework. For the description of mobile peers PPDL-profiles are used. The profiles are the trigger for the process of coordination and therefore are of high importance for the way communication is established.

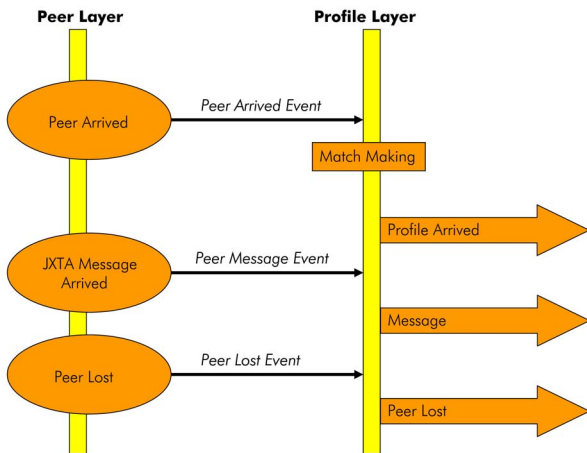


Figure 4. Profile layer integration

Within the peer-to-peer coordination framework (P2PCF) four types of profiles can be distinguished:

- i) **Private Profile** This type contains all profile data that is stored at each peer locally. Profile editors, used at the local peers, can manipulate this data in order to change the preferences of the peer.
- ii) **Public Profile** They are the particular part of a private profile which can be transmitted to other peers in the adhoc network. A public profile therefore describes a subset of the corresponding private profile. If it is appreciated not to hand-over comprehensive profile information (for example because of security concerns [7]), the private profile can be confined to a public profile.
- iii) **Interest Profile** This type of profile is used to describe an interest filter for incoming target public profiles. An interest profile contains prioritised preferences and interests regarding the structure and content of other peer’s profiles. A profile matching algorithm usually takes as input the local interest profile and the target public profile for calculating whether communication is established. The interest profile hence serves as communication barrier which only lets adequate peers join the peer-to-peer network.
- iv) **Role Exchange Profile** Profile-oriented coordination suggests communication of peers only if the peers are matching with regard to their roles. Roles describe the affiliation to a specific group of persons, the affiliation to a particular branche or a simple application. In order to inform other peers of the own embraced roles a specific type of profile is introduced, the so called role exchange profile.

3.2 Expressing Peer Interest

An essential requirement for the PPDL was to offer a means for describing preferences and interests of an entity. This feature allows to assess the subjective utility of other entities profiles. The publishing of a peer’s own interests can influence the way communication is done, hence the coordination of mobile peers is triggered by preferences. An overlapping of interests often can lead to further cooperation of mobile peers. The PPDL includes language constructs for the description of preferences. The static definition of own preferences as part of the profile allows the comparison of profiles without interaction of the user. The mobile peers can do the comparison and the profile matching, relying on the preferences, fully autonomously. A profile matching algorithm examines the subjective utility of the other profile by evaluating the fulfillment of preferences.

But the static declaration of preferences is not sufficient. It is also necessary to dynamically relate preferences to the context of their use. For this type context-aware preferences are introduced, which are evaluated depending on a user-defined context constraint. The PPDL makes no limitations regarding the number of context dimensions. All possible context information can influence the description of profiles. A context-dependent preference could for example express the preferred means of transport to work depending on the time, by choosing the taxi late at night and taking a walk the rest of the time. For the description of preferences, simple but expressive language constructs are part of the pervasive profile description language.

In cases where it is only important that a particular domain element is covered by a target profile it will suffice to declare a static structure preference. The *text* element allows to express a preference for a particular value for a domain element. For a flexible way of text comparison the placing of wildcards ("?", " *") is supported, as shown in the example below. The desired degree of *similarity* is stated as a value in the range between 0 and 1, where 0 means that no similarity is needed and 1 requires complete similarity of the elements.

```

...
<preferences>
  <match element="LastName" similarity="0.76"/>
  <match element="FirstName" text="Mil*"/>
</preferences>
...

