NEXUS – Distributed Data Management Concepts for Location Aware Applications

Steffen Volz[1] and Monika Sester[1]

[1] Institute for Photogrammetry, University of Stuttgart, Geschwister-Scholl-Str. 24D, 70174 Stuttgart, Germany
{firstname.lastname}@ifp.uni-stuttgart.de
http://www.ifp.uni-stuttgart.de

Abstract. Nowadays, mobile computers like subnotebooks or personal digital assistants, as well as cellular phones can not only communicate wirelessly, but they can also determine their position via appropriate sensors like DGPS. So-called location aware applications take advantage of this fact and structure information according to the position of their users. In order to be able to assign data to a certain location, these information systems have to refer to spatial computer models. The NEXUS[1] project, which is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), aims at the development of a generic infrastructure that serves as a basis for location aware applications. The central task of this platform deals with the data management. Since it is intended to establish a scalable middleware that is capable of handling huge amounts of data, the system has to be organized within a distributed environment. Moreover, the data that are made available by NEXUS have to fit to the needs of all kinds of location aware applications and thus multiple representations of data sets - from highly detailed to coarsely resoluted ones - have to be provided.

1 Introduction - The Basic Ideas of the NEXUS Project

In times of increasing possibilities to access information of all kinds, it is sometimes difficult to find quickly what has been searched for. Thus, a growing need for efficient methods to structure information according to ones individual needs can be observed. Location aware applications offer an approach for this purpose: they know the position of their users and provide information that is of special importance at their current location. For this reason, these information systems have to refer to spatial models that allow the assignment of data to a certain position. The central part of the NEXUS platform consists of the management of these spatial models which represent the physical world. They provide the basis for the NEXUS infrastructure and not only have to serve for particular, but for location aware applications in general.

---

[1] The project is being carried out within the scope of a research group in cooperation between the Institute for Photogrammetry, the Institute of Parallel and Distributed High Performance Systems and the Institute of Communication Networks and Computer Engineering of the University of Stuttgart.

To extend the scope of information that can be delivered by NEXUS, external information sources like digital libraries or especially the World Wide Web must be accessible via the platform as well. Therefore, the concept of virtual objects (VOs) was introduced. The virtual objects augment the NEXUS world and can have different functions. Virtual Information Towers (VITs), for example, spatially structure the data stored within other information systems. Each VIT has a certain location and - according to its significance - a predefined area of visibility (see Fig. 1).

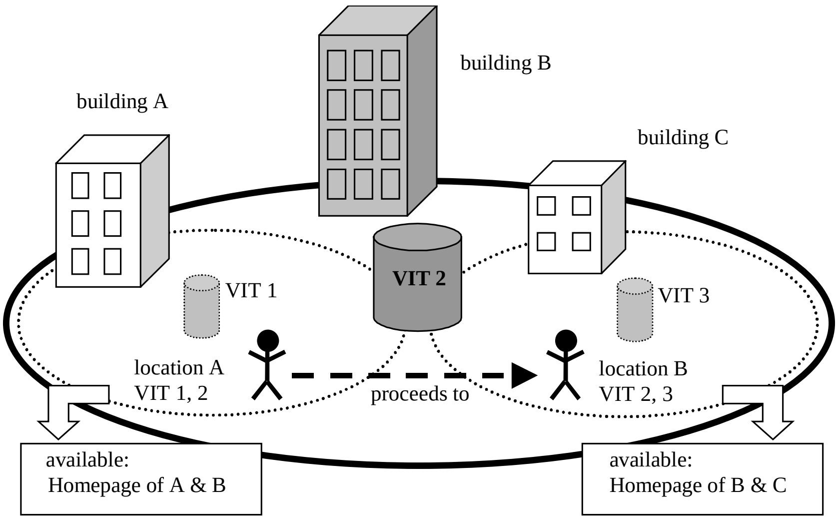

Fig. 1. Virtual Information Towers and their different areas of visibility. According to the position of a user, different information contents are provided since different VITs are visible.

Due to the fact that the spatial object models include virtual representations, they are called Augmented Area Models (see Fig. 2). Each Augmented Area Model (AAM) describes only a part of the whole NEXUS or Augmented World, respectively, e.g. an area of a city or only the detailed interior of a room. As an essential requirement, the models – which might overlap or include each other – must have a transparent structure so that their interoperability can be guaranteed.

