Workshop on Transient Anomalous Strain Detection: Summary of findings and next steps

Palm Springs, California
September 11, 2010

Jessica R. Murray-Moraleda and Rowena Lohman

Introduction

The SCEC Geodetic Transient Detection Exercise is an ongoing project targeting the SCEC Science Priority Objective "Develop a geodetic network processing system that will detect anomalous strain transients". Three phases of testing have been completed. During each phase participants were presented with test data consisting of GPS time series contaminated with a variety of noise sources in addition to, in some cases, signals due to transient fault slip. In all three phases synthetic data were used, and in Phase III one set of real data from southern California GPS sites was also included. Participants applied the transient detection methods they are developing to these data in an attempt to detect the fault slip signals in the data. Further details on the mechanics of the exercise can be found in our report on the 2009 workshop (Murray-Moraleda and Lohman, 2010; http://www.scec.org/meetings/2009am/transient.html).

At the 2010 workshop, eleven groups presented the transient detection methodology they are developing, Duncan Agnew presented improvements to his *Fakenet* code which was used to generate the synthetic test data, Tom Herring described the steps used in generating the real data that were also provided for Phase III, and Rowena Lohman then presented a summary of the results from participants in Phase III. The approaches being developed, the test data sets, and the Phase III results are detailed in the workshop slides for individual presentations which are available on the group website (http://groups.google.com/group/SCECtransient). This report will focus on the discussion which took place at the workshop, next steps for the project and progress since the workshop. The two main issues that arose in the discussion centered on further assessing what signals algorithms can reliably detect and initial testing of algorithms in a continuously-running near real-time basis.

Further Phases of the Transient Detection Exercise

The Phase III synthetic test data contained substantially more subtle and complex signals than did the Phase I and II datasets. As a result, the detections reported by the various participants

presented a more nuanced picture of the current level of detection capability among these algorithms and highlighted the approaches' strengths and weaknesses. It was generally agreed that developers must now assess their own algorithms' performance in order to improve functionality. More systematic analysis of the methods' sensitivity and false alarm rates as a function of station distribution and source characteristics is also required. The *Fakenet* code is a valuable resource for this purpose, allowing users to generate as many synthetic datasets as they like with a range of characteristics for internal testing.

At the same time, workshop attendees recognized the need for additional "blind testing" in order to obtain more objective tests of algorithm capability, allow comparison of the algorithms' strengths and complementary features, and continue to foster a community of researchers targeting the goal of transient detection. Continued use of synthetic data has a role for testing specific source characteristics and for identifying the range of signals that algorithms just cannot detect.

However, there was general agreement that the format of the test phases thus far could be improved in the following ways. First, although many algorithms can efficiently ingest one new position estimate per station daily, applying the algorithms retrospectively to ten years of data is time consuming. As a result, few participants were able to analyze all datasets as carefully as needed, and the results may be an inaccurate representation of the algorithms' capabilities. Second, presenting the summary of participants' results at the workshop allows no time for developers to assess their performance and, in the case of missed detections, analyze why the algorithm failed. Finally, in some ways the "blind test" exercise has been too blind in that participants had little input in the types of transient signals represented in the data. It is likely that participants will feel more invested in the project and be more motivated to take part if they feel they have a greater voice in debating and choosing the sorts of signals that should be included in the test data.

Based on this feedback, we recommend the following format for the next phase of the test exercise. We hope that this format will create greater buy-in among participants and lead to sustained improvement of the methods under development.

1) Developers use the *Fakenet* code (or other tools if they so choose) to generate their own data for internal testing to improve specific functionality as appropriate for their algorithms.
2) A single dataset (as opposed to 4 - 12 as in previous phases) consisting of time series for southern California GPS stations that may contain one or more transient slip signals

will be released quarterly. Participants will apply their algorithms to this dataset and upload results to the website as normal. Then, a summary of results will be posted to the website for review and online discussion. As of April, 2011, we have now made 2 datasets available to the group and have assimilated one set of successful detections into the new wiki page.

3) Prior to each quarterly release of test data, participants will propose features for the upcoming dataset that will be designed to test specific capabilities (e.g., the ability to detect multiple transients that overlap in time, a transient at the end of the time series, or a transient that occurs at the edge of the network). The proposed tests can and should include tests with real data.

Continuously running near-real-time detection algorithms

While workshop participants generally agreed that further work on improving algorithms could be fruitfully pursued for some time, there was extensive discussion of how to move the project closer to realizing the SCEC goal of operational deployment of one or more detection algorithms. In particular, the SCEC leadership expressed intense interest in having one or more detection algorithms running continuously in an automated fashion at the start of SCEC4 in February 2012.

There was discussion of who would receive (and presumably act on) reports of detected transients were an operational system to be running. Several workshop participants expressed concern that at the current stage of development, most algorithms could not produce results on an ongoing basis without a level of developer intervention that would make doing so prohibitively time-consuming. Moreover, there was universal agreement that the results of any continuously-operating detection algorithms should not be made publically available during the testing phases, especially since most algorithms still lack a rigorous detection criterion. It was also suggested that if several algorithms identified the same transient signal in real data with community consensus regarding the source, that this would be a more exciting result than pressing for lower latency at this stage.

However, many participants acknowledged that if we do not push for continuous, near real-time deployment of algorithms that it will not happen on its own in a timely manner. With one or more algorithms running in an ongoing way, other detection exercise participants who are at an

earlier stage of development could compare their results to those from the automatically running algorithms. Moreover, running a large number of tests, for example to statistically assess the sensitivity and specificity of the algorithms, can be done more efficiently if algorithms can be run with little user-intervention.

