The Waveform Correlation Event Detection System Project: Issues in System Refinement, Tuning, and Operation

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

Sandia National Laboratories

Mitchell M. Withers, Richard C. Aster

New Mexico Institute of Mining and Technology

Sponsored by U.S. Department of Energy

Comprehensive Test Ban Treaty Research and Development Program, ST485D

RECEIVED
JUL 25 1996
OSTI

ABSTRACT

The goal of the Waveform Correlation Event Detection System (WCEDS) Project at Sandia Labs has been to develop a prototype of a full-waveform correlation based seismic event detection system which could be used to assess potential usefulness for CTBT monitoring. The current seismic event detection system in use at the IDC is very sophisticated and provides good results but there is still significant room for improvement, particularly in reducing the number of false events (currently being nearly equal to the number of real events). Our first prototype was developed last year and since then we have used it for extensive testing from which we have gained considerable insight. The original prototype was based on a long-period detector designed by Shearer (1994), but it has been heavily modified to address problems encountered in application to a data set from the Incorporated Research Institutes for Seismology (IRIS) broadband global network. Important modifications include capabilities for event masking and iterative event detection, continuous near-real time execution, improved Master Image creation, and individualized station pre-processing. All have been shown to improve bulletin quality. In some cases the system has detected marginal events which may not be detectable by traditional detection systems, but definitive conclusions cannot be made without direct comparisons. For this reason future work will focus on using the system to process GSETT3 data for comparison with current event detection systems at the IDC.

Keywords: event detection, waveform correlation, event masking, master image

DISCLAIMER

This work was supported by the United States Department of Energy under Contract DE-AC04-94AL83000.

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, 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 or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

DISCLAIMER

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

OBJECTIVES

The Waveform Correlation Event Detection System (WCEDS) Project at Sandia was initiated to assess the viability of developing an improved seismic event detection system for CTBT monitoring based on the correlation of full waveforms. This type of project is warranted because although the current system (Global Association or GA) is highly sophisticated and is a vital part of the current monitoring pipeline, the quality of the bulletin produced shows substantial room for improvement. The number of *missed* events (i.e. those which must be associated by an analyst) has continued to decrease and is now approaching 10% of the final total of real events, but the number of *false* events (i.e. those that are not real and which must be disassociated by an analyst) is still large, being nearly equal to the final total of real events (pers. comm., R. Le Bras, June, 1996). This leads to a significant loss of efficiency in the system because analysts must screen all of the events to find those that are potentially important. Many approaches to improving event detection could be investigated, of course, but to create the most sensitive detector possible a full waveform approach is especially attractive.

INTRODUCTION

The WCEDS Project was initiated in the Spring of 1995 with the goal of developing a prototype which could be used to assess viability for CTBT use. By May of 1995 a rudimentary prototype was available, which was then refined based on the results of repeated testing with a set of continuous data from the Incorporated Research Institutions for Seismology (IRIS) broadband global network. Details of the initial prototype development and testing were given in a presentation at the 17th Phillips Lab Symposium (Young et al., 1995).

In the ensuing year, we have made several important refinements to the prototype and learned a great deal about the strengths and weaknesses of a waveform correlation-based system. In this paper we will summarize some of the most important new features of the present version of the detector, discuss ongoing and planned work, and provide an assessment of the potential value of the system based on what we have learned so far.

REVIEW

The detector is a grid-based automatic system which examines continuous seismic data streams for matches with some portion of a master set of waveforms which we refer to as the *Master Image* (MI). The complete grid covers the entire surface of the Earth. To understand how the detector works, think of it as creating a waveform profile for each grid point at a given time and comparing this observed profile with the expected profile if an event had indeed happened at that point and time. Prior to forming the profiles, the waveforms are processed with filtering and an STA/LTA algorithm to enhance phases relative to noise. We refer to this as *pre-processing* because it occurs prior to execution of the detector itself. The detector then runs continuously at some specified time discretization interval.

To improve the efficiency of the algorithm, at each potential origin time all possible distance correlations for each processed waveform are pre-calculated and stored in an array which we call the C matrix:

$$C_{ij} = \sum_{k=1}^{N_i} M_{ik} D_{kj} \quad (\text{EQ 1})$$

Where D is the data matrix whose columns are the observed waveforms, M is the MI matrix whose rows are the expected waveforms, and N_t is the number of time points. The detector output at each grid point is then determined by summing elements in C :

$$O = \sum_{j=1}^{N_s} C_{ij} \quad (\text{EQ 2})$$

Another way to think of this is that to calculate the detector output for each grid point, the detector follows a particular summation path through the columns of C .

