APR 21 1998

SANDIA REPORT

SAND98-0569 • UC-903

Unlimited Release

Printed April 1998

RECEIVED

APR 28 1998

OSTI

The Waveform Correlation Event Detection System Project, Phase II: Testing with the IDC Primary Network

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Chris J. Young, Judy I. Beiriger, Susan G. Moore, Julian R. Trujillo, J. Mark Harris

MASTER

Prepared by

Sandia National Laboratories

Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

DISCLAIMER

Portions of this document may be illegible electronic image products. Images are produced from the best available original document.

SAND98-0569
Unlimited Release
Printed April 1998

Distribution
Category UC-903

The Waveform Correlation Event Detection System Project, Phase II: Testing with the IDC Primary Network

Chris J. Young
Geophysical Technology Department

Judy I. Beiriger, Susan G. Moore, Julian R. Trujillo
Decision Support Systems Department

J. Mark Harris
Cooperative Monitoring Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0750

Abstract

Further improvements to the Waveform Correlation Event Detection System (WCEDS) developed by Sandia Laboratory have made it possible to test the system on the accepted Comprehensive Test Ban Treaty (CTBT) seismic monitoring network. For our test interval we selected a 24-hour period from December 1996, and chose to use the Reviewed Event Bulletin (REB) produced by the Prototype International Data Center (PIDC) as ground truth for evaluating the results. The network is heterogeneous, consisting of array and three-component sites, and as a result requires more flexible waveform processing algorithms than were available in the first version of the system. For simplicity and superior performance, we opted to use the spatial coherency algorithm of Wagner and Owens (1996) for both types of sites. Preliminary tests indicated that the existing version of WCEDS, which ignored directional information, could not achieve satisfactory detection or location performance for many of the smaller events in the REB, particularly those in the south Pacific where the network coverage is unusually sparse. To achieve an acceptable level of performance, we made modifications to include directional consistency checks for the correlations, making the regions of high correlation much less ambiguous. These checks require the production of continuous azimuth and slowness streams for each station, which is accomplished by means of FK processing for the arrays and power polarization processing for the three-component sites. In addition, we added the capability to use multiple frequency-banded data streams for each site to increase sensitivity to phases whose frequency content changes as a function of distance.

With these improvements we processed the entire test interval. Even with the directional information, however, the performance of WCEDS was uneven. After conservative screening of the WCEDS results to eliminate the worst events, we found that 81% of the WCEDS events correspond to REB events but we also found that 51% of the REB events were not in the WCEDS bulletin. Examination of the events established that performance was directly linked to the quality of the phase information available in the waveforms for each event. WCEDS detects and locates well when many phases are observable but does much more poorly when there are few. Unfortunately, few of the events recorded by the primary network have observable later phases (43% of the REB events have 5 or fewer defining detections) due to the sparseness of the sites, the predominance of short-period instrumentation, and the typically small size of the majority of the events. Given these factors, we concluded that there is little reason to believe that WCEDS could offer a substantial improvement for global monitoring of the primary network, though we are confident that it could achieve results comparable to those produced by the current PIDC automatic detection system given a commensurate level of tuning. As the current system is already tuned and operational, however, we believe that it would be imprudent to pursue further development of the WCEDS system for global monitoring of the primary network

Using the system to monitor local/regional networks, shows much more promise. Processing data from a network in south-central New Mexico, Withers et al. (1998), found that WCEDS can automatically locate events to within 3 km, and has detection capabilities approaching those of a trained analyst. This marked improvement in performance relative to the global system is directly related to the richer later phase information available in regional waveforms. If analogous CTBT monitoring scenarios can be found, WCEDS may yet offer significant advantages over current methods.

Acknowledgments

The fundamental concept upon which WCEDS is based comes from Peter Shearer's 1994 paper; we have made numerous modifications but the core of the process remains essentially the same as he described. Helpful reviews of this report were provided by Marianne Walck and Dorthie Carr. This work was supported by the United States Department of Energy under Contract DE-AC04-94AL85000. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy.

Contents

Abstract.....	3
Acknowledgments.....	5
Contents	6
Figures.....	7
Introduction.....	8
Review of the First Report	8
The IDC Primary Network.....	10
System Description	11
Setup	11
Event Processing.....	13
Testing with the IDC Primary Network	16
Data Set.....	16
WCEDS Processing Parameters	16
Output Examples.....	18
Test Results	22
Discussion.....	25
Conclusions.....	28
References.....	30
Appendix A	
Spatial Coherency Calculations	42
Appendix B	
Three-component Power Calculation.....	44
Distribution	47

Figures

Figure 1.	WCEDS operational model.....	32
Figure 2.	The IDC primary network.....	33
Figure 3.	REB events vs WCEDS events for 24 hour test period	34
Figure 4.	WCEDS detection and screening Master Images	35
Figure 5.	An example processed waveform profile for a large event	36
Figure 6.	Jordan-Syria event without and with azimuthal checking	37
Figure 7.	Nicaragua event without and with azimuthal checking	38
Figure 8.	Kermadec event without and with azimuthal checking	39
Figure 9.	Timor event without and with azimuthal checking.....	40
Figure 10.	All WCEDS events for 24 hour test period.....	41

Introduction

This report succeeds the SAND report SAND96-1916 "The Waveform Correlation Event Detection System Project, Phase I: Issues in Prototype Development and Testing". Hereafter we will refer to that document as Report 1. In that report, we described the development and testing of WCEDS, a new type of seismic event detection system based on waveform correlation, which had been developed specifically for application to Comprehensive Test Ban Treaty (CTBT) monitoring. In this report we will attempt to assess the potential utility of WCEDS for CTBT monitoring. In order to do this we tested the system on a 24 hour data set from the designated CTBT seismic monitoring network, the International Data Center (IDC) primary network. To evaluate the quality of the bulletin produced by WCEDS, we compare it with the high-quality Reviewed Event Bulletin (REB) from the Prototype IDC (PIDC), which is produced by an automatic event detection system (Global Association or "GA" -- see e.g. Le Bras, et al., 1994) and then edited by analysts to correct any errors.

We treat the REB for our test interval as "ground truth" because it is the best bulletin available. We have good reason to believe that this is a reasonable assumption; reviewing numerous randomly selected events from the REB, we found the quality of the events to be excellent. The quality of the bulletin produced directly by the automatic system at the PIDC -- the Automatic Event List (AEL) -- is much less robust, and that is the prime motivation for testing alternative automatic event detection methods such as WCEDS. Though the quality of the AEL continues to improve, many events are still missed or built improperly, and the overall level of confidence in the automatic system is low enough that all events must be reviewed to insure a high-quality final bulletin. Hence, despite the tremendous amount of development and tuning that has gone in to the GA system, the level of analyst effort required to produce an acceptable result is still high. So long as this is the case, viable alternative automatic processing methods should be considered.

The basic format of our report is as follows. First, we briefly review the system presented in the first report. Next we describe the primary network and discuss the challenges presented by it, which required us to make some important modifications to WCEDS. We then describe in detail the new version of WCEDS, highlighting the differences between the current system and the one described in the first report. Then we present the results of testing the new system with the primary network, and follow this with a critical discussion of the performance of the system and how it might be improved. Finally, we state our conclusions and make recommendations for future research.

Review of the First Report

The basic concept upon which WCEDS is based is the idea that a seismic event occurring at a specified location can be detected by collecting waveforms from the monitoring network, processing them to smooth and emphasize phases, forming a time vs. epicentral distance profile for that location, and correlating this with the profile that would be expected if an event had actually occurred. (Note that by correlating the profiles we mean correlation of the respective pairs of waveforms and summation of the results). By repeating the waveform profile correlation at regular intervals in time, we can continuously monitor this location for seismic activity. Shearer

(1994) was able to develop a global automated event detection system based on this concept by gridding up the Earth's surface to define many monitoring locations and by automating the correlation process. Shearer demonstrated that the technique can detect events missed by conventional methods, suggesting that it has a different sensitivity than do approaches relying on explicit phase detection and identification. The potential relevance of that study to CTBT monitoring, however, was not clear for a variety of reasons (different types of data, different objectives), hence the WCEDS project was initiated.

Before proceeding further, we should clarify our use of the term "waveform correlation". First, our system does not employ true waveform correlation (see e.g. Harris, 1991), wherein unfiltered or lightly filtered waveforms are correlated, but rather correlation of heavily processed, envelope-like waveforms. A system based on true waveform correlation can yield tremendous sensitivity and high-precision locations, but is unsuitable to generalized non-proliferation monitoring because it can only be used to search for events which occur where previously recorded events have occurred. By instead using heavily processed (i.e. generalized) waveforms, we can monitor all locations without any dependence on previously occurring events. The only requirement is that our model (i.e. the Master Image, which is described below) applies to any location.

