

Nexus – An Open Global Infrastructure for Spatial-Aware Applications

Fritz Hohl, Uwe Kubach, Alexander Leonhardi
Kurt Rothermel, Markus Schwehm

January 25, 1999

Abstract

Due to the lack of a generic platform for location- and spatial-aware systems, many basic services have to be reimplemented in each application that uses spatial-awareness. A cooperation among different applications is also difficult to achieve without a common platform. In this paper we present a platform that solves these problems. It provides an infrastructure that is based on computer models of regions of the physical world, which are augmented by virtual objects. We show how virtual objects make the integration of existing information systems and services in spatial-aware systems easier. Furthermore, our platform supports interactions between the computer models and the real world and integrates single models in a global “Augmented World”.

Contents

1	Introduction	3
2	General Idea	4
2.1	Augmented Areas	4
2.2	Augmented World	6
3	Example Scenario	6
4	Requirements	7
5	Related Approaches	8
6	Architecture of the Nexus Platform	9
7	Conclusion and Future Work	11

1 Introduction

When we look around us, there is a lot of information that is bound to or associated with spatial actualities or that is organized in a geographical manner. Posters are hanging on walls, signs stand at crossings, electric doors open when people approach, speeches can be heard only on certain locations and so on. The appearance of a physical object and the possibility to directly manipulate it are bound to the location of that object. As a very basic matter of nature, this fact is so fundamental that we have even problems to think of alternative possibilites. As everybody experiences this unity of information and location from birth, we can handle information that is organized in this natural way without further learning.

Up to now, only a few information system applications take into account spatial data of its users and other objects or organize information in a spatial way. These existing information system applications are currently centered around the areas of car navigation, fleet management, and map-oriented geographical information systems. Applications like this and early prototypes of augmented reality applications or such that are based on ubiquitous or wearable computers show a need for spatial aware applications. This need is limited only by the lack of a spatial sensor infrastructure on the one hand and a widespread and scalable platform for spatial information systems on the other hand. The first is subject of ongoing research with technologies like GPS that are getting cheaper and smaller. The development of the latter is the subject of this article.

Existing spatial aware applications develop all the spatial functionality and integrate the sensors that are needed by the application area. Neither can different applications share components nor are they able to interact by using the same set or at least format of spatial information. What we envision is an open global platform for spatial aware applications. This platform will not so much differ in functionality from a superset of the common mechanisms of the existing applications, but in its orientation towards a “middleware” for these and future applications that are able to interact. The intended environment for such a platform are small mobile devices carried by users. Embedded into the devices are spatial and other sensors, and wireless communication facilities. Normally the devices are also employed as user interfaces to the applications, but, as we will see, the platform can be used from any other computer, and by any other user or program, too.

The requirements of such a platform are similar to the attributes that made the WWW platform a big success: the offered service has to be simple, yet powerful, everyone should be able to offer information, and the system has to scale well for a huge number of users and information. If the research community is able to provide such a technology, a similar success like WWW can be expected since the possible applications are useful for a large audience, spectacular and can be used by even inexperienced users.

The rest of the article is organized as follows: Section 2 presents the general ideas, namely Augmented Areas and an Augmented World. In Section 3, an example scenario illustrates the advantages of a platform that supports spatial aware applications. Section

4 states the requirements of such a platform, Section 5 gives an overview of existing approaches that are related to the field of spatial information access. Section 6 outlines the architecture of a platform that fulfills the requirements. Finally, Section 7 concludes and gives a short overview of future work.

2 General Idea

The main objective of the Nexus platform is to provide an infrastructure to support spatial-aware applications. This infrastructure is based on computer representations of certain regions of the physical world, which can be augmented by virtual objects. These regions together with the virtual objects placed within them can be experienced through the Nexus platform and are called *Augmented Areas*. Technically, an Augmented Area consists of the objects in a geographical region, an installation of the Nexus system and the appropriate information content. The computer internal representation of an Augmented Area is referred to as *Augmented Area Model*. As we will show these models have a lot of advantages over conventional virtual world models.

Moreover, we do not think about the Augmented Areas as closed worlds, but integrate them into one global “Augmented World” of interconnected Augmented Areas.

