THE GNEM R&E SPATIAL DATA INTERFACE

Lisa K. Wilkening, Matthew N. Chown, and Bion J. Merchant

Sandia National Laboratories

Sponsored by National Nuclear Security Administration
Office of Nonproliferation Research and Engineering
Office of Defense Nuclear Nonproliferation

Contract No. DE-AC04-94AL85000

ABSTRACT

In the course of developing contextual data products for the Ground-based Nuclear Explosion Monitoring Research & Engineering (GNEMR&E) program, a need for rapid spatial analysis of seismic events has been realized. To this end, Sandia National Laboratories is working to develop a Spatial Data Interface (SDI) application, which will assist in the identification of these events. The SDI is supported by contextual data and reference seismic information produced by the GNEM R&E program.

The Spatial Data Interface is designed to be a generic application, which can be accessed from any of the GNEM monitoring tools (ArcView, the MatSeis package and associated regional seismic analysis tools, etc.), and which could be accessed by processes within an automated event monitoring pipeline. Contextual data is stored in a relational database, Oracle, and is accessed via an Environmental Systems Research Institute (ESRI) software package, ArcSDE. A suspicious seismic event may be submitted to the SDI, which in turn evaluates the event against a set of criteria using reference contextual data. These criteria consist of a set of spatial queries, which are the basis upon which an event may be tagged for further analysis or dismissed.

OBJECTIVE

The objective of the Sandia National Laboratories (SNL) Spatial Data Interface (SDI) is to provide a generic interface designed to enhance the processing of seismic events, which can be used either in automated or interactive processing pipelines for treaty monitoring. An application using SDI will process these events through a series of spatial queries, specifically formulated to assist in their identification and analysis. SDI will be developed using an existing programming interface for ArcSDE, an Environmental Systems Research Institute (ESRI) product. This product facilitates rapid access and query functions for spatial data stored in an Oracle database. It can take advantage of the availability of contextual and reference information properly stored within a database, and complement the capabilities of existing tools, data sets, and infrastructure.

RESEARCH ACCOMPLISHED

Introduction

Meeting the GNEM R&E program monitoring goals implies the detection, location, identification, and characterization of possibly hundreds of seismic events daily. Automated processing procedures are in place to perform this duty. As an event moves through the automated system, various software applications are accessed, which draw upon reference information stored in Oracle databases and flat files on the operating system.

Occasionally, a suspicious event is detected. An automated monitoring system may have embedded logic with which it can determine the priority of this event. This logic is based on select criteria, coupled with existing reference data sets. An event may then be dismissed as irrelevant to the monitoring mission, or tagged for further seismic analysis.

Development of the SDI is intended to enhance the automated system processing of questionable events by referring to spatially referenced datasets. Most applications support special event identification, but are somewhat limited. By incorporating access to additional detailed spatial data and enhanced query logic, monitoring analysis may be more efficiently accomplished, providing robust results for event prioritization and reporting efforts.

Spatial Data Interface Application Development

The Spatial Data Interface will be prototyped using ArcSDE 8.2 (formerly SDE, Spatial Database Engine) from ESRI. ArcSDE is an interface, which resides on top of any supported relational database (in this case Oracle), and is capable of serving spatial data to different applications. Via ArcSDE, spatial data is loaded into Oracle where it is stored and managed as tables. This application will also take advantage of the Oracle Spatial extension for spatial data types. ArcSDE can either use these data types or provide its own mechanisms for managing the feature geometry (West, 2001). In addition, ArcSDE provides fast, flexible querying capabilities by spatially indexing the data features. These indexes can be fine-tuned for maximum performance in concert with the Oracle Spatial tuning guidelines.

The SDI application itself will be designed and implemented using ArcSDE's open, high-level C programming language API. The C API allows a developer to build a custom application for ArcSDE, incorporating many advanced geographic information system (GIS) functions for spatial querying (West, 2001). The C API is also the means by which all applicable ESRI software products access ArcSDE. This flexibility would allow the SDI application to be coupled with other customized GIS applications along the monitoring pipeline, such as ArcView GIS for graphic display of SDI results. The ArcView GIS application is customized using its own object-oriented programming language, Avenue, which also allows access to ArcSDE's functionality.

