A Knowledge-Based System to Support Nuclear Test Ban Treaty Verification

Gregory Beall, Fred Dashiell, Steven Bratt, and Henry Swanger

The major technical obstacle to the signing of nuclear test ban treaties is the issue of verifying compliance. Since the banning of atmospheric and oceanic testing (Limited Test Ban Treaty of 1963) pushed testing underground, seismic monitoring has been one of the most important technologies available for monitoring compliance with test ban treaties. The goal of these treaties is to progressively reduce the allowed yield of nuclear tests (Threshold Test Ban Treaty) and, possibly, eliminate testing completely (Comprehensive Test Ban Treaty). Increasing the sensitivity of the seismic monitoring requires processing rapidly increasing volumes of naturally occurring background seismicity (earthquakes, mine blasts, and so on) that must be discriminated from nuclear tests. Scientists in the nuclear monitoring community realized that the key to improving the implied search was to bring to bear the rich body of knowledge used by human seismic analysts. This chapter describes a knowledge-based system that was built to apply this approach to the problem.

To address the problem of verification, the Defense Advanced Research Projects Agency (DARPA) has been engaged in an ongoing effort to extend U.S. capabilities to monitor the testing of nuclear weapons throughout the world by analyzing remotely sensed seismic signals. As part of this effort, in 1989, DARPA funded the Nuclear Monitoring Research and Development (NMRD) project for the development of a major system for collecting, organizing, analyzing, and archiving seismological data from many geographically dispersed seismic stations. Science Applications International Corporation (SAIC) is the prime contractor.

Seismic events as energetic as nuclear explosions generate waves in the earth that are frequently detectable beyond national boundaries. A variety of physical processes degrade the information content of the waves as they propagate, however, and the information received by a single station is generally inadequate to firmly identify and locate an event unless the station is close to the event (which is not always possible without the cooperation of the monitored country). It is frequently possible, however, to combine the information received at several dispersed stations to determine the location of an event.

The task of associating signals from a network of seismic stations was originally done manually by human analysts looking at waveforms and features extracted from the waves by signal-processing algorithms. The value of computer assistance to search through the many possible combinations of detections was quickly realized and led to automatic association programs (for example, Elvers 1980; Goncz 1980; Jeppson 1980; Slunga 1980). These programs, however, were not sufficiently flexible to readily incorporate the new expertise that has been uncovered as detection techniques improve. This chapter describes the knowledge-based system called ESAL (expert system for association and location), which, as an integral part of the NMRD project, performs this association task and identifies and locates seismic events from data received by a network of seismic stations.

System Overview

The NMRD project encompasses three separate subsystems with overlapping functions but distinct requirements: the Washington international data center (IDC), the intelligent monitoring system (IMS), and the research and development test bed (RDTB). ESAL is an integral component of the shared functions.

Washington International Data Center

The Washington IDC is one of the four international data centers participating in the United Nations Conference on Disarmament Group of Scientific Experts Second Technical Test (UN/CD GSETT-2). It is located at the Center for Seismic Studies (CSS) in Washington, D.C., and is administered by SAIC.

The Ad Hoc Group of Scientific Experts (GSE) was established by the U.N. Conference on Disarmament to develop and test new concepts for an international system of seismic data exchange for monitoring nuclear explosion testing. In July 1986, GSE, functioning under the aegis of UN/CD, proposed the design and testing of a modern international system based on the rapid exchange of seismic waveform data from a modern global seismic network and the processing of these data at international data centers. The proposed system included national data centers in each participating country that would be responsible for collecting data from designated national seismic stations, processing and analyzing the data, and transmitting results of the analysis and waveform segments to a group of international data centers located in Moscow; Stockholm; Canberra, Australia; and Washington, D.C. Each international data center would be responsible for collecting, processing, and reinterpreting the parameter and waveform data received from the national data centers to produce an optimized global seismic event list, called a *bulletin*. Analysis of the data received from the national data centers involves agreed procedures and is designed to locate as many events as possible. Currently, 28 participating nations provide data from over 50 seismic stations.

GSETT-2 has been conducted in phases, starting in mid-1988, and will conclude with a report to UN/CD on the functional capabilities and relative performance of the different technical approaches in mid- to late 1991. During this period, there has been a series of tests followed by analysis of the results and specification for modifications. A typical day of processing during a test identifies about 25 events out of some 750 detections.

Intelligent Monitoring System

The emphasis in IMS is on automating the collection and interpretation of data. It is structurally similar to IDC but has different data sources and somewhat different processing requirements (Bache et al. 1990).

First, the data currently come from a small set of high-frequency arrays[1] that are tuned to detect regional signals (unlike GSETT, which focuses on signals at global distances).[2] The physics of wave propagation through the earth is different at short (regional) distances than at greater (teleseismic) distances because regional waves largely travel through the earth's crust, and teleseismic waves primarily travel through the mantle. These differences are significant and require the use of different heuristics by IMS than those used by IDC.

