DEVELOPMENT OF A NETWORK DATA SET FOR EVALUATING DETECTION AND NETWORK PROCESSING PERFORMANCE

Benjamin C. Kohl, Theron J. Bennett, István K. Bondár, Brian W. Barker, Walter C. Nagy, Colin Reasoner, Jeffrey A. Hanson

Science Applications International Corporation

Sponsored by Army Space and Missile Defense Command

Contract No. DASG60-03-C-0009

ABSTRACT

The assessment of signal detectors and/or network based event detection, association and characterization algorithms is typically quantified in terms of probability of detection and false alarm rates or other similar metrics. This relies on the use of a data set in which both the desired "actual" signals and presumed false "noise" detections have been identified and characterized.

We have been developing a framework for constructing a data set comprised of waveforms from well-recorded high signal-to-noise ratio (SNR) events, scaled down to various sizes and embedded in clean background noise. The noise is being carefully constructed to be devoid of signals, based on the execution of standard signal detectors and by excluding time-windows of phase arrivals predicted from global, regional and local bulletins. Available noise segments are being merged into continuous blocks of several days of noise data while maintaining the basic attributes of target time period (overall noise level, seasonal, weekly and diurnal variations) and consistency with historically established noise characteristics for each station. The construction of the noise data set is iterative with interactive review of the resulting waveforms.

To provide a utility for assessing the performance of algorithms to properly detect and characterize low SNR signals and small events (e.g. sub-kiloton tamped explosions or fully decoupled cavity explosions), we are theoretically scaling data observed from a variety of large events to approximate the signals expected from smaller events. We are using a frequency-dependent scaling methodology based on the spectral ratio of the source function of the small event divided by the large event. For explosions the source spectral ratio is initially based on the Mueller/Murphy explosion source model. For earthquakes, we are developing similar methods utilizing simplified earthquake scaling models and regionally varying empirically derived sets of source spectral ratios, based on observations from earthquake mainshock/aftershock sequences. We are embedding the scaled earthquake and explosion signals into the clean background noise with distributions of local, regional and teleseismic earthquakes designed to simulate the overall detection rate observed at each of the stations. The framework we have developed allows for the embedding of events in a customizable manner to allow for the construction of various scenarios, e.g., the detection of a small explosion during an earthquake aftershock sequence. Initially we are emphasizing the construction of data sets of events and stations in southern Asia. The resulting data sets are expected to provide a valuable resource for testing and quantifying sensitiveness of signal detection and network processing at low magnitude levels and for evaluating monitoring scenarios that cannot be assessed using currently available network data.

OBJECTIVE

Practical implementations of new or improved monitoring technologies such as signal detectors, network phase association algorithms, location and event identification methods rely on quantitative assessments of performance such as probabilities of detection and false alarm rates. These types of performance metrics are typically obtained through experiments using data sets constructed from archival data records. Such experimental data sets implicitly contain signal and event recordings from numerous unknown sources (e.g. small earthquakes not listed in published local, regional or teleseismic bulletins), potentially contaminating the data set and complicating the interpretation of processing results. Furthermore, they are only representative of the event characteristics and network and station configurations in use at the time of archival data recordings. In particular, nuclear explosions from only a few isolated locations have been recorded on historical sensor networks.

Our objective is to develop an experimental network data set in which target signal and event detections are known, as well as having realistic distributions of false or "clutter" detections and background noise characteristics. We are utilizing a variety of actual nuclear explosion recordings and developing scaling and transformation algorithms to yield surrogate data for nuclear explosions originating at new locations in southern Asia. This will allow for experiments involving low threshold nuclear explosion waveform records under realistic background noise and seismicity conditions.

RESEARCH ACCOMPLISHED

The basic framework for the comprehensive network data set is depicted in Figure 1. It depends upon taking waveforms from well-recorded high signal-to-noise ratio (SNR), scaling them down to various sizes and embedding them in clean background noise. We are embedding scaled earthquake signals with distributions of local, regional and teleseismic earthquakes designed to simulate the overall seismicity and detection rates observed at each of the stations. To facilitate realistic assessments of new or improved monitoring technologies we are scaling and embedding nuclear explosion waveform recordings at low thresholds.