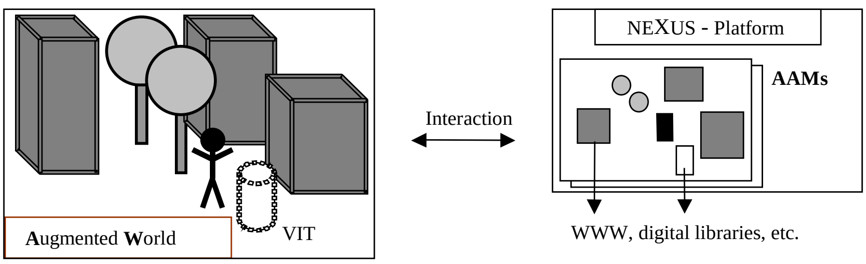

Fig. 2. The Augmented Reality and its computer representation within NEXUS

Since different applications need different resolutions of data sets, the spatial models have to manage data in multiple representations – from coarse to highly detailed levels. Furthermore, huge amounts of data must be stored and handled. For these reasons it is inevitable to use large and distributed databases as Spatial Model Servers (SMS). They have to function as a scalable middleware layer that can be accessed by multiple users. The functionalities of SMS like spatial selections or network analysis facilitate the support of the users on their way through the NEXUS world. With respect to Computer Aided Facility Management Systems (CAFM), not only the ability to perform geographical queries should be supported, but the interactive control of processes has to be realized, too. Therefore, an interesting objective comprises the implementation of active functionalities into the system in a way that objects of the real world can be controlled by their corresponding representations in the computer models and vice versa. Moreover, participants of the NEXUS system should be able to identify features of their surrounding by just pointing at them. For this reason, special sensors have to be integrated into the user's mobile device that allow the determination of directions.

The technical environment of NEXUS is assumed to be composed of so-called NEXUS stations – mobile, handheld computers like Notebooks, Personal Digital Assistants (PDAs) or Wearable Computers [18] which additionally dispose of other facilities, especially for positioning and mobile communication. In the future, appropriate sensor systems to acquire information about aspects concerning the state of the environment of the user like temperature, lighting, etc. will also be attached to the mobile devices. To determine a user‘s location in outdoor areas, the Differential Global Positioning System (DGPS) is used in combination with several supporting sensors. Indoor positioning requires different methods based on infrared signals (e.g. “Active Badge” systems), radio networks or indirect positioning techniques like image interpretation. The positioning information that is obtained from the sensor systems will be managed within special, also distributed databases for mobile objects, forming a so-called Location Service [11], [8]. The NEXUS station needs to be able to connect to the platform and external information sources like the Internet using wireless communication. For a wide area network (WAN) it can use the data service of a mobile telephone system like GSM [5] or of future systems which will be more suitable for data transmission like GPRS or UMTS. Inside of buildings, a Wireless LAN - e.g. according to the IEEE 802.11 standard - can be employed.

Within the presented publication, firstly an overview about location aware applications and their requirements concerning spatially related functionalities will be given (section 2). Following these illustrations, the global architecture of the platform is described (section 3). Section 4 focuses on the concepts that have been developed for distributed data handling and demonstrates the approaches which are applied in order to guarantee interoperability within NEXUS. Furthermore, the characteristics of the server component that is responsible for spatial data handling as well as some prototypical approaches are presented, before section 5 finally concludes the paper.

2 Characteristics of Location Aware Applications

Location aware applications know the current position of mobile users since they receive this information either directly from sensor systems or indirectly via a Location Service. Thus, they are able to determine user locations within spatial models and can therefore provide any information which is important at the current position of a user immediately. Furthermore, location aware applications offer special services for their users, particularly concerning the solution of typical GIS tasks like navigation or spatial queries.

Generally, there are different kinds of applications dealing with location awareness. More global approaches like city-, traffic- or emergency information systems stand opposite to rather local applications like museum-, exhibition- or department store information systems. With regard to the requirements of the NEXUS infrastructure, it is a precondition that the users can switch between the different applications.

To convey an impression of the functionalities that the NEXUS platform has to offer, this section will illustrate the main ideas by means of an application scenario. We assume that a fictitious businessman wants to attend a conference in a city he has never been to before and tries to profit from the services of NEXUS. He arrives at the main station and tries to get to the conference center. Via his handheld computer device, he is able to access the user interface provided to him by NEXUS-based applications. The user specifies his destination and utilizes the navigation system that leads him to the location where the conference takes place. Maps and photographs of significant points as well as synthesized voice output support him on his way and ease the orientation. Having arrived at the conference center, the NEXUS client finds himself located within the area of visibility of a Virtual Information Tower especially built up for the event. It allows the access to information concerning the schedule of talks, the conference building and its facilities, companies presenting themselves and links to their WWW sites, other participants, etc. To get into contact with persons working in the same field of activity he does, the businessman leaves a message at the virtual blackboard of the VIT.

After the conference, he applies NEXUS to search for restaurants located in the vicinity in order to have dinner. The result of the query shows a selection of objects and via their internet links, the user can inform himself about the menus of the restaurants. He chooses his favourite one and lets NEXUS guide him there. On the way, the NEXUS participant passes by a famous building which he wants to know more about. The telepointing device included in his PDA allows the identification of the object within the computer model and so any information available about the building can be displayed on his computer screen. He proceeds, but after a short while an audio signal calls his attention to another attraction close to his current position, a museum of modern art. Since the businessman specified his interest for museums when NEXUS generated his user profile, he was registered for spatial services that notify him if he is in the surrounding of such objects. The user decides to visit the museum after dinner.

