MULTI-SCALE DATA SETS IN DISTRIBUTED ENVIRONMENTS

Steffen VOLZ*, Monika SESTER*, Dieter FRITSCH*, Alexander LEONHARDI**
*Institute for Photogrammetry (ifp), Stuttgart University, Germany
Geschwister-Scholl-Strasse 24D, D-70174 Stuttgart
firstname.lastname@ifp.uni-stuttgart.de
**Institute of Parallel and Distributed High-Performance Systems (IPVR), Stuttgart University, Germany
Breitwiesenstrasse 20-22, D-70565 Stuttgart
Alexander.Leonhardi@orgel.informatik.uni-stuttgart.de

Technical IV/1

KEY WORDS: Location Awareness, Distributed Databases, Multiple Representations, Information Infrastructure.

ABSTRACT

Within the scope of the research project „NEXUS“ that is being worked on at the University of Stuttgart in cooperation between the Institute for Photogrammetry, the Institute of Parallel and Distributed High-Performance Systems and the Insitute of Communication Networks and Computer Engineering, a generic platform supporting location aware applications with mobile users will be developed. The central component of the infrastructure consists of dynamic spatial models. Since huge amounts of data have to be handled by the platform, they have to be distributed on multiple servers. Furthermore, the spatial data have to be delivered in multiple representations in order to provide the adequate level of detail for each application.

The spatially aware applications which are using the spatial models are available for the users on small handheld computers. These computers are amongst others equipped with sensors for positioning and components for mobile communication. Therefore, the applications know the position of the users within the spatial models – they are „location aware“. Moreover, the users can receive spatial data from distributed model servers by means of wireless communication. Thus, they can not only move through the real but also through the virtual world. In order to support the users on their way, GIS functionality has to be provided: they have to be able to perform spatial queries on their environment and to navigate themselves through the virtual models and the real world, respectively.

The paper will give an introduction to the general concepts of the project. It presents the approaches concerning multiple representations and describes the strategies for distributed data handling that will be applied within the platform, including some remarks on prototypical solutions.

1 INTRODUCTION

In times of the internet and growing computer networks, the issue of distributed data and processing is strongly influencing the field of computer science. This development especially affects Geographic Information Systems since they have to deal with huge amounts of data and apply computationally demanding operations. Additionally, various GIS applications have to make use of geographical data from heterogeneous resources or institutions, respectively. Thus, a growing need for interoperability between the different systems and data formats can be observed. Besides that, the applications not only have a demand for completeness of spatial information in a distributed environment, but they also need data in appropriate representations to guarantee an efficient data analysis and a clear visualization.

Considering these topics, the NEXUS project focuses on the development of a data management component that serves as a basis for so-called location aware applications. Spatial applications of this type know the current position of their 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.

As a data basis for location aware applications, spatial computer models will be provided. But, in order 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, structure the data stored within other information systems spatially. Each VIT has a certain location and - according to its significance - a predefined area of visibility (see figure 1).

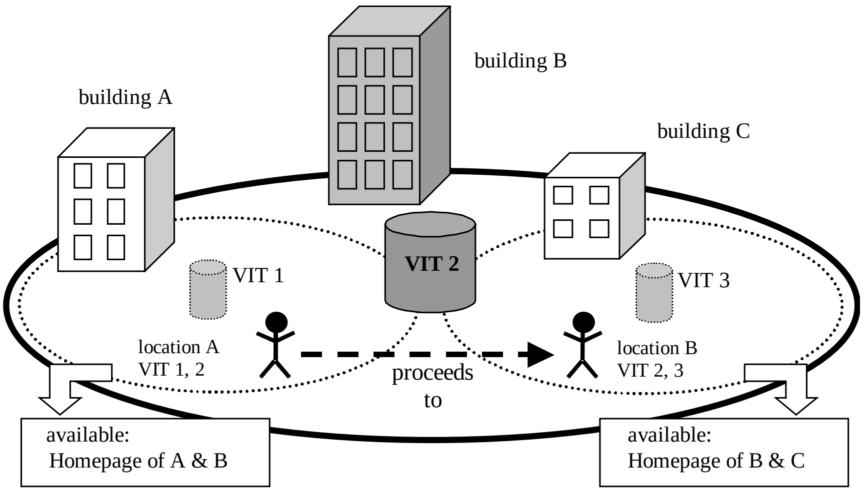

figure 1. Virtual Information Towers and their different areas of visibility.

