Seismo-Surfer: A Prototype for Collecting, Querying, and Mining Seismic Data

Yannis Theodoridis

Dept. of Informatics,
University of Piraeus
GR-18534 Piraeus, Greece
http://thalis.cs.unipi.gr/~ytheod

Abstract. Earthquake phenomena constitute a rich source of information over the years. Typically, the frequency of earthquakes worldwide is one every second. Collecting and querying seismic data is a procedure useful for local authorities in order to keep citizens informed as well as for specialized scientists, such as seismologists, physicists etc., in order to study the phenomenon in its detail. A seismic data management system should meet certain requirements implied by the nature of seismic data. This kind of data is not solely characterized by alphanumeric attributes but also from a spatial and a temporal dimension (the epicenter and the time of earthquake realization, for example). Moreover, visualizing areas of interest, monitoring seismicity, finding hidden regularities or irregularities, and assisting to the understanding of regional historic seismic profiles are essential capabilities of such a system. Thus, a spatiotemporal database system, a set of data analysis and knowledge discovery techniques and a user-friendly visualization interface, compose the Seismo-Surfer, a prototype that, further to the above, aims to integrate seismic data repositories available over the WWW.

1 Introduction

Since the ancient years, humans have been feeling, recording and studying earthquake phenomena. Taking into account that at least one earthquake of magnitude $M < 3$ ($M > 3$) occurs every one second (every ten minutes, respectively) worldwide, the seismic data collection is huge and, unfortunately, ever increasing. Scientists, such as seismologists and physicists, record this information in order to describe and study tectonic activity. For a certain time period, tectonic activity can be described by recording geographic information, i.e. epicenter and disaster areas, together with attributes like magnitude, depth, etc.

Desirable components of a seismic data management application include tools for quick and easy data exploration and inspection, algorithms for generating historic profiles of specific geographic areas and time periods, techniques providing the association of seismic data with other geophysical parameters of interest such as soil profile, geographic and perhaps specialized maps (e.g. topological and climatological) for the presentation of data to the user and, topline, visualization components supporting sophisticated user interaction.

In particular, we distinguish three user profiles that such an application should support:

– Researchers of geophysical sciences, who could, for example, be interested in constructing and visualizing seismic profiles of certain regions during specific time periods or in discovering regions of similar seismic behavior.
– Key personnel in public administration, usually asking for information such as distances between epicenters and other geographical entities like schools and heavy industries.
– Citizens (web surfers), who are interested in general about the phenomenon, and might query the system for seismic properties of general interest, e.g. for finding all epicenters of earthquakes in distance no more than 50Km from Athens.

Management of seismic data, due to their spatiotemporal nature, demands more than a relational or an object-relational Database Management System (DBMS) and a Geographical Information System (GIS) on top of the DBMS. Recent advances in the areas of non-traditional Databases, Knowledge Discovery in Databases (KDD) and Data Visualization allow better approaches for the efficient storage, retrieval and analysis of seismic data. For example, commercial DBMS's have been already providing tools for the management of spatial (2D, 3D, 4D) data, see e.g. [14], while KDD techniques have shown their potential through a wide range of applications. Additionally, several research prototypes have applied these technologies and have utilized certain benefits.

Our proposal, called Seismo-Surfer[1], consists not only of a spatiotemporal DBMS, but also of a data mining module and a graphical interface with GIS features. The development of this prototype aims to the better understanding of seismic phenomena and of the relationships between seismic parameters themselves or between those and other factors like subsoil properties, weather conditions during earthquake realization etc. We envisage that Seismo-Surfer could be a useful tool for geophysical scientists who would be mostly interested in high level concepts rather than in plain collections of raw observational data.

Moreover, it is our intention to face the challenge of exploiting the available seismic data repositories over the web. Assuming that the user would periodically, but not very frequently, check these repositories for newly available data, load them into the local database (namely, in a batch update fashion) and taking into consideration the very large size and heterogeneity of these repositories, it seems natural to extend the database of the tool with a Data Warehouse (DW) for storing aggregate information about remote data.

The rest of the paper is organized as follows. In the next section we survey available technologies and research trends that we argue a so-called Seismic Data Management and Mining System (SDMMS) should take into consideration. Section 3 presents such an SDMMS, the Seismo-Surfer prototype under development, and describes its architecture, functionality and current status of implementation. In section 4, we discuss related work and present several research prototypes developed for the management of spatiotemporal and geophysical data. Finally, section 5 concludes with directions for future work.