The SDI interface is intended to be generic and independent. Therefore, it could theoretically reside anywhere on the system and be called from any application providing that a valid interprocess communication interface is in place. The SDI will operate with minimal input from the user beyond some key information to initiate processing. This information can either be provided by the user via a GUI or can be contained in a message delivered to the SDI from another application. This information can include the error ellipse information for the event, point location information for the event, and a dataset of interest against which the spatial analysis will be made. A set of canned queries will be accessible for processing. The user will also have the ability to customize the query set or iterate the process as necessary. A user may also elect to perform spatial analysis between any of the contextual datasets available, or against a spatial feature stored in seismic database holdings (such as stations, arrival paths, etc.).

The output of the SDI will be passed back to the calling application, which can be archived, displayed (Figure 1) or used to trigger further processing. For example, a message may be passed to ArcView for viewing the selected event on a map, or to MatSeis to view associated waveform information, or to a web-based bulletin generation tool for event reporting. The goal is to provide an output format that can be easily incorporated into any monitoring task.

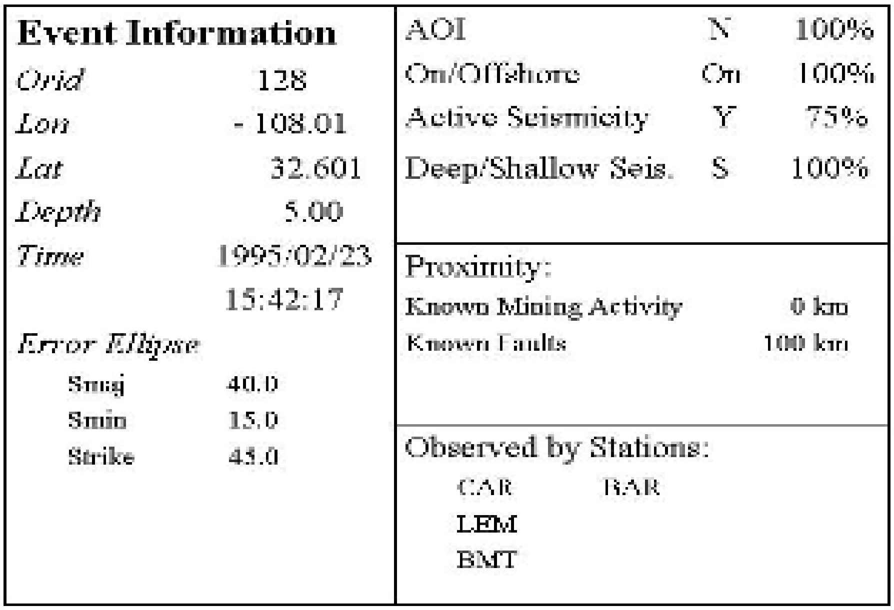

Figure 1. Example of an event identification application, which has used SDI to characterize an event.

<u>Data Sets</u>

A given monitoring system architecture will generally incorporate software and interfaces to perform seismic analysis and data visualization, either for automated or interactive processing components. Depending on the sources, the architecture will likely manage data as a mix of relational databases and flat files stored on the operating system. Therefore, the calling application must be able to access and incorporate either data storage method, or have an interface, which formats the data accordingly.

As noted above, the customized SDI interface utilizes data stored in Oracle tables. Seismic database holdings include event data and associated attribute information. Generally speaking, events have an implied spatial attribute associated with them (a latitude and a longitude), which will form the basis for many of the spatial queries. However, the remainder of contextual data holdings is stored in flat file format at present, and must be loaded into Oracle.

A wide variety of reference information may be used in order to analyze a seismic event. In order to account for the nature and location of an event anywhere across the globe, the contextual data holdings must be diverse. Spatial contextual information may include:

- Political Boundaries
- Coastlines
- Continental shelf
- Roads/railroads
- Topography/bathymetry contours
- Mines

- Seismic stations
- Regions of seismicity
- Faults
- Volcanoes

The ArcSDE software also allows spatial data to be stored with internal behaviors or relationships. The SDI can take advantage of these relationships, and other non-spatial information submitted to it, when evaluating a spatial query against these data.