Due to the fact that the spatial object models include virtual representations, they are called Augmented Area Models. 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.

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 – as well as data of different origin and thematic. 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. 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 (Starner et al. 1998) which additionally contain 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, forming a so-called Location Service (Leonhardi & Kubach 1999, Hohl et al. 1999). 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 (Eberspächer & Vögel 1998) 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.

2 MULTIPLE REPRESENTATIONS WITHIN NEXUS

Due to the various applications of NEXUS, ranging from a city information and guidance system, via information systems of department stores or airports, up to an information system of an exhibition, a huge variety of data has to be available and managed in an integrated way. Possible data sources are: scanned map sheets, images, city maps, digital

3D-city models, plans of the interior of buildings in arbitrary detail - up to a 1:1-representation of a specific product in a department store.

Thus, the data sources cover a variety of topics, have different spatial extents, and also contain information in various levels of detail. These different levels of detail are not only useful to highlight different properties of the objects. Another important aspect concerns the volume of the data. In order to work with it in a reasonable way, the data has to be available at an appropriate level of detail (or resolution) for each application. Consider e.g. the application of getting an overview of a whole city or a department store: in this case only coarse information is needed, since all the details are only disturbing the visual impression and hinder to see the important information. Only when getting closer more and more details should appear - however related to smaller and smaller extent. This means, that the information density is always the same on all the levels of detail. This coarse-to-fine-approach of visual inspection helps to manage the huge information content, but it also supports the analysis of the data, which works best and most efficient on the level of detail that corresponds to a given application.

An ideal could be to harmonize all these data sources in one single data set and extract all the necessary information for a certain application from it. Such a scenario would require operations for the extraction of different thematic components, namely automatic analysis and interpretation strategies. In order to go from one scale to the next, generalization operations are needed (Mackaness 1997). However, such generalization operations are not yet available: although there are promising and well established algorithms for dedicated problems, an integrated, global solution does not seem to be at hand in the near future. Most of the problems deal with the geometric generalization - however the question of how to aggregate the attributes is not yet solved, since it is very much application dependent. Furthermore, generalization algorithms are well suited to transfer dedicated scale ranges, however it is doubtful, that the whole possible scale range from micro to macro can be described by generalization operations. The gaps in the representations that occur at certain scales sometimes are not predictable.

Thus, NEXUS has to integrate data sets of different thematic and geometric resolution and spatial extent. Since data sets may overlap thematically and geometrically, their interoperability must be guaranteed by appropriate approaches (Laurini 1998). Such problems are being tackled by the OpenGIS initiative (OpenGIS, 1999), striving for solutions of integrating data and operations. Concerning data structures, the specifications are already given and this problem will sooner or later become obsolete. The more important problem is currently on the research agenda of OGIS, namely the specification and matching of thematic content. (Bishr 1998).

When it comes to the integration of heterogeneous data sets, one faces the problem that, although the same physical entities are modeled, the resulting data sets will not necessarily coincide. Finding corresponding objects in heterogeneous data sets is called homogenization or conflation (Saalfeld 1988, Walter & Fritsch 1999). The approach is therefore to integrate existing data sets by establishing the relations (links) between them. The relations between the data sets can either be implicit - using common geometry, or explicit, by directly linking corresponding objects in the data sets (Sester et al. 1998). Then, this knowledge can be exploited to make the transitions from one data set to the next. The links can also be described in terms of generalization operations, that can be applied for the transition between dedicated scale ranges. Ideally, this integration should occur on-the-fly, thus avoiding redundancy in data storage.

