AFRL-RV-PS-
TP-2015-0013

PIPELINE PROCESSING WITH AN ITERATIVE, CONTEXT-BASED DETECTION MODEL
Annual Report

T. Kværna, et al.

NORSAR
PO Box 53
N-2027 Kjeller, Norway

19 April 2015

Technical Report

AIR FORCE RESEARCH LABORATORY
Space Vehicles Directorate
3550 Aberdeen Ave SE
AIR FORCE MATERIEL COMMAND
KIRTLAND AIR FORCE BASE, NM 87117-5776

DTIC COPY

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 377 ABW Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RV-PS-TP-2015-0013 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

//SIGNED//	//SIGNED//
Dr. Robert Raistrick	Glenn M. Vaughan, Colonel, USAF
Project Manager, AFRL/RVBYE	Chief, Battlespace Environment Division

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
19-04-2015	Technical Report	19 Mar 2014 to 18 Mar 2015

4. TITLE AND SUBTITLE

PIPELINE PROCESSING WITH AN ITERATIVE, CONTEXT-BASED DETECTION MODEL
Annual Report

5a. CONTRACT NUMBER
FA9453-13-C-0270

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
62601F

6. AUTHOR(S)
T. Kværna, D. A. Dodge[1], D. B. Harris[2], and S. J. Gibbons

5d. PROJECT NUMBER
1010

5e. TASK NUMBER
PPM00018850

5f. WORK UNIT NUMBER
EF122183

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
NORSAR
PO Box 53
N-2027 Kjeller, Norway

[1]Lawrence Livermore National Laboratory
7000 East Avenue
Livermore, CA 94550

[2]Deschutes Signal Processing, LLC
81211 E. Wapinitia Road
Maupin, OR 97037

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory
Space Vehicles Directorate
3550 Aberdeen Avenue SE
Kirtland AFB, NM 87117-5776

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RVBYE

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-RV-PS-TP-2015-0013

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited. (377ABW-2015-0442 dtd 08 Jun 2015)

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Under existing detection pipelines, seismic event hypotheses are formed from a parametric description of the waveform data obtained from a single pass over the incoming data stream. The full potential of signal processing algorithms is not being exploited due to simplistic assumptions made about the background against which signals are being detected. A vast improvement in the available computational resources allows the possibility of more sensitive and more robust context-based detection pipelines which glean progressively more information from multiple passes over the data. In the first year of this two year contract we designed and implemented several extensions to an existing prototype detection framework to demonstrate the feasibility of improving performance from a systematic reprocessing of the raw data: signal cancellation for stripping the incoming data stream of repeating and irrelevant signals, adaptive beamforming and matched field processing for suppressing background signals and aftershock sequences, and the testing of event hypotheses by evaluating detection probabilities for both detecting and non-detecting stations, followed by optimized beamforming. In this second year of the contract, we have evaluated and enhanced significantly the signal cancellation procedures, revised the detection framework architecture in order that the procedure can be distributed and scaled efficiently, and developed a procedure for optimal detection and location of aftershocks from a target source region as a component of an iterative pipeline. This last procedure exploits significantly the probability of detection work performed in the first year of the contract.

15. SUBJECT TERMS
aftershock sequences, repeating explosions, detection framework, pattern detectors, correlation detectors, subspace detectors, matched field detectors, nuclear explosion monitoring

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Dr. Robert Raistrick
Unclassified	Unclassified	Unclassified	Unlimited	74	**19b. TELEPHONE NUMBER** *(include area code)*

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

This page is intentionally left blank.

Table of Contents

List of Figures

List of Tables

1 SUMMARY

Under existing detection pipelines, seismic event hypotheses are formed from a parametric description of the waveform data obtained from a single pass over the incoming data stream. The full potential of signal processing algorithms is not being exploited due to simplistic assumptions made about the background against which signals are being detected. A vast improvement in the available computational resources allows the possibility of more sensitive and more robust context-based detection pipelines which glean progressively more information from multiple passes over the data. In the second year of this two year contract we have continued the design and implementation of extensions to an existing prototype detection framework to demonstrate the feasibility of improving performance from a systematic reprocessing of the raw data.