A second difference involves the initial identification of phase types, which are essential to the analysis of events. Seismic waves from an event can follow a multitude of paths from the source to the detecting station, for example, crustal, shortest path through the mantle, reflected waves, and diffracted waves. The path of a detected signal, when it can be determined, is denoted by a phase ID, such as Pn or P or PKP. In GSETT, these phases are determined (generally by human analysts) as part of feature extraction performed at the station's national data center. In IMS, this information is not provided, so the task is instead performed by a knowledge-based subsystem of ESAL, referred to as Station Processing.

The third difference is that unlike IDC, which is currently only run during specific GSETT tests and with substantial human interaction, IMS is a fully deployed system running continuously and automatically on data as they arrive at CSS. IMS typically processes 500 to 1,500 detections a day and identifies 20 to 100 events.

Research and Development Test Bed

Despite the significant work done to date in this field, the analysis of seismic signals for nuclear monitoring continues to be an area of active research throughout the world. The goal of a Comprehensive Test Ban Treaty requires ever more sensitive and accurate systems. The data in the CSS database are available to researchers throughout the seismological community who want to examine the merits of different methods of processing the data. Therefore, it is necessary to provide a seismologist-oriented easy-to-use interface to support these researchers, allowing them to graphically alter the heuristics, rerun stored data, and review results.

ESAL Overview

NMRD is a large system that runs on a local area network of Sun workstations at CSS. It includes 7 major modules and 33 major processes plus a relational database. Figure 1 shows the system architecture for IMS, which is representative of the general NMRD system. Waveform data from seismic activity are accumulated at a national or regional data center, where they are signal processed (in IMS at the NORSAR Data Analysis Center [NDAC] in Norway). This processing identifies *detections*[3] in the data and extracts seismologically significant features, such as arrival time, *azimuth* (that is, direction of the incoming wave), phase velocity, amplitude, and frequency content. The features and support-

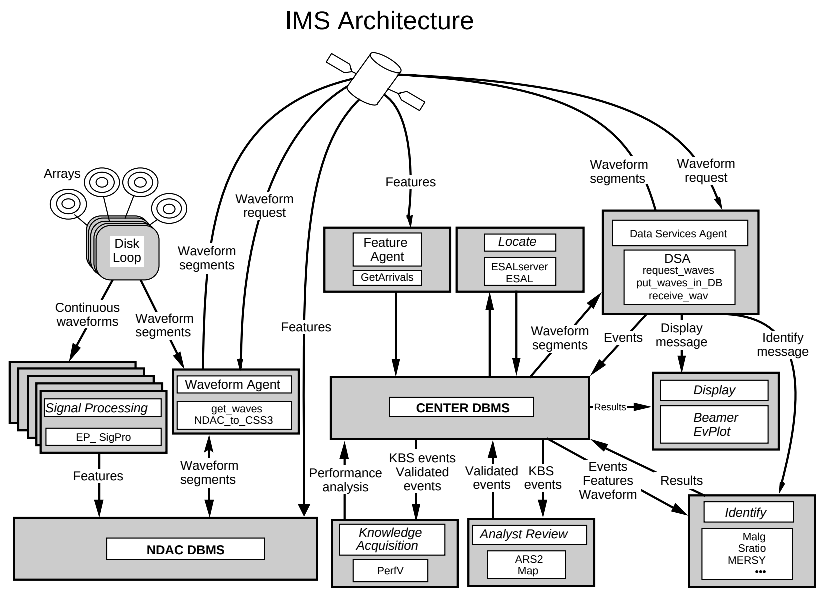

Figure 1. IMS Overview.
The architecture of IMS is shown by the major groups of processes in NDAC (left) and the center's local area networks, which are bridged by a wide-area network connection over a satellite link. The function of the agent processes is to control the data flow. The arrays and disk loops are external to IMS.

ing waveform data are transferred by satellite link and stored in the CSS database. The knowledge-based system ESAL then attempts to group related detections to identify and locate events. This process provides the essential information of location, depth, and magnitude that is needed for event identification. The results and an audit trail of the analysis are stored in the database and are subsequently reviewed by human analysts who correct the data and conclusions as necessary using the *analyst review station*.[4] A separate knowledge-acquisition tool is being developed to examine the changes made by the analysts, plus ESAL's audit trail, to help identify weaknesses in the knowledge base.

Functional Overview

ESAL is a knowledge-based system that is integrated into the much larger context of the NMRD system. ESAL can run either as part of a real-time pipeline, analyzing data as they become available in the database, or interactively, rerunning historical data from the database on demand. The functional goal of ESAL is to determine and locate all locatable events using data from a network of seismic stations. Once tentative events are identified, humans analysts can review and refine them, drawing on information not available to ESAL. The task of ESAL is to sift through the exponentially large number of potential combinations of detections to find probable groupings, a task that is extremely difficult for human analysts, especially when looking for small events hidden among other events.

