

Ubiquitous Computing – More than Computing Anytime Anyplace?

Kurt Rothermel, Dominique Dudkowski, Frank Dürr, Martin Bauer, Christian Becker
Institute of Parallel and Distributed Systems (IPVS),
University of Stuttgart

ABSTRACT

The rapid development of mobile communication networks and the availability of miniaturized and inexpensive sensor technology allow for the development of multifunctional mobile devices that are able to collect and communicate context information. We also currently see a clear trend towards the integration of embedded systems including sensors into “smart” everyday things, which is also referred to as Ubiquitous Computing. The deployment of sensor technology allows for so-called context-aware applications, going far beyond of providing mobile access to information and computing resources. Context-aware applications exploit captured sensor information to adapt to the current situation of mobile users. For instance, the information offered to a user may depend on his current location, and the way how it is presented to him may depend on the speed he moves, whether he is on its own or with others, and so on.

In this paper, we characterize context-aware applications and present the vision of a digital world model as the foundation for this class of applications. Further, the paper discusses the benefits and research challenges associated with digital world models, which allow applications to take into account the state of the real world and react to state changes.

1. INTRODUCTION

The rapid development and deployment of mobile communication networks gives rise to a broad spectrum of novel applications. Today’s 2nd generation mobile communication networks will be supplemented with more powerful UMTS networks by the end of 2003. Wireless local area networks with data rates up to 54 Mbit/s are already available today. For personal area networks, Bluetooth is the dominant technology, providing data rates up to 1 Mbit/s. Bluetooth technology is designed for cheap and low-power connectivity, and hence can be integrated into everyday things to make them “smart” and to let them communicate with other *smart things*.

We see a clear trend to multifunctional mobile devices integrating communication, computing and sensor technology into a single device. We already see tiny wrist Personal Digital Assistants (PDAs) that include powerful communication and computing technologies. Mobile phones allow for communication, computing and sometimes even for sensing: some phones already include positioning systems, such as GPS, A-GPS or Cell-Id. GPS receivers are rapidly getting smaller and cheaper so that we can expect them to be integrated into many devices or even smart things in the near future.

As for GPS, we can see an enormous miniaturization of other sensor technology as well. A great variety of sensors exist that can be used to capture context information, such as temperature, atmospheric pressure, acceleration, orientation or other objects in the vicinity. Transponder technologies, so-called *smart labels* or radio frequency tags (RFIDs), allow the attachment of tiny and cheap electronic tags to everyday things. These tags can be used to uniquely identify real world things and associate them with information.

Another trend is the development of so-called *sensor platforms* integrating communication, computing and sensor technologies. These sensor platforms may be integrated into everyday things as embedded systems (e.g. the Smart-It (M. Beigl, 2002)), or they may act as nodes of a so-called *sensor network* (e.g. the Smart Dust project (B. Warneke, 2001)). Sensor nodes will be cheap and have the size of a grain of sand, so that they can be deployed, for instance, from an aircraft in large quantities. Once deployed these sensor nodes will form a network, which allows them to cooperatively aggregate and communicate sensed information.

Many of today's mobile applications are "classical" client/server applications which allow the client to be mobile. Mobile users have access to their information and computing resources "anytime, anyplace", which offers a lot of benefits but causes new challenges as well. Nevertheless, the question arises whether the technology trends sketched above allow for a totally novel class of applications. Currently, there are two technology trends that allow applications that go far beyond conventional client/server applications. First, by integrating computing technology into everyday things, a large number of invisible computers will exist in our physical environment. These computers make things smart and allow them to communicate with other things and computers. This has been called the era of *Ubiquitous Computing* (M. Weiser, 1991), where everybody has many computers, in contrast to the PC era, where (almost) everybody had one computer. Second, the integration of sensor technology into devices and things allows for so-called *context-aware applications*. These applications exploit context information captured by sensors to adapt to the actual situation of the user and/or application. For example, the information and services offered to a user may be selected depending on the user's context. The offer may depend on the user's location, the time of day, how fast the user moves, or whether he or she is on a train or driving a car. Also the presentation of the selected information can be context-dependent. The presentation of the same information will certainly differ for a pedestrian walking in a town and a user steering a car moving at high speed on the Autobahn.