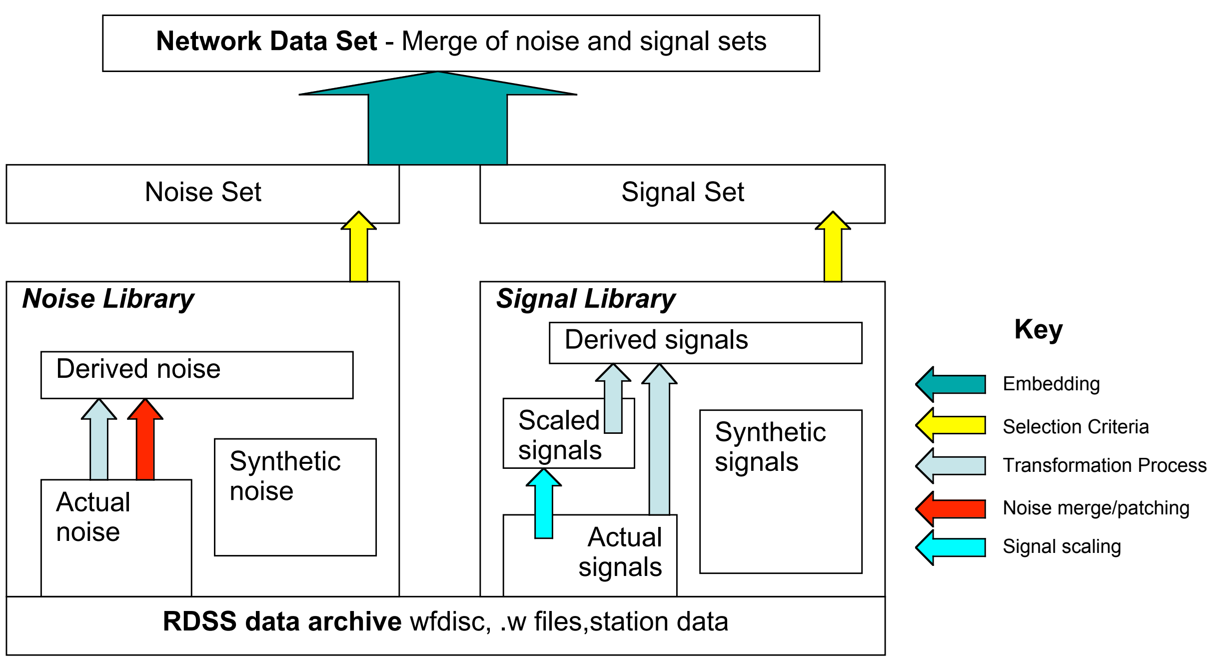

Figure 1. Basic framework for constructing a network data set suitable for systematic testing and evaluation of new or improved monitoring technologies. The data set builds upon the extensive archive of waveform data available through SMDC's Research and Development Support Services (RDSS) database. Scaled signals from a variety of sources (potentially including synthetics) are being embedded into clean background noise to construct a data set in which target and clutter detections have been characterized.

Initially we are emphasizing the construction of a data set of events and stations in southern Asia. Figure 2 shows the target area and key regional stations. To provide utility for assessing the performance of algorithms to properly detect and characterize low-SNR signal and small events, we are scaling and embedding nuclear explosion waveform data to simulate low-yield explosions from four sites in southern Asia marked in Figure 2 with blue squares. Two of the sites are known nuclear test sites (the Indian test site, and Lop Nor China) while the remaining two sites are illustrative locations for assembling surrogate nuclear explosion events. The basic methodology being used for constructing the low-yield explosion surrogate data at the sites marked in Figure 2 is readily applicable to other sites or regions.

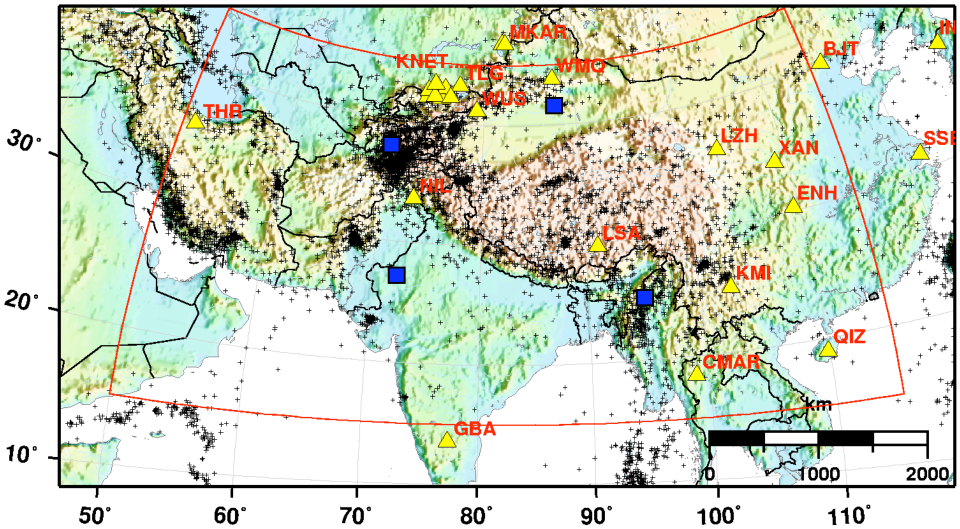

Figure 2. Area and stations in represented in network data set. While the framework we are developing is applicable to other areas, to date we have focused on southern Asia. The background seismicity is shown with "+" symbols, while the locations of hypothetical scaled low yield nuclear explosions are marked with blue squares.

Background Noise Library

The simulation of scenarios of nuclear monitoring interest requires relatively long, continuous noise waveforms spanning several days. However, because of the background seismicity of the earth, such long quiet periods do not normally occur. Figure 3 shows that for an 8-day period in 1998, detection-free intervals rarely exceeded 10 minutes at the CMAR array.