Spatial Queries

Spatial data is generally stored in one of two models, raster and vector. Raster data are represented as cells of a uniform size, each containing a single numeric value. They are often used for modeling surfaces, such as elevation or depth to basement. For the purposes of this prototype, raster data are not utilized, but may be incorporated in future iterations of this application.

The contextual data sets listed previously are stored as vector data in the SDI's prototype monitoring system. Vector data is made up of a series of features with explicit geometry. Attributes, behavioral, and relationship information may also be stored for each feature. Queries can be based on the spatial relationship of features' geometries as easily as on specific attributes or behaviors. For example, a query may select for a political boundary, which contains features of a global seismic activity data set. The query can then be modified to subselect a particular political boundary that contains only shallow seismic activity and return the results (see Figure 2).

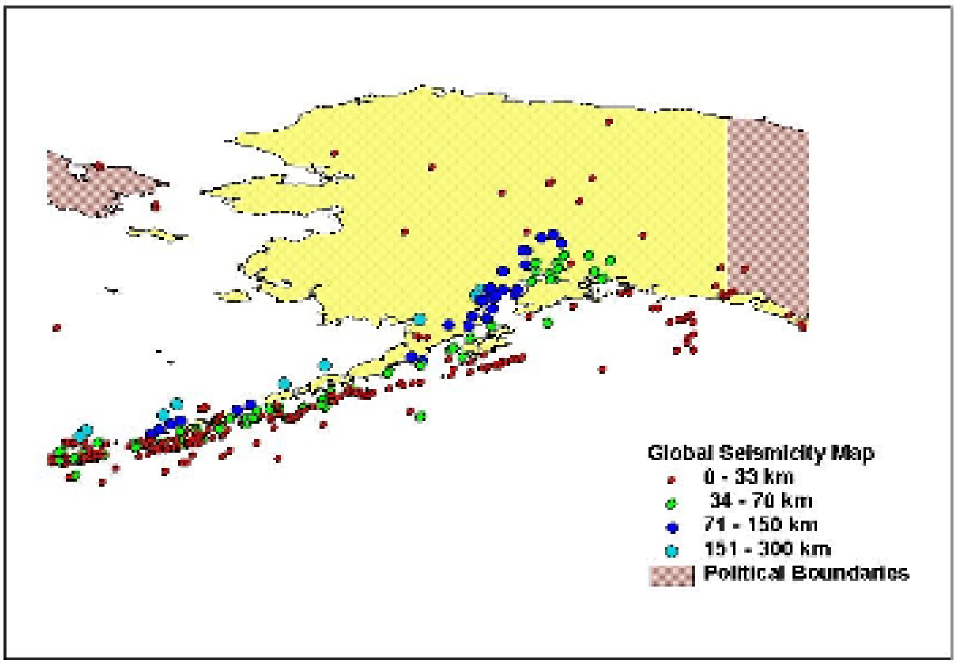

Figure 2. Graphic display of a query depicting a shallow depth seismic activity features within a political boundary.

The SDI application must be able to analyze spatial relationships between a target feature and any feature of a comparison data set. It may do this by performing any one of a basic set of built-in spatial querying capabilities that test a spatial relationship (see Figure 3), using the C API functions. SDI will also take advantage of basic GIS functions such as finding the distance and bearing between features, buffering features or generating new data sets from query results.