In the near future billions of sensors will be available in our physical environment, collecting a huge amount of context information that will be available to a wide spectrum of applications. Therefore, we assume that most of the future applications will be context-aware. Also, most of the applications discussed in the field of ubiquitous computing are context-aware.

The subject of this paper is context-aware systems, their applications, and the scientific challenges resulting from their realization. The remainder of this paper is structured as follows. The next section gives a precise definition of the term context and a classification of different kinds of context information, before characteristics of context-aware applications are introduced and explained by using examples of existing systems. Section 3 deals with spatial models which are needed when location is included as a context parameter. Section 4 presents the vision of a global digital world model, while Section 5 discusses the scientific challenges associated with this vision. The paper concludes with a short summary.

2. CONTEXT-AWARE SYSTEMS

2.1. Context and Context-Aware Applications

Context-aware systems are characterized by their exploitation of context information, like the location of users and objects. There are many different definitions of context (A. K. Dey, 1999; G. Chen, 2000; K. Mitchell, 2002), which makes a clarification of this term necessary:

Definition: Context. Context is the information which can be used to characterize the situation of an entity. Entities are persons, locations, or objects which are considered to be relevant for the behavior of an application. The entity itself is regarded as part of its context.

One or more entities may be relevant for an application's behavior. This definition is based on the notion of context as introduced by (A. K. Dey, 1999). Context-aware applications access and use context information:

Definition: Context-Aware Applications. An application is context-aware if its behavior is influenced by context information.

We will consider two examples, a fleet management and a vehicle navigation system. Locations defined by the underlying road maps and mobile objects representing vehicles are relevant for both applications. While the context of the fleet management system is composed of all vehicles operated by the corresponding carrier, their position and further information like cargo and date of delivery, the navigation system considers only a single vehicle, which is the one to be navigated to a certain destination. In the navigation application, the context includes the vehicle's current position, its destination, the road network and possibly additional information like the current traffic situation. These examples show that different applications may have different sets of relevant entities, and the relevant context of these entities may differ as well.

Two kinds of context can be distinguished, *primary context* and *secondary context*. The first type of context is defined to be location, identity, and time, while the latter can be derived from further attributes of entities, such as speed or acceleration. Primary context is used as an index to access context information, and typically, combinations of location, time and identity are used.

In the fleet management scenario described above, a vehicle's identity can be used to determine its current location, while the location index can be used to find out about all vehicles in a certain area, like a road or a town. If we combine time with location and identity, we can determine where vehicles were located in the past and predict which vehicles will be in a certain area in the future, for instance. Secondary context can be accessed through primary context. For instance, the speed of a certain vehicle can be accessed by using the vehicle's identity, while the average speed of the vehicles in a certain area can be determined by using the location index to find the individual vehicles in this area.

2.2. Characteristics of Context-Aware Systems

Applications can adapt their behavior based on context information in three different ways:

Context-Based Selection: Context can be exploited for the selection of services and information. The location of services and information and their distance to a user are important criteria to classify them and make the selection based upon this classification. In particular, the location is of high relevance because users are often interested in entities in their vicinity, e.g. restaurants, taxis, printers, etc. For example, the VIT system (A. Leonhardi, 1999) is based on the metaphor of so-called *Virtual Information Towers* (VITs), each of which has a position and range of visibility. The posters (HTML pages) attached to a VIT are only "visible" for a user if he is located within the VIT's range of visibility. Additionally, both VITs and posters can be associated with time periods that define when they are visible. Therefore, both location and time are relevant and impact the selection of information.

Context-Based Presentation: The way how an application presents information to its users depends on context. Presentations of the same piece of information may for example differ in the level of detail, the presentation metaphors, or even in the presentation media used.