[1] General information is available at http://thalis.cs.unipi.gr/~ytheod/software/seismo/.

2 Requirements for Management and Mining of seismic data

A combination of three state-of-the-art database technologies is required for the efficient handling of seismic data, namely, spatiotemporal databases, data warehouses and data mining techniques.

2.1 Spatiotemporal Databases

Modelling the real world for seismic data applications requires the use of spatiotemporal concepts like snapshots, changes of objects and maps, motion and phenomena [16]. In particular, we are concerned with the following concepts:

– *Spatial objects in time points.* It is a simple spatiotemporal concept where we record spatial objects in time points, or, in other words, we take snapshots of them. This concept is used, for example, when we are dealing with records including position (latitude and longitude of earthquake epicenter) and time of earthquake realization together with attributes like magnitude, depth of epicenter, and so on.
– *Spatial objects in time intervals.* This could be the case when we intend to capture the evolution of spatial objects over time, for example when, additionally to the attributes mentioned previously, we are interested in recording the duration of an earthquake and how certain parameters of the phenomenon vary throughout the time interval of its duration.
– *Layers in time points.* Layers correspond to thematic maps showing the spatial distribution of certain attributes in the database. The combination of layers and time points results into snapshots of a layer. For example, this kind of modelling is used when we are interested in magnitude thematic maps of earthquakes realized during a specific day inside a specific area.
– *Layers in time intervals.* This is the most complex spatiotemporal concept we are interested in for the modelling of earthquake phenomena. For example, modelling the whole sequence of earthquakes, including the smaller in magnitude that precede or follow the main earthquake, uses the notion of layers in time intervals.

It is clear that we are mostly interested in the spatiotemporal attributes of earthquake data. For example, typical queries that involve the spatial and the temporal dimension of data are the following:

– *Find the ten epicenters of earthquakes realized during the past four months, which reside more closely to a given location.*
– *Find all epicenters of earthquakes residing in a certain region, with a magnitude $M > 5$ and a realization time in the past four months.*
– *(Assuming multiple layers of information, e.g. corresponding to main cities' coordinates and population) find the five strongest quakes occurred in a distance of less than 100Km from cities of population over 1 million during the 20th century.*

In order to support the above data models and queries, novel data types [4] have been proposed in the literature. Seismic data are multi-dimensional and, as a consequence, require different techniques for their efficient storage and retrieval than those traditionally used for alphanumeric information. A great deal of effort has been also spent for the development of efficient Spatial Access Methods (SAM) and some (R-trees, Quadtrees) have been integrated into commercial DBMS's (Oracle, Informix, DB2). Recently, SAMs that take the dimension of time into consideration have been also proposed [15, 9, 20]

2.2 Data warehouses

Additional to the spatiotemporal DBMS, for the reasons stated in section 1, a data warehouse approach can be adopted for the integration of the available sources of seismic data over the Web and the utilization of on-line analytical processing (OLAP) technology. A data warehouse is usually defined as a subject-oriented, integrated, time-variant, nonvolatile collection of data in support of management decision making process [5]. We illustrate the benefits obtained by such an approach with two examples of operations supported by spatial data warehouse and OLAP technologies:

- A user may ask to view part of the historical seismic profile, i.e. the ten most destructive quakes in the past twenty years, over Europe, and, moreover, he/she can easily view the same information over Greece (more detailed view, formally a *drill-down* operation) or worldwide (more summarized view, formally a *roll-up* operation).
- Given the existence of multiple thematic maps, perhaps one for quake magnitude and one for another, non-geophysical parameter such as the resulting damage, they could be overlayed for the exploration of possible relationships, such as finding regions of high, though non-destructive, seismicity and vice versa.

A data warehouse is based on a multidimensional data model which views data in the form of a data cube [1]. A data cube allows data to be modelled and viewed in multiple dimensions and is typically implemented by adopting a star schema model, where the data warehouse contains a *fact table* related with a set of *dimensional tables*, e.g. *quake(quake id, quake type, ...)*, *geography(geography id, region, country, continent)* and *time(time id, date, month, year, century)*. Fact table contains measures on seismic data, such as *number of earthquakes*, and keys to each of the related dimension tables (figure 1). Especially for seismic data, where multidimensional information is involved in the dimensional tables as well as in the fact table, spatial data cubes are also of interest [19].