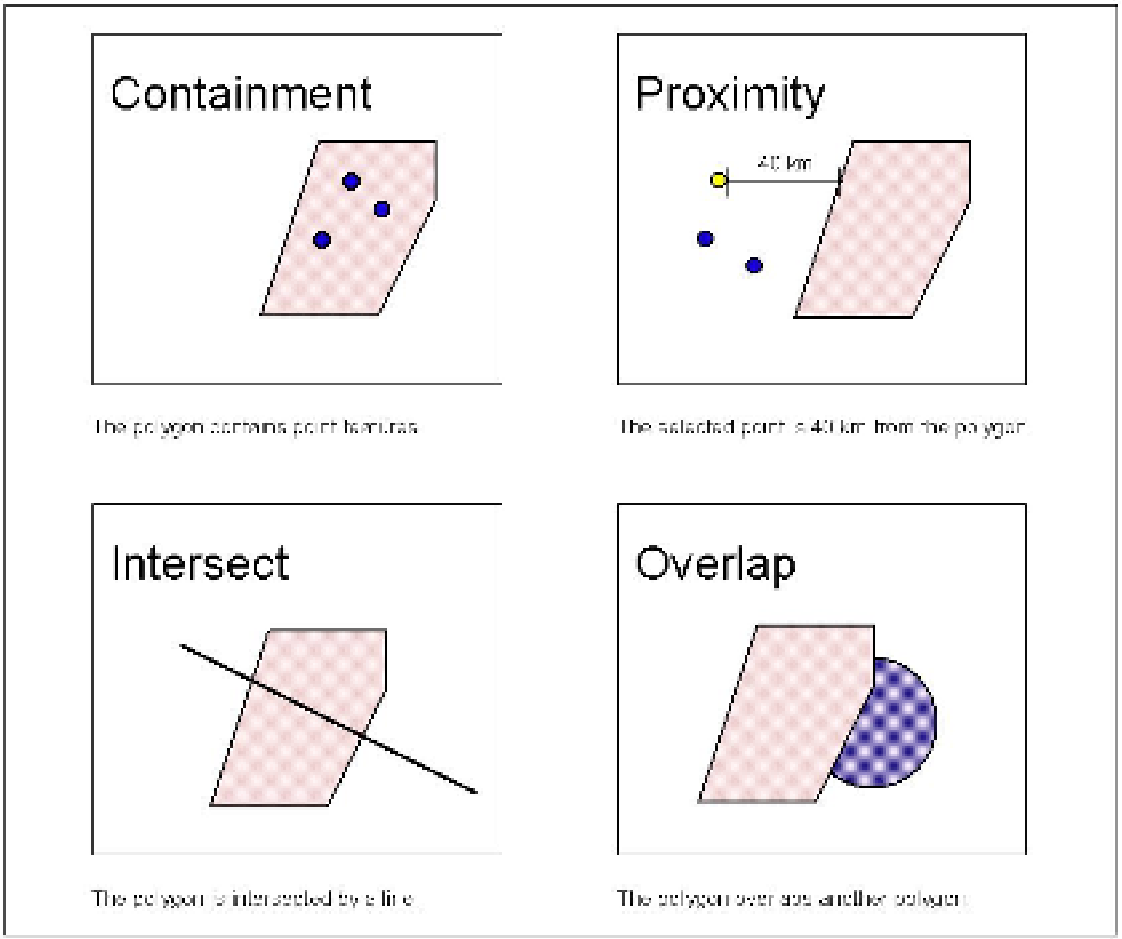

Figure 3. Examples of common spatial relationships between features.

While an automated procedure may have existing logic for defining priorities of seismic activity monitoring, performing spatial analysis on suspicious events can enhance the process by finding relationships between contextual features. Figure 1 gave an example of simple, first-cut analysis. Many times, this may be all that is required to dismiss an event. However, for those suspicious events, which are not immediately dismissed, the SDI can add another dimension to the process, strengthening confidence in the event analysis. The utility of this application is most apparent when placed in the context of normal event processing. An example is depicted in Figure 4.

Figure 4 models an automated processing system within which decisions are made regarding a suspicious event. In this case, an event, which occurs beyond realistic testing depths, will be dismissed immediately. This decision will have been made outside the SDI application. If the answer is yes, a spatial query is submitted to the SDI, which evaluates the overlap of the event's error ellipse with the predefined area of interest data set. The result is returned to the calling application. The next call may be to check the event's error ellipse against seismic data sets of the region, or its proximity to dams or roadways. Specific contextual data, such as depth values ("Deep" in Figure 4), may be queried repeatedly as the event is more accurately defined and characterized. This process can iterate as necessary through the automated analysis pipeline, querying the event against a series of contextual data sets until either a decision is reached, or the event is kicked out to an interactive processing pipeline for special event analysis. SDI analysis may then be reapplied.