ESAL contains two separate knowledge-based subsystems: Station Processing and Network Processing. Figure 2 shows the top-level screen of ESAL's user interface, which reflects the control flow in ESAL. Detection, station, and earth-model data for a seismologically meaningful time period[5] are loaded into ESAL from the NMRD database. The number and types of the stations are unrestricted. The data are optionally processed by the *Station Processing subsystem*, which determines the phase identifications of the detections and provides tentative event groups of detections from a single station. The heart of ESAL is the *Network Processing subsystem*, which examines a set of detections from a network of seismic stations and attempts to determine which groups of detections were generated by the same events. The time, location, magnitude, and composition of these events is the output of ESAL.

In addition, ESAL writes an audit trail of its major decisions, which is stored in the database. When human analysts correct the results produced by ESAL, these changes are also stored in the database. They will subsequently be analyzed by a separate knowledge-based system,[6] which will make deductions about the reason for the corrections and use the audit trail to indicate flaws in ESAL's reasoning. This tool will be important for performance validation and knowledge acquisition.

User Interface

ESAL has a seismologist-oriented graphic user interface that allows a user to graphically examine and modify heuristics, examine the content of solutions, and review details of the reasoning process and evolution of the solutions. Over 200 parameters are available to the user to control ESAL's processing. These parameters and related rules are organized and accessed by module, as seen in figure 3, where the user is examining some of the parameters controlling the Special technique for event hypothesis generation.[7] Figure 4 shows part of the interface available for examining ESAL's solutions. The upper area shows detections along a time line and events (the rectangles) that were constructed from them. The lower area shows high-level information about one of the events and the detections that were considered during its evolution. Selection of an individual detection displays detailed information.