An algorithm for stripping the incoming data stream of repeating and irrelevant signals prior to running primary detectors has been modified and evaluated; this procedure has been documented in a draft manuscript for submission to a scientific journal. The detection framework architecture has been revised in order that the procedure can be distributed and scaled efficiently. The revised architecture is also more easily maintained since the data buffers are no longer shared among multiple objects, and the system including the detection cancellation algorithm has been successfully evaluated. Extensive aftershock sequences can result in a significant deterioration of automatic event bulletins due to incorrect association of phase detections. We describe the key elements of a targeted event detection and location scheme for identifying events in a limited source region using specially optimized context-based detectors such that all corresponding phase detections can be removed from the data stream.

2 INTRODUCTION

The traditional processing flow in monitoring pipelines consists of a transformation of continuous waveform data into a stream of parametric information describing discrete detected arrivals. Critical event-building functions (association, location and identification) are performed subsequently on the parametric data. This processing architecture is a legacy of computing resources available during the 1980s and 1990s, for which very intensive signal processing operations were computationally prohibitive. Expensive signal processing operations were minimized in such systems by performing only a single pass over the waveform data.

One consequence of this architecture is that only very simple assumptions are possible about ambient background conditions in the selection and development of signal processing algorithms. The key assumptions – clearly false in important and not-infrequent circumstances – are that events occur in isolation and that the background consists of white noise, usually uncorrelated and uniform in power among array sensors. Pipelines do have procedures to disentangle overlapped and interleaved seismic arrivals, but these usually are confined to association algorithms operating on the parametric representation of detected arrivals. Except to allow analysts to verify whether events have been built properly, and build new ones as necessary, monitoring systems generally do not revisit waveform data once it has been reduced to its parametric description.

This single-pass architecture imposes substantial limitations on the signal processing algorithms that are applied for detection and phase characterization. Among these is the inability to adapt algorithms to suppress or cancel interfering signals with known characteristics. Future pipelines could allow waveform data to be revisited by detection and characterization algorithms that exploit some understanding or model of the context in which they operate.

In contemplating the signal processing front-end of future pipelines, we assume the overall systems would have characteristics or functions not present in current implementations. For example, we imagine that central data collection systems would have autonomous components to characterize local sources of interfering signals near each station and remove or suppress those signals from the data stream before it is passed to the main pipeline for detection and subsequent processing. Such a subsystem would generalize existing masking functions that handle dropouts and data spikes to handle interfering, propagating transients from local nuisance sources. Examples of such sources abound: glaciers near SPITS producing icequakes, cultural sources near KSRS producing lone Rg phases. At times, these stations can be plagued with thousands of nuisance transients. A key element of this project is to build and test a prototype detection and cancellation component.

Our second assumption about future pipelines is that they would maintain constantly-refined models (i.e. world views) of seismic activity around the globe, and use those models to adapt the signal processing functions of the pipeline to optimize detection and characterization functions. In such architectures, model state might be determined initially by pipeline processes not very different from those used currently. However, waveform data could be extensively reprocessed with algorithms conditioned on hypotheses about the current state of seismic activity (e.g. large aftershock sequence in progress, mid-day mining explosions anticipated, etc.). A mature picture of activity may emerge after several iterations of model-

driven signal processing on the continuous data stream. It is not that this picture of activity is of direct monitoring interest, but rather that it may allow more effective screening of interference in the continuous data stream.