```

In order to express logical combinations of preferences *logical constructors* (eg. *AND*, *OR*, etc.) are provided. Inside of a logical constructor preferences have equal priorities; there are no restrictions regarding the depth of nesting.

In general, context is characterised by *n* dimensions (e.g. time, place, environmental conditions, emotional states, etc.). For being open for every occurrence of context, the PPDL makes no assumptions/restrictions regarding the particular context dimensions, i.e. context constraints can be defined/implemented dynamically. Figure 5 illustrates an example, where the specific preferences are only valid if a particular context constraint is satisfied; overlapping and repetitive patterns are allowed. Given a *n*-dimensional context space, the preferences diagram also gets *n*-dimensional.

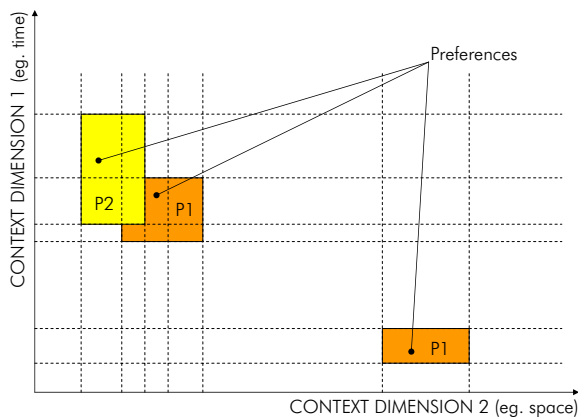


Figure 5. Time-Space Preferences Diagram

3.2.1 Context Constraints

Context constraints are a means to confine the multi-dimensional context space by creating smaller subspaces. Such a subspace, regarding the two dimensions time and place, could for example be a certain context subspace that is determined by the place *metro station* and by the time *February 2nd, 2006*. An evaluation of this context constraint is only valid if all confined dimensions are satisfied concerning the current context situation. In this document context is illustrated in the form of *n*-dimensional diagrams, which represent context constraints as geometrical shapes. The above mentioned example would be illustrated as the subspace represented by the yellow rectangle in the diagram on the right-hand side of Figure 6. The description of

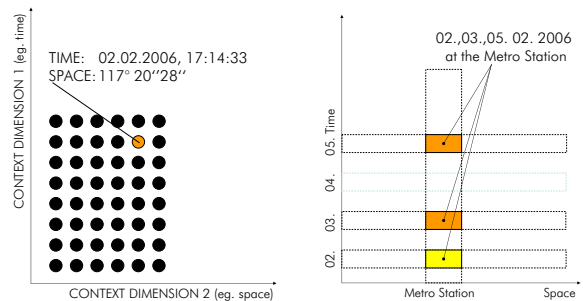


Figure 6. Context Atoms and Aggregates

more complex shapes than a single point in the context space (as shown in the diagram on the left-hand side of Figure 6) is possible by the creation of context aggregates which can be utilized by a geometrical set of operations (eg. union, intersection, etc.) with the possibility of creating multilevel definitions of context constraints.

3.3 Profile Matching

Peers need to know how their profile compares to the profiles of target mobile peers, and to which degree they match [18]. The comparison of two profiles and the subjective interpretation whether the pair of profiles matches is a hard and non-trivial problem. Each mobile peer can have its own preferences that affect the result of a profile match, hence profile matching has to be done individually for each mobile peer. Results of profile matching must be comprehensible and expressive, therefore the results of profile matching algorithms have to be scaled within the range [0, 1]. Figure 7 below illustrates the considerations made for the implementation of the profile matching process.