A second potential cause for misunderstanding comes from the correlation process itself. In fact, we use a simple dot product of two time series:

$$= \sum_i X_i Y_i \quad (\text{EQ 1})$$

rather than a true mathematical correlation, wherein the dot product would be normalized by the powers in each of the series:

$$= \frac{\sum_i X_i Y_i}{\left(\sum_i X_i\right)\left(\sum_i Y_i\right)} \quad (\text{EQ 2})$$

We eschew the true correlation due to the common occurrence of intermingled phases from multiple events on waveforms from global networks. In such cases, regardless of the value of the dot product in the numerator for the event of interest, normalizing by the power of the station will lead to a low correlation value due to the significant power in the non-matching phases from the other events. For example, if a phase from a local event happens to arrive during the recording of a teleseismic event, regardless of the number of correlatable phases in the teleseismic event, the local phase may completely dominate the observed power and thus lead to a deceptively low correlation for the teleseismic event. Another reason to avoid true correlation is that if any weighting is used for the Master Image, normalization will remove it. For these reasons, for our "correlation" we use the un-normalized dot product instead of the true mathematical definition.

We began our prototype development with Shearer's system, but had to make several significant modifications to suit the requirements of the CTBT monitoring environment. Shearer processed the waveforms to simplify and improve the correlations, and we followed suit, using

slightly modified algorithms better suited to our data. The basic way in which the correlation for a grid point is calculated, however, had to be radically changed. CTBT monitoring requires a much finer grid with significantly more grid points, which led to poor performance using the correlation method described above. We were able to solve this problem by re-designing the process. Forming an observed profile for each grid point and correlating with the predicted profile leads to a huge number of redundant correlations when applied to the global grid of points, because the correlation of a given station's waveform and the predicted waveform for a given distance will be calculated as many times as there are grid points at that same distance from the station. (Note that this is only true if both the predicted and the observed waveforms are azimuthally invariant -- we assume that this is true.) This redundancy can be avoided by computing, for each station, all of the correlations of the observed waveform with the predicted waveform at each possible distance, and storing the results in a correlation matrix (C), which has a column for each station and a row for each distance. The overall correlation value at each grid point can then be found by summing the appropriate values in C. Using this method, we were able to achieve a tremendous increase in speed for large grids (greater than two orders of magnitude for a 1 degree grid spacing), making the monitoring of large grids more feasible.

WCEDS finds events by searching a given interval of origin times for the grid point and origin time point that have the maximum detector output. If the maximum output exceeds an empirically derived detection threshold, an event is declared. Multiple events with inter-mingled phases are detected by an iterative event masking process. After an event has been detected, the corresponding correlations (true and false) in C are masked using the masking matrix, X. The maximum grid point/time point search is then repeated with the masked C to find other events, which are in turn masked (i.e. the information in X is updated). This process is repeated until the maximum output does not exceed the detection threshold or a specified maximum number of events has been built.

A simplified diagram of the entire WCEDS operational model is shown in Figure 1. New functionality added specifically to process the primary network data (and described later in this report) is indicated with *italics*.

The system was tested on an hour of data from the IRIS (Incorporated Research Institutions for Seismology) broadband, global, three-component, seismic network. Examination of the data indicated that 4 of the 5 events listed in the PDE (Preliminary Determination of Epicenters, produced by the U.S. National Earthquake Information Center) during that hour were observable by the IRIS network. After tuning to optimize performance, we processed the interval and WCEDS built all 4 of these events. This test on a small interval of data was far from conclusive, but was encouraging enough to motivate us to continue the project and test the system with primary network data.

The IDC Primary Network

International seismic monitoring for the CTBT will be done using the IDC primary network, as agreed upon in the text of the CTBT. This is a heterogeneous network consisting of two general types of sites: arrays of single vertical component short-period instruments, and sites with a single set of three-component broadband instruments. The arrays are especially effective at detecting

high-frequency P phases, particularly first arrivals from an event, but offer little sensitivity to lower-frequency coda phases due to their band-limited instrumentation. The three-component sites offer better low-frequency information but do not provide the noise-cancellation capabilities of the arrays and so typically have much higher detection thresholds.

At the time of our test, there were 35 sites in the network: 22 three-component stations and 13 arrays (Figure 2). Ultimately there will be approximately 50 total sites, with the majority being arrays. Though the primary network has been designed with the intent of providing as uniform global coverage as possible, in fact, the coverage is uneven due both to physical constraints (all the sites are located on land) and political constraints (not all countries have been willing to provide the same level of access). As a result, some areas of the world are probably covered better than is warranted (e.g. Scandinavia where 3 arrays are at local-regional distances), while other seismically active or politically interesting areas are poorly covered (e.g. Tonga-Fiji is more than 40 degrees away from the nearest primary station).

The quality of the installation at each primary network site can be expected to be good, but regardless, this uneven coverage will present a significant challenge for any type of system that is to be used to monitor the primary network. In our initial testing we found that for WCEDS to produce a bulletin of quality even approaching that of the REB, it was necessary to modify the system presented in Report 1 to use directional information. This violated the basic structure of the system presented in Report 1, but could not be avoided. Only by using azimuth and slowness (or equivalently, angle of incidence) to verify compatibility of an observed signal with a theoretical signal, could we reduce the number of false events built and the mislocations of true events to acceptable levels. We also found cases, however, where these checks could degrade performance: e.g. an obvious signal clearly belonging to a given event had inconsistent directional information and therefore could not be associated. Thus, imposing this type of consistency check, while essential for processing the primary network, is a delicate matter. In our implementation, tolerances are specified separately for each station; ideally, they should be further specified by source region, ala the GA system at the PIDC.

System Description

WCEDS consists of two types of processes: setup and event processing. The setup processes are those which must be run one time to set up static data files needed for the event processing. The event processing processes are those which are run continuously to produce the event bulletin. The descriptions provided here are brief -- the interested reader should refer to Beiriger et al. (1998) for a much more detailed description of the WCEDS software design.

Setup

There are two setup (one-time) processes which must be run prior to any processing of data by WCEDS: Master Image creation and Grid Map creation. Once these processes are run, data can be processed continuously by WCEDS without re-running the setup processes, unless the parameters have changed.

Master Image Creation

The Master Image (MI) defines the expected patterns of arrivals in the data that indicate that an event has occurred. Literally, it is a collection of the predicted generalized waveforms for a set of discrete distances spanning the range of possible epicentral distances (0 to 180 degrees for a global network). The correlation value calculated for each grid point is found by summing the correlations of the processed observed waveform for each station with the MI waveform at the corresponding epicentral distance. The MI can be created either empirically from stacked processed waveforms (e.g. Shearer, 1991) or from a set of travel time curves (see Report 1). For various reasons discussed in the first report, we chose to use a travel time-based MI then and continue to do so for the primary network testing.

One distinct advantage of using a travel-time based MI is that it is much easier to control the weighting of the phases. Each phase within the MI can be weighted independently and the weight for a given phase can be made to vary with distance. Optimizing the weighting is not important for large, well-recorded events with many obvious phases, but for poorly-recorded events with few clearly identifiable phases, it can be very important. For these types of events, several locations may have comparable correlation values. Proper weighting can cause the system to select the most logical solution. Without the use of directional information, in fact, weighting is the best method we have found to improve the quality of the WCEDS bulletin. Unfortunately, we have not found a simple method to choose the optimal weighting. We use trial and error, performing repeated processing of selected intervals containing events which define the lower limit of what we feel is detectable, and adjusting the weights to get the best results.

In the Shearer study, a single MI was used for all of the processing. As described in Report 1, we have found it to be essential to use two separate MI's: one for detection and one for phase screening. The detection MI should include only the phases that are most commonly seen, because including additional phases will decrease the sensitivity of the detector. This is because both the processed waveforms and the MI's are always one-sided (positive) and hence any additional phases which are included in the detection MI will correlate with noise, leading to some non-zero contribution to the overall correlation products. Thus, unless a phase is typically observed, it is better to omit it from the detection MI. The screening MI, on the other hand, must include any phase which has an appreciable likelihood of being observed. This is because any phases that are present for a given event but that are not screened will lead to false correlations with the detection MI, which could lead to spurious events.

The basic structure of the MI's used here is the same as that described in the first report, except that they have now been augmented to include the slowness for each phase for directional consistency checking. Our MI's are based on the IASP91 travel time curves (Kennett and Engdahl, 1991), with the width of the various phase branches being determined by the grid spacing, MI time and distance discretization, origin time processing interval, and range of possible event depths. See Report 1 for details.

Grid Map Creation

The Grid Map (GM) is an important feature of the system that was not explicitly described in Report 1. The GM consists of a set of structures -- one per grid point -- which contain the information needed to accurately and efficiently calculate the overall correlation for each grid point. In

the system described in the first report, the information for each grid point consisted of a list of the stations and for each, the position (row) in the C matrix corresponding to the epicentral distance from each grid point. Thus, to calculate the correlation for a particular grid point for a particular origin time, all that is needed is to use the GM information for that grid point to sum the appropriate rows for all of the columns of the C matrix for that origin time. To implement directional consistency checking for the primary network system, we augmented the GM structure to include the theoretical azimuth to each station (obtained from spherical geometry). Note that we do not include theoretical slowness values in the structure because this information is better stored in the MI.

Event Processing

Once the MI and GM have been created, event processing can be initiated. There are three parts to this: waveform processing, correlation, and event detection/location.