NEW RESEARCH ACCOMPLISHED

Event Masking and Iterative Event Detection

As might be expected, fundamentally changing the basis of event detection creates a new set of problems. One of the most basic is the problem of the sidelobes produced by the correlation. If the observed signal and the master signal each have multiple phases, then the cross correlation function will show sidelobes in addition to the true correlation peak. We refer to these sidelobes as *false correlations* because they are due to correlations of an observed phase with the wrong master phase. To understand how false correlations occur, consider a simple hypothetical system with three phases: A, B, and C. For a data set we will use a waveform from one station at which all three phases are observable and distinct.

FIGURE 1. True and false correlations, 3 phase system: fixed origin time

As Figure 1 shows, in addition to the true correlation at the correct origin time at distance D_3 , there are lesser but significant false correlations at D_2 and D_1 . For any given distance, there will also be false correlations both before and after the true origin time (Figure 2). This means that there will be certain *incorrect* grid and time points for which false correlations will occur, and if the output at any of these points exceeds a detection threshold a false event will be declared at that

point. Perhaps more importantly, it will be very difficult if not impossible, to detect any real events of smaller magnitude which may have occurred nearby in time.

FIGURE 2. True and false correlations, 3 phase system: changing origin time

To solve this problem it is necessary to remove all the effects of a given event, both true and false, once it has been detected. If this can be accomplished then events with intermingled phases could be detected through an iterative process of detection and masking. The most obvious solution would be to zero the waveforms where phases have been observed for a detected event. This does imply removing data, but it is important to realize that this is really no different than what goes on in trigger-based systems such as GA. Virtually no system will allow the phase observed at a given time to be associated with more than one event: a choice must be made (this is the topic of “conflict resolution”, e.g. Beall et al., 1995). Once a trigger has been associated with a given event it is effectively removed from further consideration which is equivalent to zeroing the waveform from which it came.

We chose another approach, however, because in our system the censored waveforms would have to be re-correlated with the MI and this is a computationally expensive process. Instead, we elected to mask the correlations in the C matrix that are associated with a detected event and then recompute the output (O; Equation 2) at the grid points. We wish to note however, that as we have continued to refine the code we have greatly increased the speed at which the C matrix can be computed and iterative calculation of C may now be viable.

In order to mask properly, it is necessary to determine which stations and which phases at those stations contributed to the event detection. Fortunately, this is straightforward in the WCEDS scheme (details are provided in Young et al., 1996). Once this information is available, all possible correlations between observed phases and MI phases are calculated and stored in masking matrices, referred to as the X matrices. There will be an X matrix for each origin time, just as there is a C matrix for each origin time. The grid is then searched again over the specified time interval, but in this case the correlation sums are calculated from the C matrices masked by the X matrices. If any other grid points are found to exceed the detection threshold then the largest output point will be found, an event will be declared, X information will be determined, and so on. Using this iterative technique, we have conducted several tests in which multiple events with intermingled phases were successfully detected within the same time interval.

Continuous Near Real Time Execution

While the above results are encouraging, bulletin completeness is not the only criteria which must be considered for a monitoring system: near real time response is also essential. For this reason event detection systems used for continuous monitoring divide data into small segments to insure that events are detected as soon as possible. The minimum length of the segment which can be processed is constrained by the requirements of the detection system. In our case, to run the detector for even one potential origin time, we must have data from that origin time through the time span of the MI or all of the potentially available phases may not contribute to the correlation. Presumably systems that need very quick response times would have to use shorter MIs.

Let us consider how the processing of segments could occur by considering an 8 hour interval of data with 4 events in it.

FIGURE 3. Continuous execution models

If we process the data as a single long interval (Figure 3a) then all of the events present should be built properly, but the time delay in declaring the events is too large, particularly for an event occurring early in the interval. An obvious solution is to divide the interval into smaller segments and process these as they become available; in this case we could choose four 2 hour segments (Figure 3b). The difficulty with such a system is caused by events which occur near the end of a segment, and this is a common occurrence for a global network when segment lengths are short.

Such an event may have phases which correlate in more than one segment and consequently the event may not be built properly. In this case two of the real events have correlations in multiple segments and as a result we declare two extra false events. Unless there is some communication between the processing of the segments, we cannot recognize these as false events.

To deal with this problem we chose to implement the concept of overlap or *look back*. The process is illustrated in Figure 3c. The MI has a correlation time length associated with it (the number of columns in the MI defines the time span), so it is apparent that we can trust any detected events whose origin times fall more than that time length before the end of the current segment. These events will have secondary peaks due to false correlations that will be examined as the detector runs through the current time segment, but the true correlation will also be examined and will be preferred because it must have a greater value than any of the false correlation points. Hence the true point will be found, and all of its correlations will be correctly masked and the detector will be able to find nearby smaller events. The difficulty comes in situations when it is not possible to guarantee that the detector will have had the opportunity to check the true correlation point and prefer it to one of the false correlation points. Any events detected in the segment following the trusted event segment fall in this category. As soon as it is possible to correlate with any phase corresponding to an event whose true origin time lies outside of the segment currently being processed, events built must be labelled as suspect. Note however, that these events still must be built and masked because whether or not they are real, the correlations related to them may corrupt the association process in the current time period.