Within the project, we have defined the spatially related requirements of location aware clients (that have partly been illustrated by the application scenario) in a Use-Case model. It specifies the actions that the platform must be able to perform (see Fig. 3). Between the different Use-Cases, generalization relationships can exist: the generalized Use-Case is contained in the superordinated one as a black box. Use-Cases are called abstract, if they only exist to be used by other Use-Cases; otherwise they are concrete [9].

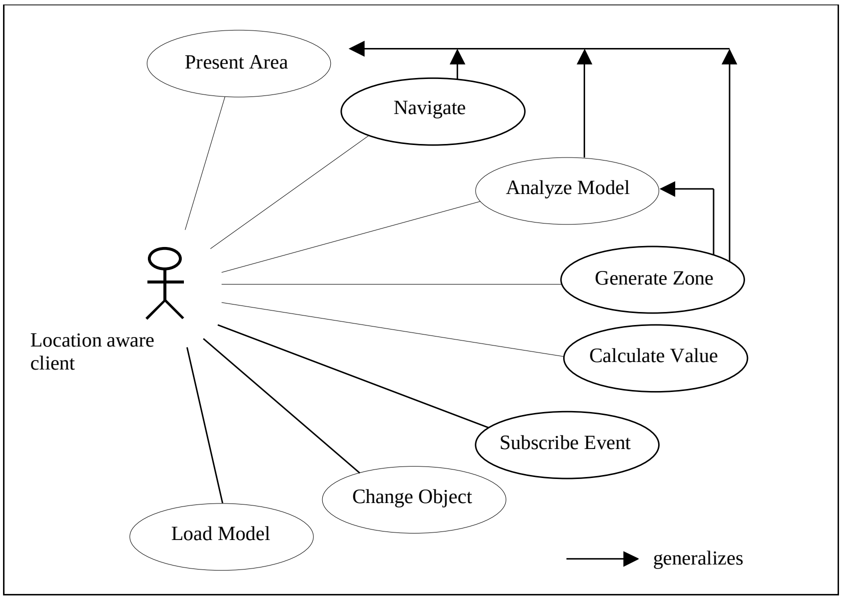

Fig. 3. Use-Case diagram showing the spatially related services that are demanded by location aware applications

The Use-Case *Present Area* allows location aware clients to provide a map for their users, showing their immediate environment in 2D or 3D. It is generalized by *Navigate* (function: find shortest/best paths), *Analyze Model* (function: perform queries on attributive or spatial properties) and *Generate Zone* (function: generate 2D or 3D buffers around spatial objects). *Calculate Values* is responsible for measurements and statistical computations. Users of location aware applications can furthermore utilize the *Subscribe Event* Use-Case in order to register themselves for receiving notifications if a spatially related event occurs (e.g. if they access a predefined area as it was described in the application scenario). If users are authorized to manipulate the geometry or attributes of spatial objects, *Change Object* will be applied. *Load Model* is a Use-Case that has been determined for mighty clients that can carry out the processing on their own and simply download the data they need from the platform.

3 Architecture of the NEXUS Platform

The NEXUS platform is designed to be a generic infrastructure that supports all kinds of location aware applications. To guarantee the functionalities that the platform claims to provide, the system is composed of four units which have to cooperate: the user interface, the sensor systems, the communication and the information or data management (see Fig. 4). The latter will be described in more detail in the next section, whereas the other components are presented in the following.

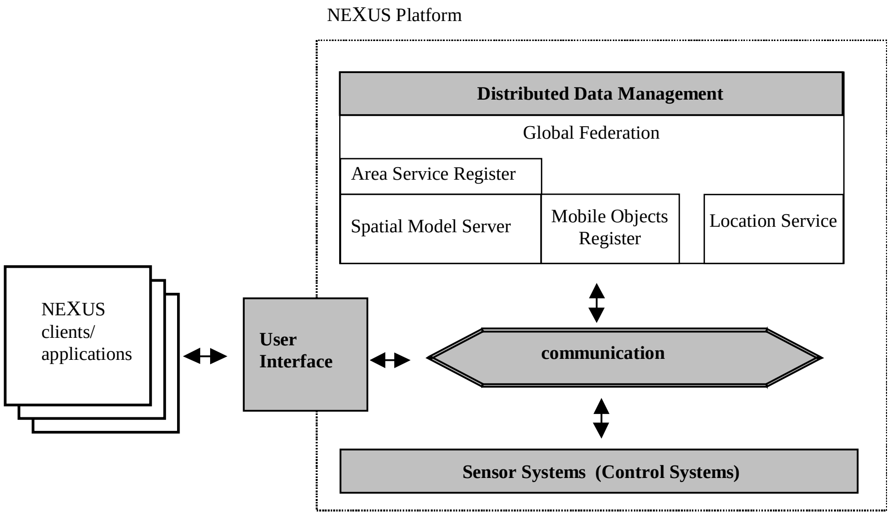

Fig. 4. Architecture of the NEXUS Platform

The NEXUS clients or location aware applications, respectively, access the platform via a standardized *user interface* which is running on the mobile device carried by the user. It has to take into account the individual demands of the various information systems and must facilitate a seamless interaction with the infrastructure for the user, i.e. the performance of queries has to function in a straightforward manner and the visualization of results must be intelligible as well. Furthermore, the interface has to adapt to the different kinds of NEXUS stations, especially concerning the different levels of computing power, the different amounts of memory, the different levels of network connection or the different displays. A PDA will e.g. require a completely different user interface than a Wearable Computer. Of course, the NEXUS system has to guarantee platform independent usage and so the different operating systems must also be reflected by the user interface.