Waveform Processing

In order to process data from the primary network arrays in a manner consistent with our azimuthally-independent C matrix approach, we required a method to produce a single optimal data stream for each array to correlate with the MI. Planar wavefront beamforming is conceptually the simplest approach to processing array data, but it would lead to a large set of beams spanning the set of directions corresponding to the grid points, and consequently would be incompatible with our requirement of azimuthal independence. Instead, we chose to use the spatial coherency method of Wagner and Owens (1996) which produces a single output stream for any multi-channelled site (three-component, array, three-component array). For a given time window, the algorithm transforms each of the channels to the frequency domain and then performs an eigenvalue decomposition of the covariance matrix. The output for each window is the principal eigenvalue, and has been shown to be comparable in sensitivity to a beam aimed at the known source. A detailed description of the algorithm is given in Appendix A. A distinct advantage of the spatial coherency method is that it automatically adapts to changes in the dominant direction of energy, and hence always provides the maximum observed coherency. It is also more robust than planar beamforming because it does not assume a planar wavefront. We found this algorithm to be so effective and simple that we chose to use it for processing all of the primary network sites, both arrays and three-component.

As mentioned above, in order to adequately monitor the globe using the primary network, we found it necessary to use the directional information available from the three-component and array sites. Spatial coherency does not provide directional information, so we had to use other methods to obtain it. In the case of the arrays we use an FK transform on a moving data window and find the peak value: the corresponding azimuth and slowness are the outputted values. For three-component data, for a moving window we compute the polarized power over a slowness grid similar to that used for FK and, again, take the azimuth and slowness values from the slowness grid direction with the greatest power (see Appendix B for details).

To improve sensitivity, we also modified WCEDS to allow the production and use of multiple processed streams for each site. For a global system, each station may be detecting events at dis-

tances ranging from local to teleseismic. The signals from these different regimes have very different frequency content, making it difficult to define a single processed stream which has good sensitivity to all of them. Instead, we allow the user to specify several different streams for each site, with the stream parameters tuned to the particular targeted signal regime. For a global network, we have found it to be effective to create three different types of streams: low frequency, mid frequency, and high frequency. Each is correlated with the detection MI separately, as if it came from a different station. Not only does the use of the multiple streams increase the frequency sensitivity of the system, but because we provide a means to control the distances at which each of the streams are correlated, it also provides a means to avoid spurious correlations: e.g. a high frequency stream added to improve regional sensitivity can be prevented from being used for correlations beyond regional distances. Without directional information available, this correlation range control can be critical: e.g. if we were processing a network of single component sites. With directional information, however, slowness constraints provide a more precise means to accomplish the same thing.

Correlation

The basic correlation technique is essentially as described in Report 1, except that we correlate separately for each frequency band stream and for each phase in the detection MI and store the results in separate C matrices. These modifications were needed to implement the directional consistency checking, which is done on a stream-by-stream and phase-by-phase basis. The separate C matrices allow us to control not only which stations contribute to the correlation value at a grid point, but also which streams and which phases from each station. In the new method, the correlation values for each of the C matrices are generated as before, except that the detection MI used to calculate each contains only the appropriate phase. Both modifications improve system sensitivity, but they also considerably increase memory usage and this can become prohibitive if the number of detection phases and/or processed streams per station is large.

The correlation step is also the point at which we implemented the slowness consistency check. As a result, we in effect guarantee that we never load slowness-incompatible correlations in the C matrices. With the new phase-by-phase correlation, the check is straightforward: if the slowness stream associated with the data stream is not consistent (within a prescribed tolerance) with the slowness of the MI phase against which the data is being correlated, the resulting correlation value will not be added to C. The check is intended to prevent observed and theoretical phases with grossly incompatible slownesses from correlating with each other (e.g. Pn with PKP). As mentioned previously, this type of check is essential for processing data from a sparse network like the primary network, but it can also degrade the sensitivity of the system if not used carefully. Imposing overly tight tolerances on the match between observed and predicted slownesses for some stations can prevent many useful arrivals from being properly associated with the grid point corresponding to the true location.

Event Detection and Location

Detection and location are accomplished by searching over a specified time interval (e.g. 2 hours) for the grid point and time point with the highest correlation value: if that value exceeds a specified threshold, an event is declared. The correlation value for each grid point for a given origin time is found by following the appropriate summation path through the C matrices for that origin time, using the information stored in the GM. The new system differs slightly from that

described in Report 1 in two ways. First, the summation must be done over all of the detection phases and all of the frequency bands due to the new phase-separated and stream-separated correlations. Second, for each phase correlation that exceeds a prescribed phase significance threshold, the observed azimuth is checked against the theoretical azimuth recorded in the GM: if the two do not match within a prescribed tolerance, the phase correlation will not be added to the total value for the summation path. Note that separate checks are made for each frequency band stream, so if any of them are azimuthally consistent, the correlation from that station-phase will be allowed. Thus, it is entirely possible that the list of contributing stations and phases for a given grid point will differ for the three frequency band streams.

As with the slowness check discussed above, the azimuthal consistency check is delicate. Directional consistency checking must be used to adequately build many of the smaller events commonly listed in the REB, but poor choices of tolerances can degrade performance.

Event Masking

Once an event is built, it must be masked to allow the detector to find other events within the same time interval. To properly mask an event, it is necessary to identify all of the phases which are present, not just those in the detection MI, hence the use of the phase screening MI. Masking is performed for every combination of observed phases with detection MI phases. The only significant difference in the new system is that because there are separate C matrices for each detection phase, there must be masking for each. While this does imply separate X matrices for each detection MI phase, the same is not true for each frequency band stream; for a given phase, if an observable phase is found on any of the streams, it is assumed to be present on the other streams as well, and hence the same masking can be used for all of them.

Though the separate C matrix per detection phase approach has some disadvantages -- increased memory usage and system complexity -- it makes the masking significantly more precise. Consequently it should lead to fewer inadvertent "blind spots" wherein an event could be missed. For example, in the previous system, because the correlations for all of the detection MI phases were summed into the same C matrix, masking {distance, station} positions in that matrix meant that in effect we were masking all of the detection MI phases for that station at that distance, regardless of whether or not they had actually been observed. In the new system the correlations for each detection MI phase are stored and masked separately. If only one of the phases is actually observed, then only the correlations with that phase would be masked, which should prevent us from inadvertently masking critical phases from another event. Note that without the use of a directional consistency check there would be no way to tell whether the phase belongs to one event or another, so one might just as well use the single, combined C matrix approach.

Continuous Operation

WCEDS was intended to be compatible with the IDC/USNDC automatic data processing pipeline, and hence we designed it to be able to process data continuously. A model for continuous processing was described in Report 1, but had not yet been implemented when that report was published. The model has since been thoroughly tested, and we used it for our processing of primary network data. For a detailed description of the design, we refer the reader to Report 1. The key points are:

- WCEDS breaks the data flow up into finite chunks (e.g. two hours), which it then processes iteratively, building as many events as possible until a minimum detection threshold is reached or a maximum number of events has been built.
- The chunks of data are further divided into trusted and untrusted intervals, according to whether all of the data necessary to perform the correlation with the MI is available in the chunk. The detected events and the corresponding X matrices are kept separately for the trusted and untrusted intervals.
- When WCEDS proceeds to the next chunk of data, it overlaps the previous chunk by an amount equal to the length of the untrusted interval (i.e. the entire untrusted interval from the previous chunk will be re-processed as a trusted interval in the new chunk).
- Trusted events are put in the final bulletin; untrusted events are not.
- Only the masking from the trusted events in the previous chunk is carried over to process the new chunk to prevent trusted events from being rebuilt. Untrusted events should be rebuilt if they are legitimate events.

Testing with the IDC Primary Network

Data Set

For our testing we selected a day's worth of data spanning 00:00:00 GMT on 12/24/1996 to 00:00:00 GMT on 12/25/1996. The primary network configuration for which we were able to obtain data consisted of 13 arrays and 22 three-component sites (Figure 2). We chose to exclude data from the stations CMAR (Chang Mai Array, Thailand) and PDY (Peleduy, Russia) due to consistent noise problems, but used all the other stations. To assess the WCEDS results, we compare with the REB from the PIDC. This is the best bulletin available for the primary network, and should include essentially all of the events that could be detected and located with this data set, though we note that the PIDC used data from a few additional stations which were not available to us.

The 24 hour interval we used was selected essentially at random, but nevertheless posed several excellent tests for WCEDS. The REB lists 79 events, with mb values ranging from 2.5 to 4.7 (Figure 3, top). The number of defining stations used for location (ndef) ranges from 3 to 53. Throughout our discussion, we prefer to use ndef rather than magnitude when discussing the events because the former is a more reliable indicator of how "detectable" an event is.

WCEDS Processing Parameters

For simplicity and quicker processing time, we used a single, uniformly spaced surface grid only. The grid spacing was 2 degrees, leading to a total of 10357 grid points. Although WCEDS can easily be extended to process grids at several depths (see Withers et al., 1998), we chose not to do so because it would considerably slow the processing time. (Note that the problem is not related to the number of grid points -- relatively few are needed to monitor the limited areas of the

Earth which have deep seismicity -- but to the fact that separate C matrices must be calculated for each depth.) Given the relatively small number of deep events in the test interval, using deep grids was unlikely to have a significant effect; of the 79 total events in the REB, 51 were assigned $d = 0$ km, and 12 more were assigned depths between 0 and 100 km, leaving only 16 with assigned depths greater than 100 km.