This scenario then allows to integrate new data sets, as they become available, which is in the spirit of having an open system comparable to the WWW, where any kind of information can easily be integrated. First of all, their geometric and thematic properties have to be extracted. These specifications then lead to the integration into the system.

3 ARCHITECTURE OF THE 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 has 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 figure 3). The latter will be described in more detail in the next section, whereas the other components are presented in the following.

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.

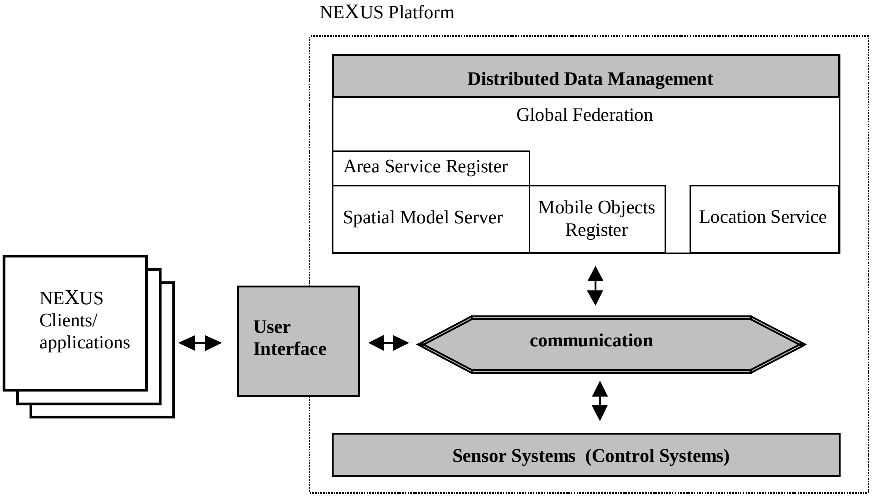

figure 2. Architecture of the NEXUS-Platform

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 (e.g. via a CAFM system).

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 DATA MANAGEMENT WITHIN NEXUS

As the Augmented World Model contains a large amount of information, a centralized server will not be able to cope with the necessary amount of queries, calculations and updates. Instead, we propose a distributed architecture, where local servers are responsible for certain areas. They are ordered hierarchically, in a way that servers at a higher layer are covering a larger area with a lower accuracy (see figure 3). Thus, each type of query can operate on the appropriate data resolution.

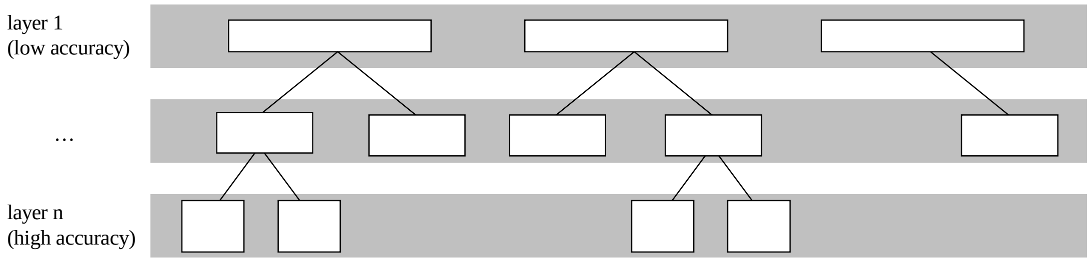

figure 3. A possible multi-tiered architecture of spatial data servers.

The requirement to distribute the data across many servers leads to the necessity to establish a Global Federation. Its functionality is described in section 4.1. 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 (see section 4.3). 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 Area Service Register can be organized in terms of a spatial indexing mechanism, like an R-Tree. In addition to the geometry, also the thematic information content of data sets should however be refered to here. 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 section 4.2).

4.1 General query strategies

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 solution back to the Federation which propagates it to the client.