Further to the operations of roll-up and drill-down described above, typical data cube operations include *slice* and *dice*, for selecting parts of a data cube by imposing conditions on one or more cube dimensions, respectively, and *pivot*, which provides alternative presentations of the data to the user.

2.3 Data mining

The integration of data analysis and mining techniques into an SDMMS ultimately aims to the discovery of interesting, implicit and previously unknown knowledge. We study the integration of three basic techniques for this purpose: methods for finding

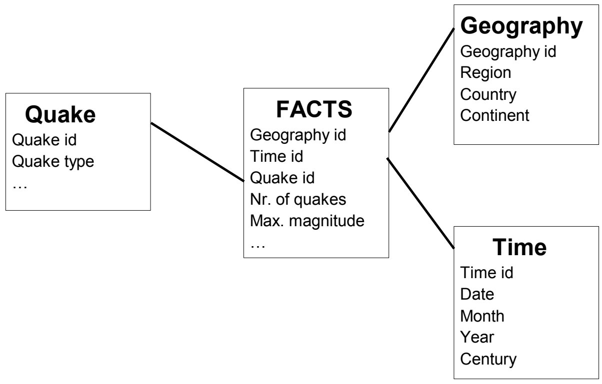

Fig. 1. A possible star schema for an SDMMS.

association rules, clustering algorithms and classification methods. Recently, there have been proposals that expand the application of knowledge discovery methods on multi-dimensional data [10, 11].

- Association rules have are implications of the form $A \Rightarrow B$, $A \subset J$, $B \subset J$ where A, B and J are sets of items and are characterized by two numbers, s, which is called *support* of the rule and expresses the probability that a transaction in a database contains both A and B, and c, which is called *confidence* of the rule and expresses the conditional probability that a transaction containing A also contains B. For example, an association rule could be that during a specific time period, earthquakes of $M > 4$ (itemset A) occurred in a certain region in Greece (itemset B), with confidence 10% and support 80%.
- Clustering algorithms [8, 5] group sets of objects into classes of similar objects. Possible applications on seismic data could be for the purpose of finding densely populated regions according to the Euclidean distance between the epicenters, and, hence, locating regions of high seismic frequency or dividing the area of a country into a set of seismicity zones (e.g. low / medium / high seismic load) .
- Data classification is a two-step process [5]. In the first step a classification model is built using a *training data set* consisting of database tuples that it is known to belong in a certain class (or, in many cases, an attribute of the tuples denotes the corresponding class) and a proper supervised learning method, e.g. decision trees [5]. In the second step, this model is used for the classification of tuples not included in the training set. For example, we could classify earthquake data according to magnitude, location of epicenter or their combination.

Visualization techniques can be used either for the purpose of presenting query results or for assisting the user in the formulation of queries and allowing visual feedback to the Knowledge Discovery in Databases (KDD) process. For example, spatial regions can be selected graphically and queries concerning these regions could be subsequently formulated or the selection of variables on which classification will be performed could be decided after the examination of a parallel coordinate plot [7]. Furthermore, visual interaction tools allow quick and easy inspection of spatiotemporal relationships as well as evaluation and interpretation of data mining results.

Widely used visualization techniques include the usage of geographic maps for the visualization of the spatial and temporal distribution of data attributes, clusters and classes and tree views, scatter plots and various types of charts for the visualization of mining results. Examples of mining results visualization include, among others, tree-views for the results of the application of a decision tree learning algorithm and scatter plots for the visualization of relationships between variables of association rules.

3 The Seismo-Surfer prototype

For an SDMMS following the above requirements, we propose the architecture presented in figure 2. A number of filters perform Extract-Transform-Load (ETL) operations for integrating data from external data sources (e.g. web sites) into the local database. Hence, a new filter has to be implemented into Seismo-Surfer each time we would like to connect a new source. The purpose of the data load manager is twofold: (i) it loads filtered data into the local database and (ii) it calculates aggregations and cube operations and updates the data warehouse.