Context-driven signal processing operations might include beamforming algorithms that search for events of interest among strong transient interference. Current beamforming algorithms are designed for very benign, white, uncorrelated noise background conditions. They are not well suited to detection among events in an aftershock sequence, for example. However, it should be possible to configure beamforming and related (for example, matched field processing) operations in second passes over data to make use of knowledge of interfering events gleaned from an initial pass. Specifically, we will investigate a data-adaptive, empirical matched field processor that substantially suppresses aftershock signals while remaining sensitive to signals from a known test site. The capability of this algorithm to suppress interference is significantly beyond the ability of a conventional beamforming algorithm because it is specifically designed to reject events with known characteristics.

The model-building supervisory function of future pipeline architectures also could instigate signal processing operations to support or refute hypotheses being formed or tested. For example, having formed a marginally-supported event hypothesis with a minimum of phase detections, the supervisor could direct special beams to search for phases predicted to exist at stations without detections. Current architectures have some capability to perform this function. For example, the IDC system performs associations, forms and locates events based on early-arriving seismic and hydroacoustic data (forming the SEL1 bulletin). The system automatically requests data from auxiliary stations based upon event hypotheses reported in that bulletin, and performs detections and measurements on those new data. However, this function can be substantially enhanced to direct re-examination of all waveform data, both to search for missing phases as noted and to confirm detections in the context of an hypothesized event (with location and magnitude estimates).

The general objective of this project is to examine whether a processing approach that allows waveform data to be revisited repeatedly can substantially improve detection performance in a monitoring pipeline. The specific objectives are to determine whether:

1) information about transients from local sources surrounding an array can be systematically discovered by an autonomous system and exploited to reduce interference with beamforming operations,
2) an advanced pipeline could configure adaptive beams to reject aftershocks more effectively than conventional recipe beams,
3) catalog events could be more effectively built by reprocessing data with beams

specifically designed to search for phases predicted to exist under event hypotheses but which were not detected on a first pass over the data with conventional recipe beams,

4) false associations can be discovered with specialized beams designed to check phase detections made with conventional recipe beams.

The specialized beams or detection algorithms we have in mind for reprocessing the data in pursuit of objectives 3 and 4 are likely to be adaptive beams configured to search for specific phases in correlated background noise or among waveforms from competing events. Since they would be applied in a second or third pass over the data, they would be designed to detect phases with specified characteristics while rejecting signals already detected and characterized in a previous pass over the data.

3 TECHNICAL APPROACH

The introduction outlines a number of specific procedures that a context-based detection model can incorporate to produce a higher fidelity description of relevant sources of seismicity on local, regional, and global scales. The previous annual report described three specific approaches: signal cancellation for problematic transients, adaptive beamforming with matched field detection for aftershock suppression, and the testing of seismic event hypotheses through the reprocessing of raw data. Here, we describe in detail:

(1) the modified signal cancellation procedures, including mathematical formulation of the methods,
(2) the revised detection framework architecture,
(3) the current status of a context-based procedure for detecting and locating events in a limited source region, using specifically targeted detectors.

3.1 Cancellation of Repetitive Transients in Seismic Array Data Streams

Abstract

Highly repetitive local seismic interference often complicates operations at key stations in earthquake and verification monitoring networks. In some applications it may be desirable to remove interference transients from the continuous data streams of seismic arrays as a preprocessing step before normal detection, phase identification and event formation functions. We present several algorithms for transient interference suppression based on empirical models for interference sources. The general approach these embody is a dual of the traditional Widrow-Hoff noise canceller. Cancellers assume that a separate recording of an interference source time history is available and that this time history is correlated with components of unwanted noise in a primary data stream. The objective of traditional cancellers is to estimate a transfer function – in seismic applications, a Green's function – relating the continuous source to interference in the primary stream, allowing the interference to be estimated and removed from the stream. By contrast, our application assumes that the Green's function is known – through past observations of repetitive transients – leaving the source time history to be estimated. This is a more ambitious, but feasible, cancellation operation. We examine two algorithms, one which makes a continuous estimate of the interference source time history and a second which models the source as producing a relatively sparse set of discrete, impulsive events. In the latter case, correlation

detectors find these discrete excitations to build an impulse-train representation for the source.