Sensor systems on the one hand have to measure global indicators of the real world like temperature or illumination. On the other hand, they must detect rather object related information, especially the position of the NEXUS participants. Therefore, the NEXUS station has to be equipped with appropriate sensors. They record the different data and forward them to the data management component to keep the Augmented World Model in a highly current state. But often, the various data have to be combined to one final value. In the case of positioning data, for example, the different sensors like DGPS, pedometers, digital compasses, etc. yield individual measurements which have to be aggregated to one coordinate pair on the server side so that this information can be used at all.

Furthermore, the sensor systems are also related to the functionalities concerning the control of objects of the real world. Authorized users will have the opportunity to change the properties of objects in the Augmented World Model and to propagate these changes to the appropriate server via their NEXUS station. By modifying the computer representation, the corresponding real world object can be induced to perform the specified event. For example, a user could select a certain room on the display of his computing device, change its attributes concerning the conditions of illumination and send the modifications to the platform. The data management component would carry out the commands and thereby initiate a communication process that triggers the light switch and changes the illumination in the real environment.

The *communication* unit is responsible for the data exchange between the different components of the infrastructure. Generally, the NEXUS environment consists of different networks and so the handover between them has to be solved. Moreover, the efficiency of data transmission must be increased, e.g. by means of caching or hoarding techniques. Furthermore, the communication component has to realize the adaptivity of the system in order to prevent it from breakdown if parameters like the bandwidth change.

4 Distributed Management of Heterogeneous Spatial Data

The NEXUS platform provides information and services for its users. Since huge amounts of data and multiple customers have to be handled by the infrastructure, a centralized solution is inadequate if the system lays a claim to scalability. Thus, the data must be organized in a distributed environment. But not only their storage, also the processing has to be shared amongst different servers to optimize response time. Moreover, the spatial data must be offered in multiple representations in order to meet the individual needs of the different applications. Hence, appropriate algorithms to deduce all the necessary levels of detail have to be implemented. Besides that, the interoperability of the different data sets must be guaranteed. Another vital aspect concerning the acceptance of NEXUS deals with the privacy of data and the protection against illegal data manipulation. The preconditions mentioned above imply various problems that have to be solved within the project by appropriate approaches. Some of them will be discussed in this section.

4.1 General Concepts for Distributed Data Handling

The requirement to distribute the data across many servers leads to the necessity to establish a Global Federation. Since the management of user positions demands for other optimization procedures than the management of rather static objects (as they are usually dealt with in GIS), the storage and processing of NEXUS-related data has to be divided into separate parts. With regard to mobile objects, user specifications (like name, user profile, access rights) can be handled in conventional databases within an Object Register, whereas user positions have to be updated frequently and thus must be administered separately by a Location Service suitable for that purpose.

Concerning the geographical data, a so-called Area Service Register documents the assignment of the more static Augmented Area Models to the Spatial Model Servers (SMS). The Spatial Model Servers themselves store and manage the various computer models (including the virtual objects and their connections to external information spaces) and provide the necessary GIS functionalities (see Fig. 5).