The detection MI consisted of P (i.e. the first arriving IASP91 phase at every distance) and S only. No other phase was observed often enough to warrant including in the detection MI. In fact, even S was marginal, but our tuning indicated that when present it was useful enough to offset any problems related to false correlations. P was given a weight of 2 at all distances; S was weighted as 1. The screening MI consisted of P, PcP, PP, PPP, S, ScP, SS, SSS, Lg, PKPab, PKPbc, PKPdf, and PKiKP. These phases were selected by reviewing the REB to find the most commonly observed phases. In keeping with the grid, only surface (zero depth) MI's were used, though the phases were broadened slightly (about 5 seconds) to help detect deeper events. The MI distance span was 0 to 180 degrees with a discretization of 1 degree, and the time span was 0 to 2400 seconds with a discretization of 1 second. Both the detection and screening MI's are shown in Figure 4.

Whenever possible, three sets of processed streams were created for each site: low frequency (0.5 - 1.5 Hz), mid frequency (1.5 - 2.5 Hz), and high frequency (2.5 - 3.5 Hz). If the site had only short-period data, only mid and high streams were created; if the site had broadband data, a low stream was also created. We selected these three pass-bands after extensive analysis of selected REB events as the bands which consistently showed the best signals. In some cases using additional bands (e.g. lower frequency for S arrivals) could lead to a few additional phase contributions, but not without increasing the overall noise level for event detection. Also, for the primary network, additional phases seem to be observed only for larger events, for which there are generally enough P arrivals to lead to a good location regardless. Note that this conclusion might not apply to a different network; the primary network design emphasizes P phases (i.e. first arrivals), and consequently the lack of later phase information observed may be misleading.

We allowed each of the frequency streams to correlate at all distances (0 to 180 degrees) to take advantage of the surprisingly high-frequency P phases which we observed for some events at great distances (e.g. PKP phases). Without using slowness consistency checking, this would be impractical due to the potential false association of local phases.

To create the streams to correlate for each band, the waveform processing consisted of the following sequence: bandpass filtering (3 pole, Butterworth), spatial coherency processing, STA/LTA processing (STA length = 3 seconds, LTA length = 27 seconds), and zeroing below a specified threshold (to decrease noise correlations). An example of a processed data profile is shown in Figure 5. It is worth noting that the processed waveforms for this event, which is one of the better recorded events in the test interval ($n_{\text{def}} = 53$ in the REB), are dominated by the first arrivals. This is an important result which we will return to when evaluating the WCEDS performance. Despite a very flexible signal processing interface and extensive tuning, we were unable to bring out significant later phase information, even for the best events.

For the spatial coherency calculation, the processing window length was 10 seconds, and the windows were overlapped by 90% to create output streams of one sample per second. To create the azimuth and slowness streams for the directional consistency checks, the processing consisted of the same bandpass filtering followed by the FK or polarization calculations described previously. Again, the processing window length was 10 seconds and the windows were overlapped by 90% to get a 1 sample per second output. Azimuthal consistency tolerances were generally 10 degrees at the arrays and 20 degrees at the three-component sites, though two of the three-component sites (KBZ and ZAL) had to have azimuthal checking turned off entirely due to consistently poor estimates even on high SNR signals. Slowness tolerances were 4 seconds/degree, a fairly loose constraint. We experimented with much tighter slowness constraints but chose to abandon them due to uneven results. For some events with poor azimuthal station coverage, tight slowness constraints (e.g. 0.5 seconds/degree) led to tremendous improvements in location, but in other cases the same constraints prevented events with good station coverage from being built correctly or at all, presumably due to heterogeneities along the ray path which cause the ray to diffract significantly and therefore to have a large mismatch with the predicted slowness. If detailed geographically varying slowness information were available for the stations, it could be incorporated into the Grid Map, ala GA, allowing us to use a much tighter slowness constraint.

The event detection correlation sum threshold was set to 250, a number which was determined empirically. We ran WCEDS on selected intervals within the test set which contain what we judged to be the smallest events which we could reasonably expect to detect, and noted the detector output for the nearest grid point/time point for these events. The smallest value (250) was selected as our detection threshold. We also implemented a parameter to control the maximum number of events which could be built in a given interval to prevent the system from spending too much time on one interval. This parameter was set to 16 events. For the origin time discretization, we used 4 seconds based on a series of tests to find the largest discretization which could be used without missing events. Data were processed in 1 hr 40 min chunks, with a trusted interval of 1 hour and an untrusted interval of 40 minutes.

Output Examples

Before discussing the overall performance of WCEDS, in order to better understand how WCEDS works with real data, we briefly examine four events from the test interval. Though we have chosen to impose azimuthal consistency checking in the actual processing of the test interval, for illustrative purposes we show results with and without the checking. Note that all four of these events were detected by WCEDS. In all but the last case, the epicentral information listed at the start of each section comes from the REB.

Jordan-Syria Border, 22:16:28.7, $d=28$ km, $mb = 4.5$, $ndef = 53$

This event, which was shown in Figure 5, was the most widely recorded event in the interval and so could be expected to be an example of WCEDS at its best. As discussed above, however, relatively few phases other than the first arrival are recorded in the REB, which lists: P or Pn (30), Pg (2), Sn (1), Lg (1), PcP (2), and pP (19). The lack of later phase information is typical of the REB events and probably reflects the bias of the primary network design towards first arrivals.

The most sensitive stations are the arrays, and these consist of short-period vertical instruments, so S-wave phases and other lower frequency phases can be expected to be under-represented.

For our test, pP is a nuisance phase because we are using a surface MI only, which cannot match this phase. Fortunately, for such a shallow depth the broadening introduced by the processing merges P and pP anyway, so we had no trouble with spurious correlations caused by the unmatched pP.

The WCEDS output for this event without azimuthal consistency checking is shown in the top half of Figure 6. We show the network correlation values plotted at each grid point on the globe for a series of processed origin times relative to the REB origin time: -4 minutes, -2 minutes, 0 minutes, +2 minutes. These pictures illustrate typical features of WCEDS output. The panel for each time shows a set of intersecting circles centered around the primary network stations. (Compare with Figure 2 for the station locations). In some cases, the rings are more than 90 epicentral degrees from a station and so appear to be centered on the corresponding antipode. Each ring represents the correlation of some feature in the one or more of the processed streams for the corresponding station with one of the phases in the detection MI (in this case either P or S). If the feature is in fact one of the phases included in the detection MI, then it will have a true correlation ring corresponding to the observed phase correlating with the correct detection MI phase, and one or more false correlation rings corresponding to the observed phase correlating with the other detection MI phases (in this case there is only one other phase). This concept is discussed in much greater detail in Report 1. The important points are that the correlations observed for each phase will show up as rings, and that only one of the rings will correspond to the true correlation. At the correct origin time and only at that time, the true correlation rings for each observed phase (in this case P_{obs} with P_{theo} and S_{obs} with S_{theo}) will be co-located and the values added, making the correlation value on that ring greater than on any other ring. Further, at that time and only at that time, the rings for every correlating station will intersect at the location of the event. Again, we refer to the earlier report for more detailed discussion.

With these ideas in mind, the WCEDS output sequence can be easily understood. In the -4 min panel, the highest correlation rings are centered around the primary network arrays in Europe (ARCES, ESDC, FINES, GERES, and NORES all correlate well). These are the P_{obs} with P_{theo} correlations, and there seem to be at least three prominent ones. There is also a fainter ring corresponding to the African station BGCA. The European station rings intersect already, but not all at one well-defined location because the origin time is incorrect. The P_{obs} with S_{theo} rings have much smaller radii and are much fainter, but can also be seen. There are no observable S phases (see Figure 5) so there are no S_{obs} with S_{theo} rings.

As the time advances to -2 min, the rings draw in closer to the stations, creating a better defined region of intersection in the Middle East, though there is still not one clear point of intersection. The smaller circles make it easier to find the corresponding station for each: the ring corresponding to ESDC, the array in Spain, is particularly clear. Note that the BGCA ring is smaller also, but has not yet intersected with the rings from the European stations.

At 0 min, the European rings have contracted further and now the region of intersection is well-defined and corresponds closely to the REB location. The event actually detected by

WCEDS (which used an azimuthal constraint) had a location difference of 1.9 degrees and an origin time difference of 12.7 seconds (early) relative to the REB values. The BGCA ring is still present though difficult to see, and now it passes through the same location. In fact, this location is intersected by several more rings than are visible here; when we examine the WCEDS bulletin we find that in addition to the stations listed above, significant correlations were also found for ILAR, KBZ, and YKA. This panel should also make it clear that not all of the rings that are visible are related to this event. The processed waveforms for each station can be expected to have many occasional above-average amplitude features, and these will always correlate with the detection MI phases at some combination of origin time and distance, leading to rings. It is the intersection of rings from many stations that define events.

At +2 min, the rings have contracted yet further and no longer intersect at one location. Note that the BGCA ring is now clear again. Based on the sequence of panels one might conclude that the ultimate fate of the correlation rings is to collapse inward to each of the corresponding stations, and this is indeed the case. Viewing an animation of the WCEDS output sequence would be like watching a reversed film sequence of pebbles dropped into a pond; ripples propagate inward toward the source points until they reach them and disappear. Along these lines, those familiar with the techniques used in the petroleum industry might view WCEDS as a type of migration processing (e.g. Yilmaz, 1987). In this case, we assume the Earth structure to be known via the MI, and seek the epicentral parameters. Because the origin time is also unknown, we must migrate for a set of possible origin times and choose the one which yields the best answer.