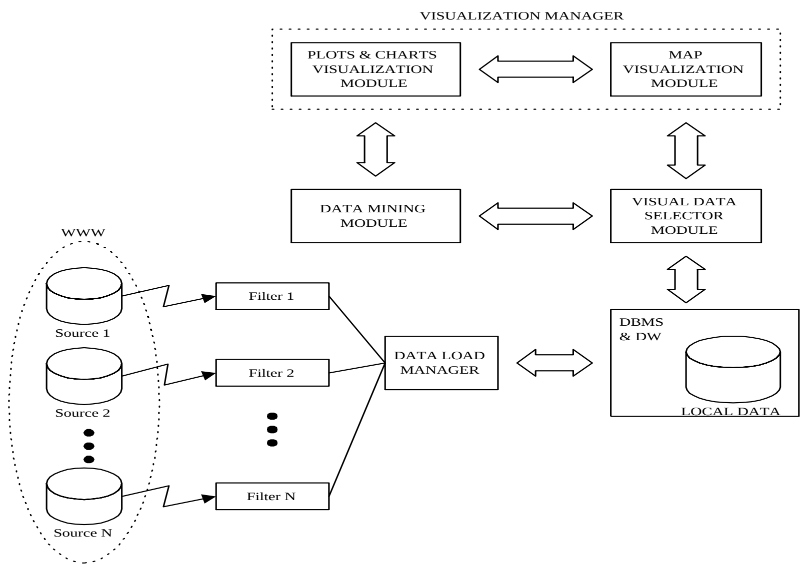

Fig. 2. The Seismo-Surfer architecture.

The visual data selector module allows the selection of subsets of data, e.g. by visually zooming into an area of interest or by setting constraints on the values of desirable attributes (figure 3), and the selection of the abstraction level at which the data are studied, e.g. by performing DW operations like drill-down, roll-up etc.

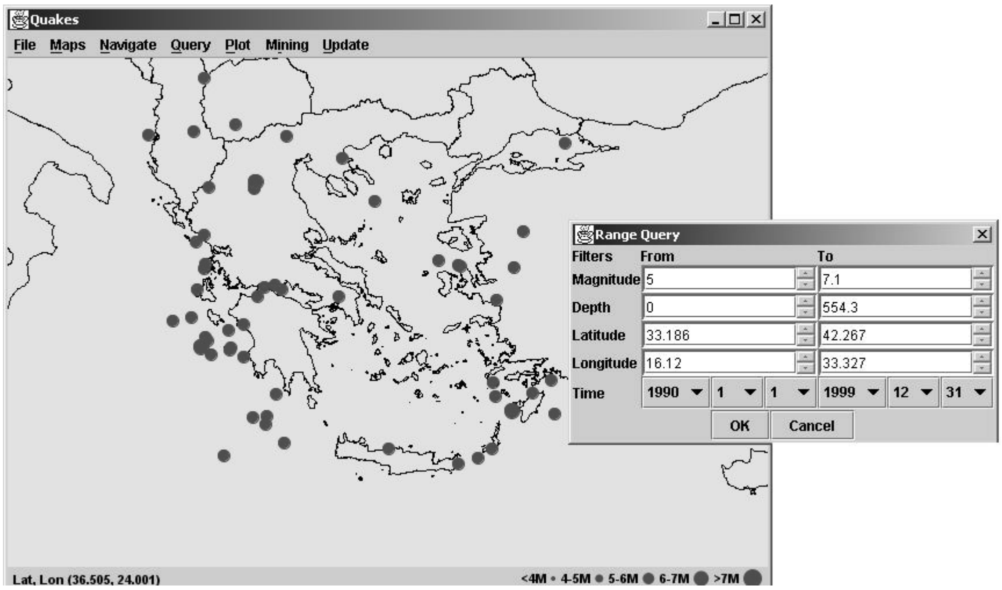

Fig. 3. Screen-shot of our prototype illustrating the spatial distribution of earthquake epicenters in Greece during 90's, corresponding to $M_L \geq 5$.

On the other hand, the data mining module includes a set of data mining algorithms providing the main functionality, i.e. classification, association rules and clustering (figure 4). Both off-the-shelf algorithms provided together with the DBMS and from scratch implementations of novel techniques can be exploited.

After the user has selected the data of interest, either by posing queries or after performing a data mining task on the selected data, a data visualization technique can be applied by the visualization manager. The visualization manager consists of two modules; the first generates plots and charts and the second performs visualization of data over maps.