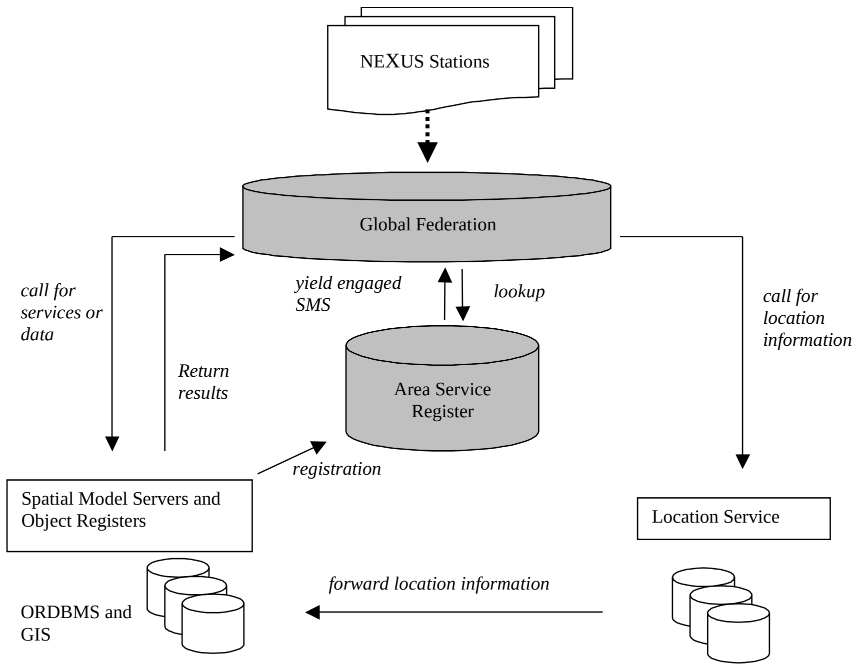

Fig. 5. Distributed data management and query strategies within NEXUS

Any query or update request that is posed on the NEXUS platform is first sent to the Global Federation component. The Federation then informs itself using the Object Register and the Location Service about the specifications and the current position of the client. Subsequently, it queries the Area Service Register in order to find out, which Augmented Area Models or Spatial Model Servers, respectively, have to be involved in the processing of the demanded task. It then causes the Location Service to forward the user position to the relevant SMS and calls these SMS to perform the required services. After having carried out the necessary steps, the selected SMS instances yield their partial results back to the Federation which aggregates them to one solution and finally propagates it to the client. For some tasks that are not decomposable (like shortest path analyses), though, this strategy is inadequate since it would not produce a reasonable result. If such a case occurs, the Global Federation must have the ability to load the necessary data from the servers involved and to process the query itself. For this reason, it has to possess GIS functionality as well.

4.2 Interoperability of Heterogeneous Data Sets

The Augmented World Model has to represent all the features of the physical world that could be of importance for any location aware application. Thus, the data sets have to be as detailed as possible. On the other hand, not each application needs data of high resolution, but rather prefers coarser representations that only contain the most important features. In consideration of efficient data analysis and visualization as well as low data transmission rates, it is very important to be able to provide the adequate level of detail for each application task. For this reason, it is a vital prerequisite for the spatial models that they can manage multiscale representations.

Concerning the interoperability of different spatial models, the specifications of the OpenGIS Consortium have to be taken into account. Moreover, since the infrastructure deals with heterogeneous data in multiple representations, the co-processing of differently resoluted data sets must be possible as well. For this reason, a strategy to homogenize data in multiple representations has to be developed. The following sections present the approaches pursued to guarantee the interoperability of heterogeneous data and describe methods for the derivation of multiscale data sets.

Interoperability Approaches. The applications must be able to perceive the combination of different Augmented Area Models as one big single model on which they can operate, i.e. the complexity of the distributed management of heterogeneous data has to be hidden from them. In order to guarantee this interoperability between the various spatial data sets, not only their data formats have to be exchangeable, but also the different representations of data have to be harmonized so that they can be processed together. This requires the definition of the relationships between spatial data sets which generally can be manifold. Augmented Area Models can share any geometric-topologic relation like overlap, adjacent, include, etc. [6]. Of equal importance are semantic relations which will be determined in terms of refinement or reduction, e.g. concerning geometry or attributes. The following example that describes the relations between two Augmented Area Models might clarify this approach:

Model A: generalized city model
Model B: detailed building model

Spatial relation:	B is completely contained in A (inclusion)
Semantic relation:	B refines the geometry AND the attributes of the building representation in model A

To allow the co-processing of differently resoluted data sets, the relations or links between the different representations, as well as the links between the individual objects not only have to be known basically, but they have to be specified in detail. Many difficulties are involved in these homogenization techniques. Some of the possible approaches to find a solution for this task are presented in the following section.

The topic also affects the query strategy that has been developed for the platform. Generally, the representation of a data set which is used to process a query depends on the kind of application and on the type of query. If multiple spatial models are needed to perform a procedure, their possibly varying levels of detail have to be adjusted so that a common usage can be guaranteed. The homogenization has to be carried out by the Spatial Model Servers if they do not provide the suitable representation already. Reflecting this approach, the scheme that is displayed in Figure 5 has to be extended. Within the course of a query, the Global Federation not only has to ask the Area Service Register where the suitable data are located, it also has to specify which representation is needed. Thereupon, the Area Service Register finds out by checking the meta information of the data sets involved if the required resolution can be derived. Only then the Federation initiates the communication with the appropriate Spatial Model Servers and calls them to carry out the demanded task.

In the context of interoperability, the specifications of the OpenGIS Consortium, Inc. [14] also play an important role for the NEXUS project. Especially the interfaces for the implementation of vector objects, attributes and spatial references as well as the specifications for queries and feature modifications will form a basis for interoperability within NEXUS. With respect to the Area Service Register, the concepts for Catalog Services which establish an infrastructure for "spatial search engines" and deal with standards for metadata information, must be integrated. The OpenGIS presentation specifications that aim at providing geospatial technologies for mobile devices as well as the feature identity and relationships specification [13] are also highly coinciding with the demands of the NEXUS platform.

Linking Heterogeneous Data Sets. Data sets may differ considerably, even if they represent the same physical entities. Besides differences in data formats, there are discrepancies concerning thematics and geometry. They mainly result from different underlying data models, different acquisition time, different operators capturing the data, etc. Another source of differences are gross errors in the data. Also, the nature of the objects can result in differences, e.g. natural objects that have a boundary that is not clearly defined [3]. Thus, beside the management of thematic and geometric differences, also imprecise or missing data has to be taken care of. The integrate use of all available data sources has to provide mechanisms for identifying and linking corresponding objects [10]. These links have to be either explicit or implicit - based on common geometry [17].