The key to continuous processing in our system is to treat the trusted events and the suspect events differently when the detector moves on to the next time period. Except for this modification the processing is as discussed above: events are built and stripped away until no grid point in the segment exceeds the detection threshold. When the detector is ready to move on, the next segment to be processed must overlap the previous segment by the correlation time length of the MI, i.e. by the length of the suspect event interval. As we begin processing the next interval, we start with an X matrix created from the previously detected trusted events, but we discard the suspect events and the related masking. Now we begin processing as above, except that what was the suspect interval in the previous segment is now part of what will be the trusted interval in the current segment. Thus, any of the discarded events which were legitimate will be rebuilt while false events will be rejected in favor of the corresponding true events. In this manner the process continues one segment at a time, perhaps building false events but correcting them before they are added to the final bulletin.

Master Image Refinements

In the initial stages of the project, we chose to use travel-time based MIs rather than empirically derived images (stacks) a la Shearer because of the difficulty in controlling the number, width, and weighting of the phases in the latter. To form the travel-time based MIs we convolved the IASPEI travel time curves (Kennet and Engdahl, 1991) with a shaping function to give them finite correlation widths chosen to match the observed widths in the processed data. For the most part, MIs formed in this manner yielded satisfactory results, but in fact the MIs were not truly correct. In addition to compensating for processed phase width, the width of the phases used in the MI should also be wide enough to 1) compensate for limited origin time sampling, 2) compensate for limited origin depth sampling, and 3) compensate for inaccuracies in the MI distance against which the waveform for a given station is correlated (due to both the discrete sampling provided by the grid and to the discretization interval of the MI). Origin time and depth effects are straight-

forward. To correct for origin time discretization, one simply widens all intervals by half of the time discretization at which the correlations are to be calculated. The depth discretization is a more complicated effect but can be approximated by a similar uniform widening dependent on the depth range to be spanned by each MI.

The distance discretization effects require a more complex phase and distance dependent adjustment of the correlation widths.

FIGURE 4. Master image correlation widths and distance discretization

Unless an event happens to occur exactly at a grid point and stations happen to lie at integer multiples of the MI distance discretization from that point, observed waveforms will be correlated with MI waveforms for *incorrect* distances. As shown in Figure 4a, to compensate for these effects the MI correlation width should be widened. Note that the necessary discretization width is related to the slope (slowness) of the phase: the greater the slope of the phase, the wider the dis-

cretization width. For a uniformly spaced grid of spacing L the largest possible distance shift due to grid discretization, Y , is:

$$Y = L/(\sqrt{2}) \quad (\text{EQ 3})$$

As is shown in Figure 4b. On top of this one must apply the effect of the distance discretization in the MI; all grid point to station distances must be rounded to the nearest distance represented in the MI before correlation. To understand these effects, consider an example where a station is 72.6 degrees from an event and the grid spacing is 2 degrees ($Y \sim 1.4$). The possible range of distances to the nearest grid point are 71.2 degrees to 74.0 degrees. However, if the discretization interval of the MI is 1 degree, then the range of distances to the nearest grid point for purposes of correlation is 71 degrees to 74 degrees. Given a set of travel time curves, the grid spacing, and the MI distance discretization, one can generate the required correlation widths for stations at every distance. The process is illustrated in Figure 4c.

Station-specific Processing

As mentioned above, prior to correlation the data streams are pre-processed to enhance phases relative to noise. In our original version of the code, all data streams were pre-processed in the same manner. This led to performance problems for some time segments due to the varying quality of the data at each station. We have now implemented a pre-processing control file which allows individualized control over the processing and weighting for each station:

