

Workshop Report

Southern California Earthquake Center Workshop on Transient Anomalous Strain Detection: Blind Test Results and Future Planning

September 12-13, 2009

Palm Springs, CA

Jessica R. Murray-Moraleda and Rowena Lohman

Introduction

As an outgrowth of SCEC's central role in the development of the Southern California Integrated GPS Network (SCIGN), one of SCEC's science priority objectives is to "develop a geodetic network processing system that will detect anomalous strain transients". Given the growing number of continuously-recording geodetic networks consisting of hundreds of stations it is no longer feasible to inspect all station-component time series by eye to monitor for anomalous behavior. A means for systematically searching geodetic data for transient signals has applications for network operations, hazard monitoring, and event response, and may enable identifying events that would otherwise go (or have gone) unnoticed. We began the SCEC Transient Detection Test Exercise in 2008 to foster an active community of researchers working on this problem, explore promising methodology, and combine effective approaches in novel ways. This report summarizes a workshop held September 12 – 13, 2009 in Palm Springs, California to assess what we have learned to-date and discuss areas of focus as the project moves forward.

The test exercise is an opportunity for researchers to develop and test methods for geodetic transient detection. Through the exercise, test data are distributed to participants who apply their detection algorithms and report back on any transient signals they find. Phase I of the exercise ran from January 15 to March 15, 2009 and Phase II from June 16 to September 7, 2009. Thus far test data have been synthetic Global Positioning System (GPS) observations with realistic white and colored noise and seasonal signals (common-mode and otherwise). The workshop focused on the results of Phase II which used twelve datasets, each a collection of time series mirroring the spatiotemporal coverage of continuous GPS data in southern California since 2000. Four datasets had no transient tectonic signal added while the rest included signals obtained from models of temporally-varying slip on known faults in the region.

At the workshop, groups developing detection algorithms presented their approaches, the true signals in the test data were revealed, and Phase II results (submitted before the workshop) were summarized. The workshop concluded with a group discussion of directions to pursue for Phase III. The workshop agenda, slides from presentations, a summary of Phase II test data signals, and detailed information on source-time history and noise components of the Phases I and II test datasets may be found on the Google Groups website for the project (<http://groups.google.com/group/SCECtransient>).

Methodologies under development

Eleven groups presented their approaches to transient detection at the workshop. All methods under development rely in some way on using the spatial coherence of a transient deformation source to separate it from time varying signals due to other processes. These approaches

included time series analysis that assessed the significance and space-time coherence of rate changes in the data, principal component analysis with adaptations to address data gaps and propagating signals, assessment of the degree of correlation among time series for certain regions or time periods, application of Kalman filtering to estimate time-dependent strain and deformation signals, visualization of strain rate changes, and multi-resolution analysis of spatially-coherent signals in varying-length portions of the time series. Many groups used combinations of these approaches in developing their detection algorithms.

The test data

The test data consisted of time series that included contributions from secular motion (using the SOPAC velocity solution for southern California), time-dependent noise (flicker, random walk, and seasonal noise), common mode signals with distance-dependent amplitude, and deformation due to slip on known faults (based on the SCEC Community Fault Model). The slip signals could be spatially-variable, could have time histories that followed a variety of functional forms, and could propagate spatially over time. The twelve datasets were divided into two groups of six, labeled A and B.

Phase II Group A data included two datasets with relatively large deformation signals, two with more subtle transient deformation, and two with no deformation signal. The deformation sources for the large signal cases were 1) propagating slip on the San Gabriel fault with ~3.7 meters cumulative slip occurring over ~6 months and an equivalent Mw 6.5, and 2) two transients on adjacent parts of the Compton fault, the first a propagating slip event lasting 6 months and the second a transient in which the majority of slip occurred over ~10 days, with a combined equivalent Mw 6.0. The more subtle transients were 1) a propagating slip event on the Oakridge fault lasting ~1.25 years with cumulative slip of 7 cm and equivalent Mw 5.5, and 2) an increase in slip rate beginning in early 2002 and lasting for seven years on the Temecula fault with an equivalent Mw 5.85.