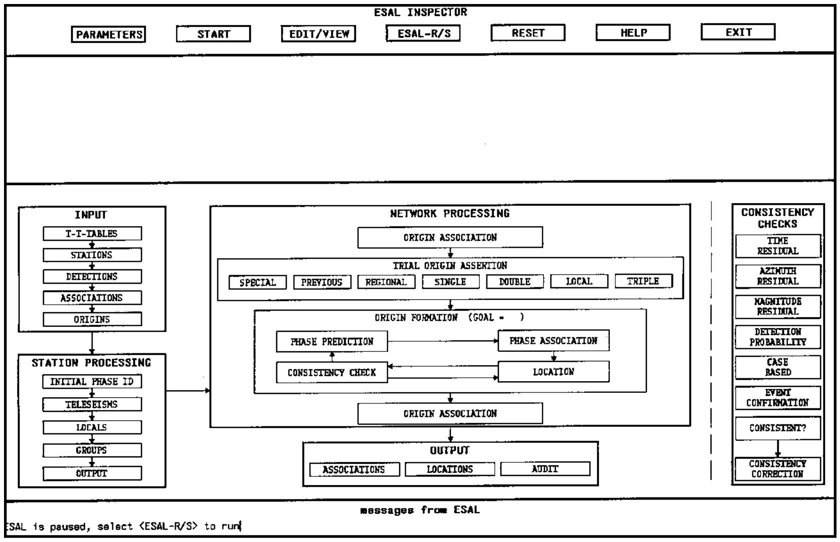

Figure 2. Top-Level ESAL *User Interface.*

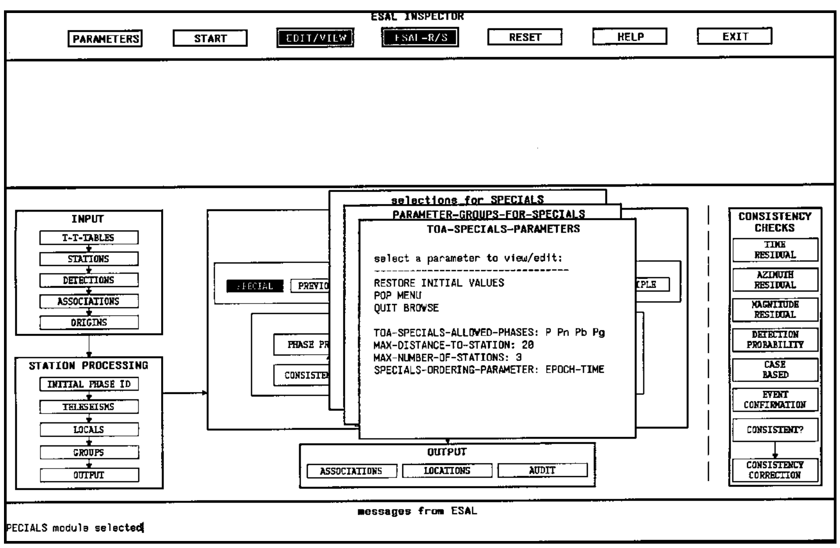

Figure 3. Access to ESAL *Heuristics.*

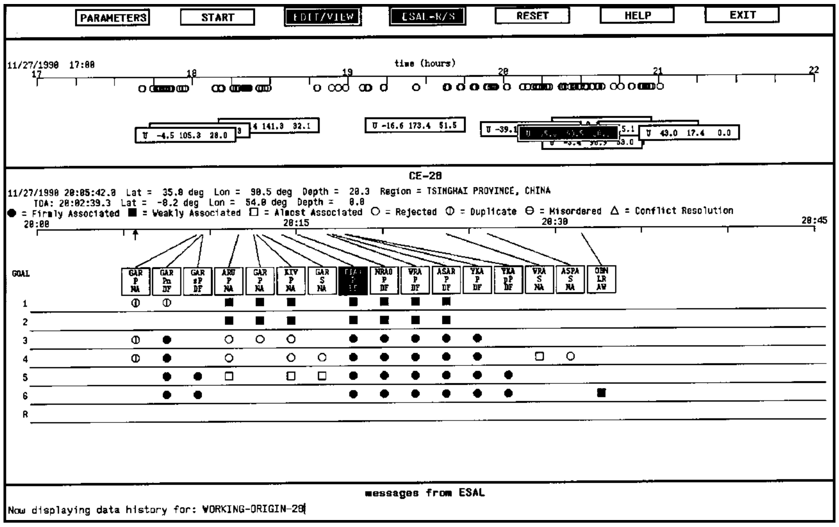

Figure 4. Inspection of ESAL *Results.*

This type of interface is a significant advance over traditional automatic association programs and provides major support for use of the tool for research and development. A user can examine results in detail, adjust heuristics, and then rerun the data. This kind of window into the workings of an automatic association program has not existed before and should provide new insights into its complex behavior.

Design Considerations

Flexibility is the dominant design consideration for ESAL for a number of reasons: First, GSE specifies a number of the seismological heuristics to be used when running GSETT data, which must be modified when processing the regional data in IMS; it might also be desirable to modify them even when processing teleseismic data from other sources.

Second, it is frequently desirable to adjust the search for efficiency depending on the nature of the data, for example, regional or teleseismic. It is also desirable to adjust the thoroughness of the search because the density of the data varies; it is generally desirable to process the data as thoroughly as possible in the available time.

Third, the entire NMRD system is undergoing rapid evolution. ESAL must be able to accommodate data from new stations with new characteristics, new regionally specific earth-model data, and new regionally specific expertise associated with these stations.

Fourth, the field of automatic association itself is an area of active research, especially for regional data. It is essential that seismologists be able to easily examine alternative heuristics.

ESAL must be able to run in a continuous pipeline to support IMS but must also allow the graphically driven interactive reprocessing of data to support RDTB.

ESAL Architecture

ESAL was implemented in ART, Common Lisp, Fortran, and C and is deployed on a Sun workstation. This workstation is part of a local area network of Sun workstations that run the full NMRD system (including the database, telecommunications, signal processing, and analyst review station). The hardware was chosen for smooth integration with the overall system architecture. ART was chosen as the primary programming language after previous experience by SAIC showed that another popular shell was too slow and that a custom inferencing system required too much programming effort and too much focus on the inferencing strategy, distracting effort from the seismological issues. ART was known to have a full set of features, superior performance, and a strong record of successfully deployed systems.

Station Processing

An essential piece of information used by the Network Processing subsystem is the phase ID for each detection. Unlike other data features used by Network Processing, phase IDs cannot be extracted by signal processing alone. During GSETT, phase IDs are determined by human analysts at the various national data centers. In IMS, however, these elements are not provided and are determined instead by Station Processing.

The Station Processing module of ESAL analyzes signals from one station at a time. The features of interest, extracted by Signal Processing before submission to ESAL, include phase velocity, direction of source, wave amplitude, frequency content, and polarization. The principal objectives of this module are to derive the most likely phase type (for example, Pn, Sn, Noise) for each detection and associate detections into groups that are likely to have been produced by the same event to the extent that such conclusions are reliable using information from just one station.

These phase types and groupings are produced for each station being analyzed and are stored in the database. They can then be passed on together to the Network Processing module for more complete location processing. For flexibility, Station Processing can also be used as a standalone module, without also running the Network Processing module.

Currently, Station Processing is most useful for signals produced by events relatively close to the station (local and regional events) because the heuristics are best understood for the closer distances. However, the architecture is designed so that as new stations with more sensitive instruments come online, and new heuristics are developed, single-station treatment of teleseismic data can easily be added to the rules of the module.