This event would have been detected and located very accurately by WCEDS without azimuthal consistency checking, but this would not be true for some of the other events in the 24-hour interval, hence the need to implement the directional consistency checking. The same sequence of panels for this event with azimuthal consistency checking is shown in the lower half of Figure 6. In this case, correlations from a given station-phase pair have been disallowed if the directions from the corresponding azimuth and slowness streams are not consistent with the expected values plus or minus the tolerance value specified for that station. Now we see that most of the portions of the rings are gone because they are inconsistent with the theoretical directional information. The resulting sequence of panels is much less cluttered, though perhaps also a little less easy to understand. Nevertheless, with the understanding we have gained from our detailed interpretation of the sequence in the top half of Figure 6, this sequence should make sense. In this case, we see arcs of high correlation closing in on the known location and coalescing at that point at the correct origin time.

Nicaragua, 14:34:16.1, $d=24.3$ km, $mb = 4.0$, $ndef = 12$

This event is included as another example of an event which was well-handled by WCEDS. The event was not seen as widely as the Jordan-Syria event, due to the smaller number of primary network stations nearby, but it was still relatively well-recorded. Most of the picks in the REB are first arrivals, but there were weaker PcP or pPcP's picked at 4 stations.

The sequence of panels without azimuthal checking is shown in the top half of Figure 7. Throughout the sequence, the dominant feature is the P_{obs} with P_{theo} ring associated with the southwest US array TXAR. The smaller P_{obs} with S_{theo} ring for this same station is also easily visible. One can also find a strong ring for the array YKA in northern Canada. There are other sta-

tions which contributed to the detection of this event (PDAR, ULM, SCHQ, ASAR, and WRA) but their correlation values are too low to yield visible correlation rings in this figure. The location, as defined by the intersection of the rings in the 0 min plot, is not as sharp as for the previous event, but is still fairly good. Interestingly, the agreement with the REB event information was actually better: the true maximum correlation event built by WCEDS had a location difference of 0.8 degrees and an origin time difference of 4.1 seconds (early) relative to the REB values. Apparently, though the contributions of the stations whose correlation rings are not visible here are small, they are significant and the overall good distribution of stations leads to an accurate location, though the elongated shape of the region of high correlation values in the 0 min panel indicates that it is not as well-constrained as was the previous event.

The sequence of panels with the azimuthal checking is shown in the bottom half of Figure 7. Again, the directional consistency checking is not needed to detect this event, but it makes the preferred location all the more apparent. The dominant feature in each panel is the arc of high correlation from TXAR.

Kermadec Islands, 05:11:07.3, $d=0.0$ km, $mb = 4.3$, $ndef = 5$

This event represents a considerable step down in quality from the previous two, though it is intermediate in magnitude. This is due to the relatively sparse network coverage provided by the primary network for the south Pacific (and the southern hemisphere in general, for that matter). The complete REB information consists of a first arrival (P or PKP) at five stations: CTA, ASAR, WRA, BGCA, and FINES. Note that three of these are in Australia, and one of those (CTA) was not available in our data set.

The sequence of panels without azimuthal checking is shown in the top half of Figure 8. The most prominent features by far are the rings associated with the two Australian arrays ASAR and WRA. For our purposes these arrays are very nearly co-located, hence their correlation rings almost perfectly overlap. The highest correlation rings are P_{obs} with P_{theo} while the smaller, lower correlation rings are P_{obs} with S_{theo} . The separation between the rings for the two arrays is just barely discernible, and is probably most easily seen in the -2 min panel. There are several other fairly strong correlation rings, but most of these are not related to the event.

At the REB origin time, the REB location (just north of the North Island of New Zealand) seems to be intersected by not only the rings from the Australian arrays but also rings which come from stations in very different locations. One of these is from BGCA in Africa, while the other appears to be from FINES in Finland. WCEDS also includes a very weak LPAZ correlation which is not readily apparent in this figure and is not included in the REB.

In this case, the true maximum correlation event built by WCEDS had a location difference of 2.2 degrees and an origin time difference of 4.7 seconds (late) relative to the REB values. These are quite respectable values for this event and they are, perhaps, deceptively good. Without the presence of the BGCA correlation the location would be poorly defined, making the directional consistency checking essential.

The sequence with directional checking is shown in the bottom half of Figure 8. The tremendous value of the checking is evident. We can still see a problem with resolving the correct dis-

tance of the event from the Australian arrays due to the large slowness tolerance used, but never the less the area of compatible locations has been greatly reduced. Clearly a tighter tolerance on slowness could constrain this event better, but we found that for several other events, the observed slownesses for the P phases showed much larger disagreements with the predicted values. This was particularly true of subduction zone events (e.g. near Japan), perhaps due to path diffractions caused by interactions with subducting slab structure. Applying a tight slowness tolerance to these phases would prevent them from being associated properly, hence we chose to use the loose 4 sec/deg value. Clearly, being able to use a different slowness tolerance for each grid point would offer the best solution, if such information were available.

Timor, 00:54:00.0, depth = 0 km, mb = ?, ndef = 2 (WCEDS information)

The information from this event comes from WCEDS because the event was not included in the REB. Nevertheless, examination of the waveforms revealed that it is without question a legitimate event. WCEDS built the event using only the two Australian arrays, but for one of these (ASAR) it was able to correlate S as well as P, and that helped constrain the location. The event provides an excellent illustration of the importance of directional consistency checking.

The sequence of panels without azimuthal consistency checking is shown in the top half of Figure 9. In this case, the situation is a little different because there is an observed S phase, leading to another set of correlation rings. Again, the prominent features are the rings associated with the Australian arrays. The highest correlation rings are P_{obs} with P_{theo} and S_{obs} with S_{theo} . These rings are all so close together that it is difficult to identify them separately. The lower correlation rings are P_{obs} with S_{theo} (within the high rings) and S_{obs} with P_{theo} (outside of the high rings). It is clear in this case that the true correlation rings intersect to the north of Australia. Without azimuthal checking, however, the bearing is ambiguous and a large mislocation is possible. Apparently, there is very little activity elsewhere on the network as few rings from non-Australian stations are visible.

The sequence with the checking is shown in the bottom half of Figure 9. The improvement is dramatic; the preferred location near Timor is clear, though again we note the ambiguity in range due to the loose slowness check.

Test Results

Performance

In order to give the reader some sense of the operational practicality of WCEDS, we present performance metrics. It should be noted, however, that WCEDS is a prototype system and has not been tuned to optimize performance. We feel confident that should an operational system be built, performance could be considerably improved. There are many ways to accomplish this, and we point some of these out in our discussion. With these caveats in mind then, we present the metrics. Note that all processing was done on a single Sun Ultra Sparc 1 workstation with 128 Mb of RAM.

Pre-processing the data took about 12 hours per hour of data for the entire network, which comes to about 12 days to process the complete interval. Clearly this is an unacceptable figure for

an operational automatic monitoring system which is supposed to keep up with continuous data flow, but it is misleading. In operational automatic monitoring systems where huge amounts of data must be continuously processed such as the PIDC, the same problem would occur if all of the processing were done on a single machine. Instead, the processing is distributed across a pool of dedicated workstations (or alternately, across processors if multi-processor machines are used), and the bottleneck is removed. For our test we used only one processor for the entire 35 stations, each of which had at least three channels, hence the unimpressive performance. Based on our results, to keep up with the data flow would require 12 processors instead of 1, but this number could probably be reduced significantly by optimizing our code.

Once the processed streams were available, running the 24-hours of data through the correlator and locator took about 4 days, or a factor of 4 times the length of time processed. This figure is much better than the waveform processing, but would still not be acceptable for an operational system. Again, though, we note that performance could be easily improved by parallelizing the detection and location processes, in this case by using different machines to work with different parts of the grid (this is a distinct advantage of grid-based approaches).

Examining the correlation/location performance in greater detail, we find that a total of 137 events were built, which works out to an average of 42 minutes per event. As would be expected, the amount of time needed to process each 1 hr 40 min chunk of data was not equal; it depended on the number of events built. The shortest time interval was 2:07 (4 events) and the longest was 6:46 (16 events -- the maximum number allowable). As we discuss below, the number of events which would pass simple screening criteria was much lower, only 35% of the total. If a means could be found to prevent the other events from being processed, then the total processing time would be 31 hours. While this is not an achievable goal -- any detection system will produce spurious events -- it is important to note that the less time an automatic system wastes in producing unwanted answers, the more efficient it becomes. In other words, tuning not only improves the quality of the bulletin but also improves the efficiency of the system.

Bulletin

To get rid of the lowest quality events, we imposed a set of simple screening criteria on the WCEDS output: we pass only those events with $ndef \geq 3$ and correlation value > 300 . This gets rid of 89 events, leaving us with a total of 48. The remaining events are shown in the bottom half of Figure 3. By carefully examining each of the events we determined that 39 of them (81%) correspond to REB events. The mislocations relative to REB range from 1 degree to 43 degrees. The worst are the Tonga-Fiji events due to the poor network coverage in that area and the loose slowness tolerance we used. The mislocations for these events are almost entirely along a great circle path connecting the true location with the Australian arrays. This suggests that if regionalized directional information were available for grid points in the south Pacific, much tighter slowness constraints could be applied and the mislocations could be significantly reduced.