Because most participants focused their efforts on the Group A data, discussion did not address the Group B dataset in detail. Like Group A, four of the six Group B datasets included transient signals due to fault slip. The character of the signals was similar to those in Group A, although none were as large as the Group A San Gabriel event.

Phase II Results

The two largest Group A displacement signals were visually apparent in the time series. All participants successfully detected these events, although not all groups detected the propagation of the first transient on the Compton fault. None of the participating groups detected the two more subtle Group A transients. In addition to the lower magnitude of the signals, the relatively long duration of these events and the fact that, in the Oakridge case, only a small portion of the fault was slipping at any given time contributed to the difficulty in detecting these signals. It was argued during the workshop that in the Temecula case it was the first ~2 years of the dataset that were anomalous rather than the following 7 years, highlighting the difficulty in identifying transient signals that have a long temporal extent relative to the length of existing time series within the network. We also discussed the idea that an eventual systematic transient detector may need to focus on transients that fall within a particular range of magnitudes and durations in order to be robust.

Developing metrics for reporting results has not been straightforward since different approaches to transient detection focus on different aspects of transient behavior and return different information. In particular, although the transient source of interest in the test exercise is fault slip, not all methods incorporate a fault model, let alone estimate the space-time slip function. Therefore, we aimed to require the most basic information from all participants in order to facilitate comparison of findings among groups. Specifically, the spatial extent is to be described by the centroid location and an ellipse based on the spatial extent of the surface displacement signal attributed to transient deformation. The temporal signature is to be defined by the centroid time and duration. These metrics are detailed on the group website.

In Phase I these metrics were defined fairly qualitatively in the instructions to participants, and we found that we needed to specify the requirements in greater detail for Phase II. Even with the more detailed guidelines, however, the way in which different groups interpreted the metrics appeared to vary. For instance whether a group used a 95% confidence ellipse for the spatial extent (indicating that 95% of the transient signal occurred within the area defined by the ellipse) or some other confidence level, and whether they reported the semi-major and semi-minor axes of the ellipse or the major and minor axes may be responsible for some of the scatter seen in the reported location and spatial extent of the deformation signals for the two large-signal Group A datasets. Likewise, there was a noticeable amount of scatter in centroid time and duration estimated by the participating groups. Some of this may be due to different interpretations of centroid time. For instance, some groups may have defined this quantity as the midpoint of the time period during which transient deformation occurred while others may have defined it as the time at which half of the total transient deformation signal had accumulated. One way to address this problem would be to require results to be submitted through a web form linked to an application that would generate maps of the solution on the fly, allowing groups to identify whether their results were being interpreted correctly. We are investigating the feasibility of this technology further. We will also continue to refine the metrics in future phases of the test exercise.

Discussion

Following the presentation of test data and results, we discussed the most important areas on which to focus future test phases in order to achieve the objective of developing an automated detection system for anomalous strain transients.

We began with the question of what we have learned so far. It is evident that algorithms can retrospectively detect signals that are already visually apparent in the time series, but we seem to have made little progress on more subtle signals and have not yet assessed the real-time capabilities of the methodologies currently under development. Participants agreed that test data with large signals (e.g., the Group A San Gabriel and Compton fault transients described above) are useful for software validation but, given their high SNR, do not drive further development of algorithms to new levels of sophistication.

As one means for assessing our progress it will be important to establish the currently achievable detection thresholds as a function of signal magnitude, spatial extent, duration, and network configuration, as well as to quantify the false alarm rate. This will require establishing even more specific metrics for comparing results. It was generally agreed that participants should

report error bounds on their estimated spatial and temporal parameters, although there was less consensus as to whether these bounds should be strictly quantitative or whether more qualitative assessments (e.g., assigning a level of confidence on a 5-point scale to each detection) would also be useful. We propose to allow either type of confidence assessment for Phase III test data, with the idea that we will require quantitative bounds in the future if they prove to be feasible.