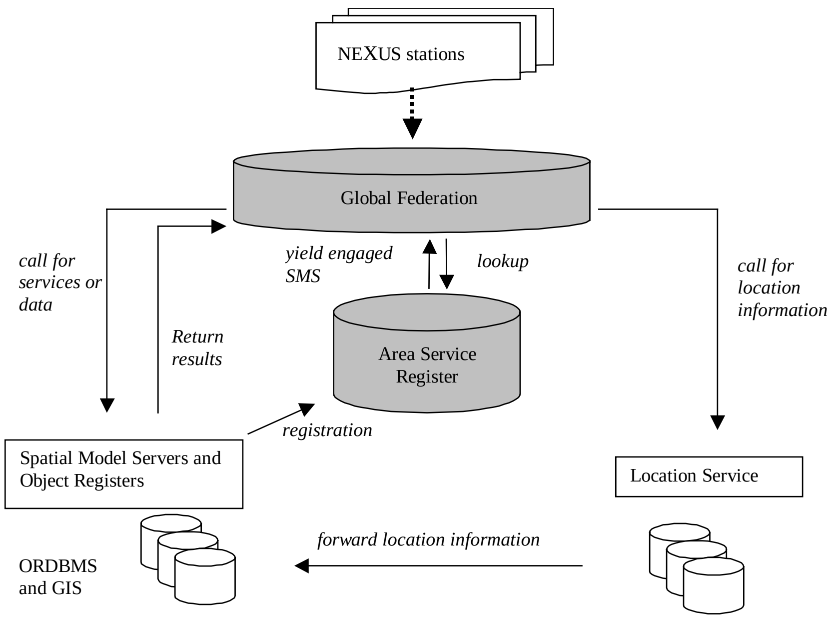

figure 4. Distributed data management and query strategies within NEXUS.

4.2 Distributed Spatial Data Handling

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, e.g. network analyses can not be carried out efficiently within the DBMS yet. 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 (WANG 1999). Therefore, the general concepts for the course of queries presented in chapter 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.

Besides these strategies for the improvement of response time by optimizing operations according to the type of query and the amount of effort that a procedure demands, another factor plays an important role for the processing in a distributed environment: the type of client. 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 User-API that is located on the client side and makes the GIS functionalities of the server available (see section 4.4). 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.

The way queries are handled is also decisively influenced by the distribution of the data. The easiest case occurs, of course, if one Spatial Model Server stores all the necessary spatial data for a demanded query (e.g. a spatial selection of objects on the basis of a buffer; see figure 5, Query 1). But if data from multiple servers have to be involved in a processing task (see figure 5, Query 2), different strategies are possible. As one solution, the server that holds the biggest part of the data needed (i.e. server B in Query 2) acts as the main server which receives the remainder of the data from the other server(s) (i.e. server A in Query 2) and processes the query itself. But if heavy operations have to be carried out, this concept does not perform in an optimal manner. Thus, in terms of query decomposition and optimization of response time, another solution is more appropriate: each server yields partial results which are aggregated by the Federation component. Yet, there are other restrictions. In case of network analyses (see figure 5, Query 3), this strategy does not guarantee a reasonable result since this type of query can not be decomposed. Hence, another approach suggests to forward all the necessary data to the Federation which aggregates them and processes the query itself.