For instance, the REAL project (J. Baus, 2002) supports adaptive navigation for pedestrians. The user interface is adapted according to the current user situation (e.g., speed of movement) and capabilities of his mobile device (e.g., screen size, voice output). For example, the user's speed determines whether he is navigated by means of a detailed map or just by arrows indicating the direction to move.

Context-Based Action: Actions are initiated when conditions defined on context information become true, i.e., changes in context may trigger actions. An example is the forwarding of a phone call to an answering machine if the callee is in an important meeting at the moment and does not

want to be disturbed (R. Want, 1992). The ComMotion system (N. Marmasse, 2000) sends reminders in form of voice messages to users who enter a certain geographic area, e.g., when a user comes close to a shoe shop the system reminds him that he wanted to buy some shoes.

3. SPATIAL CONTEXT MODELS

Location information is interpreted based on spatial models. Those models typically contain objects that represent physical entities (e.g. items and persons) of the real world. They can be augmented by so-called virtual objects, which are anchor objects of links to existing information spaces, such as the VIT mentioned above. Each object in the spatial model has a unique identity and a location. Further, object attributes model other context information of these objects. This way, different attributes represent different aspects of the context of an object.

Spatial models may have substantially different complexities. Both topological and topographical models are possible, which span a spectrum ranging from simple coordinate systems to detailed 3D models of the real world. Locations can be determined by geometric or symbolic addresses. For example, the location of a person can be determined by using geographic coordinates (e.g. GPS) or a symbolic address (e.g. the Active Badge System (R. Want, 1992)). The coordinates provided by a positioning system are the basis for determining the location of an object. If other spatial relationships are required in addition to position information, for example to determine regions or neighborhood relations, additional information about the relations between two coordinates are required. If geographic coordinates are available, regions and neighborhoods can be determined by using the geometric distance function. In the case of symbolic coordinates, additional information is required. Topological models map spatial relationships onto coordinates and allow for instance the modeling of neighborhood relationships also for symbolic coordinates.

Context-based applications are able to select those objects that are relevant to them from the spatial model by using the indices [location, time] and [identity, time], where the primary context time may implicitly refer to the present time in many cases. Additional context information can be realized based on these objects by restrictions on the objects' attributes. For example, in fleet management, one can assign a new freight item to a vehicle by first selecting all vehicles that are close to the item's origin and then choose those vehicles that are available. In this case, the time of inquiry and the location of the item are the primary context types, while the property "is available" is modeled by a secondary context type attribute.

The interaction of context-based applications and the context model can be realized in different ways. Besides a query type of interaction also event-based interactions are possible. Applications can describe particular events by predicates and will be notified when these events occur. *Spatial events*, describing the occurrence of a certain geographic constellation of mobile and static objects are of major interest. For instance, if two objects move within spatial proximity, or if an object enters a location, e.g. a building or a room, a notification which is due to the occurrence of the event can be emitted.

Spatial context models can be classified according to the covered spatial area, the complexity of the model abstractions, and the dynamics of the model, as depicted in Figure 1. This figure also classifies some existing applications.




Figure 1: Classification of Spatial Context Models

Spatial Scope: Spatial models describe a part of the real world. In the extreme case, the model can be global, meaning that it covers the entire world. In contrast, local models only cover more or less limited portions of the real world. An example of a spatially scoped model is the model of a building that includes floors, rooms, people, etc. of that building. The spatial context models underlying current *location-based services* are mostly designed for global use. Global models are highly desirable as they allow context-based applications to offer their services at arbitrary locations.

Complexity of Model Abstraction: The variety of applications accessing spatial information requires different levels of abstraction. The abstractions supported by spatial models range from point coordinates, over 2D projections and 2.5D models up to complex 3D models of the reality. Clearly, the more complex the model abstractions, the more information can be extracted from the model allowing for a wider spectrum of applications. For example, in a coordinate based system only the location of objects and their geometric distance to other objects can be determined, which is not sufficient for a navigation application, which needs information about road network, such as provided by the Geographic Data File Format (GDF). By providing complex 3D models of the reality, not only realistic visualizations can be generated, but also the orientation and position of users can be derived by image recognition based on the stored model information. Current global models are mostly simple point coordinate models, while some spatially limited models provide for more complex abstractions.