	stream	distance range	weight	filter parameters	STA/LTA parameters
##USA ---					
#Eastern U.S. & Caribbean					
stream=CCM	1	0.0 4.0	0.0	0.6 ~2.0 10.0 3 BP causal~	1 ~3.0 27.0 1 30.0~
stream=CCM	2	0.0 50.0	0.0	0.8 ~0.8 2.0 3 BP causal~	1 ~3.0 27.0 1 30.0~
stream=CCM	3	0.0 180.0	0.0	1.0 ~0.1 1.5 3 BP causal~	1 ~3.0 27.0 1 30.0~
stream=HRV	1	0.0 4.0	0.0	0.6 ~2.0 10.0 3 BP causal~	1 ~3.0 27.0 1 30.0~
stream=HRV	2	0.0 50.0	0.0	0.8 ~0.8 2.0 3 BP causal~	1 ~3.0 27.0 1 30.0~
stream=HRV	3	0.0 180.0	0.0	1.0 ~0.1 1.5 3 BP causal~	1 ~3.0 27.0 1 30.0~
# San Juan has some sort of bizarre hi freq. noise source, so omit					
# Local hi. freq. stream:					
stream=SJG	1	0.0 50.0	0.8	1.0 ~0.8 2.0 3 BP causal~	1 ~3.0 27.0 1 30.0~
stream=SJG	2	0.0 180.0	1.0	1.0 ~0.1 1.5 3 BP causal~	1 ~3.0 27.0 1 30.0~
#Western U.S.					
stream=ANMO	1	0.0 4.0	0.0	0.6 ~2.0 10.0 3 BP causal~	1 ~3.0 27.0 1 30.0~
stream=ANMO	2	0.0 50.0	0.0	0.8 ~0.8 2.0 3 BP causal~	1 ~3.0 27.0 1 30.0~
stream=ANMO	3	0.0 180.0	0.0	1.0 ~0.1 1.5 3 BP causal~	1 ~3.0 27.0 1 30.0~

FIGURE 5. Pre-processing control file

With the new structure, the user has full control over the number of processed streams which will be formed for the correlation for each station, the distance range over which they will be used for correlation with the MI, the weighting for each stream, the filter parameters, and the STA/LTA parameters. The user can use these controls to preferentially filter out known noise sources, down-

weight less-reliable stations, etc., which can greatly improve the performance of the system. Changes in frequency content with distance can be taken advantage of by using sets of band-pass filters and specifying the precise distances at which the filtered streams will be correlated (note that more than one stream can be correlated at a given distance).

As is shown in Figure 5, the new structure also provides the user with the opportunity to take a regional approach to the problem. Groups of stations from similar regions can be grouped together for parameter control and comparison. The groups of stations can then be down- or up-weighted to control their effect on the overall detector output.

CONCLUSIONS AND RECOMMENDATIONS

Much was accomplished in the first year of the WCEDS Project. We developed a sophisticated detector prototype which we used to run an extensive series of tests on waveform data from the IRIS broadband global network. The problems we encountered during these tests led us to add important new capabilities to the original code including event masking and iterative detection, continuous execution with lookback, creation of better Master Images, and individualized station pre-processing. Each of these capabilities improved the quality and completeness of the bulletin produced by the system.

The potential value of the system for CTBT monitoring, however, is still not clear. In some situations, the detector has successfully built small events which might not be detectable by trigger-based systems such as GA, but in others it has missed events which should be detectable by many existing systems. Perhaps the single greatest reason for the lack of a clear conclusion about the usefulness of WCEDS has been our inability to make direct comparisons with existing systems. To do this we must be able to process the data from GSETT3, and this has only recently become available to us. Once the necessary modifications have been made to process this data we should be able to make direct comparisons with GA on the same data set and this should make the value provided by a full waveform system should become much clearer. Our primary goal in the current year is to make this comparison.

Access to GSETT3 data should also provide us with the means to address another important idea: the possible use of waveform correlation for event screening. As mentioned above, one of the major problems facing the current system is that it generates a very large number of false events whose validity cannot be determined with the information available in the database. Yet these events can easily be determined to be false when the waveforms are viewed by analysts, suggesting that waveform correlation might provide a means to improve automated screening. We will pursue this idea once GSETT3 data is available to us by computing correlation products for all of the events (true and false) in the IDC bulletin and looking for an obvious difference. This is in fact a much easier problem than has been addressed in the WCEDS project because many of the difficult issues which we have had to solve (e.g. false correlations) are not relevant.

REFERENCES

- Beall, G., R. Le Bras, W. Nagy, T. Sereno and H. Swanger (1995). Global association system phase 2: conflict resolution, *SAIC document #SAIC-95/1029*, 46 pp.
- Kennet, B. L. N. and E. R. Engdahl (1991). Traveltimes for global earthquake location and phase identification, *Geophys. J. Int.*, 105, 429-465.

- Shearer, P. M. (1994). Global seismic event detection using a matched filter on long-period seismograms, *J. Geophys. Res.*, **99**, 13713-13725.
- Young, C., J. Beiriger, J. Trujillo, M. Withers, R. Aster, L. Astiz, and P. M. Shearer (1995). WCEDS: a waveform correlation event detection system, *17th Seismic Research Symposium on Monitoring a CTBT*.
- Young, C. J., M. Harris, J. Beiriger, S. Moore, J. Trujillo, M. Withers, R. Aster (1996). The waveform correlation event detection system project, phase I: issues in prototype development and testing, *SAND report*, in press.