In an attempt to outline the functionality of Seismo-Surfer, we divide the operations supported in three main categories:

– *Spatiotemporal Queries.* The tool provides the capability of performing queries concerning the spatiotemporal as well as traditional one-dimensional properties of seismic data and their relationships. Thus, the user can study seismicity and by isolating regions and time periods of interest, execute nearest neighbor queries (e.g. *"find all epicenters of earthquakes in distance no more than 50Km from Athens"*), find earthquakes occurred in inhabited areas (range queries), calculate distances between epicenters and other geographical entities like schools and heavy industries, frequency of disastrous earthquakes occurrence and other statistics. Input of queries is performed with the assistance of visual controls and the results can be presented in form of tables or charts and over maps.

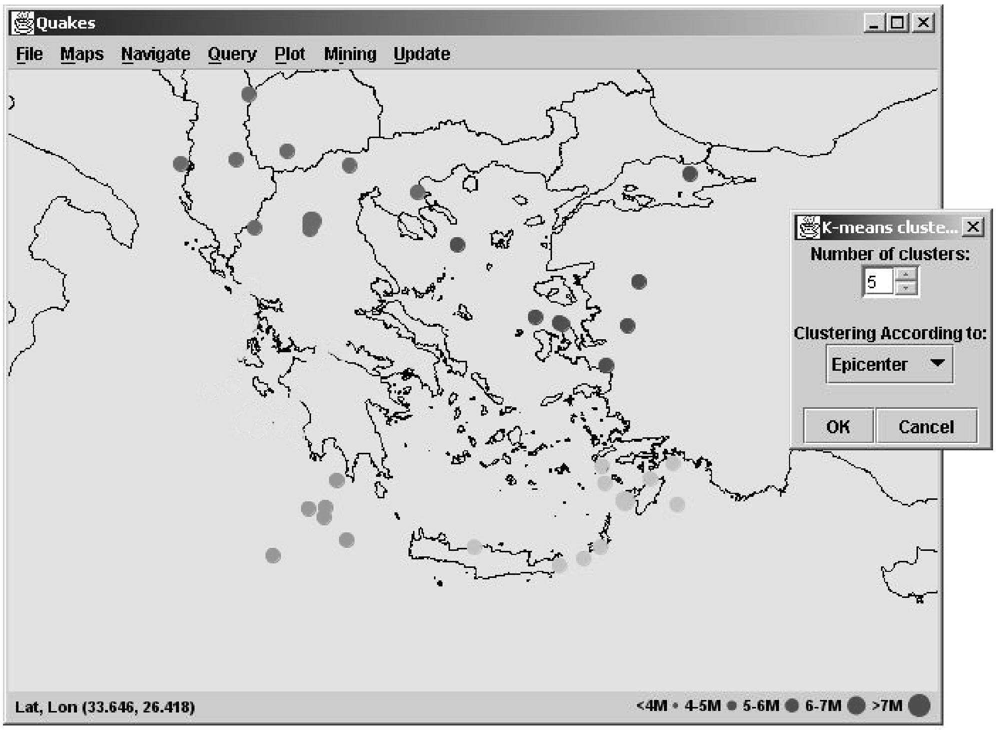

Fig. 4. Clustering of earthquake epicenters in the Aegean and surrounding area using the k-means algorithm. Five clusters are visualized, composing the well known to seismologists "Hellenic arc".

- *Data Cube Operations.* The prototype can deal with the very large size of the observational data sets by supporting summarized views of data in different levels of abstraction of the spatial (e.g. province, country, continent), temporal (e.g. month, year, ten year period) or any other dimension characterizing the data.
- *Remote data sources management functionality.* The proposed architecture provides a flexible implementation framework. For example, as an option, only summaries of seismic data could be stored locally and in case the user requests a detailed data view which is not available, additional data can be loaded, from the remote (web) source, on demand.
- *Simple Mining Operations.* By means of clustering, classification and association rules, SEISMO-SURFER provides capabilities of constructing seismic profiles of certain areas and time periods, of discovering regions of similar seismic behavior, of detecting time recurrent phenomena and of relating seismic parameters between themselves and to information like damages, population of areas, proximity of epicenter to cities etc.
- *Phenomena extraction.* The data mining module is also used for the automatic extraction of semantics from stored data, such as the characterization of the mainshock and possible intensive aftershocks in shock sequences (for example, see figure

5, where M_L values, i.e. the local magnitudes[2] of a shock sequence, are depicted). Moreover, one could search for similar sequences of earthquake magnitudes in time or in space, i.e. time-patterns that occurred in several places or space-patterns that occurred in several time periods.