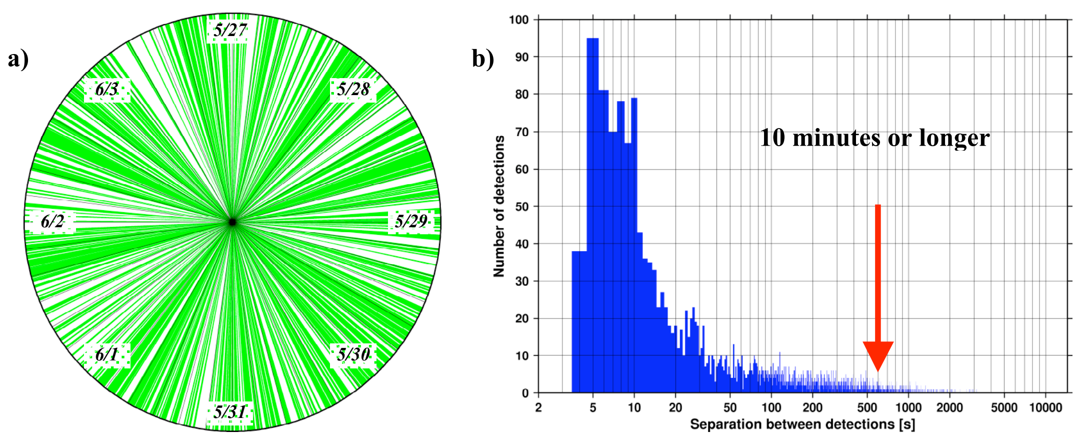

Figure 3. a) Detection-free segments (green) of 10 minutes or longer at CMAR between 5/27/1998 and 6/4/1998. b) Distribution of the lengths of detection-free waveform segments. Long continuous blocks of detection free segments are not readily available; hence we are developing a merging technique that allows numerous detection free segments to be combined into a single continuous trace.

Figure 4 illustrates the procedure of merging noise segments to produce a continuous waveform, devoid of any signals. The interpolation and tapering between the noise segments is carried out in the frequency domain to avoid abrupt changes in the frequency content of the resulting waveform. The final derived waveform is subject to further quality assurance procedures (e.g. no spurious detections are generated at segment boundaries, the spectrogram does not show discontinuities, etc.) before it is used as background noise for embedding events

To provide a realistic background for the embedding of scaled signal data, the merging approach must accommodate the variability of noise conditions observed at each station. We are currently engaged in a number of complementary efforts to characterize seismic and infrasonic noise conditions (e.g. Bowman et al., 2004). Figure 5 illustrates that noise conditions at a station can dramatically change from season to season. The waveform merging procedures ensure the preservation of seasonal and diurnal characteristics in the target, derived noise waveform, as well as the noise levels relative to high and low noise conditions. The low, high and average noise models for a given station build upon the noise characterization efforts of Bahavar and North (2002).

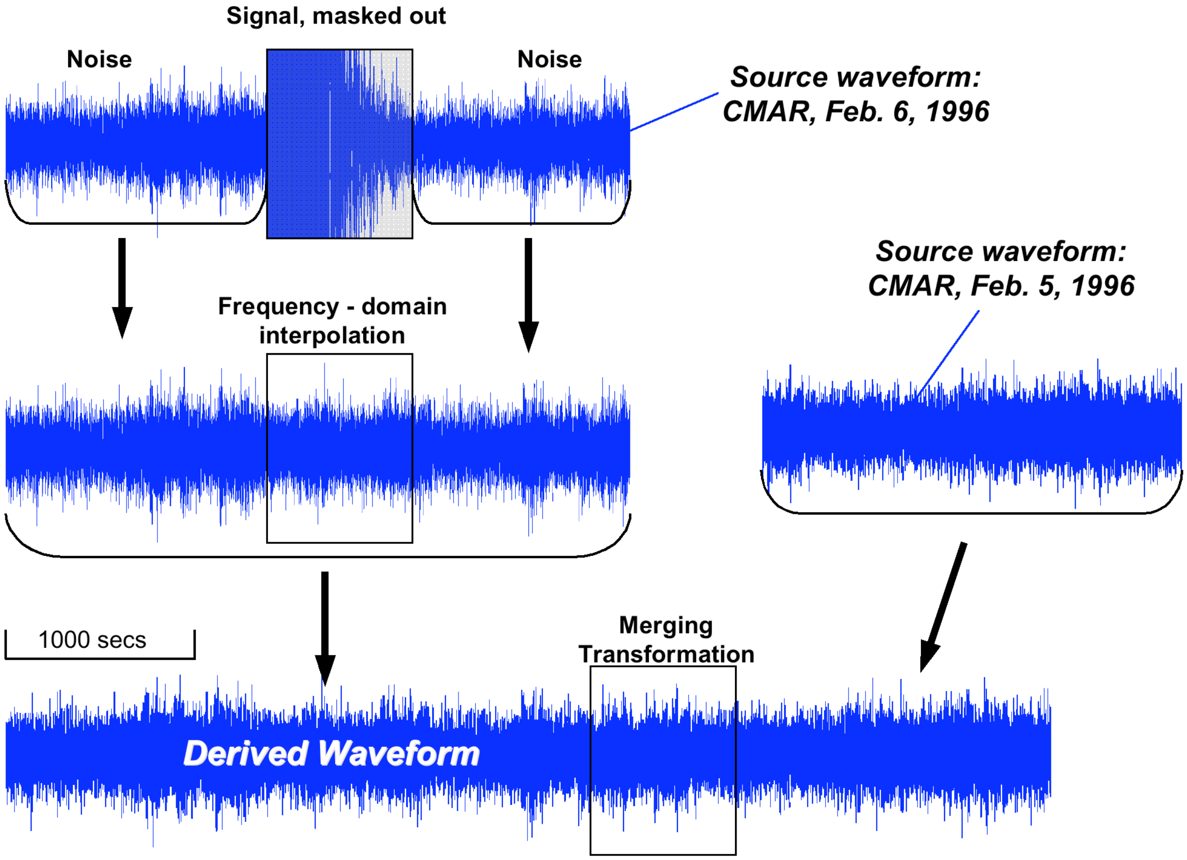

Figure 4. Illustration of the procedure of merging noise segment to produce detection free continuous background noise waveforms. In this example, two waveforms segments and two transformation operators are utilized to form a single continuous waveform trace devoid of signals. In the first waveform (upper-left) an observed signal is first masked and then replaced with a noise segment interpolated from the neighboring noise records. A second segment (upper-right), from a different day but with compatible spectral characteristics is merged with the main segment to extend the resulting derived waveform. These types of operators are repeatedly utilized to form a detection free background noise data set, suitable for embedding scaled signals.