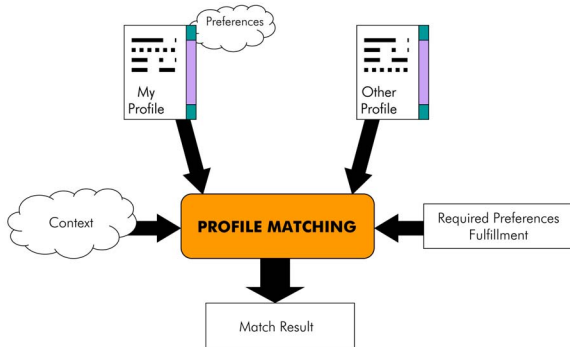


Figure 7. Profile Matching Approach

3.4 Similarity Analysis

A Peer's location is characterised by a m -dimensional vector (points in an m -dimensional euclidean space). The similarity between two peers is a measure of how closely they resemble each other (eg. the shorter distance between objects in the space, the more similar objects are). Similarity and dissimilarity have an inverse relationship: the greater the similarity, the smaller the dissimilarity and vice versa; both have a monotonic relationship with distance. $\forall i, j, k \in X$ the similarity σ of objects $i = (i_1, i_2, \dots, i_m)$, $j = (j_1, j_2, \dots, j_m)$ and $k = (k_1, k_2, \dots, k_m)$ is defined by the following set of relations, which has to be fulfilled in order to guarantee a fundamental, algebraic based profile matching process.

$$0 \leq \sigma(i, j) \leq 1 \quad (1)$$

$$\sigma(i, i) = 1 \quad (2)$$

$$\sigma(i, j) = \sigma(j, i) \quad (3)$$

$$\sigma(i, k) \leq \sigma(i, j) + \sigma(j, k) \quad (4)$$

Similarity search in many cases is done by the application of known similarity metrics like the euclidean distance or the minkowski metric [10], the latter has been used for our experimental setting. All metrics calculate the similarity of features or entire feature vectors. Although semistructured data can be transformed to some kind of feature vector the results of above mentioned metrics for similarity search in semistructured data will not be satisfying because the distance calculation of flat feature vectors does not consider the possible hierarchical structure of the semistructured data.

Other solutions have to be employed, but similarity search in semi-structured data is an extremely difficult task because of the possible complexity of the document structure and because of the fact, that a

document schema must not necessarily exist. Several approaches already exist for this problem [1], [16], [12], [19], but the most are not well suited, eg. IBM XML Diff [9] and DeltaXML [15], that simply compares XML-trees and their leafs. The research area of web query languages is also active in the domain of similarity search, but the research results are only employable for the application area of web information retrieval.

4 Context-Aware Peers in FMS

To demonstrate the potential of context-aware profiles, a flexible manufacturing system (FMS) demonstrator has been built on top of the P2PCF. Considering a FMS scenario where –as opposed to the usual centralized control approach– goods are processed in a totally distributed way, stationary and mobile peers are implementing the whole production process. In the demonstrator, both active (like assembly machines or carrying carts) and passive (like goods or parts of goods) are represented by peers, both in the physical, as well as in the logical sense. A physical object, like a welding robot, however, needs to be integrated into the landscape of FMS objects, all interacting according to the P2P paradigm. The physical activity of the robot needs to be a direct effect of a control issued by the corresponding (software) process.

The architectural concept of the FMS demonstrator is to associate with each and every physical object one instance of a peer, executing the P2PCF. To this end, we have developed a miniaturized hardware platform, the Peer-it platform, serving as a stick-on solution for physical objects. Peer-it's integrate hardware and software into a single, self contained P2P solution. Sticked




Figure 8. Welding-Robot, Processing-Peer

onto a physical object, e.g. a manufacturing robot, and linked with the controls and drivers of it, the Peer-it with its instance of the P2PCF makes the physical peer a logical node in the P2P application. Carrying carts, robots, machinery, etc., when equipped with stick on Peer-it's thus set up a spontaneous network of peers, controlled by the P2PCF. For example sticking a Peer-it computer onto a welding robot makes a "processing