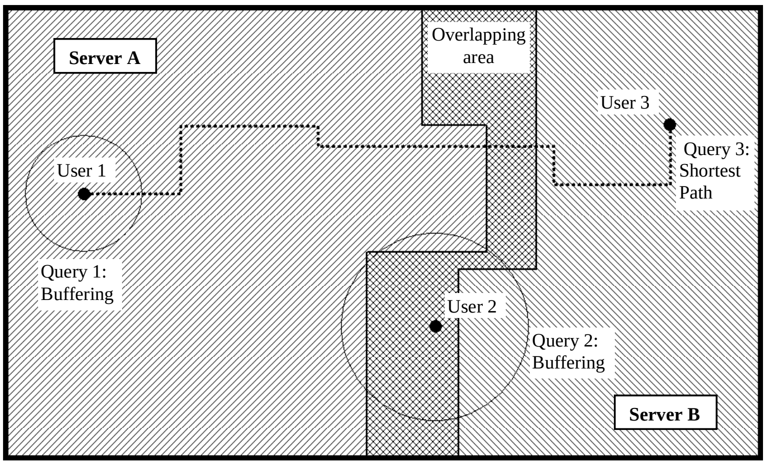

figure 5. Single and Multi-Server query processing

An important topic within the context of distributed data management deals with the interoperability of spatial data. On the one hand, the general interoperability of data in different formats has to be taken into account. For this reason, the specifications of the OpenGIS Consortium have to be reflected within the NEXUS project. On the other hand, also the co-processing of differently resoluted data sets must be realized (see section 2).

4.3 Management of the Location Service

A very important aspect of the Augmented World Model is the knowledge about the exact location of the mobile objects, especially of its users. The location information for the mobile objects is gathered from different types of sensor systems. As mentioned above, this can for example be a GPS sensor attached to the NEXUS-Station a user is equipped with. As the NEXUS platform requires location information with a high accuracy, this information is highly dynamic, depending on the movement characteristics of the object. Therefore, special mechanisms and a special component, the Location Service, have to be established for handling the location information. This service will have to deal with a high number of updates. To reduce them as much as possible, special update protocols will be used (see below).

In Leonhardi & Kubach (1999), an architecture for the Location Service has been proposed. Its functionality is as follows:

- A position query (locOfObj) returns the current location of a mobile object.
- A range query (objAtLoc) returns the identifiers of all objects inside a given area.
- An event mechanism notifies an application, if certain situations occur. For example, if a user enters or leaves an area (onEnters, onLeaves) or if a certain group of mobile users meets (onMeeting).
- The Location Service can also periodically update a certain item of location information with a given accuracy, which may be important for keeping an Augmented Area Model up to date.

As with the Spatial Model Servers, the location information for all mobile objects inside a given area are stored on one certain *Location Server*, which is associated with this area. All sensor systems that acquire location information for objects in this area report it to this Location Server. The Global Federation sends its queries to the nearest server, which answers it if the information is available locally. Otherwise, it forwards the query to the appropriate server(s), which it determines through the Object Register, in case of a position query, or the Area Service Register, in case of a range query. From the Object Register the central attributes of a mobile object (e.g., its maximum and average speed) and the knowledge, on which location servers the location information for a certain mobile object is currently stored, is obtained. It is used to find the appropriate Location Server for a position query. The Area Service Register is queried to find the Location Server(s) which are responsible for the area of a range query.

For the transmission of the location information, either from a mobile device with a GPS sensor or within the Location Service, special update protocols are used to reduce the necessary number of updates. In Leonhardi & Rothermel (2000), basic update protocols for transmitting location information efficiently have been compared. In order to improve the efficiency, different types of dead-reckoning protocols will be applied, where the server estimates the current location of the mobile object. The client calculates a similar estimation and sends an update when the current position differs from the estimated one by more than a certain degree. In contrast to Wolfson et al. (1999), the mechanism we plan to develop will not assume prior knowledge about the future route of the mobile object.

4.4 Prototypical solutions

So far, a prototypical Spatial Model Server has been established that uses the functionalities of ArcView GIS and its Internet Map Server extension. To access the server, a Java library has been developed as a User-API. It is located on the client side and shields the applications from internal details of the GIS. By the use of this interface, a client can carry out spatial selections (*objectsInArea, closestObjects*, etc.), navigation tasks (*findShortestPath*), queries on attributes (*getAttribute*) and graphical presentations (*getCurrentView*) (Schützner 1999).

To demonstrate the capabilities of the SMS prototype, a sample application has been developed that utilizes the User-API and provides some of the implemented services (e.g. navigation or also graphical presentations) via a GUI. The application forms an URL that denotes the demanded services along with their arguments each time a user wants to perform a query. At first, the HTTP string is sent to a web server machine where the so-called ESRIMap Web Server Extension (WSE) is located. This component is employed by ArcView in order to accept service invocations from the client. 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. Thus, the URL is forwarded by the WSE to the appropriate SMS node where it is dispatched by ArcView's Internet Map Server extension. This triggers the execution of commands within the GIS that are finally producing the query result. Depending on the type of query, the result is propagated to the client either as plain text or as an image that shows the generated map.