Furthermore, by providing links between them, the two visualization modules can be used in parallel. For example, rules could be presented in the form of a chart in one window, while, in another window, a map could be showing the spatial distribution of the data for which these rules hold true.

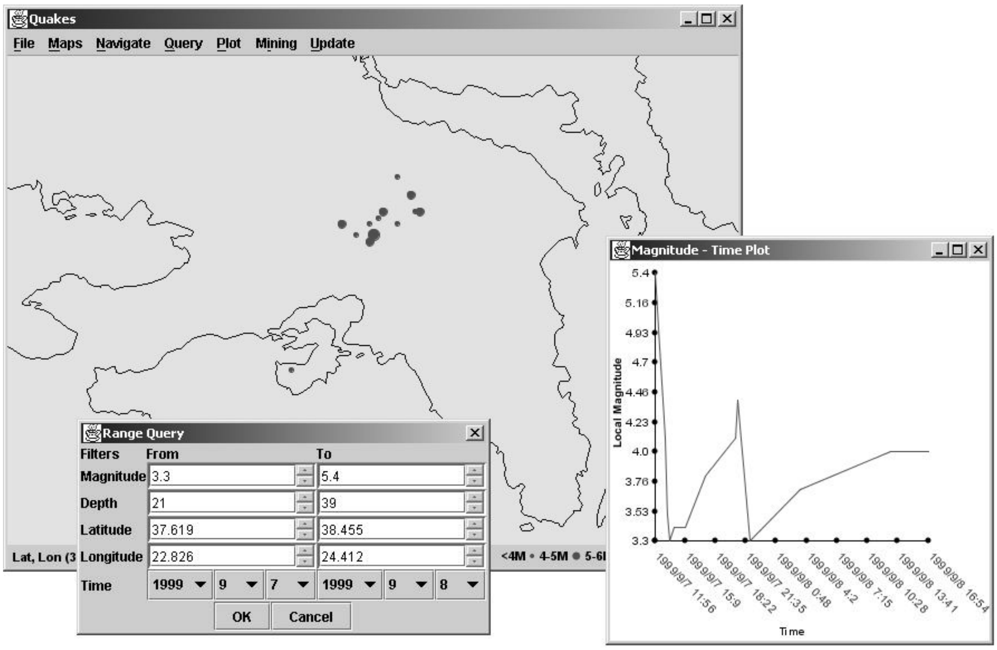

Fig. 5. A series of shocks occurred on and the day after the 7th of September 1999 in Athens, Greece. The greatest peak refers to the disastrous mainshock of local magnitude $M_L = 5.4$, while it is clear that at least one intensive aftershock followed a few hours later.

Currently, the prototype is at an early stage of development. In particular:

– Two remote sources have been integrated: one from the Institute of Geodynamics, National Observatory of Athens, with earthquakes in Greece since 1964 and one

[2] M_L (local magnitude), developed by Richter in 1935, is only one of earthquake magnitude scales used by seismologists. Other scales also include M_S (surface wave magnitude) and M_W (moment magnitude). For a discussion on these and other scales, see http://www.seismo.berkeley.edu/seismo/faq/magnitude.html.

for recent earthquakes worldwide from the US Geological Survey[3]. Users can check for newly arrived data and, if they exist, the database is updated accordingly.
- Queries are formulated using visual controls for setting conditions on both the alphanumeric and spatiotemporal attributes of data, retrieving and visually presenting qualifying items.
- A data mining algorithm for clustering epicenters, the well known k-means, has been implemented and integrated into the system.
- Basic map functionality (zoom in/out and shift operations) is also provided, while pairs of attributes of the earthquakes whose epicenters are shown on the map can be further explored by means of 2D plots.

Next stages of development include the implementation of more sophisticated visualization techniques, of more data mining algorithms, of the data warehouse and of more filters for the integration of other web data sources. The underlying database technology used is a modern Object-Relational DBMS, namely Oracle *9i*. For our selection, we considered two issues: first, the special data types, predicates and advanced indexing mechanisms for spatial data incorporated (R-trees and Quadtrees), allowing integration of the models described in section 2 into the database system, and, second, the off-the-shelf data mining algorithms included in this version (naive bayes and decision trees for classification, association rules, clustering using k-means). The data warehouse is also planned to be implemented by taking advantage of the ready to use functionality provided by Oracle *9i*, including support for constructing data cubes, for performing aggregation queries etc.