The integration of data sets of similar geometry and semantics can be realized with conflation techniques, relying either on geometry alone, or also including context in terms of relations [19]. Data sets of different resolution have to be harmonized with respect to semantics and geometry. Multiscale representation can in principle be generated from a highly detailed data set. This has the favorable effect that the links between the objects are produced directly. This generalization problem is tackled by cartographers as well as by computer graphics and computer vision specialists. Depending on the data type, different generalization operators are possible. In image processing, raster maps are transferred to image pyramids simply by smoothing operations [1]. 3D-surfaces are tackled with mesh simplification algorithms [4], [15]. Such approaches mainly take geometry into account, assuming that the biggest objects are also the important ones. In order to include also semantics, other approaches have to be used. Here, model based or knowledge based techniques can reflect the importance and the behaviour of the objects [12].

The following example (Fig. 6) shows a multiscale data set of a test area of the city of Stuttgart. The highest level of detail (LOD 3) represents buildings in a complex 3D structure. It was acquired by a method that combines digital ground plans with height information from laser scanner data and is able to capture detailed building geometries with different roof types automatically [7]. The second level of detail simplifies the buildings to single entities in 2.5D. This data set was generated by assigning the maximum height of the building ground plan as the z-value. The coarsest level represents only aggregates of buildings in terms of a city block structure. It could be derived by integrating the city block boundaries from the municipal information system (cadastral data set) and the average heights of the buildings inside a city block. Thus, the scale-dependent representation was on the one hand generalized from the data set with the highest level of detail, however at one step (transfer from LOD2 to LOD1) combined with other information (in this case the city block boundaries). This process guarantees that the links between the objects in the different LOD's are known and can be exploited by the system.

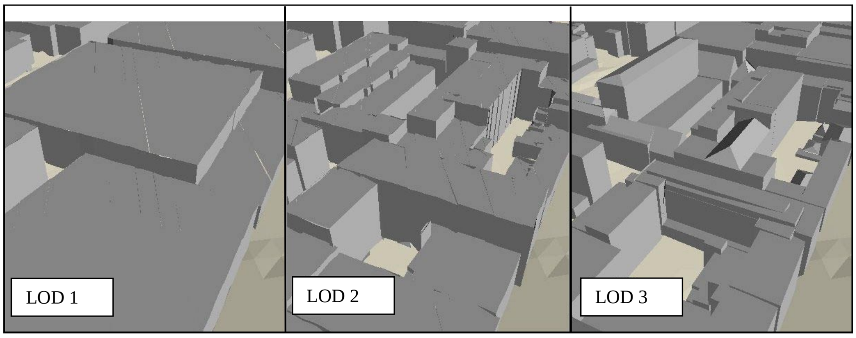

Fig. 6. A Scene of the Virtual City Model of Stuttgart in different levels of detail

4.3 Storage and Processing of Spatial Data within SMS

For the management of the various Augmented Area Models, special servers have to be established that support the storage and processing of spatial data. Typically, object relational DBMS like Oracle, DB2 or Informix in combination with spatial data blades like the Spatial Database Engine by ESRI are used for that purpose. Although SQL meanwhile allows geographical queries, it does not cover the complete functionalities of a GIS. Thus, geographic information systems have to be integrated into the distributed SMS environment as well. In comparison with centralized versions, distributed GIS additionally offer a more sophisticated support for data sharing and also improve the reliability and the opportunities for system growth. Another important advantage concerns the increased efficiency that can be achieved by distributing costly operations to multiple sites for parallel processing [20]. Therefore, the general concepts for data handling presented in section 4.1 have to be extended by query optimization strategies. With regard to distributed DBMS, there are four basic steps that have to be carried out in order to reduce response time: query decomposition, data localization, global optimization and local optimization. The query decomposition which divides a query into subqueries, each of which can be executed at a site, and the global optimization that is responsible for developing a strategy for processing the subqueries will be solved by the Global Federation. The Area Service register has to deal with the data localization and the Spatial Model Servers themselves must optimize the subqueries at their individual sites.

Interfaces of SMS. Generally, the Spatial Model Servers should consider that different clients with varying computing power are accessing the platform and so an interface must be provided for each of them. Mighty clients only need the feasibility to receive complete spatial models and process the data on their own. Usually, the different kinds of location aware applications employ a customer API that is located on the client side and makes the GIS functionalities of the server available (see following section). Furthermore, also simple clients have to be supported by an HTML- or WML-based interface. For system administrators, a management API has to be prepared (Fig. 7).