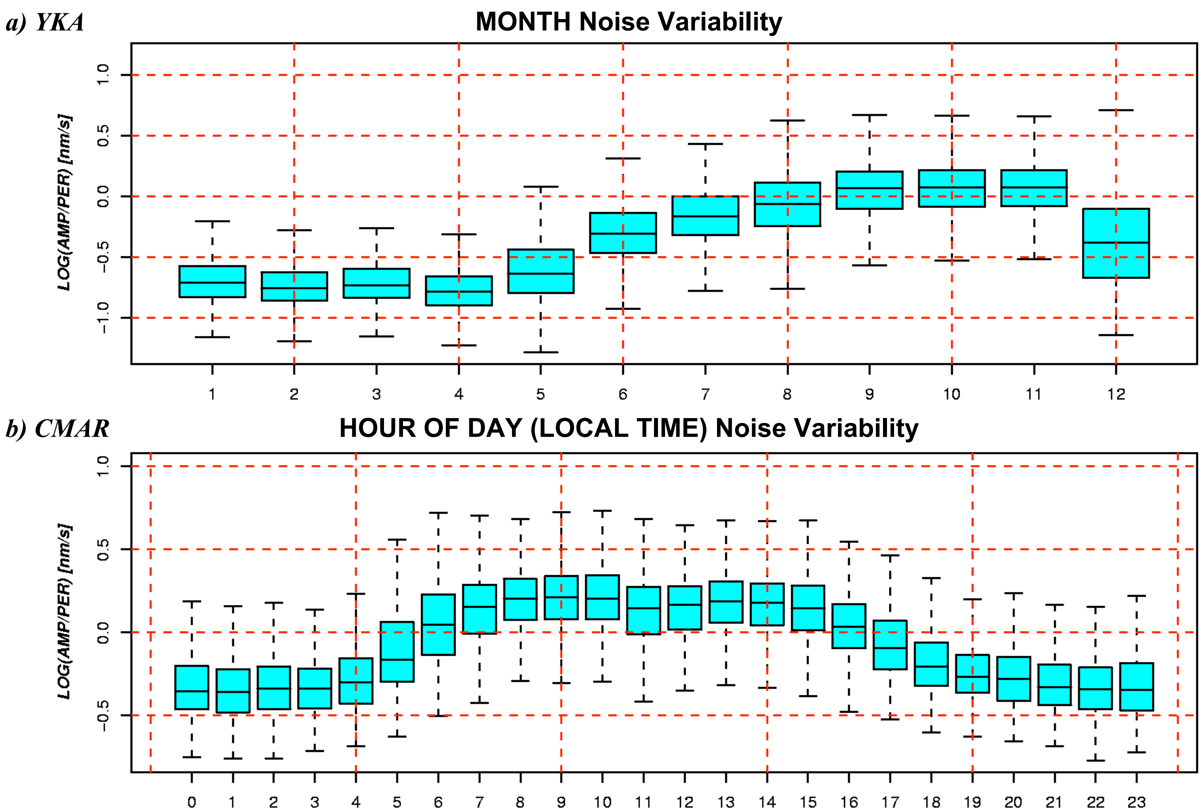

Figure 5. Seasonal variations of noise levels at YKA and diurnal variations at CMAR. The boxes show the median (horizontal line) and outline the upper and lower quartiles of the noise levels. The whiskers extend about 2.5 times the standard deviation. The noise level at YKA (panel a) dramatically decreases during the winter months. Noise levels at CMAR (panel b) show increased levels during the day and decreased levels at night. The construction of the clean background noise set retains these basic features for all stations.

Explosion Sources and Scaling

The expected ground motion time history for a small explosion can be determined from the observed ground motion from a large explosion at the same station if the spectral ratio of the two seismic source functions can be estimated. This approach was successfully applied to scale a number of large nuclear explosions at Lop Nor (Kohl et al. 2002), where we approximated the source spectral ratio using the well-documented Mueller/Murphy explosion source model (Mueller and Murphy, 1971). As a test of the applicability of this methodology, we first theoretically scaled selected observed regional data from the larger of two reference explosions (i.e. May 15, 1995, m_b(REB) = 5.73) to the source conditions of the smaller reference explosion (i.e. July 29, 1996, m_b(REB) = 4.71) and compared the resulting synthetics with the corresponding observed data from the smaller explosion. The results for station MAK ($\Delta \approx 7°$) are shown in Figure 6, where it can be seen that the theoretical scaling very accurately accounts for the observed frequency dependent changes in waveform characteristics, as evidenced by the pronounced change in the L_g/P amplitude ratio between the larger and smaller explosions.