4 Related work

Taking into consideration the availability of massive spatiotemporal data sets, the need for high level management, analysis and mining has already been recognized by several researchers as shown by the development of several research prototypes. In the sequel, we briefly describe the purpose and functionality of the most important, to our knowledge so far, that have appeared in the literature. We distinguish two categories of interest. The first includes prototypes that perform one or more of the operations of managing, mining and visualizing either spatial or spatiotemporal data, while the second includes relevant prototypes - applications for scientific and geophysical data.

Geo-miner [6] is a data mining system for spatial data that includes modules for mining (rules, classifications, clusters), for spatial data cube construction and for spatial OLAP. These modules are accompanied with a geo-mining query language and visualization tools.

In [2], the building of an integrated KDD environment for spatial data is proposed, based on the prototypes Kepler, for data mining, and Descartes, for interactive visual analysis. This work is primarily focused on visualization techniques that increase the effectiveness of the KDD process by means of sophisticated user interaction.

[3] Available at http://www.gein.noa.gr/services/cat.html and http://neic.usgs.gov/neis/bulletin/, respectively.

In [3], five case studies are presented in order to bring up critical issues and to show the potential contribution of applying data mining techniques for scientific data while application prototypes for several scientific fields (SKICAT for determining whether observed sky objects are stars or galaxies, JarTool for searching for features of interest on the surface of planets, a system for mining bio-sequence databases, Quakefinder for detecting tectonic activity from satellite data and CONQUEST for the analysis of atmospheric data) are presented. This study is mainly concerned with the issues of dealing with the very large size of observational data sets and of how KDD techniques can be used for extracting phenomena on high conceptual level from the low-level data.

Finally, commonGIS [12] is a web-based system for the visualization and analysis of spatially-related statistical data based on the previously mentioned Descartes prototype and on a geo-spatial database. CommonGIS has been used for the construction of a database of earthquake events registered within the European Countries between 500 AC and 1984 DC.

We intend to combine some interesting functionalities of the above systems and integrate them into Seismo-Surfer. What most distinguishes our prototype from related work is that it proposes a integrated environment for managing and mining spatial data and for the exploitation of heterogeneous data sources.

5 Conclusion

In this paper, we have proposed a novel SDMMS architecture, described its functionality and outlined the potential benefits, by providing extended examples, for potential users: researchers of geophysical sciences, key personnel in public administration as well as people who are just interested on querying or viewing seismic data (web surfers). Main issues discussed include spatiotemporal concepts necessary for the modelling of seismic activity and for efficient storage and retrieval of seismic data, KDD and DW technologies for dealing with the very large amount of available seismic parameters observations and for their effective analysis, and visualization techniques that empower the user to fully exploit the capabilities of an SDMMS by linking it with KDD operations.

Additionally, we have presented Seismo-Surfer, an SDMMS prototype under development, which fulfills the requirements of the proposed SDMMS architecture. Current implementation includes the functionalities of (a) importing data from a remote web source into the local database, (b) asking queries on both the alphanumeric and the spatiotemporal attributes of data, (c) clustering spatial attributes, such as epicenters, and (d) providing map operations (zoom in/out, shift) and 2D plots. Next development stages include, among others, the data warehouse facility, more data mining algorithms, and filters for new remote data sources.

The current version of Seismo-Surfer is available at desktop mode. Apart from the above mentioned further development issues, future steps include a web interface, a generator of tectonic activity scenarios and simulator in the line of the GSTD generator for spatiotemporal data [21], as well as a study of how KDD techniques can be applied on seismic data. For example, the exploitation of already proposed clustering algorithms satisfying the special requirement of spatial databases [17, 18] could also be a task for future work.

Acknowledgements

The author wishes to thank the anonymous referees for their comments. This work was partially supported by Directorate-General Environment of the European Commission under grant PEADAB (Post EArthquake Damage and usability Assessment of Buildings)