2.1 Augmented Areas

Each Augmented Area is limited spatially, e.g to the area of a certain city or the inside of a building, and consists of physical objects and virtual objects. While virtual objects can only be perceived by a user through the Nexus platform, physical world objects exist in the physical world. Each object is located at a spatial position, which may change over time. As depicted in Figure 1 Augmented Areas may overlap or include each other.

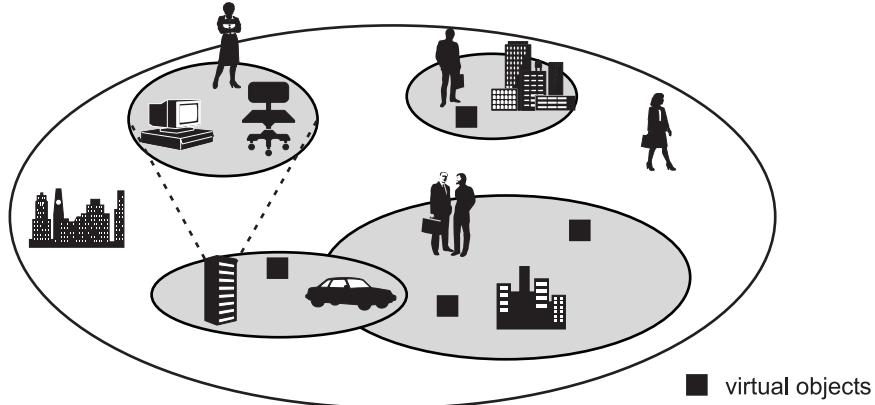


Figure 1: Augmented Areas (shaded) within the physical world

For every object of interest in an Augmented Area there is a datastructure in the according model that represents the object. Note that this typically includes the user

and other persons of interest who are located in the Augmented Area. An important attribute of each datastructure is the spatial position of the according object. In the case of mobile objects the position may be determined by locating systems, such as the Global Positioning System (GPS) (Hofmann-Wellenhof et al., 1997) or the Active Badge System (Want et al., 1991). The Augmented Area Models allow Nexus-based applications not only to profit from the user's current position, but also to take into account the position of all other objects, e.g. to detect when two objects meet each other. The user's current position is, for example, important for spatial-aware information systems, such as the Cyberguide (Long et al., 1996), whereas the position of a number of objects within an Augmented Area is of interest for Intelligent Vehicle Highway Systems.

Another feature of the Augmented Area Models is that they are some kind of "living models". By the term "living model" we mean, that changes in the Augmented Area, noticed by the Nexus platform, are automatically propagated to the model. The other way round, changes of the model state might trigger actions in the Augmented Area. An example is the manipulation of a virtual control board (see example scenario in Section 3).

Existing information systems and services, e.g. the World Wide Web (WWW) or on-line banking, can easily be integrated into the Nexus infrastructure via virtual objects. For this purpose virtual objects are used as an anchor in the Augmented Areas to access external systems and services. An example for this kind of virtual objects is given in the example scenario in Section 3.

Figure 2 illustrates the central role of the Augmented Area Models. They are the connecting link between client applications, existing information systems/services and the Augmented Areas.

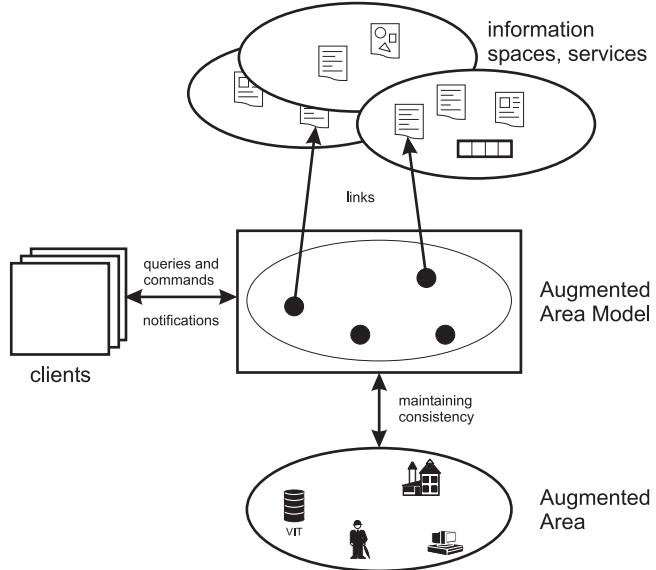


Figure 2: Augmented Area Model

2.2 Augmented World

In addition to the services needed by isolated spatial-aware applications, the Nexus platform also offers the integration of single Augmented Areas into one global “Augmented World”. To make this integration easier the Augmented Area Models are described in a uniform way and are accessible through the same application interface offered by the Nexus platform.