We should also point out that our locations are from a 2 degree grid, while the REB locations come from a more traditional fitting of the travel times from picked arrivals (see Bratt and Bache, 1988 for details on the location algorithm), which is not limited to a grid of locations. Even for the AEL, GA outputs hypothetical locations from its grid (currently set at about 3 degree spacing),

but these are then automatically refined. Consequently, on average our locations cannot possibly be expected to be as precise as those from the PIDC, regardless of the quality of the data.

Comparing the events in Figure 3, it is apparent that for the REB events that were widely observed (i.e. have large *ndef*'s), the WCEDS events match well. As the number of observed phases decreases, so does the correspondence between the events. WCEDS did miss a few of the large *ndef* events, but in almost every case, examination of the processed waveforms revealed that the signals were marginal, suggesting that these events had been aided considerably by the human analysts involved in producing the REB. The one notable exception was a large mid-Atlantic Ridge event (*ndef*=30) with good signals which was missed by WCEDS. However, we found that this event was part of a series and had occurred nearly co-located to and just 40 seconds prior to another, even larger event (*ndef*=40), which WCEDS did build. Apparently the separation between the phases of the events was too small to be distinct for our method due to the broadening of the phases in the processed waveforms and in the MI.

Further comparison reveals that 40 events in the REB (51% of the total) were not seen by WCEDS. This is a large number that warrants investigation. In examining all of the REB events, we found that 48 of them (61%) have no reported phases other than the first arriving P-type (i.e. Pn, P, or PKP) and that 34 (43%) have 5 or less defining detections. Note that these are the results *after* analyst review wherein considerable effort was made to find all available arrivals. Given these figures, it is perhaps understandable that WCEDS missed such a large number of the REB events. WCEDS relies on the information that is in the processed waveforms alone to detect and locate events. The system works best with large amounts of data, even if some of the data is of lower quality; it works much more poorly with sparse data sets, regardless of how good the data might be. Unfortunately, the latter seems to be much more typical of primary network events.

If we do not use any screening criteria, a total of 137 events were built by WCEDS (Figure 10). This is not an unreasonable number for a global system, but a quick glance at the locations suggests that there are significant problems with many of the events. Note the large number of events in aseismic regions. Not surprisingly, investigation of these events established that most (if not all) were built incorrectly. Typically, these events consist of a single good local or regional arrival at one station combined with correlations of noisy waveforms at other stations to produce the spurious location. Recall that correlating with a single station provides a ring of possible source locations. Even the proper radius of the ring is not clear unless more than one phase has correlated (e.g. P and S); it can trade off with origin time. Thus, without any other true correlations from other stations, any random noise correlation can lead to a slightly higher correlation value at a certain grid point and determine the location. The directional consistency checks considerably reduced the number of these events, but clearly many are still present.

There are also unusually large numbers of events located in the vicinity of Australia and Alaska. Examination of these events reveals the same problem: each typically consists of a single good local arrival at either ILAR for the Alaskan events, or WRA or ASAR for the Australian events, combined with noise from some other stations. These are legitimate events but they cannot be well located by WCEDS without observations from other stations. This is not a problem limited to WCEDS: most event location systems locate single-station events poorly, so generally such events are screened out so that they will not appear in the bulletins produced. For example, the GA

processing at the PIDC is currently prohibited from producing single-station events. Thus, our screening criteria are not unreasonable.

Discussion

Comparing the output of a prototype automatic system with the REB can be useful, but only if the proper perspective is taken. First, it is important to realize that it is an unrealistic goal for any crudely-tuned prototype to match the overall performance of a soundly designed, mature, extensively-tuned system. Second, one should also recognize that even if the prototype system were tuned nearly perfectly, it would still be unlikely to match the REB because that bulletin has been reviewed by analysts and corrected for any mistakes made by the automatic processing. One might well argue that a fairer comparison would be with the un-reviewed AEL (automatic event list) produced directly by the PIDC's automatic event detection system, but we chose not to do this because it would force us to compare one subjective bulletin with another, leading to confusion about which is correct. Instead, we chose to compare to the closest thing we could find to "ground truth" for our test set, i.e. the REB. The ultimate goal of any automatic system should be to approach the quality of the reviewed bulletin, hence comparison to the REB is a useful process.

In terms of general reliability and system behavior, we found the performance of WCEDS to be good for a prototype system. Very few problems were encountered with the system behaving in a manner that was inconsistent with the expected behavior. If a suitable number (generally at least 4 or 5) of the processed waveforms for an event showed good correlatable features matching the phases in the detection MI, the event was always detected and was generally well-located, depending on the geographical distribution of the correlating stations and the strengths of the signals. Prior to the use of the directional consistency checking, notable exceptions were events in the south Pacific, where the poor primary network station coverage handicapped the system. The use of directional consistency checking considerably improved the detection and location of these events, though, as we have seen, their locations were often still poorly determined due to our loose slowness constraint.

Our experience with the primary network suggests that processing without using directional information will not work due to the sparse, uneven coverage. Our initial attempts at processing the test set without using directional information led to large numbers of false events and to huge mislocations of true events. For smaller events, the coverage of the primary network in many areas is simply not good enough to give WCEDS the resolution needed to locate them without using directional information. Fortunately, employing directional information can dramatically improve the situation, and directional information is available for the entire network (either from array or three-components). Using directional information is not, however, without pitfalls. We chose to use simple azimuth and slowness consistency checks (observed vs. theoretical) to control the correlations of observed data with azimuthally independent MI's. This method proved effective in many cases but also led to situations where well-observed events were either not detected or not located properly due to larger than expected discrepancies between the observed and theoretical values, which forced to adopt large tolerances. In general, we found arrays to be more reliable than three-component sites, and azimuth to be better behaved than slowness, but exceptions occurred for each case. Allowing station-dependent azimuth and slowness tolerances makes the

checks more robust, but more precise source location-dependent values are probably ultimately needed. As the PIDC continues to add to their archive of events, such corrections may soon become available (e.g. Koch and Kradolfer, 1997).

Even using directional information, however, if we examine the quality of the WCEDS on an event-by-event basis, the results are somewhat disappointing. WCEDS detected and located virtually all of the REB events with a large number of defining stations ($n_{\text{def}} > 12$), but performed very unevenly on the low n_{def} events. In a few cases our system built events which were not included in the REB, but in many more it missed REB events entirely. We examined both types of events in great detail to understand why this was the case and found that the answer always lies in the processed waveforms themselves. None of the REB events that were missed had readily apparent arrivals in our processed waveforms; if the events had appreciable arrivals corresponding to the phases predicted in the detection MI, they were built. Thus, we found that the problem is not in the WCEDS design but rather in the data itself: many of the smaller events simply do not have enough information in their processed primary network waveforms to allow WCEDS to detect and/or locate them properly. It is possible that better waveform processing algorithms could change this conclusion, but the minimal number of later phases recorded in the analyst-aided REB suggests that this is unlikely. Unless the amount of phase information available for primary network data can be improved, WCEDS can be expected to have little advantage over systems that work exclusively on first arrivals. This is not to say that WCEDS cannot be made to yield comparable results with a system like GA. We are quite confident that similar performance could be achieved given a commensurate level of effort towards tuning and refinement. Our goal, however, has been to assess the potential to *improve* the current monitoring capability, and we see no reason to believe that WCEDS can do this for global monitoring of the primary network.

One of the original appeals of WCEDS was its relative simplicity compared to other automatic systems, but this does not come without a price: because WCEDS contains very little of the complicated, Bayesian-style logic that is embedded in the other network event detection codes (e.g. Le Bras et al., 1994), it is all the more dependent on the processed waveforms themselves to uniquely characterize an event. Ironically, we have found that for the primary network, such phases are rare except for very large events -- the type of events that nearly all methods can detect and locate well. The smaller events in the REB typically have no later phases, giving WCEDS little if any advantage over traditional methods. This is because most of these events occur beyond regional distance from the nearest primary network station.

It is possible that for another type of global network with better sensitivity to coda phases, WCEDS could do significantly better. Using more long-period data could improve sensitivity, but probably only for larger events, which are generally not the events of interest for CTBT monitoring. Smaller events typically do not generate significant amounts of low frequency energy. Other means to improve later phase contributions would be to use three-component arrays to take advantage of energy on horizontal components, or to try to increase the number of stations in the network to get more regional waveforms, which often have obvious coda arrivals.

The benefits of this latter approach have been demonstrated in another study which we have conducted using data from a local network in central New Mexico (Withers et al., 1998). In that study, we found that WCEDS could deliver excellent automatic locations for events within the

network (within 3 km) and detection capability approaching that of highly trained analysts. This despite the fact that the version of the system used for the local network is much less sophisticated. The improvement is clearly related to the quality of the regional data processed. Unlike the teleseismic waveforms from the primary network, the regional waveforms for the New Mexico network often have prominent secondary phases, and these give WCEDS the improved sensitivity for which it was designed. The implications of this study for the CTBT monitoring problem are not clear, however. WCEDS may prove valuable, but perhaps only if the monitoring agencies become interested in working with data from much denser networks than the primary network.