While operational detection methods are the long-term goal, the target for February 2012 would necessarily be modest and certainly would not be described as "operational". There was consensus that implementing one or more algorithms to automatically report detections on a weekly basis was a good first target, and this would mesh well with the weekly availability of the final orbits used in the GPS processing. Algorithm developers would be notified of these detections for review, similar to the way in which automatic earthquake location and magnitude reports are reviewed by a seismologist. Initially only participants in the project, SCEC leadership as appropriate, and perhaps GPS network operators would have access to the results. Once detection algorithms are operational, USGS would likely have the responsibility, given their mandate, to respond to potential transient fault slip events. The output of detection algorithms will also be useful to network operators for monitoring data integrity.

Implementation of one or more continuously-operating algorithms will require IT infrastructure to provide a steady stream of input (i.e., daily GPS positions) to the algorithms, ingest the output of algorithms to display detection information in map and time-line view, and alert developers when their algorithm returns a detection so. There are several sources of continually-updated daily processed GPS data that might be used as input. Two applicable examples are those from the Plate Boundary Observatory (PBO) processing centers and the NASA MEaSURES project, both of which cover southern California. We currently anticipate using PBO data, in part because of Tom Herring's existing history of working with this project. It may be possible to draw upon parts of the CSEP infrastructure for running the algorithms and reporting results, or at least to use its development as a model. In CSEP, participants provide an executable of their software as well as source code. Their software is run within the CSEP testing architecture using standard input and output without participant intervention thus reducing the burden for developers of maintaining continuous operation of their code and result reporting

Publication of the results thus far

Following the workshop it was suggested that the results of the transient detection blind test exercise thus far be published in order to summarize and make available to the broader community what we have learned. Such a report would have two main components: a description of each approach and its pros and cons, and a discussion of what combination of features from similar methodologies might provide the best results. We envision the report to be brief, focusing on outcomes and next steps. Thus, description of methodologies should refer to published work where possible with necessary details for yet-to-be published algorithms given in an online supplement. Because the detection methods that have been employed in the test exercise generally fall into several broad categories (for example, methods based on PCA, Kalman filtering, or strain rate mapping), the report should devote ample attention to the ways in which complementary approaches may be used in tandem and what the next steps are for implementing ongoing transient monitoring.

Action items include the following

1) Contact all participants in the blind test exercise to solicit their participation.
2) Establish the scope and format of the document and identify an appropriate journal to which to submit the report.
3) Devise a schedule for compilation of participant contributions to the report, editing, and review.

Incorporation of other data types

Discussion of extending the test exercise to other data types was limited as attention during the workshop focused primarily on the issue of continuously-operating algorithms. However, it should be noted that one of the future SCEC activities described in the SCEC4 proposal is the generation of a Community Geodetic Model which would incorporate both GPS and InSAR time series. Thus, adapting or developing transient detection algorithms to use InSAR, either alone or in combination with GPS data, is an important target for future work. Indeed, the U.C. Riverside group presented results at the workshop that demonstrated the application of a PCA-based approach to removing seasonal signals from InSAR data. Such tools are potentially valuable for enabling the identification of other time-varying signals in the data as well.

The next Transient Detection workshop

Annual workshops like those that we have held at the 2009 and 2010 SCEC meetings are an important means of maintaining project momentum. These workshops attracted a broad audience and provided a forum for test exercise participants to describe their methodology. As we move forward, however, some modifications are required to better reflect the evolving needs of the project and address drawbacks of the current format. The envisioned 2011 workshop will be more targeted, thus helping to advance the state of the art more efficiently.

We propose a two-part workshop aimed at those groups who are actively participating in the test exercise and/or extending their methodology to continuous operation. The first half of the workshop will focus on advances to methodology based on results from the periodic blind tests that will be held throughout the year. By making the true signals available directly after each blind test, participants will have had time to investigate any failures of their algorithms in time for the workshop, and discussion can focus on the strengths of individual approaches and ways in which multiple methods might be used together. The second half of the workshop will be dedicated to planning for the implementation of continuously-operational algorithms, building on discussions that will be occurring throughout 2011 with the groups who plan to participate in this aspect. One difficulty that some participants have faced in previous years is that SCEC funding has arrived sufficiently late in the year for groups to complete all the work they proposed in time for

workshops held at the SCEC annual meeting. Therefore, we propose to hold this one-day workshop directly after the AGU Fall Meeting in San Francisco.

Conclusion

The SCEC Transient Detection Exercise has succeeded in raising awareness of possible approaches to the detection problem, bringing together a group of researchers actively working on this problem, and spurring the development of a diverse set of algorithms. It is now time to move to the next stage in which we identify the most promising combination of approaches and strive to implement these tools in a continuously-operating manner. SCEC Leadership has now committed to providing support for the development of the necessary IT infrastructure over the timescales discussed above. For now, our proposed set of tasks looking ahead to the future includes the following:

1) Conduct periodic blind tests designed with greater participant input and involving a more manageable amount of test data.
3) Provide a framework for the extension of one or more algorithms to continuous operation.
4) By the time of the SCEC meeting, demonstrate a simplistic algorithm on real data within this framework
5) Publish the results to-date as a record of what has been learned and a foundation for further development.
6) Encourage work that will enable extension of transient detection to other data types.
7) Hold a targeted workshop for active participants which will focus on the application of methodological advances and concrete next steps for continuous operation.

References

Murray-Moraleda, J. and R. Lohman (2010), Workshop targets development of geodetic transient detection methods, *Eos Trans. AGU*, *91*(6), 58.