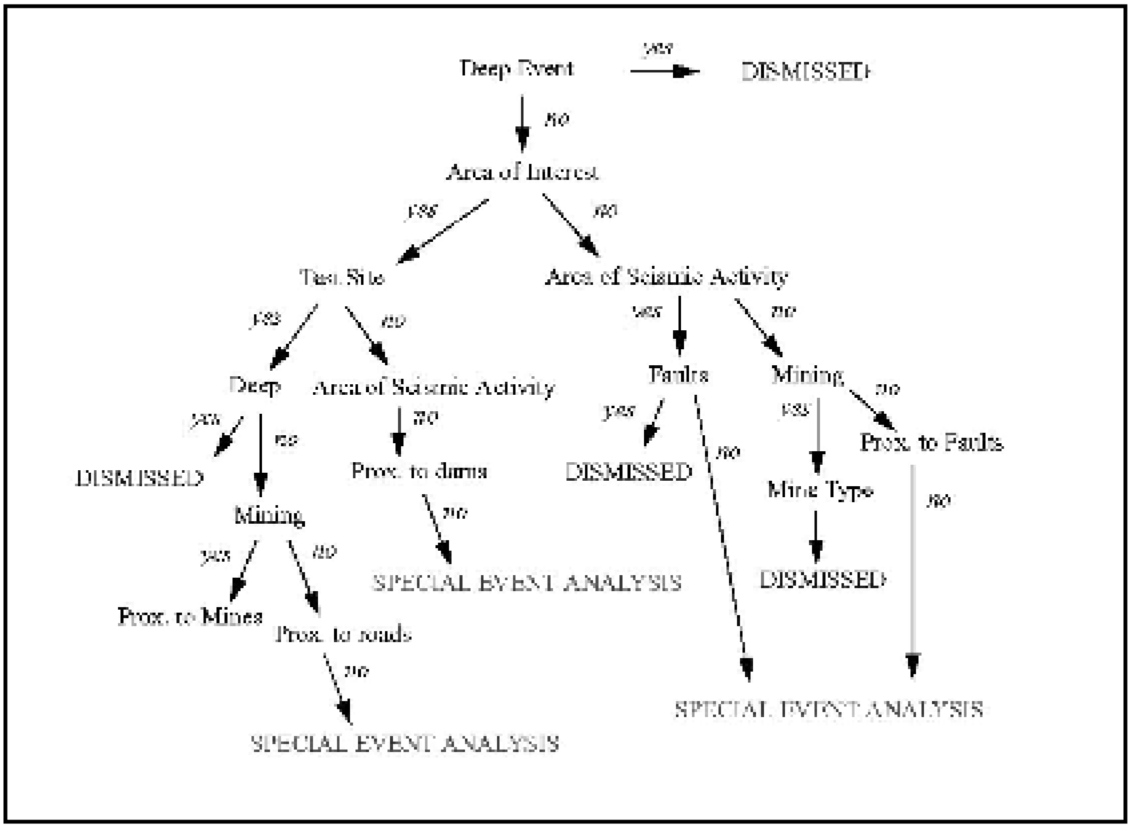

Figure 4. Example of an automated processing decision tree. Event processing moves through the tree, interfacing with SDI for spatial analysis where applicable.

Interactive event processing may require entirely new spatial querying capability. For example, the automated procedure will have established the initial key criteria that an event is shallow, in an area of active deep seismicity, and close to, but not inside, an area of interest. The SDI may then be called upon to answer spatial questions, which provide the analyst with context about the geographic and geophysical attributes of the region. These may include questions about the predominant geologic features in the region, topographic elevations, proximity to faults, types of materials mined in the area, or proximity to human habitation and other man-made features.

CONCLUSIONS AND RECOMMENDATIONS

The application developers recognize that there may be several iterations of the SDI before it is considered a “complete” tool. Future developments will likely focus on creating the “front-end” application to structure the processing rule-set and handle information passing between monitoring system components. As its use grows, new datasets may be developed and new spatial relationships defined. The questions will continue to define the system requirements. Given the availability of and confidence in the spatial contextual data, the SDI should prove to be both an efficient and effective analysis tool, and it should be considered for use as part of an automated event monitoring pipeline.

REFERENCES

West, R. (2001), Understanding ArcSDE, Environmental Systems Research Institute, Inc. (ESRI), Redlands, CA.