peer” in the sense of the FMS out of it. Technically, the

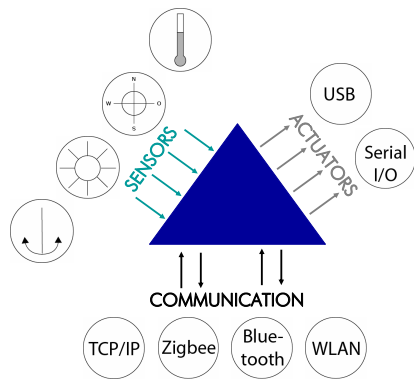


Figure 9. Three channels of a P2P-system

Peer-it concept (see Figure 9) is an "all-in-one" sensor-actuator-communication platform consisting of the execution platform (CPU, memory, standard interfaces), a sensor array and a collection of actuators. Wireless communication (based on protocols IEEE 802.11 and IEEE 802.15) serves to support communication in the nearby proximity. Sensors are dedicated to collect data characterising the situation of a peer with respect to environmental conditions (like temperature, light, humidity, noise, time, place, etc.). Actuators respond to the control triggers resulting from the application into the mechanical means of the respective peer. Communication among such Peer-it's is based on the concept of proximity detection and the mechanisms implemented in the P2PCF. As such, the Peer-it hardware platform, together with the P2PCF executing on this platform, can be considered an autonomous self contained and highly portable P2P computing platform.

Based on the Peer-it computing platform, mobile peers occur in the roles of transport-peers, able to move goods from one space to another; Processing-peers able to assemble or manufacture goods; Artefacts, representing the good per se, and monitoring-peers, able to inspect the configuration and states of all other peers. All these peers interact, once they come into spatial proximity to each other and based on the exchange of their role profiles related to their particular situation (see Figures 8, 10).

By using the P2PCF, the transport-peer is automatically aware of each processing-peer within the manufacturing cell (by periodically advertising of each peer and the process of profile matching) and each processing-peer is automatically aware of the transport-peer, configuring available processing- and transport entities is not required. Based on the autonomy of processing-peers and due to the fact that they can communicate with other machines, a

peer-to-peer enabled machine can collect the main parts of the required software and configuration from the environment. The machine configures itself through the awareness of its environment, moreover, each other entity of such a FMS becomes aware of the new machine and can automatically configure itself accordingly.




Figure 10. Demonstrator setting

The P2PCF, which is the software basis of Peer-it's, allows for fully decentralized ad-hoc organization and coordination based on standard compliant components (OSGi), named a "shared-nothing" technology. The hardware-part is based on the industry standard PC104/+, allowing easy extension, device heterogeneity and transparent integration of resourceless peers (object peers, artefacts).

5 Conclusions and Future work

To support the spontaneous interaction among mobile peers, we have developed context aware profiles, encoding the self description of one peer's abilities and preferences in a standardized, XML-based metadata format. A performance efficient profile hierarchy is proposed, allowing mobile peers to quickly check for the compatibility of their semantic role within an application and to relate their self descriptions to the situation, i.e. the context of their "getting together". The PDDL language is proposed to cope with the dynamic variability of self descriptions of mobile peers, in that an n-dimensional context model is coded with the self describing attributes of a peer within the same profile. A modular, hierarchical context information structuring mechanism allows to build syntactic context constraints, which, with the support of the peer hardware sensors, can be verified at run-time. One peer's profile data, therefore, can always be related to the current context constraints, and thus implement context awareness. The exchange of situational profile data among co-located peers, and the individual analysis of

similarity among the received and one peer's own profile hence allows for an efficient implementation of context aware behaviour in peer-to-peer applications per se.

Both, the mechanism of context constraints in PDDL structured profile data, and the mechanisms for the similarity analysis have been implemented in the peer-to-peer coordination (software) framework, P2PCF [6], [8]. On the other hand, an autonomous hardware platform integrating multiple sensors, actuator arrays and wireless communication technologies has been built, namely the Peer-it platform. Peer-it's executes the P2PCF. With this setting, the capabilities and advantages of context aware profiles and the efficiency of the profile hierarchy could be demonstrated in a FMS scenario, where autonomous products in a totally distributed manner control their own production process by spontaneously interacting with transport peers (like e.g. carrying carts) or processing peers (like machines in assembly lines).