Our ability to quantify network monitoring capability in the region of interest is limited because past nuclear explosions have occurred at only a few sites and have been recorded at just a few regional monitoring stations. In order to develop a more representative data set for the region, we are investigating the possibility of using "surrogate" explosions and associated "surrogate" regional recordings to simulate what might be expected if an explosion occurred at a new, uncalibrated source location and was recorded by existing, or possibly future, regional recording stations. For example, if a nuclear explosion were to occur in the Pamir/Hindu Kush region and was recorded by stations at regional distances, it might be expected that the characteristics of those recordings would be

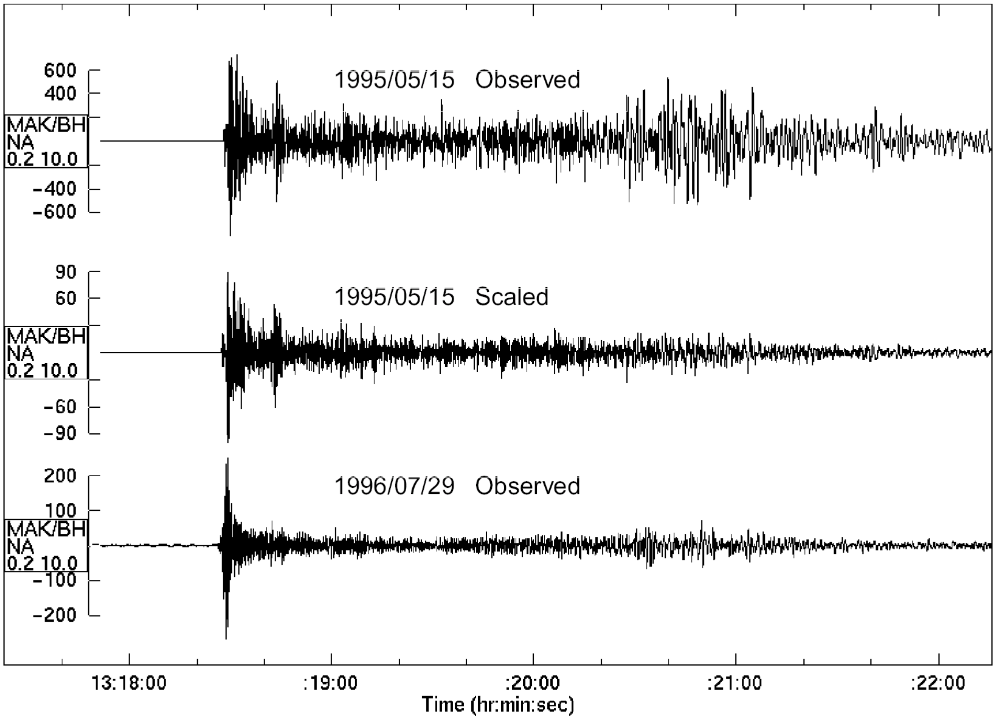

Figure 6. Comparison of the station MAK seismogram obtained by scaling the observed data from the larger May 1, 1995 (m_b= 5.73) reference Lop Nor explosion (top) to the source conditions of the smaller July 29, 1996 (m_b= 4.71) reference Lop Nor explosion (center) with the corresponding observed July 29, 1996 data (bottom).

similar in some ways to the regional signals that have been recorded from previous explosions that have occurred in the region (i.e. explosions at the Lop Nor, India and Pakistan test sites or Soviet PNE's). If we accept some uncertainty, then we can use these prior regional explosion recordings to assess the detection, location and identification performance of a regional network in monitoring a hypothetical future explosion.

Figure 7 shows a sample of recorded regional and far-regional signals from 5 previous explosions of similar magnitude that have occurred in the region, plotted in order of increasing distance. These waveforms represent the general characteristics to be expected at these distances from explosions occurring in the region. At the closer stations, there is abundant broadband S-wave energy which in some cases is equal in size to the P-wave energy. At greater distances, the P-wave energy begins to dominate until the point where the broadband S-wave energy disappears. In a general sense there does not seem to be any particular dependence of these characteristics on the individual source regions. So, if we place our surrogate explosion of an assumed yield at a given location, we can select from this list of recordings, scale the waveforms according to our explosion scaling laws, and "create" the network that we want to have record this simulated explosion.

We are also exploring ways of simulating small station distance changes by applying time and frequency corrections, in order to produce waveforms at distances where none currently exist. If we can accept even more uncertainty, the library of regional explosion signals could be expanded to include explosions recorded from other, more diverse test sites (e.g. STS, NZ, NTS...). It should be noted that other researchers (e.g. Fisk et al., 2000) have, in the course of developing regionalized discriminants, compared nuclear explosion waveform data from a variety of test sites. Clear differences due to variability in attenuation models and other factors have been observed. This suggests that the development of surrogate waveform explosion data at new sites can at best be considered representative of an event recorded in an uncalibrated region. It should not be considered a simulation of precisely what would be expected to be recorded from that site.