There are two major steps in the Station Processing module (figure 2): initial phase classification and event-group association with phase identification. This latter step is divided into special rules for teleseisms, local events, and more generic regional events.

The first step classifies each detection into one of several mutually exclusive categories: regional P-type, regional S-type, teleseism (that is, from a distant event), or noise. This initial phase identification for a detection is based on feature values of the signal without consideration of the relationship it might have with the other detections being considered. The heuristics for this classification are implemented using rules with patterns describing the various categories. The value ranges that delimit the categories in the patterns are parameters that can be modified by the user.

Final phase identification and association into likely event groups is done by considering the detections at the station in time order. The earliest arriving unassociated P-type wave is used as the foundation of a new event group, and the list is searched for following detections that are to be included in the new group. Here, the interrelationships of the detections are used to derive the most likely specific final phase ID (for example, Pn, Pg, Sn, Lg) and decide whether a given detection is consistent with the developing group. The context of a detection within a group, that is, its relationship to other detections in the group, is combined with empirically derived probabilities using Bayesian reasoning to derive a most likely final phase identification.

The reasoning in the Station Processing module is monotonic using a constrained depth-first search. That is, an event group is completely formed, and the phase IDs are affirmatively established and not subsequently revised until additional information is brought to bear during the Network Processing phase. The rules are forward chaining. A major feature is the ability to easily add special heuristics for particular stations or geographic regions to the reasoning.

Network Processing

Network Processing examines detections from a network of seismic stations and attempts to determine which groups of detections were generated by the same events. It consists of an event hypothesis generator

(Trial Origin Assertion in figure 2) and a multilevel loop (Origin Formation) that iteratively adds corroborating detections, relocates the event, and then checks for self-consistency. If this process identifies sufficient corroborating data, the hypothesis is confirmed; otherwise, it is abandoned. The input data used by Network Processing are detection features (for example, arrival time, azimuth, velocity, amplitude, phase ID), station data (for example, location, sensitivity), earth-model information (travel time and amplitude-distance data for the various phases), and user parameters that control heuristics regarding the nature of the events to be formed and the manner in which they are to be formed. It can also include previously formed events and tentative single station groupings that are to be reprocessed using new data or different heuristics. The output are event characterizations (time, location, magnitude) and the set of associated detections. Both are written to the CSS database for subsequent analyst review.

The heuristics used in Network Processing fall into two classes: those controlling the composition of the events that are formed, for example, restrictions on what detections can be used to define an event and the minimum number and character of defining detections, and those controlling how the events are to be identified and formed, that is, control of the search process.

The heuristics in the first class are tightly specified by GSE in IDC. Many of these heuristics are relaxed or modified, however, when running IMS. The heuristics controlling the search have more overlap between IDC and IMS but vary with the type and source of the data.

Network Processing Architecture. Network Processing is implemented as a rule-based forward-chaining system. Detection and station data are stored in ART schemas (also called frames) that allow an object-oriented structuring of the data, pattern matching by rules for the expression of complex conditions, and the use of procedural access and active values for efficient processing when pattern matching is not required. Heuristics are expressed in rules, schemas, and procedures. Some heuristics are encoded in Lisp for efficiency. Many heuristics, however, require a rich pattern-matching language; such as the example in figure 5.

To maintain flexibility, rules are made as generic as possible and frequently represent only the structure of the heuristic. The specifics of the heuristic, such as thresholds, phase restrictions, and minimum requirements for confirmation, are represented in a declarative form that is easily modified by the user. This approach emphasizes the separation of the data from the inferencing mechanism and is a crucial factor in allowing the system to handle multiple seismic domains. Fortran

```
IF there is ?detection-1 associated-with ?event
   which was detected at ?time-1 at ?station,
       which is at ?distance from ?event
   and which has associated-phase ?phase-1,
       which is a member of ?P-type-phases
   AND there is a second ?detection-2 associated-with ?event
   which was detected at ?time-2 at ?station
   and which has associated-phase ?phase-2,
       which is a member of ?S-type-phases
   AND ( ?time-1 – ?time-2 ) < ?max-P-S-time-for-regional-event
   AND ?distance > ?max-distance-to-regional-event

THEN ?detection-1 and ?detection-2 may not be used
   to compute the location of ?event
```

Figure 5. Complex Conditions Require a Rich Pattern-Matching Language.

and C routines handle a variety of numerically intensive calculations based on the earth model, such as the computation of seismic wave point-to-point travel times.