5 CONCLUSION

Within the NEXUS project, the focus has to be put on the development of flexible and efficient techniques for the management of spatial data in order to support different kinds of location aware applications. So far, the basic concepts for a distributed data and query handling have been evaluated and specified. Furthermore, strategies for the use of multiple representations and for the interoperability of heterogeneous data have been examined. 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 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 further analyzed. A first prototype of the NEXUS infrastructure will be available in two years.

REFERENCES

Bishr, Y., 1998: Overcoming the semantic and other barriers to GIS interoperability. Int. J. Geographical Information Science, Vol.12, No. 4, pp. 299-314.

Eberspächer, J. and Vögel, H.-J., 1998: GSM Switching, Services and Protocol, B.G. Teubner, Stuttgart, Germany.

Egenhofer, M., 1991: Reasoning about Binary Topological Relations. In: Günther, O. Hrsg., "Advances in Spatial Databases - Proceedings of the 2nd Symposium on Large Spatial Databases SSD '91".

Hofmann, C., 1999: A multi tier framework for accessing distributed, heterogeneous spatial data in a federation based EIS. In: Proceedings of the 7th International Symposium on Advances in Geographic Information Systems (ACM-GIS '99), Kansas City, MO, USA.

Hohl, F., Kubach, U., Leonhardi, A., Rothermel, K. and Schwehm, M., 1999: NEXUS - an Open Global Infrastructure for Spatial-Aware Applications, Proceedings of the 5th International Conference on Mobile Computing and Networking, Seattle.

Laurini, R., 1998: Spatial multi-database topological coninuity and indexing: a step towards seamless GIS data interoperability. Int. J. Geographical Information Science, Vol.12, No. 4, pp. 373-402.

Leonhardi, A. and Kubach, U., 1999: An Architecture for a Universal, Distributed Location Service. Proceedings of the European Wireless '99 Conference, München.

Leonhardi, A. and Rothermel, K., 2000: A Comparison of Protocols for Updating Location Information. Submitted to the ACM/IEEE Int. Conference on Mobile Computing and Networking, MobiCom 2000.

Mackaness, W., Weibel, R. and Buttenfield, B., 1997: 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.

McKee, L. and Kottman, C., 1999: Inside the OpenGIS Specification. Photogrammetric Engineering & Remote Sensing Vol. 65, No. 12, pp. 1345-1359.

Open GIS Consortium, 1999: Inc. Open GIS Consortium. http://www.opengis.org

Saalfeld, A., 1988: Automated Map Compilation, International Journal of Geographical Information Systems 2(3), pp. 217 - 228.

Schmalstieg, D., 1997: Lodestar - An Octree-Based Level of Detail Generator for VRML. in: Proceedings of SIGGRAPH Symposium on VRML.

Sester, M., Anders, K.-H. and Walter, V., 1998: Linking Objects of Different Spatial Data Sets by Integration and Aggregation, GeoInformatica(4), pp. 335-358.

Schützner, J., 1999: Entwicklung einer Server-Komponente für ein räumliches Modell in NEXUS. Diploma Thesis Nr. 1768, Faculty of Computer Science, University of Stuttgart.

Walter, V. and Fritsch, D., 1999: Relational Matching of Spatial Data. International Journal of Geographical Information Science(5), pp. 445-473.

Wang, F., 1999: Query Optimization for a Distributed Geographic Information System. Photogrammetric Engineering & Remote Sensing Vol. 65, No. 12, pp. 1427-1438.

Wolfson, O., Sistla, A., P., Chamberlain, S. and Yesha, Y., 1999: Updating and Querying Databases that Track Mobile Units. Distributed and Parallel Database Journal, 7(3), pp. 1-31.