References

- [1] S. Baumann. Music similarity analysis in a p2p environment. Technical report, German Research Center for Artificial Intelligence, Erwin Schrödinger Str., 67663 Kaiserslautern, Germany.
- [2] A. Bergholz. *Querying Semistructured Data Based On Schema Matching*. PhD thesis, Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät II, 1999.
- [3] N. Chang. Data manipulation in an xml-based digital image library. Technical report, School of Library, Archive and Information Studies, University College London, London, UK, 2004.
- [4] C. Chen. Structuring and visualising the www by generalised similarity analysis. Technical report, Dept. of Computer Studies, Glasgow Caledonian University, Glasgow, G4 0BA, UK.
- [5] A. K. Dey and G. D. Abowd. Toward a better understanding of context and context-awareness. Technical report, GVU Technical Report GIT-GVU-99-22, 6 1999.
- [6] A. Ferscha, M. Hechinger, R. Mayrhofer, and R. Oberhauser. A light-weight component model for peer-to-peer applications. In *Proceedings of the 24th International Conference on Distributed Computing (MDC04)*. IEEE Computer Society Press, 3 2004.
- [7] A. Ferscha, M. Hechinger, F. Ortner, and R. Mayrhofer. Securing passive objects in mobile ad-hoc peer-to-peer networks. In *Electronic Notes in Theoretical Computer Science: SECCO 2003*, 6 2003. Elsevier Science, 3, 3, 85, June 2003.
- [8] M. Hechinger. A peer-to-peer framework for mobile ad-hoc networks. Master's thesis, Institut für Pervasive Computing, Joh. Kepler Universität Linz, 4 2005.
- [9] IBM. Xml diff and merge tool - a peer-to-peer framework for mobile ad-hoc networks. <http://www.alphaworks.ibm.com/tech/xmldiffmerge>, last visited: 26. 04. 2006.
- [10] N. Jiang, C. Schmidt, V. Matossian, and M. Parashar. Enabling applications in sensor-based pervasive environments. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 21(9), pages 871–883, 1999.
- [11] N. Jiang, C. Schmidt, V. Matossian, and M. Parashar. Enabling applications in sensor-based pervasive environments. Technical report, The Applied Software Systems Laboratory Department of Electrical and Computer Engineering Rutgers University, Piscataway NJ 08855, USA, 2004.
- [12] T. Karagiannis, M. Faloutsos, and M. Molle. A user-friendly selfsimilarity analysis tool. Technical report, Department of Computer Science & Engineering, University of California, Riverside.
- [13] G. Lau, H. Wang, K. Law, and G. Wiederhold. Relatedness analysis approach for regulation comparison and e-rulemaking applications. Technical report, Stanford University, Dept. Of Civil & Environmental Eng. And Computer Science Department, Stanford, CA 94305-4020.
- [14] G. T. Lau. *A Comparative Analysis Framework For Semistructured Documents, With Applications To Government Regulations*. PhD thesis, Stanford University, Dept. Of Civil & Environmental Eng. And Computer Science Department, Stanford, CA 94305-4020, 2004.
- [15] M. E. Ltd. Deltaxml - managing change in an xml environment. <http://www.deltaxml.com/index.html>, last visited: 26. 04. 2006.
- [16] J. McHugh, S. Abiteboul, R. Goldman, D. Quass, and J. Widom. Lore: A database management system for semistructured data. volume 26, pages 54–66. SIGMOD Record, 1997.
- [17] D. O'Sullivan, D. Lewis, and V. Wade. Enabling adaptive semantic interoperability for pervasive computing. Technical report, Knowledge and Data Engineering Group Trinity College, Dublin.
- [18] Y. Papakonstantinou, H. Garcia-Molina, and J. Widom. Object exchange across heterogeneous information sources. In *Proceedings of the International Conference on Data Engineering (ICDE)*, pages 251–260, Taipei, Taiwan, March 1995.
- [19] D. Quass, A. Rajaraman, Y. Sagiv, J. Ullman, and J. Widom. Querying semistructured heterogeneous information. In *Proceedings of the International Conference on Deductive and Object-Oriented Databases (DOOD)*, pages 319–344, Singapore, December 1995.
- [20] A. Toninelli, A. Corradi, and R. Montanari. Semantic discovery for context-aware service provisioning in mobile environments. Technical report, Dipartimento di Elettronica, Informatica e Sistemistica, Università di Bologna, Viale Risorgimento, 2-40136 Bologna, Italy.