Due to the uniform description, relationships between different areas can easily be ascertained. Examples for such relationships are the distance between areas or the inclusion of one area in another. Furthermore, the characteristics of different areas, e.g. their levels of detail or size, are comparable.

Since all models are accessible through the Nexus application interface, an application can easily switch between different models. For example, a tourist guide system may switch from the global city model to a local museum model when the user enters the museum. Furthermore, we plan to realize a hand-over mechanism between different applications, which allows to offer user-friendly integrated services instead of forcing the user to run single, isolated applications. Therefore, the city guide and the museum guide could be separate Nexus-based applications, which are integrated in the tourist guide system.

Another advantage of the uniform format of the Augmented Area Models is that new models can be integrated into the global “Augmented World” at any time. After the integration the new models can immediately be accessed by any Nexus-based application. The integration works similar to the integration of a new WWW-server into the Web: the institution providing a new model, just has to register it at the Location-Management component (see Section 6). This easy integration of new models will help the Nexus world to grow fast and to reach a widespread coverage.

3 Example Scenario

In this section we present an example scenario to demonstrate the potential of Nexus and Nexus-based applications.

Our scenario shows how a Nexus-based conference system can assist the visitors of a large conference with a big audience, many talks and workshops. The Augmented Area Model used by this system covers the whole building where the conference is held, and includes all visitors within the building, all rooms of the building as well as important items within the building, e.g. overhead projectors. The Nexus platform can control some of these devices by sending control signals to them. It determines the current position of every visitor with an indoor infrared tracking system.

The spatial-aware information component of the conference system allows the visitor to access information about the conference in a very convenient manner. It is based on the Virtual Information Tower (VIT) concept. VITs are virtual objects, which are located

at all places of interest. If a VIT becomes “visible”, i.e. the visitor gets close to it, he will be able to access the information attached to this VIT through his mobile device. For example, a VIT in the entrance hall may offer information about the conference’s agenda and which workshop or talk is held in which room. A VIT within a room may provide more specific information about the workshop held in this room. The visibility area of a VIT can be specified arbitrarily.

Now let us assume that a visitor *A* is listening to a talk. During the talk another visitor *B* asks an interesting question. Therefore, *A* wants to know more about *B*. *A* can access the information *B* is willing to offer about his person by just pointing at *B*. The pointing may be performed through the graphical user interface on *A*’s mobile device or more conveniently through a pointing device. Note that this kind of information access is not supported by the World Wide Web. In the WWW you always need a hint about the information you are looking for, e.g. *A* needs at least *B*’s name to be able to search for *B*’s homepage. If later in the day *A* wants to talk to *B*, but does not know where to find *B*, the conference system will again be able to help *A*. It just invokes a navigation application that guides *A* to the current location of *B*.

The conference system might also be used to control the technical devices within each conference room. Therefore, a virtual control board could be offered to authorized persons. The manipulation of this virtual object might control the devices located in the room, e.g. the light or the volume of the speaker system. This is possible because the Nexus platform allows changes of the Augmented Area Model to be propagated to the Augmented Area by sending control signals to physical devices.

Each of the tasks in the above scenario (information, navigation and controlling) can be performed by a generic Nexus-based application. Then, it is the common underlying Augmented Area Model that plugs them together.

Due to space limitations we do not show how different Augmented Area Models can cooperate. But, it should be clear that the model of the conference building used by the conference system could cooperate well with a model of the city where the conference is held, in order to provide information about the city to the conference visitors.

4 Requirements

Now we will identify the requirements for the Nexus platform that arise from our idea. These requirements can be classified according to the main challenges the Nexus platform has to cope with:

Mobility: First of all, Nexus has to deal with the mobility of the users and other physical world objects. Therefore, it must support mobile wireless communication and small mobile computing devices, such as Personal Digital Assistants and notebooks.