For the REB events which WCEDS did build, the mislocations ranged from less than one degree to more than 40 degrees (assuming, of course, that the REB locations are accurate). Obviously the grid spacing (2 degrees) puts an inherent limit on the location accuracy for any grid search system, but WCEDS has some additional factors which are less easily quantified. The system is based on a dot product of processed waveforms with predicted waveforms, so the preferred location will be that which yields the optimal product, i.e. the optimal alignment of the waveforms. In the case of our traveltimes-based MI's, the predicted waveforms are weighted boxcars, so the dot products are essentially weighted integrations of the processed observed waveforms over the phase intervals defined by the boxcars. Thus, to find the preferred location, WCEDS will search for the grid point and origin time that will yield the maximum integrals of the processed waveforms. There is no reason to believe that this will be the same grid point and origin time point that would best fit discrete first arrivals. In fact, only in the case of high-frequency, low-noise data should the two locations be expected to be very similar, which is just what we see when the system is used to locate local/regional events (Withers et al., 1998). For teleseismically detected events such as those in the REB, one can expect essentially random mislocations due to variations in the processed waveforms caused by source mechanism and attenuation.

How to compensate for such effects is a difficult problem with no obvious solution. We have experimented with using different shapes to better match the MI to the observed processed waveforms, but found no general shape which could be shown to work better than a boxcar; specialized shapes will match some events well but do very poorly with others. The problem is that the shape of the processed waveform is dependent on many dynamic factors (e.g. event rupture mechanism, station characteristics), so no single, specific static shape works particularly well overall. The preferred method to improve WCEDS locations may be to re-process the events, with a more precise technique. Staying completely within the WCEDS methodology, one could set up a much finer grid covering the approximate event location, collect regional data for the event, and re-run WCEDS, as is done by Withers et al. (1998), who use a coarse-grid (10 km spacing) for the entire state of New Mexico and a finer grid (1 km spacing) for central New Mexico. Alternatively, one could re-process the teleseismic data for the events found by WCEDS to precisely pick the arrival times and use a non-grid-based, traditional location method ala the PIDC to obtain a much more accurate solution.

If WCEDS locations are to be used directly, then some sort of confidence estimate is needed. Because of the fuzziness in the locations due to the use of processed waveforms, this is not a straightforward process. One can contour the global output of WCEDS at the time of detection (Figures 6-9) and gain considerable insight into the uniqueness of the event location, but this is not the same as a true confidence region, which would rigorously indicate the effects of uncer-

tainty in the model and data upon the location (e.g. Flinn, 1965; Buland, 1976). An estimate of the model error could be obtained from the derivation of the traveltime curves used for the MI (e.g. the fit of the IASPEI curves to the data used to generate them), but for teleseismic locations this contribution is likely to be far outweighed by the error due to the processing of the waveform data, which would be much more difficult to estimate. Instead of generating error estimates for discrete arrivals, we would need to generate estimates for entire waveforms and develop the theory for how this information would be propagated through to the location. This would seem to be a difficult problem which we have chosen not to address until a clear need is demonstrated.

While the gridded output is not a confidence region, as we have seen it can still provide considerable insight into the quality of the solution. A simple shape such as an ellipse could be fit to the information on the grid, but this would actually decrease the information content, in some cases potentially leading to incorrect interpretations. We feel that the raw output itself is much more useful. Further, by iteratively re-processing the events and leaving out station and/or phases, the gridded output could be used to assess the significance of each separate piece of information (station-phase pair) to the solution resolution.

Conclusions

The Sandia Labs WCEDS project has succeeded in building and testing a prototype seismic event detection system based on a type of full-network waveform correlation (Shearer, 1994), which has allowed us to assess the potential usefulness of such a system for CTBT monitoring. As a complex data processing system, WCEDS performs well, functioning essentially as designed. Processing speeds are below acceptable levels for continuous data flow, but this problem can be easily remedied by spreading the necessary processing across several workstations, as is currently done for CTBT monitoring at the PIDC. Despite the general robustness of the software however, utility for CTBT monitoring is not clear. We tested WCEDS on a 24-hour interval of primary network data for which we treated the PIDC REB as "ground truth". WCEDS does well with large events but performs unevenly for smaller events, missing many of those present in the REB. The underlying reason for the disappointing performance is easily identified as the lack of later phase information available in the waveforms recorded by the primary network, which is verified by the small number of later phases listed in the REB. The lack of later phase information can be attributed to the sparse coverage of the primary network and its low sensitivity to phases other than first arrivals, both of which make it poorly suited to the WCEDS technique.

Monitoring this type of network implicitly requires event detection and location algorithms to make as much as possible out of a few, well-determined observations (e.g. time, azimuth, and slowness for the P arrivals at two stations). This is at odds with the philosophy of WCEDS wherein events are built by pattern-matching simplified, processed, waveform profiles. There are numerous ways in which the WCEDS event bulletin could be improved for this network (e.g. different waveform processing algorithms, finer grid, use of MI's for several depths, use of regionally varying MI's, etc.), but we do not feel that any of these are likely to appreciably change our assessment of the potential utility of the technique as applied to globally monitoring the primary network. Our detailed comparison of the WCEDS bulletin and the REB suggests that while extensive tuning of WCEDS could allow it to eventually match the performance of the automatic event

detection system in use at the PIDC, it is unlikely to improve overall global monitoring of the primary network. Hence given that the PIDC system is already tuned and operational, we see no reason to pursue further development of the WCEDS system for global monitoring of the primary network.

The usefulness of WCEDS on a regional or local scale, however, seems to be a different matter. In a related paper, Withers et al. (1998) report considerable success using the technique to monitor data from a local network. Unlike teleseismic waveforms from small events, regional waveforms often have prominent secondary phases, and these seem to give WCEDS the improved sensitivity for which it was designed. The implications of this study for the CTBT monitoring problem need to be more clearly understood, but if proper application scenarios can be found, WCEDS may yet prove useful for CTBT monitoring.

References

- Beiriger, J. I., S. G. Moore, J. R. Trujillo, C. J. Young (1997), The Waveform Correlation Event Detection System Global Prototype Software Design, *SAND Report #SAND97-3179*.
- Buland, R. (1976). The mechanics of locating earthquakes, *Bull. Seismol. Soc. Amer.*, **66**, 173-187.
- Flinn, E. A. (1965). Confidence regions and error determinations for seismic event location, *Rev. of Geophys.*, **3**, 157-185.
- Harris, D. B. (1991). A waveform correlation method for identifying quarry explosions, *Bull. Seismol. Soc. Amer.*, **81**, 2395-2418.
- Jordan, T. H. and K. A. Svedrup (1981). Teleseismic location techniques and their application to earthquake clusters in the South-Central Pacific, *Bull. Seismol. Soc. Amer.*, **71**, 1105-1130.
- Kennett, B. L. N. and E. R. Engdahl (1991). Traveltimes for global earthquake location and phase identification, *Bull. Seismol. Soc. Amer.*, **105**, 429-465.
- Koch, K. and U. Kradolfer (1997). Investigation of azimuth residuals observed at stations of the GSETT-3 Alpha Network, *Bull. Seismol. Soc. Amer.*, **87**, 1576-1597.
- Le Bras, R., H. Swanger, T. Sereno, G. Beall, R. Jenkins, W. Nagy, and A. Henson (1994). Global association final report, *SAIC document #SAIC-94/1155*, 28 pp.
- Shearer, P. M. (1994). Global seismic event detection using a matched filter on long-period seismograms, *J. Geophys. Res.*, **99**, 13713-13725.
- Shearer, P. M. (1991). Imaging global body wave phases by stacking long-period seismograms, *J. Geophys. Res.*, **96**, 20353-20364.
- Wagner, G. S. and T. J. Owens (1996). Signal detection using multi-channel seismic data, *Bull. Seismol. Soc. Amer.*, **86**, 221-231.
- Withers, M. M., R. C. Aster, C. J. Young (1998). An Automated Local/Regional Seismic Event Detection and Location System Using Waveform Correlation, submitted to *Bull. Seismol. Soc. Amer.*.
- Yilmaz, O. (1987). Seismic data processing, *Soc. of Expl. Geophys.*, Tulsa, 526 pp.
- Young, C., M. Harris, J. Beiriger, S. Moore, J. Trujillo, M. Withers, and R. Aster (1996). The waveform correlation event detection system project, phase I: issues in prototype development and testing, *SAND Report #SAND96-1916*.

Figure 1. WCEDS operational model

Illustrating at a high-level the operational model for the WCEDS system. The text to the right highlights the important features of each step. New features are indicated by *italics*.

Figure 2. The IDC primary network

The IDC Primary network as of December, 1996. Diamonds represent array sites; triangles represent three-component sites.

Figure 3. REB events vs WCEDS events for 24 hour test period
 (a) The events in the REB for our 24 hour test period. The shading of each circle (i.e. event) is proportional to the number of defining detections (ndef) used in the location. The primary network stations are indicated with diamonds and triangles as in Figure 1. (b) Same, but for screened WCEDS events (corr ≥ 300 , ndef ≥ 3).