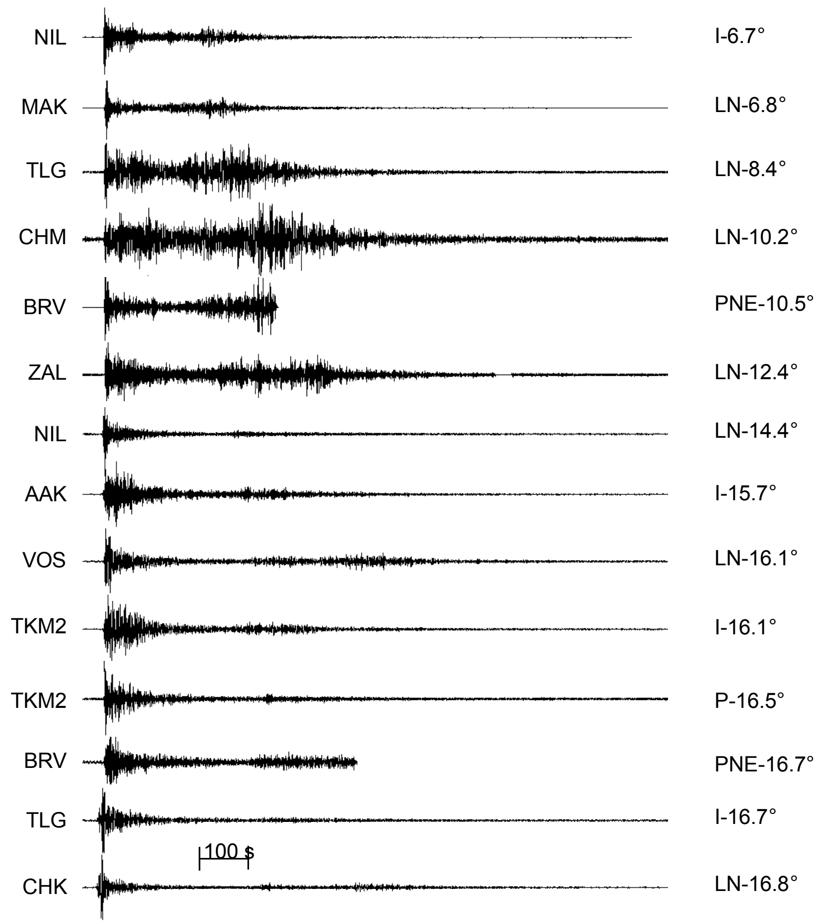

Figure 7. Regional recordings from the Pakistan explosion of 5/28/98 ('P', REB m_b = 4.8), the Indian explosion of 5/11/98 ('I', REB mb = 5.2), the Soviet PNEs of 8/15/73 and 4/11/72 ('PNE', m_b = 5.3, 4.9 respectively) and the Lop Nor explosion of 7/29/96 ('LN', m_b = 4.8).

Earthquake Sources and Scaling

To provide realistic conditions for testing and evaluating new monitoring technologies, the network data set is being constructed to retain a reasonable representation of background clutter signals originating from earthquakes. To this end, we have been reviewing the natural seismicity of the region with the intent of abstracting a representation of the relative level and frequency of natural earthquake activity that will affect seismic monitoring.

The region encompasses several rather active zones including (1) the Pamir-Hindu Kush area, (2) India-Bangladesh-Myanmar border area, (3) Yunan-Szechwan China area, (4) southern Iran, (5) southwest India (Bhuj), and (6) Turkmenistan, as well as several intervening relatively quiet areas. To analyze the natural seismicity, we binned the overall region into a series of 5°-x-5° blocks and used a combination of the Reviewed Event Bulletin (REB) results for the time period 1994/10/19-2001/09/01 as well as published information on earthquake recurrence rates for Flinn-Engdahl (FE) seismic areas (Kagan, 1997). For each bin in our regional grid, we assigned an FE designation and used the published recurrence rates to determine the slope (b) and the REB results to set the level of activity (a) for cumulative earthquake recurrence relations represented by

$$\text{Log } N = a - b\, m$$

where N is the number of earthquakes per year exceeding magnitude m. Figure 8 shows the number of earthquakes ($m_b \geq 3$) expected per year for each element in our grid covering southern Asia. The predicted seismicity levels generally agree with historical activity and published information for the region, i.e. high levels of activity in bins corresponding to most active zones.

The recurrence relations define distribution functions for natural seismicity in southern Asia. We are using these distributions to define expected event magnitudes from earthquakes in the region. That is, for each zone in Figure 8 we have identified a representative large high-SNR "designated" earthquake and have derived the projected number of events at a particular magnitude expected to occur in that zone within a given time period of interest (e.g. 2 weeks). We are constructing the network data set by scaling and embedding the representative earthquake to the projected magnitudes and distribution applicable to that zone, and repeating this for all relevant zones in Figure 8. By tracking the precise embedding times of the signals, the network data set implicitly includes a full inventory of a realistic distribution of real and clutter (e.g. small earthquake) signals. This will facilitate signal detection and network processing experiments resulting in rigorous estimates of probabilities of detection and false alarm rates.

Latitude Range	Longitude Range												
	50-55	55-60	60-65	65-70	70-75	75-80	80-85	85-90	90-95	95-100	100-105	105-110	110-115
40-45	14.7	0.38	4.16	9.54	18.6	44.6	35.4	23.9	2.95	1.62	2.57	3.63	14.4
35-40	25.1	24.5	5.37	120.	645.	281.	19.4	53.7	17.3	23.9	12.3	5.37	4.67
30-35	30.1	32.3	26.3	199.	169.	77.6	34.6	120.	51.2	41.6	24.5	5.37	0.93
25-30	158.	87.0	11.7	371.	11.4	10.7	33.8	79.4	70.7	85.1	107.	10	0.93
20-25	0.	1.14	4.16	10.7	43.6	2.18	3.80	3.80	125.	75.8	21.3	3.31	2.18
15-20	1.69	6.76	0.56	0.	5.88	2.69	0.	1.62	13.4	16.9	3.63	1.81	1.20

Figure 8. Number of events per year of $m_b \geq 3.0$ for 5°-x-5° blocks in southern Asia, based on recurrence relations. This distribution function is being used to define the target magnitudes, and how often representative earthquakes are scaled and embedded in the background noise for each zone.