Heterogeneity: Nexus must be usable in a large variety of Augmented Areas, ranging from areas covering a whole nation to areas just covering an office. This leads to a very heterogenous system, which has to support different network technologies and different

tracking systems. It must also be possible to extend Nexus without a big effort to support new technologies. Moreover, we want Nexus to adapt dynamically to changes in the user's environment, e.g. when the available network bandwidth changes. Further aspects of heterogeneity come from the different services and information systems that must be accessible through Nexus and the various input and output devices that have to be supported, e.g. the pointing device mentioned above.

Interoperability: The interoperability between different Nexus-based applications and different Augmented Area Models must be guaranteed.

Scalability: Nexus has to scale well for both a large number of users and a large number of objects to be stored in the Augmented Area Models.

Privacy: Since Nexus stores information about each user's location, it must be ensured that only those persons get information about a user who are authorized by this user to do so.

5 Related Approaches

Single aspects of the Nexus platform can be found in other systems and are of interest in several research areas. Here we present some of the most prominent concepts and discuss their relation and differences to Nexus. An important aspect of Nexus are mobile computers and communication. Its basic functions therefore also use many of the technologies developed in Mobile Computing (Satyanarayanan, 1996).

Ubiquitous Computing: In Ubiquitous Computing, as described by Weiser (Weiser, 1993), it is assumed that computers will be abundant in the near future. Small computing devices with the ability to communicate are to be integrated into objects of everyday use, enabling them to give information about themselves and to interact with their users. One of the main concerns of Ubiquitous Computing has therefore been the development of suitable hardware. In Nexus all supported objects are represented in a computer model. Thus, information can be attached to objects that do not have a computer interface. Nexus is also able to place virtual objects in the real world through which information can often be accessed more intuitively.

Location- or Context-Aware Information Systems: A number of systems can be described by the terms location- or context-aware. They use the knowledge about the position of their users to provide information about the area and nearby objects, as background information for answering queries more specifically or to trigger certain actions (Schilit et al., 1994). Further information about the environment, like temperature, etc. can also be taken into account (Beadle et al., 1997). Location-aware applications are often built around a specific positioning technology like the Active Badge System (Want et al., 1991) or the ParcTab (Schilit et al., 1993), which were developed for an office environment. In contrast to Nexus, most location-aware systems use a very specialized model of the real world which is only suitable for a single environment and a certain type of application.

Geographical Information Systems and Navigation Systems: The functionality for storing and retrieving data in association to geographic objects is provided by Geographic Information Systems. They combine a map and a database to perform spatial queries, like searching for all objects of a given type within a certain distance. If the map is connected to a device which determines the position of its user, e.g. GPS, and a program provides the algorithms to calculate the most appropriate way from one location to another, it can be used as a Navigation System. Geographical Information Systems usually provide only a certain level of detail, whereas in Nexus different models with different levels of detail and abstraction can be combined. Also, Geographical Information Systems do not support virtual or active objects and updates between the real world and its model.

Augmented Reality: The research area of Augmented Reality has as its goal the integration of a virtual and the real world. Special devices, e.g. a heads-up display, enable a user to move in this augmented world. An example for a typical application is a system which allows a worker to “see” power lines and water pipes which run in the walls of a building. Other systems also allow an interaction with the augmented world (Nagao and Rekimoto, 1996). The range of Augmented Reality applications is usually restricted by their hardware requirements, although much progress has been made using wearable computers (Starner et al., 1997). In contrast to existing Augmented Reality systems, Nexus is not limited to certain applications or areas, as different models with different levels of detail can be combined.

In summary, existing approaches for spatial information access usually collect location information themselves and create their own model of the world. A single platform like Nexus, which provides these services, allows developers of such applications to concentrate on their application functionality. Furthermore, a bigger audience can be addressed and a cooperation between different applications is possible.

6 Architecture of the Nexus Platform

The main purpose of the Nexus platform is to manage the Augmented Area Models and to present a uniform access to the combined Augmented World Model for different spatial-aware applications and services, allowing them to interact and to cooperate. It is therefore not intended to be used only in a single application, but to provide a standardized interface for basic spatial-aware functionality. This aim is reflected in the architecture of the platform. Because an Augmented Area Model can contain lots of objects, it must be possible to distribute it across a number of different nodes to ensure scalability.