Dynamics: The dynamics of a model determines how fast the content of a model may change over time, if at all. In the case of a static model, such as a digital street map on a CD-ROM, the model is not changed at all. Clearly, both the real-world aspects that can be covered by the model and the degree of consistency with the real world depend on the model dynamics. For example, the model should support dynamic information if a navigation system wants to route the vehicles based on the current road conditions.

4. VISION: GLOBAL DIGITAL WORLD MODELS

From the preceding classification of different approaches in Figure 1, we conclude that existing spatial context models are either rather simple and have global scope or are detailed and locally scoped. Typical examples of the former are telematic services like Passo (Passo, 2003) and Tegarón (Tegarón, 2003) used for car navigation. They only provide maps for the outdoor domain consisting

mainly of roads and static objects and it is difficult to extend these systems beyond their designed purpose. Therefore, they have a rather low complexity of abstraction and low to medium dynamics. The provided models span a wide area (e.g. several countries). On the other hand, we have locally scoped applications like the city guides Guide (N. Davies, 1999) and Cyberguide (G. D. Abowd, 1997). Usually, these models have medium complexity of abstraction and limited support for extensibility needed to support a wide spectrum of applications.

In Figure 1, we can identify a gap which has not been targeted so far. This gap relates to spatial context models which are highly detailed, have global scope and are highly dynamic. Therefore, the vision is a global, worldwide spatial context model, which is realized by a federation of partial models, and where the complexity of abstraction of the spatial model extends far beyond the context models used by today's applications.

These spatial context models will include representatives of a variety of stationary and mobile physical objects of the real world, such as streets, buildings, rooms and people, vehicles and smart everyday objects. Also, these models will be augmented with virtual objects (e.g. Virtual Post-its), which serve as anchors to information in existing information spaces or services, e.g. the web or digital libraries. In addition, the model may capture dynamic states of the represented real world objects, which can be accomplished by integrating sensors into the real world objects. These sensors update the model constantly, i.e. changes in the real world are reflected in the world model. It is also possible to model *interactive objects*, which are connected with actuators, allowing changes in the model to affect the state of objects in the real world.

Despite of its evident benefits, a homogenous spatial context model integrating all available sensor information is highly unrealistic. Rather than that, one should expect that a large number of heterogeneous partial context models will coexist and that federation concepts are implemented to provide a global view onto the entire context model. There are many reasons for the federation of models, such as different levels of geographic relevance of context information, a variety of entities that are to be modeled at different levels of detail, different modeling methods and last but not least organizational and security considerations.

The costs of administration of a global spatial context model will certainly be tremendous, and an important prerequisite for such a model is the amortization of these costs by a large number of applications using the model. Therefore, the context model must be as *generic* as possible to provide support for a wide spectrum of applications.

As in the case of the existing World Wide Web, one can expect that a large number of service providers will exist that will integrate particular context models into the federated model. These models can complement one another and overlap geographically, for example, a model including city map information and a detailed plan of a building in the same city. In addition, several alternative models can be provided for the same geographic region by different service providers. The federation of these models results in a global view onto the model information, where potential inconsistencies must be automatically detected and resolved.

Integrating individual spatial context models by a federation is transparent to applications, and asynchronous and synchronous access mechanisms allow uniform access to model information via the federation. Following the *pull paradigm* applications can issue queries, in particular, spatial queries, such as "Which persons are currently located in room 38.01 of building X". On the other hand it is possible to be registered for events according to the *push paradigm*, in particular, for spatial events, such as for the spatial predicate "More than five people are located in room 38.01 of building X". Predicates can have global character and they can span several partial models.