Figure 4. WCEDS detection and screening Master Images

(a) Detection MI and (b) Screening MI used for primary network test. The detection MI uses only P and S because for most events recorded by the primary network, these are the only phases observed.

Figure 5. An example processed waveform profile for a large event

An example processed waveform profile for a magnitude 4.5 event near the Jordan-Syria border. The processing consists of 1.5-2.5 Hz bandpass filtering, spatial coherency processing, and STA/LTA processing.

no
azimuthal
constraint

azimuthally
constrained

Figure 6. Jordan-Syria event without and with azimuthal checking

The shading indicates the correlation value at each grid point (red is highest, blue is lowest) at times of -4 min, -2 min, 0 min (i.e. origin time), and +2 min relative to the origin time from the REB. The contributing stations (all P only) are ARCES, BGCA, ESDC, FINES, GERES, ILAR, KBZ, NORES, SCHQ, ULM, YKA, and ZAL. For each panel, the white dot indicates the REB location.

no
azimuthal
constraint

azimuthally
constrained

Figure 7. Nicaragua event without and with azimuthal checking
As in Figure 6. The contributing stations are ASAR, PDAR, SCHQ, TXAR, ULM, WRA, and YKA.

no
azimuthal
constraint

azimuthally
constrained

Figure 8. Kermadec event without and with azimuthal checking
As in Figure 6. The contributing stations are ASAR, BGCA, LPAZ, and WRA.

Figure 9. Timor event without and with azimuthal checking
 As in Figure 6. The contributing stations are ASAR (P & S), and WRA. In this case, the white dot shows the WCEDS location.

Figure 10. All WCEDS events for 24 hour test period

All of the events detected and located by WCEDS for the test period. No screening has been applied.

Appendix A

Spatial Coherency Calculations

Assume that we want to calculate a spatial coherency stream for N_{ch} data channels, $x_i(t)$. The channels can correspond to elements of an array, different directional components, or both. Let T_{total} be the total length of time spanned by each of the channels, T_{window} be the time window over which the calculation is to be performed, and p_{ol} be the percent overlap between successive windows.

First, we will calculate the necessary parameters for the windowing. The total number of whole windows to process will be

$$n_{win} = \text{int} \left(\frac{(T_{total} - T_{window})}{\left(1 - \frac{p_{ol}}{100}\right)} \right) + 1 \quad (\text{EQ A.3})$$

where int rounds down to the nearest integer. The start and end times for each window will be

$$t_{start} = K \left(1 - \frac{p_{ol}}{100} \right) T_{window} \quad (\text{EQ A.4})$$

where $K = 0 \dots n_{win}$, and

$$t_{end} = t_{start} + T_{window} \quad (\text{EQ A.5})$$

The sampling rate of the spatial coherency output will be

$$\text{samprate} = \frac{1}{\left(1 - \frac{p_{ol}}{100}\right) T_{window}} \quad (\text{EQ A.6})$$

Now we show the spatial coherency calculation. For the K th window, we take the FFT's of each of the data channels to transform to the frequency domain:

$$x_i^K(t) \rightarrow x_i^K(\omega) \quad (\text{EQ A.7})$$

and form the covariance matrix of the transformed channels:

$$\text{cov}_K = \begin{bmatrix} \tilde{x}_1^K(\omega) \cdot \tilde{x}_1^K(\omega) & \dots & \tilde{x}_1^K(\omega) \cdot \tilde{x}_{N_{ch}}^K(\omega) \\ \dots & \dots & \dots \\ \tilde{x}_{N_{ch}}^K(\omega) \cdot \tilde{x}_1^K(\omega) & \dots & \tilde{x}_{N_{ch}}^K(\omega) \cdot \tilde{x}_{N_{ch}}^K(\omega) \end{bmatrix} \quad (\text{EQ A.8})$$

where $\tilde{x}_i^K(\omega) = x_i^K(\omega) - \bar{x}_i^K$. The spatial coherency output for the K th window is then the principal eigenvalue of the K th covariance matrix:

$$spacoh(K) = \lambda_{max}(cov_K) \quad (EQ A.9)$$

Appendix B

Three-component Power Calculation

We want to calculate azimuth and slowness streams from the direction of peak polarization for a three-component station with channels $x_i(t)$ where $i=1,2,3$. The data time windowing and new sample rate are as described in Appendix A, Equations A.1-4.

The method works by setting up a set of directional vectors to span slowness space and then projecting the information from each of the three data components onto these vectors to find the direction with the greatest polarized power. Let N_T be the number of points in each window. Then the total power for the Kth window is

$$P^K = \frac{\sqrt{\sum_{i=1}^3 \sum_{j=1}^{N_T} (x_{ij}^K)^2}}{N_T} \quad (\text{EQ B.10})$$

This value will be used to normalize the polarized powers in each direction.

We begin by defining the directional vectors associated with the three data channels. We do not need to assume that the components are in a particular order (e.g. Z, N, E), nor even that the components are perpendicular. Let the horizontal angle of rotation counterclockwise from due east for each component i be θ_i , and the angle of rotation upward from vertically down be ϕ_i . If we adopt as our coordinate system 1=east, 2=north, 3=vertically down, then the directional vectors are

$$\hat{u}_i = \begin{bmatrix} \sin(\phi_i) \cos(\theta_i) \\ \sin(\phi_i) \sin(\theta_i) \\ \cos(\phi_i) \end{bmatrix} \quad (\text{EQ B.1})$$

The direction space over which we will calculate polarization power is specified as a grid of slowness values (s_{east}, s_{north}). Let the range of possible values for each be given by

$$s_j = s_{min} + (j-1) \cdot delslow \quad (\text{EQ B.2})$$

where $j = 1 \dots numslow$, $delslow$ is the slowness increment, and $numslow$ is the number of slowness values. The total number of polarization space vectors will then be $N_S = numslow^2$. Note that typically s_{min} is negative and $numslow$ is chosen so that $s_{max} = abs(s_{min})$ so that the grid of slowness values is symmetric about zero.

For the projection, we must convert the slowness coordinates to the same angular coordinates used in Equation B.1. To do this, we must assume a local velocity, vel_{loc} whose units we will assume to be consistent with the slowness units. Then for each direction, k

$$\sin(\varphi_k) = vel_{loc} \cdot \sqrt{s_{east_k}^2 + s_{north_k}^2} \quad (\text{EQ B.3})$$

and

$$\tan(\theta_k) = \frac{s_{north_k}}{s_{east_k}} \quad (\text{EQ B.4})$$

Solving each of these equations for the angles, the vector for each direction can be written

$$\hat{v}_k = \begin{bmatrix} \sin(\varphi_k) \cos(\theta_k) \\ \sin(\varphi_k) \sin(\theta_k) \\ \cos(\varphi_k) \end{bmatrix} \quad (\text{EQ B.5})$$

We can now calculate the projection coefficients, c_{lm} , of data stream direction vectors l onto polarization space vectors m :

$$\begin{bmatrix} c_{11} & \dots & c_{1N_s} \\ c_{21} & \dots & c_{2N_s} \\ c_{31} & \dots & c_{3N_s} \end{bmatrix} = \begin{bmatrix} u_{11} & u_{12} & u_{13} \\ u_{21} & u_{22} & u_{23} \\ u_{31} & u_{32} & u_{33} \end{bmatrix} \cdot \begin{bmatrix} v_{11} & \dots & v_{N_s 1} \\ v_{12} & \dots & v_{N_s 2} \\ v_{13} & \dots & v_{N_s 3} \end{bmatrix} \quad (\text{EQ B.6})$$

Using these coefficients we can calculate the projected data streams $w_k^K(t)$ for the K th window:

$$\begin{bmatrix} w_{11}^K & \dots & w_{1N_T}^K \\ \dots & \dots & \dots \\ w_{N_s 1}^K & \dots & w_{N_s N_T}^K \end{bmatrix} = \begin{bmatrix} c_{11} & c_{21} & c_{31} \\ \dots & \dots & \dots \\ c_{1N_s} & c_{2N_s} & c_{3N_s} \end{bmatrix} \cdot \begin{bmatrix} x_{11}^K & \dots & x_{1N_T}^K \\ x_{21}^K & \dots & x_{2N_T}^K \\ x_{31}^K & \dots & x_{3N_T}^K \end{bmatrix} \quad (\text{EQ B.7})$$

Where x_{ij}^K refers to the j th time point of the i th component stream in the K th window. Now we calculate the power in each projection stream k and normalize by the total power in the K th window:

$$P_k^K = \frac{\left(\sqrt{\sum_{j=1}^{N_T} (w_{kj}^K)^2} \right) / N_T}{P^K} \quad (\text{EQ B.8})$$

Or, if we want to emphasize the peak value (e.g. for an FK-like display of the slowness grid), we can instead use

$$P_k^K = \frac{\left(\sqrt{\sum_{j=1}^{N_T} (w_{kj}^K)^2} \right) / N_T}{P^K - \left(\sqrt{\sum_{j=1}^{N_T} (w_{kj}^K)^2} \right) / N_T} \quad (\text{EQ B.9})$$

In either case, the azimuth and slowness (or equivalently, the east and north slownesses) for the K th window are taken from the projection space stream with the maximum normalized power.