The models store the current state of all their mobile, stationary and virtual objects. Updates between Augmented Area Models and the Augmented Areas are performed by the platform. Clients may query the current state of a model or subscribe to a certain event and get a notification from the Nexus platform whenever it occurs. From virtual objects links lead to specific items in an information space (e.g. WWW, a Digital Library

or X.500) or services. Authorized clients, e.g. spatial-aware services, are able to create, modify and delete virtual objects or set links to information spaces.

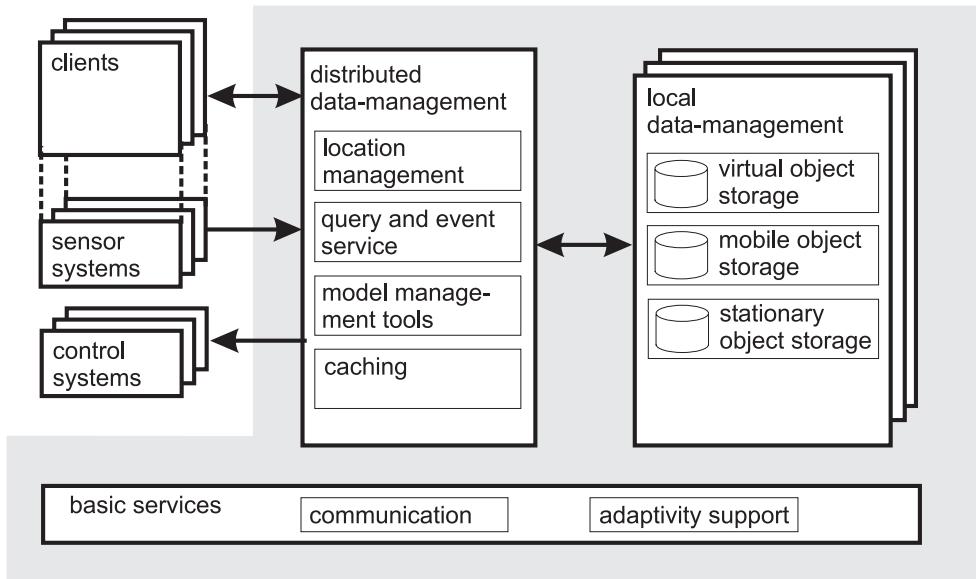


Figure 3: Architecture of the Nexus platform

The architecture of the platform is shown in Figure 3. The purpose and functionality of its components is as follows:

External components: *Clients*, i.e. spatial-aware applications and services, use the functionality of the Nexus platform by accessing the Augmented World Model through its standardized interface. The models are linked to the real world through the *sensor- and control-systems*. Sensor-systems, e.g. a GPS sensor, provide information about the state of objects in the real world. In many cases a (positioning) sensor is coupled with a client, as the client is interested in its own position and environment. Via the control-systems the platform is also able to change the state of certain real world objects. In both cases a standardized interface is used in order to allow new sensor- and control-systems to be added.

Basic services: Basic services include a support for *communication* and the *adaptivity of applications*. They are used by all other components. The former handles communication between the different components in the Nexus platform and allows them to communicate efficiently and transparently with objects on stationary as well as on mobile computers. The latter uses a flexible definition of quality-of-service to manage the available resources and notify applications of impending changes. Thus, an appropriately designed application can react to changes of the available resources.

Local data-management: Different parts of an Augmented Area Model are stored in different nodes of the Nexus platform and are managed by the components which are responsible for the local data-management. Nodes are either single computers or clusters of computers containing databases that store and retrieve spatial data efficiently.

According to their different update characteristics there are special databases for storing virtual, mobile and stationary objects. Stationary objects like buildings etc. need only few updates, while virtual objects and links are modified moderately by clients or services and mobile objects have to be updated very frequently if their position changes fast.

Distributed data-management: Distributed data-management is performed by a logically centralized set of components which control the access to the local, distributed and maybe replicated nodes, where different parts of the models are stored. A discussion of how to distribute the components of the platform across different mobile and stationary computers is not presented here. Parts of it can be found in (N. N., 1999).