The synthetic displacement time series used in Phases I and II were provided without any stated standard errors or covariance matrices. It was pointed out that although the noise spectra of the test data can be assessed using statistical methods that do not rely on knowing the data covariance, a priori information about data errors can be valuable in certain situations. For example, a failing antenna may result in decreasing SNR and increasing position uncertainties, and this information can be incorporated into detection algorithms. Simulating the data covariance structure for synthetic data will require further examination of error statistics and may be included in Phase III data, with the caveat that the reported data errors may either over- or under-estimate that true data errors, just as is found with real data.

While it was clear that in future phases test data with subtle signals are needed to establish detection limits, there are also a variety of features found in real time series that further complicate transient detection, especially in real time. These should be included in the test datasets as well. For example, offsets due to non-tectonic sources such as instrument malfunctions are commonly seen in GPS time series. In retrospective analysis these can be removed from the data ahead of time or estimated simultaneously with other quantities if the times of offsets are known a priori. However, in a real-time implementation of transient detection algorithms there must be a robust mechanism for identifying offsets without prior knowledge. Spatially-coherent non-tectonic transient signals (including spatially-coherent seasonal signals, sometimes with time-varying amplitude) are also often found in geodetic datasets and are particularly problematic if the goal is to identify transient deformation sources due to fault slip or volcanic processes. More general source types such as coseismic (both large and small) and associated postseismic signals, poroelastic response, volcanic sources, and a strain “wave” propagating across a broad region were also discussed for future test datasets. A means of identifying the character of signals that will always be difficult to deconvolve from seasonal and anthropogenic sources will be an important part of any automated detection program.

Another topic which received considerable attention was the relative merits of using synthetic versus real test data. We chose to use synthetic data for Phases I and II. With a known signal, evaluating an algorithm’s success is easier, and participants may later use the known data characteristics to improve or identify problems with their methodology. However, synthetic data embodies assumptions regarding the kinds of signals found in geodetic data, and participants might tune their algorithms to these anticipated signals. The consensus among workshop participants was that there is substantial additional source complexity yet to be added to the synthetic time series (e.g., offsets, postseismic deformation, spatially-coherent nontectonic signals) and that detection algorithms are still early in their development. Therefore, Phase III should continue to use synthetic data (or potentially a mixture of synthetic and real datasets), but the time series should contain a greater variety of signals, and the target signals for detection (e.g., those due to fault slip) should be more subtle than the large-amplitude signals seen in some Group A datasets.

Participants also recognized that other data types can provide important constraints on the transient detection problem, and we discussed expansion of the test to other observations. InSAR, with its superb spatial coverage and sensitivity to vertical deformation, and strainmeter data, which provide a higher sampling rate and sensitivity to smaller-amplitude signals than GPS, are the most likely candidates. Incorporating multiple data types that differ greatly in their spatial and temporal sampling will pose new challenges for detection algorithms. Furthermore, because these data have error characteristics and noise sources that are less well understood than those of GPS data, simulating these types of observations for synthetic test cases is not straightforward. Therefore, participants concluded that Phase III of the test exercise should continue to focus on GPS data with the goal of better establishing detection levels, assessing the significance of reported detections, and addressing a broader range of complexities representative of those found in real GPS data. In the mean time, the development of means for generating synthetic strainmeter, InSAR, or other types of test data should be supported.

Next Steps

Phase III test data will be released in early 2010 and (like Phases I and II) will consist of a collection of datasets. These will be primarily synthetic data and will incorporate the more complex deformation signals and noise sources described above, however it is possible that some real GPS datasets will be included as well.

In addition to the Phase II Group A and B datasets, a set of six Group C datasets was prepared but not yet released. We are releasing the Group C datasets simultaneously with posting of this workshop report for groups to work on between now and the start of Phase III. The deadline for submission of Phase II Group C results is February 1, 2010.

Duncan Agnew has developed software for generating the synthetic test data, and this code will be available to participants so that they can generate additional data for internal testing of their algorithms. However, before general release of the code beta testers are needed. Interested groups should contact Duncan Agnew and cc Rowena Lohman and Jessica Murray-Moraleda.

We anticipate that a future workshop will focus on establishing GPS transient detection thresholds, means of evaluating test results using real data, and generation of test datasets for InSAR and/or strainmeter observations.