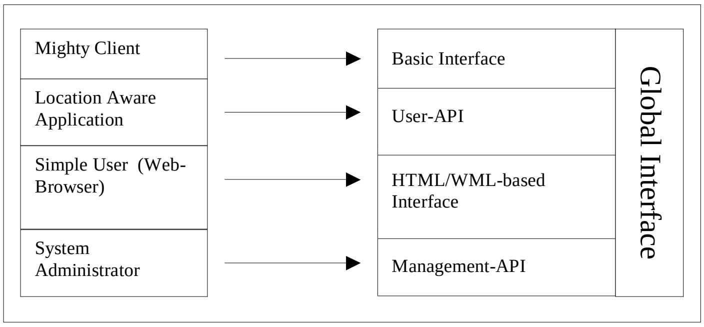

Fig. 7. The different interfaces for Spatial Model Servers

Prototypical approaches. Until now, a first version of a User-API has been developed as a Java library that uses the commercially available GIS product ArcView by ESRI together with its extension Internet Map Server (IMS) as a prototypical Spatial Model Server. It provides the services that have been defined as Use-Cases in section 2, i.e. it allows to carry out spatial selections (*objectsInArea*, *closestObjects*, etc.), navigation tasks (*findShortestPath*), attributive queries (*getAttribute*) and graphical presentations (*getCurrentView*) [16]. The Java packages that are located on the client side shield applications from internal details of the GIS. For this reason, the SMS can be accessed from virtually any platform.

By means of the User-API, the applications form URLs that denote the demanded services along with their arguments. They are forwarded via a basic Federation component to the Web Server and further to the ESRIMap Web Server Extension (WSE) that is employed by the SMS in order to accept service invocations. The WSE also balances the work load, making use of the capability of ArcView-IMS to distribute service calls roughly equally amongst identical ArcView instances. This enhances the performance of SMS. Thus, the WSE selects the SMS node that is most suitable for the demanded task and propagates the HTTP-String there. ArcView's Internet Map Server extension then parses it and triggers the execution of server scripts that are appropriate for the handling of the specified services and so the results of the queries are finally produced. The spatial data of the SMS may reside on remote ORDBMS that are capable of storing object geometries.

Prior to launch, an SMS signs itself up at the Area Service Register and identifies the area, the object classes and the levels of detail of the data sets it manages. Moreover, an SMS has to be able to receive the positioning information of the users from the Location Service. Generally, this communication could be done directly with ArcView, but the programme does not support the RMI mechanism that is applied by the other NEXUS components. For this reason, a communication-DLL has to be established that provides this method and mediates between the GIS on the one hand and Location Service and Area Service Register on the other hand (see Fig. 8).

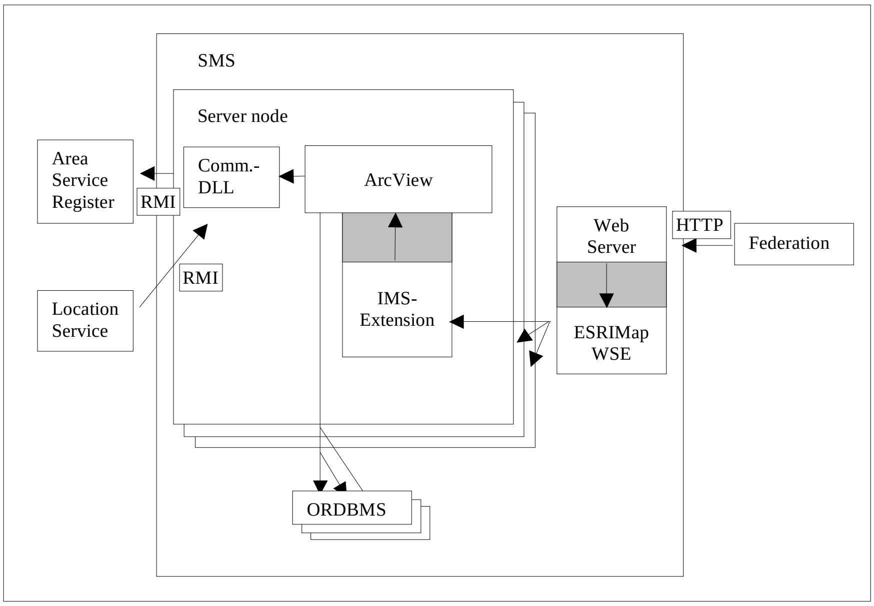

Fig. 8. The architecture of an SMS based on ArcView GIS

For an efficient communication within a network, the amount of data that has to be transferred between client and server must be minimized. For this reason, the SMS should preserve the states of its location aware clients, e.g. containing the map extent or the specification of results of previous queries. Thus, also less data has to be stored at the client side which is particularly important for weak clients. ArcView IMS, however, only allows stateless communication and therefore the management of the clients had to be integrated into the server. At first, a client has to open a session before it can invoke services from SMS. Then, the state of the client is stored within different tables that are updated by server scripts each time there are changes until the session is closed again.

To demonstrate the capabilities of an SMS, a sample application has been developed that makes use of the customer API and allows the invocation of different services via a GUI. It enables users to sign themselves up at the SMS. The SMS then integrates the user as a mobile object in the spatial model. By entering a distance in meter, the application selects and highlights all spatial objects that lie within this distance to the user. From the current set of selected objects, the closest one can also be determined. Additionally, objects can be selected by specifying the attributes that have been assigned to them in the context of the application. Besides that, the application accomplishes navigation tasks. It yields the optimal path that visits points specified by the user. The results of spatial queries are prepared as maps and sent to the client as images.