Complexities in Network Processing. The problem of associating detections with the correct event is more difficult than it might appear because of the large volume of data and the presence of noise. The signals from an event can be detected over a time period that is large compared to the interval between events.[8] Critical discriminating phases can be lost in an overlapping signal from another event. In addition, such strong discriminants as azimuth are frequently unavailable or have large uncertainties. All the detection's features, including the phase ID, carry uncertainties that ESAL must use in its reasoning. Because of these uncertainties, there can be numerous self-consistent ways to combine a set of detections into event hypotheses that ESAL must select between. Furthermore, the number of detections that must be analyzed increases rapidly with the sensitivity of the system (for example, there are roughly 10 times as many magnitude 5 events as magnitude 6 and 10 times as many magnitude 4 as 5 [Bolt 1976]; most of the events located by IMS are below magnitude 2).

Search in Network Processing. An unusual feature of this problem relative to many AI search problems is that ESAL must find as many locatable events as possible within system constraints (for example, acceptable processing time). An exhaustive search is an exponential function of the data density, so the search must be restricted, but it is done in a flexible way so that the extensiveness of the search can be balanced against the amount of data. The search can roughly be described as a *best-first search*. One event hypothesis at a time is generated and allowed to evolve. If the resultant event is considered acceptable, it is retained, and the defining detections are removed from consideration for the construction of subsequent events. If the event is judged to be unacceptable (generally because of an inadequate amount of corroborating data), then it is dissolved, and the associated detections are made available to subsequent events. This process continues until the event hypothesis generator is exhausted.[9]

Hypothesis Generation and Development. A number of techniques are available in ESAL for generating event hypotheses (Trial Origin Assertion in figure 2); the selection and order of the techniques used are under the control of the user and can vary depending on the source and nature of the data. Within each technique, additional heuristics determine the specific order of hypothesis generation and the quality checks that control the minimum requirements for the hypotheses to be considered.

The evolution of the hypothesis (Origin Formation in figure 2) is at the heart of ESAL and is the most complex part of the processing. The loop in Origin Formation has Phase Prediction-Association, Location, and Consistency Check as its main phases. Given an event hypothesis (from Trial Origin Assertion), the set of available detections is examined to identify detections that might corroborate the hypothesis (Phase Prediction-Association). Detections are checked for compatibility with both the location of the event and the composition of the event (for example, the phases detected by a single station must obey various ordering constraints). The new detections are tentatively associated with the event, and the event is relocated using the new information. The solution is then checked for seismological self-consistency (Consistency Check), and detections that are no longer consistent are removed and the event relocated. The hypothesis is then reexamined for additional corroborating detections (Phase Prediction-Association). This refining process continues until no more detections are associated, or a previously examined node in the search space is revisited. This process actually proceeds through a series of stages in an effort to stabilize the convergence of small events, initially using the most reliable data (such as teleseismic primary phases), then adding less reliable data (such as regional and secondary phases) as the solution stabilizes. The number and structure of these stages is under the control of the user.

Network Processing Internal Model. The heuristics that control the association of a detection with an event depend not only on the relationship of the detection to the event but also on the relationship of other detections to the event and of the detection to other events. As previously mentioned, the initial detection and station data are represented by schemas; event hypotheses are also represented by schemas. In addition, as events are formed and evolved, ESAL maintains a dynamic solution model of these relationships in the form of a semantic net. The links between detections and events are themselves objects (schemas) of a single type with many attributes (such as associated phase, type of association, phase and association confidences). This representation allows reasoning about complex aspects of the structure of the net, not just about individual objects. Having a single type of link greatly simplifies the expression of the rules for reasoning about the net, which is essential to the goal of keeping ESAL flexible.

All the data features that ESAL reasons about have uncertainties arising from measurement and the seismological models.[10] These uncertainties are propagated into the event location and are then used to compute confidence levels for the association of detections with events. Confidence levels are uniformly computed throughout ESAL using conventional probability theory and are the basis for determining consistency and resolving conflicts.

Search Control in Network Processing. The search in Network Processing is somewhat unusual in that it is based on the "iterate until convergence" model but must deal with the problem of a minimum discrete step size (adding or removing one detection) that is sometimes large compared to the separation of the local minima of the evaluation function. This step can produce a search that drifts and follows complex orbits and that is difficult to control in numeric models. In ESAL, it is not uncommon to find complex orbits where some detections are added, others are removed, more are added and removed, and the system returns to the same node. ESAL cannot afford to maintain a full representation of the search space, so instead, it records snapshots of significant nodes in the path that it traverses so that it can appropriately terminate these potential loops.[11] At the conclusion of Origin Formation, the event hypothesis is tested against various minimality and validity checks before it is confirmed.