The component for *location management* is responsible for continuously mapping objects and areas to the node where the appropriate information is stored. It is used by the *query and event service* to provide an API by which applications and services can transparently access the complete Augmented World Model. Queries and subscriptions for events are forwarded to the appropriate nodes.

For certain applications clients or services need to be able to modify virtual objects or to introduce new links to items of information. The *model management tools* provide the functionality to add, modify or delete virtual objects or links to information spaces. Links and virtual objects may be organized in different layers, which are used to group objects, which e.g. have similar content or belong to the same provider. For each user a specific view of the model can be created.

Distributed data-management is also responsible for *caching* because it has the knowledge of where the information is stored. The caching of Nexus should be location-aware, i.e. take into account which information is frequently accessed in a certain area. Together with suitable prefetching techniques it is able to reduce the delay in accessing information and to support a certain degree of disconnected operations.

All components of the platform have to deal with the important issue of security and privacy. A global control of access is handled by the distributed data-management.

7 Conclusion and Future Work

In this paper we have presented the general idea of a platform that provides the basic functionality for spatial-aware applications. This platform will allow such applications to be created much more easily as they can rely on a common infrastructure. The platform maintains specific models for certain areas of the real world which allow a user to access information or services by their spatial position or through real world objects. It also combines the single models into an Augmented World Model that can be used by applications to interact and to cooperate. We have outlined the architecture of this platform and discussed some related approaches.

We envision a global coverage of the Nexus-platform with an universal Augmented World, into which new models, information and services can be integrated as easily as pages into the World Wide Web. We have outlined some possible applications for

such a platform and believe that once it is available, many further attractive applications will emerge, again similar to the World Wide Web.

Currently, a first project evaluates different aspects of the Nexus idea and builds a first prototype. The prototype will be used for experiments and to measure usage characteristics and load. Based on the experiences in this project a first version of the Nexus platform will be build.

References

Beadle, H., Harper, B., Maguire Jr., G., and Judge, J. (1997). Location aware mobile computing. In *Proceedings of the IEEE/IEE International Conference on Telecommunications*, pages 1319–1324, Melbourne, Australia. IEEE.

Hofmann-Wellenhof, B., Lichtenegger, H., and Collins, J. (1997). *Global Positioning System: Theory and Practice*. Springer-Verlag TELOS.

Long, S., Kooper, R., Abowd, G., and Atkeson, C. (1996). Rapid prototyping of mobile context-aware applications: The cyberguide case study. In *Proceedings of the Second Annual International Conference on Mobile Computing and Networking (Mobicom'96)*, pages 97–107, Rye, New York, USA. ACM.

N. N. (1999). *Analysis of Distribution Schemes for the Management of Location Information*. Technical report, N. N.

Nagao, K. and Rekimoto, J. (1996). Agent augmented reality: A software agent meets the real world. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96)*, pages 228–235, Kyoto, Japan. AAAI Press.

Satyanarayanan, M. (1996). Fundamental challenges in mobile computing. In *Proceedings of the Fifteenth ACM Symposium on Principles of Distributed Computing*, pages 1–7, Philadelphia, PA, USA. School of Computer Science, Carnegie Mellon University, USA. ACM.

Schilit, B., Adams, N., Gold, R., Tso, M., and Want, R. (1993). The parctab mobile computing system. In *Proceedings of the Fourth Workshop on Workstation Operating Systems (WWOS-IV)*, pages 34–39, Napa, CA, USA. IEEE.

Schilit, B., Adams, N., and Want, R. (1994). Context-aware computing applications. In *Proceedings of the First IEEE Workshop on Mobile Computing Systems and Applications*, pages 85–90, Santa Cruz, CA, USA. IEEE.

Starner, T., Mann, S., Rhodes, B., Levine, J., Healy, J., Kirsch, D., Picard, R., and Pentland, A. (1997). Augmented reality through wearable computing. *Presence*, 6(4).

Want, R., Falcao, V., and Gibbons, J. (1991). The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102.

Weiser, M. (1993). Some computer science issues in ubiquitous computing. *Communications of the ACM*, 36(7):74–83.