5 Conclusion and Outlook

Within the NEXUS project, the focus has to be put on the development of efficient techniques for the management of spatial data in order to support different kinds of location aware applications. So far, the concepts for a distributed data and query handling have been evaluated and specified. A prototypical version of a Spatial Model Server could be implemented on the basis of a commercial GIS. In order to allow the usage of the data and services available on the server side, a User-API has been provided that can be accessed by NEXUS clients. A sample application uses this interface to interact with the SMS, being able to perform advanced GIS queries like navigation or spatial selections. Upcoming tasks will have to deal with a more sophisticated integration of the Federation component into the data management unit of the NEXUS platform and also with the implementation of strategies concerning the processing of queries in a distributed environment. Especially the problems that arise from the need for an interoperable data handling of multiple representations will have to be faced. A first prototype of the NEXUS infrastructure will be available in two years.

References

1. Ackermann, F., Hahn, M.: Image Pyramids for Digital Photogrammetry. In: Ebner, H. et al. (eds.): Digital Photogrammetric Systems, Karlsruhe (1991) 43-58
2. Anders, K.-H., Sester, M.: Methods of Data Base Interpretation – Applied to Model Generalization from Large to Medium Scale. In: Förstner, W., Plümer, L. (eds.) SMATI’97: Semantic Modelling for the Acquisition of Topographic Information from Images and Maps. Birkenhäuser (1997) 89-103
3. Burrough, P. A: Natural Objects with Indeterminate Boundaries. In Burrough, P.A., Frank, A.U. (eds.) ESF-GISDATA, Vol. 2 (1996) 3-28
4. De Floriani, L., Puppo, E.: Hierarchical Triangulation for Multiresolution Surface Description. ACM Transactions on Graphics 14(4) (1995) 363-411
5. Eberspächer, J., Vögel, H.-J.: GSM Switching, Services and Protocol. B.G. Teubner, Stuttgart Germany (1998)

6. Egenhofer, M.: Reasoning about Binary Topological Relations. In: Günther, O. Hrsg.: Advances in Spatial Databases - Proceedings of the 2nd Symposium on Large Spatial Databases SSD '91. (1991)
7. Haala, N., Brenner, C.: Rapid Production of Virtual Reality City Models. GIS 12(2), (1999) 22-28
8. Hohl, F., Kubach, U., Leonhardi, A., Rothermel, K. and Schwehm, M.: NEXUS – an Open Global Infrastructure for Spatial-Aware Applications. Proceedings of the 5th International Conference on Mobile Computing and Networking. Seattle, (1999)
9. Jacobsen, I., Booch, G., Rumbaugh, J.: The Unified Software Development Process. In: Booch, G.; Jacobsen, I., Rumbaugh, J. (eds): Object Technology Series, Reading, MA, USA (1999)
10. Laurini, R.: Spatial multi-database topological continuity and indexing: a step towards seamless GIS data interoperability. Int. J. Geographical Information Science, Vol. 12, No. 4 (1998) 373-402
11. Leonhardi, A., Kubach, U.: An Architecture for a Universal, Distributed Location Service. Proceedings of the European Wireless '99 Conference. München, (1999)
12. Mackaness, W., Weibel, R., Buttenfield, B.: Report of the 1997 ICA Workshop on Map Generalization. Technical report, 19th until 21st of June 1997. http://www.geo.unizh.ch/ICA/. Gävle Sweden (1997)
13. McKee, L., Kottman, C.: Inside the OpenGIS Specification. Photogrammetric Engineering & Remote Sensing Vol. 65, No. 12 (1999) 1345-1359
14. Open GIS Consortium: Inc. Open GIS Consortium. http://www.opengis.org (1999)
15. Schmalstieg, D.: Lodestar – An Octree-Based Level of Detail Generator for VRML. In: Proceedings of SIGGRAPH Symposium on VRML (1997)
16. Schützner, J.: Entwicklung einer Server-Komponente für ein räumliches Modell in NEXUS. Diploma Thesis Nr. 1768, Faculty of Computer Science, University of Stuttgart. (1999)
17. Sester, M., Anders, K.-H., Walter, V.: Linking Objects of Different Spatial Data Sets by Integration and Aggregation, GeoInformatica(4) (1998) 335-358
18. Starner, T., Schiele, B., Pentland, A.: Visual Context Awareness in Wearable Computing. Proceedings of the Second International Symposium on Wearable Computers (ISWC'98). Pittsburgh, PA, USA. (1998) 50-57
19. Walter, V., Fritsch, D.: Relational Matching of Spatial Data. International Journal of Geographical Information Science(5) (1999) 445-473
20. Wang, F.: Query Optimization for a Distributed Geographic Information System. Photogrammetric Engineering & Remote Sensing Vol. 65, No. 12 (1999) 1427-1438