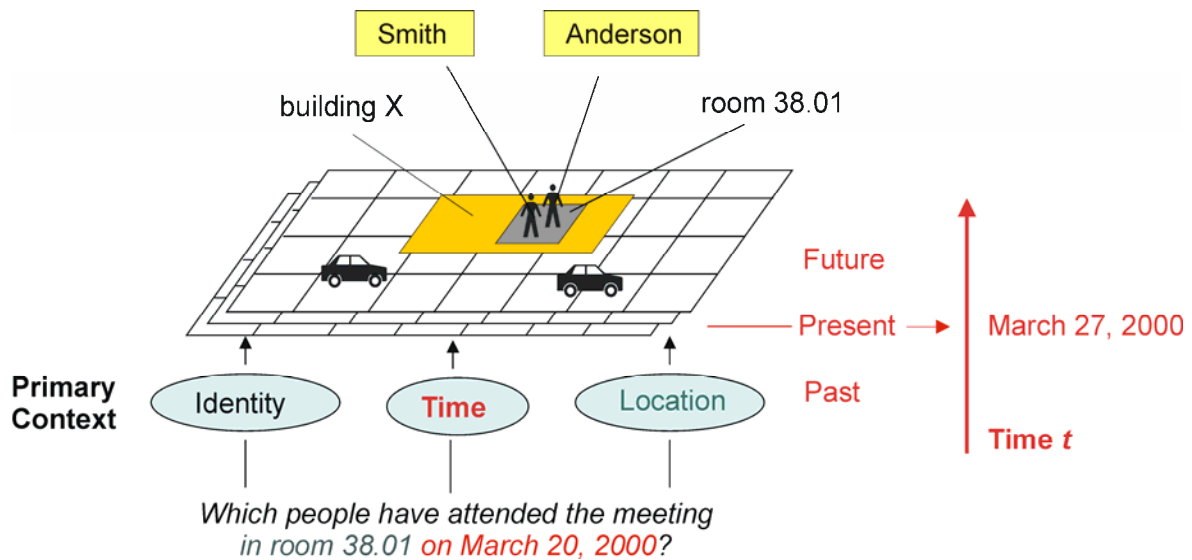


Figure 2: Spatial Model and Indices (Identity, Time, Location)

By integrating *temporal concepts*, spatial context models can reflect more than only the current state of the modeled portion of the real world. By including a temporal dimension, it becomes possible to refer to states of the past and future. For example, temporal concepts make it possible to issue queries such as “Which people have attended the meeting in room 38.01 on March 20, 2000?” Integrating such concepts allows for a number of innovative applications, such as traffic congestion analyses and projections in the area of telematics.

5. SCIENTIFIC CHALLENGES

In order to realize the previously described vision several interesting scientific challenges have to be met. In the year 2000, the Nexus project (F. Hohl, 1999) was started at the University of Stuttgart, Germany, with the goal to develop concepts and methods for the support of mobile and location-based applications. This project has been extended to the interdisciplinary Center of Excellence “World Models for Mobile Context-Aware Systems” funded by the DFG since January 2003, which deals with concepts and methods for the realization of global world models.

One premium research challenge is the modeling and management of global, detailed and highly dynamic spatial models. The Nexus world model comprises location-based data of stationary as well as mobile real-world objects plus virtual objects. In order to be able to support a wide spectrum of applications the model must be highly extensible. Further questions arising from the management of the federated world model include scalable system architectures, appropriate federation concepts as well as consistency concepts and mechanisms.

In order to access the world model, efficient and robust mechanisms enabling a mobile device to interact with the model have to be developed. This includes a seamless integration of different network technologies as well as optimized access to the world model. A model of the communication network can be used to decide whether communication should be deferred until an area with good network connectivity is reached, or data can be hoarded on the mobile device and then be used later in areas with weak network connectivity. Also innovative communication concepts and protocols which rely on information stored in a world model can be realized, like protocols for sending messages to all users in a certain geographic area. Moreover, security and privacy mechanisms are necessary to protect model data. Following the concept of multilateral security, the security needs of all parties interacting with the world model need to be balanced.

Other scientific challenges result from the integration of sensor data and the presentation of and interaction with the model. Regarding the quantity of sensor data that has to be integrated, methods

for automatic integration are necessary. What is more, inconsistencies originating from different sensors need to be resolved. Presentation of model information requires context-aware adaptation techniques that allow the presentation of model data depending on the current situation of the user and the capabilities of the user's device.