Backtracking and Conflict Resolution. In a conventional automatic association program, the construction of the event would be complete at this point, and the search for the next event would begin. A major innovation in ESAL is the introduction of nonmonotonic reasoning. Programs like that described by Slunga (1980) follow the strict rule that a defining detection cannot be reconsidered by subsequent events. This strategy is correct (and beneficial to the management of the search) if there is a high confidence in the association of the detection with the event.[12] Not all detections can be associated with high confidence, however, because of the uncertainties previously discussed. Therefore, the conventional strategy can erroneously associate a detection with an event, thereby making it unavailable when the correct event is subsequently hypothesized. ESAL introduces the concept of *weak association*, which allows certain associated detections to be considered by subsequent event hypotheses and be temporarily multiply associated. This process allows additional small-event hypotheses to be considered at the processing cost that multiple associations must be resolved before the event is confirmed. The user can control how many detections are weakly associated, thereby trading processing time for expanded search.

Development and Deployment

The overall design of the NMRD system, including a high-level description of ESAL processing, was submitted to DARPA in mid-1988. The detailed design of ESAL started in late 1988. Two Inference knowledge engineers worked full time with periodic input from initially one and eventually four expert SAIC seismologists. The first version of the Network Processing system was delivered to SAIC in the spring of 1989. Station Processing was added at a later stage, the first version being delivered in late 1989. New versions of both systems, containing requested extensions, were released at roughly quarterly intervals. ESAL was first used in GSETT-2 testing in January 1990. ESAL has been fully deployed in the IMS application since November 1990.

ESAL was initially validated by running it with historical GSETT data and having the results accepted by the seismologists on the team. As GSETT testing progressed, the available data increased. These data are now used for regression testing of new releases. ESAL's conclusions in IMS are continuously reviewed by human analysts; however, ESAL is not currently expected to do as well as seismic analysts who have access to the actual waveform data. The following counts indicate an approximate measure of the size of ESAL:

Number of ART rules:	250
Lines of ART code:	12,634
Initial ART schemas:	672
Schemas during run:	~ 1,500
Number of Lisp functions:	1,677
Lines of Lisp code:	35,061
Lines of Fortran and C code:	13,422

ESAL is currently used on a continuous, daily basis within IMS at CSS in Washington, D.C. Within the IMS processing, ESAL has processed all detection data received from the Scandinavian high-frequency arrays

NORESS, ARCESS, and FINESA since 1 November 1990. From this date through 11 January 1991, ESAL processed 64,572 automatically detected seismic phases to form a seismic bulletin of 3,889 local, regional, and teleseismic events. All ESAL solutions are reviewed by an analyst and corrected if necessary. Processing in the IMS application includes both Station Processing and Network Processing. A German array, GERESS, is being added in early 1991, and there is a plan to add two Polish three-component stations and seven three-component stations in the USSR and China by mid-1991. IMS is the only seismic processing system in the world capable of fully automated processing of local, regional, and teleseismic data.

In addition to IMS, ESAL is used in the Washington IDC during GSETT-2 testing. Nineteen days of data between 16 January 1990 and 4 December 1990 were processed as a part of GSETT-2. For these days, ESAL located 487 events using 13,878 reported detections. As in IMS, all ESAL solutions were reviewed by an analyst and corrected if necessary. The final operational exercise of GSETT begins in April 1991, covering 42 continuous days of processing. As part of the technology transfer encouraged by UN/CD and DARPA, ESAL will be installed in Norway in mid-1991 for the Norwegian Council for Scientific and Industrial Research.

NMRD is still actively evolving with the addition of new seismic stations with different detection capabilities (for example, the three-component stations in Poland, the Soviet Union, and China) and the refinement of seismological heuristics, especially in the areas of region-specific refinements and the station processing of teleseismic data. The flexible design of ESAL allows the seismologists to make many adjustments without altering any source code. Alterations that require source code changes are made by the original developers. The audit trail mentioned earlier provides automated support for knowledge acquisition, but the knowledge base update is a manual process.

Innovations

ESAL's innovation is mainly seen in comparison with other automatic association programs. Innovations include the following:

First is the flexibility to process data of diverse quality and type; support multiple seismic domains; alter the reasoning process at run time; control the search, particularly with regard to the concept of weak association and revision of previous hypotheses; and easily add new heuristics.

Second is the ability to automatically assign phase identification to regional and teleseismic phases (Station Processing). Few existing automatic association programs include this ability as an automated procedure.

Third is the ability to express much more complex and knowledge-rich consistency checks than conventional automatic association programs.

Fourth is the graphic user interface. Conventional automatic association programs have no capability that is analogous to the ESAL interface for inspecting intermediate processing results.

Fifth is the ability to run either as part of a data processing pipeline or in interactive mode.

Relative to other knowledge-based systems, ESAL is unusual in that it simultaneously supports a fully deployed real-time system and a research and development tool. The search employed in Network Processing is also unusual among typical knowledge-based applications in that the final solution should contain all viable partial solutions (although this kind of search seems to apply to a variety of problems where it is essential to explain all sources of overlapping data, such as fault monitoring in a network of sensors where errors can propagate significant distances from the source).