Key to this approach is the scaling of earthquake data. For source scaling we are using, amongst others, a Brune ω^2 source model (Brune, 1970; 1971) with some modifications including those of Hanks and Kanamori (1979) and source dimension relations determined by Wells and Coppersmith (1994) for continental earthquakes. Figure 9 (left) shows the P-wave displacement amplitude spectra predicted by these scaling models over the magnitude range from 4.8 to 2.8, and we show on the right the relative spectral amplitudes as a function of frequency for scaling 4.8 magnitude earthquakes to the lower magnitudes based on this model. Over this magnitude range, the model predicts significant differences over frequency bands from just below 1 Hz to frequencies in the range 10-20 Hz, with the two corners in the right figure representative of the magnitudes of the two earthquakes (i.e. in scaling larger events the lower corner would be shifted to the left).

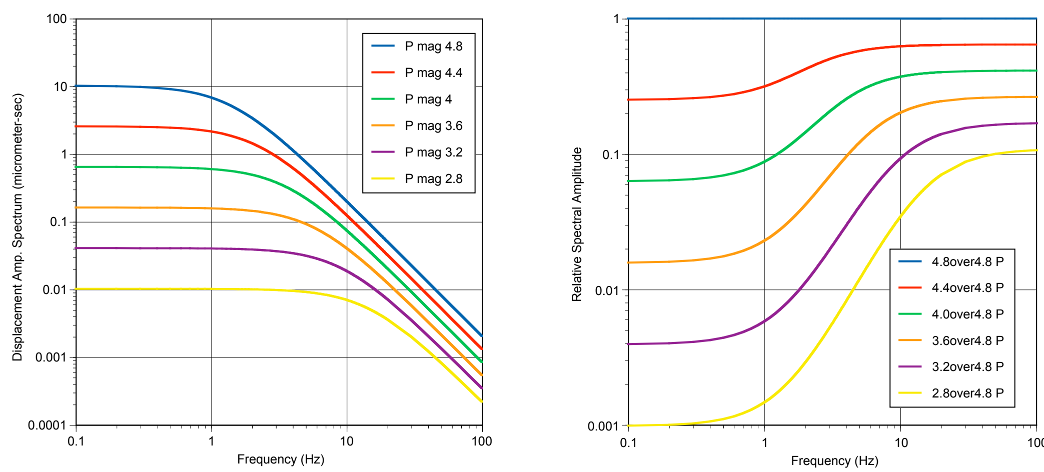

Figure 9. Frequency-dependent earthquake source scaling relationships based on Brune model and Wells-Coppersmith relations for source dimensions of continental earthquakes

Our quasi-empirical approach involves developing relative spectral amplitude scaling relations from the theoretical models and validating their applicably by comparing with empirical spectral ratios derived from sequences of well-recorded events (e.g. large aftershock/mainshock pairs). Recognizing that the detailed earthquake source characteristics depend on variables such as fault plane geometry and dynamic stress drop, in addition to the long-period source size, which is usually quantified as the earthquake moment, we expect different models to be applicable under different conditions. Hence we maintain flexibility and employ a quasi-empirical approach, including models first demonstrated for earthquakes in the Lop Nor region (Kohl et al. 2002).

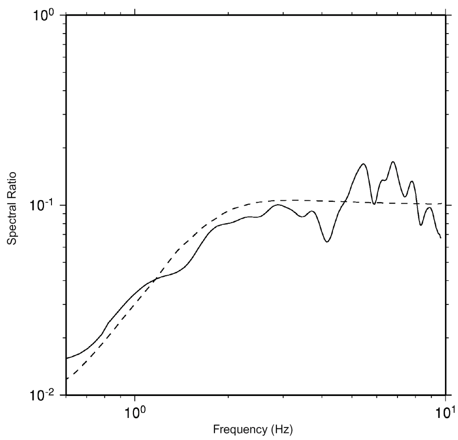

Figure 10. Average observed and model predicted source spectral ratios for the Lop Nor earthquakes of 1999/01/27 and 1999/01/30.

For example, we computed the average P-wave spectral ratio between two reference earthquakes at Lop Nor (i.e. 1999/01/30, m_b(REB) = 5.34 and 1999/01/27, m_b(REB) = 3.93) using data recorded at a number of regional stations. We then fit the average P-wave spectral ratio to a simplified one-dimensional explosion source model. The resulting fit to the average source spectral ratio for the 1999/01/27 and 1999/01/30 earthquakes is shown in Figure 10, where it has been assumed that the ratio of the source corner frequencies scales as the cube-root of the ratio of the two earthquake moments. The simple source model fits the observed spectral ratio quite well over the frequency range extending from about 0.5 to 10 Hz.