In order to derive requirements for our research and to evaluate our research results we will build advanced context-aware applications based on our world model. Those applications include a navigation system for blind people and a Smart Factory. Finally, social aspects of context-aware systems and their application have to be examined to make sure that those systems will be accepted by a wide user community.

6. SUMMARY

The integration of computer, communications and sensor technology in everyday things and mobile devices allows for the collection and communication of context information. This gives rise to a new class of applications, which are able to adapt their behavior according to their relevant context. In this paper, we have given a general classification of context and context models and introduced the vision of global digital world models as basis for context-aware applications. We have pointed out many research problems that have to be solved to make this vision a reality. These challenges are in the field of communications, security, data management and modeling, presentation of model information, sensor integration, applications as well as social aspects. The variety of challenges requires an interdisciplinary approach like the one pursued by the Nexus project.

7. REFERENCES

- G. D. Abowd, C. G. Atkeson, J. Hong, S. Long, R. Kooper, M. Pinkerton (1997): Cyberguide: A Mobile Context-Aware Tour Guide. *Wireless Networks* 3(5), pp. 421-433
- J. Baus, A. Krüger, W. Wahlster (2002): A Resource-Adaptive Mobile Navigation System. In: *Proceedings of the International Conference on Intelligent User Interfaces*, San Francisco, ACM Press, pp. 15-22
- M. Beigl, H. Gellersen (2003): Smart-Its: An Embedded Platform for Smart Objects, *Smart Objects Conference (sOc)*, Grenoble, France
- G. Chen, D. Kotz (2000): A Survey of Context-Aware Mobile Computing Research. Dartmouth College, Technical Report TR2000-381
- N. Davies, K. Cheverst, K. Mitchell, A. Friday (1999): Caches in the Air: Disseminating Tourist Information in the GUIDE System. In *Proceedings of the Second IEEE Workshop on Mobile Computing Systems and Applications*, New Orleans, Louisiana, USA, February 1999
- A. K. Dey, G. D. Abowd (1999): Towards a Better Understanding of Context and Context-Awareness. Georgia Institute of Technology, College of Computing, Technical Report GIT-GVU-99-22
- A. Harter, A. Hopper, P. Steggles, A. Ward, P. Webster (1999): The Anatomy of a Context-Aware Application. In: *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networking*, Seattle, WA, ACM Press, pp. 59-68
- F. Hohl, U. Kubach, A. Leonhardi, K. Rothermel, M. Schwehm (1999): Next Century Challenges: Nexus – An Open Global Infrastructure for Spatial-Aware Applications. In: *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '99)*, pp. 249-255

- A. Leonhardi, U. Kubach, K. Rothermel (1999): Virtual Information Towers – A Metaphor for Intuitive, Location-Aware Information Access in Mobile Environment. In: Proceedings of the Third International Symposium on Wearable Computers, San Francisco, CA, pp. 15-20
- N. Marmasse, C. Schmandt (2000): Location-Aware Information Delivery with ComMotion. In: Proceedings of the Second International Symposium on Handheld and Ubiquitous Computing (HUC), Bristol, UK, Springer, pp. 157-171
- K. Mitchell (2002): A Survey of Context-Awareness. Lancaster University, Internal Technical Report. <http://www.comp.lancs.ac.uk/~km/papers/ContextAwarenessSurvey.pdf>
- Passo (2003): Vodafone Passo, <http://www.passo.de>
- Tegaron (2003): T-Traffic, www.tegaron.de
- R. Want, A. Hopper, V. Falcao, J. Gibbons (1992): The Active Badge Location System. ACM Transactions on Information Systems, 10(1): 91-102
- B. Warneke, M. Last, B. Liebowitz, and K. S. J. Pister (2001): Smart Dust: Communicating with a Cubic-Millimeter Computer. IEEE Computer, 34(1), pp. 44-51
- M. Weiser (1991): The Computer of the 21st Century. Scientific American, 265(3), pp. 94-104