References

1. S. Agarwal, R. Agrawal, P. Deshpande, A. Gupta, J. Naughton, R. Ramakrishnan, and S. Sarawagi. On the computation of multidimensional aggregates. In *Proceedings of the 22nd International Conference on Very Large Databases, VLDB*, pages 506–521, Bombay, India, September 1996.
2. G. Andrienko and N. Andrienko. Knowledge-based visualization to support spatial data mining. In *Proceedings of the 3rd Symposium on Intelligent Data Analysis*, pages 149–160, Amsterdam, The Netherlands, August 1999.
3. U. Fayyad, D. Haussler, and P. Stolorz. KDD for science data analysis: Issues and examples. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining*, pages 50–56, Portland, Oregon, USA, 1996.
4. R. Guting, M. Bohlen, M. Erwig, C. Jensen, N. Lorentzos, M. Schneiderand, and M. Vazirgiannis. A foundation for representing and quering moving objects. *ACM Transactions on Database Systems (TODS)*, 25(1):1–42, March 2000.
5. J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann, 2000.
6. J. Han, K. Koperski, and N. Stefanovic. GeoMiner: a system prototype for spatial data mining. In *Proceedings of ACM SIGMOD International Conference on Management of Data*, pages 553–556, Tucson, Arizona, USA, May 1997.
7. A. Inselberg. Multidimensional detective. In *Proceedings of IEEE Visualization Conference*, pages 100–107, Phoenix, AZ, October 1997.
8. A. Jain, M. Murty, and P. Flynn. Data Clustering: A Review. *ACM Computing Surveys*, Volume 31(3):264–323, September 1999.
9. G. Kollios, V.J. Tsotras, D. Gunopulos, A. Delis, and M. Hadjieleftheriou. Indexing Animated Objects Using Spatiotemporal Access Methods. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 13(5):758-777, September/October 2001.
10. K. Koperski and J. Han. Discovery of spatial association rules in geographic information databases. In *Proceedings of the 4th International Symposium on Large in Spatial Databases, SSD*, pages 47–66, Portland, Maine, August 1995.
11. K. Koperski, J. Han, and J. Adhikary. Mining knowledge in geographical data. *Communications of the ACM*, 26(1):65–74, March 1998.
12. U. Kretschmer and E. Roccatagliata. CommonGIS: a European project for an easy access to geo-data. In *Proceedings of the 2nd European GIS Education Seminar, EUGISES*, Budapest, Hungary, September 2000.
13. M. Nascimento and J. Silva. Towards Historical R-Trees. In *Proceedings of ACM symposium on Applied Computing*, pages 235–240, Atlanta, GA, USA, February/March 1998.

14. Oracle Corporation. Oracle R Locator - Location-based Services for Oracle 9i. *Oracle Technology Network data sheet*, 2002. Available at: http://otn.oracle.com/products/oracle9i/datasheets/spatial/9iR2_locator_ds.html.
15. D. Pfoser, C. Jensen, and Y. Theodoridis. Novel approaches in query processing for moving object trajectories. In *Proceedings of the 26th International Conference on Very Large Data Bases, VLDB 2000*, pages 395–406, Cairo, Egypt, September 2000.
16. D. Pfoser and N. Tryfona. Requirements, definitions, and notations for spatiotemporal application environments. In *Proceedings of the 6th international symposium on Advances in Geographic Information Systems, ACM-GIS '98*, pages 124–130, Washington, DC, USA, November 1998.
17. J. Sander, M. Ester, H.-P. Kriegel, X. Xu. "A Density-Based Algorithm in Spatial Databases: The Algorithm DBSCAN and Its Applications", Data Mining and Knowledge Discovery, pp.169-194, 1998.
18. C. Sheikholeslami, S. Chatterjee, A. Zhang. "WaveCluster: A-MultiResolution Clustering Approach for Very Large Spatial Database", Proceedings of 24th VLDB Conference, New York, USA, 1998.
19. N. Stefanovic, J. Han, and K. Koperski. Object-Based Selective Materialization for Efficient Implementation of Spatial Data Cubes. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 12(6):938-958, November/December 2000.
20. Y. Tao and D. Papadias. MV3R-Tree: A Spatio-Temporal Access Method for Timestamp and Interval Queries. In *Proceedings of the 27th International Conference on Very Large Data Bases, VLDB 2001*, pages 431-440, Roma, Italy, September 2001.
21. Y. Theodoridis and M. Nascimento. Spatiotemporal datasets on the WWW. *SIGMOD Record*, 29(3):39–43, September 2000.
22. Y. Theodoridis, M. Vazirgiannis, and T. Sellis. Spatio-temporal indexing for large multimedia applications. In *Proceedings of the 3rd IEEE Conference on Multimedia Computing and Systems*, pages 441–448, Hiroshima, Japan, 1996.