The scaling model was evaluated by scaling the observed data from the larger earthquake (1999/01/30) to the source conditions of the smaller earthquake (1999/01/27) and then comparing the results with corresponding observed broadband data at selected regional stations. The results for station MAK are shown in Figure 11 where it can be seen that the adopted scaling model accurately accounts for the observed frequency-dependent changes in waveform characteristics, as evidenced by the pronounced change in the Pn/Lg amplitude ratio between the larger and smaller earthquakes. In this case, where the range in amplitude scaling is about a factor of 50, the broadband amplitude level of the theoretical synthetic matches that of the corresponding observed seismogram to within about a factor of 1.5 at this station. We conclude that this model adequately accounts for frequency-dependent effects of source size on the seismic data observed from these Lop Nor earthquakes.

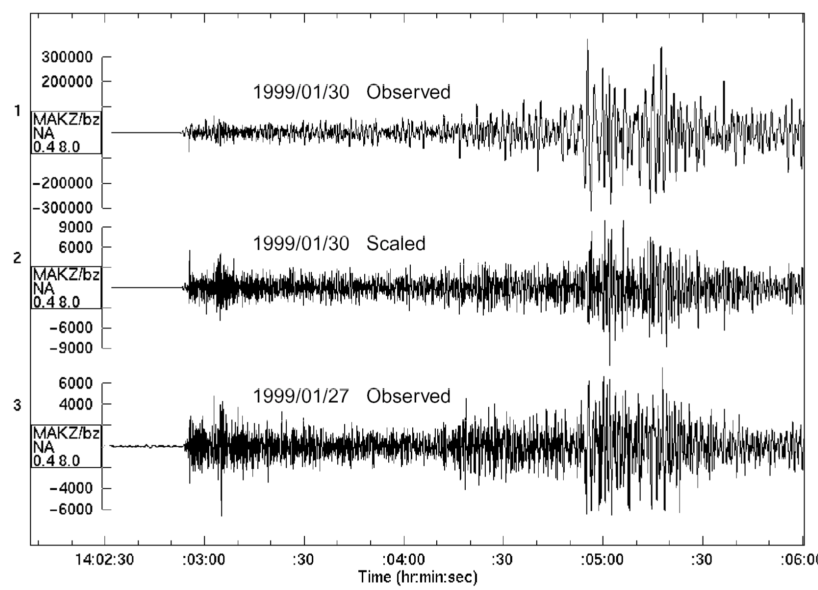

Figure 11. Observed MAK data from a 1999/01/30 (m_b = 5.34) Lop Nor earthquake (top) compared to the earthquake scaled to m_b = 3.93 (center), and compared to observed MAK recording from the 1999/01/27 (m_b = 3.93) event (bottom).

CONCLUSIONS AND RECOMMENDATIONS

We are developing a network data set comprised of actual and surrogate nuclear explosion waveform data to simulate the realistic monitoring conditions expected for southern Asia. By using an approach that involves embedding scaled nuclear explosion waveform data, as well as scaled earthquakes, the data set includes a complete inventory of event and signal attributes over a wide range of magnitude thresholds. Through the incorporation of surrogate nuclear explosion waveform data, the data set will be useful for evaluating monitoring scenarios involving new locations and the latest sensor systems, independent of the availability of historical nuclear explosion data records.

This data set will support a variety of experiments of monitoring technologies focused on improved signal detection, network association and identification capabilities. Through the use of the network data set, with its controlled content including only known scaled explosions and earthquake signals embedded in detection-free clean background noise, rigorous performance metrics will be readily measured. Thus the resulting data sets are expected to be a valuable resource for testing and quantifying the performance of signal detection and network processing at low magnitude levels and for evaluating monitoring scenarios that cannot be assessed using currently available network data.

REFERENCES

Bahavar, M. and R. North (2002), Estimation of background noise for International Monitoring System seismic stations, *Pure Appl. Geophys.,* 159, 911-944.

Bowman, R., E. Baker and M. Bahavar (2004), Infrasound station ambient noise estimates, *in Proceedings of the 26th Seismic Research Review: Trends in Nuclear Explosion Monitoring*, Orlando, FL (Sep. 21-23, 2004).

Brune, J. (1970), Tectonic stress and the spectra of seismic shear waves from earthquakes, *J. Geophys. Res.* 75, 4997-5009.

Brune, J. (1971), Correction*, J. Geophys. Res.* 76, 5002.

Fisk, M., S. Bottone, R. Carlson, H. Gray and G. McCartor (2000), Event Characterization Development and Analysis at the Prototype IDC, *DTRA Technical Report*, DTRA-TR-99-5.

Hanks, T. C., and H. Kanamori (1979). A moment magnitude scale, *J. Geophys. Res*. 84, 2348-2350.

Kagan, Y. Y. (1997). Seismic moment-frequency relation for shallow earthquakes: Regional comparison, *J. Geophys. Res.* 102, 2835-2852.

Kohl, B., R. North, J. R. Murphy, M. Fisk and G. Beall (2002), "Demonstration of Advanced Concepts of Nuclear Test Monitoring Applied to the Nuclear Test Site at Lop Nor, China", *in Proceedings of the 24th Seismic Research Review – Nuclear Explosion Monitoring: Innovation and Integration,* LA-UR-02-5048, Vol. 1, pp 302 – 312.

Mueller, R. A. and J. R. Murphy (1971), Seismic characteristics of underground nuclear detonations. Part I. Seismic spectral scaling, *Bull. Seism. Soc. Am*. 61, 1975.

Wells, D. L., and K. J. Coppersmith (1994), New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement, *Bull. Seism. Soc. Am*. 84, 974-1002.