Summary

NMRD is one of the most sophisticated systems in the world today for the seismic detection of nuclear testing. It was built to test new ideas about the integration and processing of data from a global network of seismic sensors, and it was designed to use knowledge-based system technology to provide the flexibility to evolve along with the seismological understanding of such a system. This flexibility was demonstrated in late 1990 when, in less than a month, the system was modified to process previously unseen data provided by a different branch of the Department of Defense.

ESAL represents a new generation of automatic association programs that take advantage of knowledge-based technology to incorporate the rich and evolving body of knowledge used by human seismic analysts. ESAL is fully deployed in a continuously running real-time system and is also used as a research and development tool to improve the understanding of such systems.

As for acceptance by the customer, DARPA has identified NMRD as one of the outstanding programs within its funding. The value of this system was not envisioned in hours or dollars saved but in the confidence it could provide to policy makers that a nuclear test ban treaty can indeed be verified. In the world's changing political environment, its greatest value might end up being the ability to monitor regions of the world where treaty negotiations have failed.

Acknowledgments

ESAL was developed as part of the Nuclear Monitoring Research and Development (NMRD) project funded by the Defense Advance Research Projects Agency (contract number MDA972-88-C-0024). The top-level architecture and heuristics for ESAL were specified by Science Applications International Corporation (SAIC), the NMRD prime contractor. SAIC also provided the numeric routines for earth-model calculations. The ESAL knowledge base was designed and implemented by Inference Corporation as a subcontractor to SAIC. The authors are indebted to the staff members of the Geophysics Division at SAIC, San Diego, including Thomas C. Bache, director, Jeff Given, and Ethan Brown, and to Bradley P. Allen, director of advanced technology, Inference Corporation.

Notes

1. A *seismic array* is a set of seismic sensors dispersed over a relatively small geographic area whose signals are combined and processed together to produce improved signal-to-noise ratios and improve the extraction of certain features, such as signal azimuth. A variety of new data sources and data types will be added to the IMS network in the near future.

2. Regional distances are within about 2,000 kilometers; teleseismic distances are larger.

3. A detection is determined by a significant increase in the signal over background noise and is interpreted to be an arriving wave front from a seismic event.

4. The *analyst review station* provides the analyst with an extensive graphic and procedural interface to review ESAL's conclusions. This interface includes waveform data, regional maps, and satellite photos that were not available to ESAL. Analysis of the corrections that are made reveals that analysts make use of information in the waveform that is not contained in the current feature set. This analysis might lead to an improved feature set for future processing.

5. It is desirable to have all the detection data for an event, which implies a minimum time interval of about 1 hour. In practice, the interval varies between 3 and 24 hours depending on the density of the data.

6. A performance validator, PerfV, is currently under development.

7. The user typically loads an initializing set of parameters that characterize the data to be run, for example, GSETT data, and then optionally modifies a few of them interactively.

8. Distinct events can occur minutes apart, yet the various waves received by a station from a single event can arrive over a period of tens of minutes because of the different paths taken through the earth.

9. One might expect the search to end when all the detections are associated with an event. Experience has shown, however, that generally about half of the detections are never associated with an event and represent either noise or disturbances too small to be detected at multiple stations (detection at multiple stations is generally a requirement for confirming an event hypothesis).

10. Some of the uncertainties arising from the seismological model might eventually be overcome with sufficient observational experience; measurement uncertainties caused by noise in the seismic signals are unavoidable.

11. It is not adequate to simply record the node. Because reaching a previous node does not indicate an invalid solution, it is also necessary to record information about how to proceed from the node.

12. A detection can only have come from one event; consequently, it is not physically meaningful to have one detection associated with two events.

References

Bache, T. C.; Bratt, S. R.; Wang, J.; Fung, R. M.; Kobryn, C.; and Given, J. W. 1990. The Intelligent Monitoring System. *Bulletin of the Seismological Society of America* 80B: 1833–1851.

Bolt, B. 1976. *Nuclear Explosions and Earthquakes.* New York: W. H. Freeman.

Elvers, E. 1980. International Seismological Data Center: Procedures to Check Events through Dynamic Information and to Estimate Magnitudes, FOA Report C20368-T1, National Defense Research Institute, Stockholm, Sweden.

Goncz, J. H. 1980. Present Status and Dynamic Planning for Automatic Association Programs, Seismic Data Analysis Center Report SDAC-TR-80-2, Teledyne Geotech Co., Arlington, Virginia.

Jeppson, I. 1980. International Seismological Data Center: Automatic Association of Long Period Surface Wave Data, FOA Report, C 20385-T1, National Defense Research Institute, Stockholm, Sweden.

Slunga, R. 1980. International Seismological Data Center: An Algorithm for Associating Reported Arrivals to a Global Seismic Network into Groups of Defining Seismic Events, FOA Report, C20386-T1, National Defense Research Institute, Stockholm, Sweden.