We give an example of icequake cancellation at the SPITS array in Svalbard and find that the continuous-source estimator is highly effective at removing interference transients at the cost of distorting desired signals. Distortion is reduced as the number of array channels is increased, but would require an unrealistic number of channels for high-fidelity preservation of desired signals. The intermittent source algorithm largely avoids distortion of desired signals but is not completely effective in removing interference. It is sufficiently effective, as we demonstrate, for significant improvements to pipeline operations. We also developed a post-processing correction to the continuous-source canceller that appears to be very effective in mitigating the distortion problem. The corrected continuous canceller would be the algorithm of choice for this application, were it not for its very high computational cost.

I. Introduction

Often it is the case that key seismic stations, otherwise ideally placed for some monitoring function, are located close to sources of seismic interference. This observation is true, in particular, for stations located in polar regions, in which the movement of glaciers and sea ice can generate thousands of small "icequake" events in a day. These events can cause major difficulties for monitoring pipelines, particularly in the detection and association phases of operation. A particularly clear example is seen in the SPITS array located in the European arctic (Figure 1). SPITS is a key station for monitoring the Arctic region, including the former nuclear test sites on Novaya Zemlya. For events in the Novaya Zemlya region, SPITS has, during normal background noise conditions, the capability to detect events down to magnitude 2.0 [Gibbons et al., 2011].

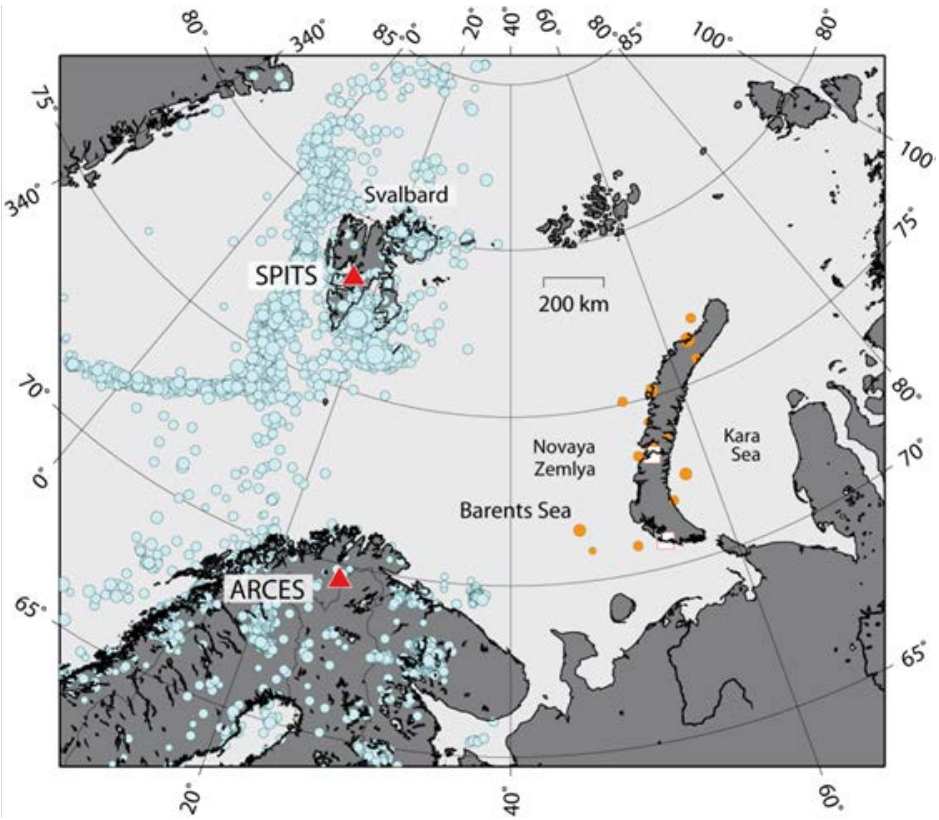

Figure 1. Location of the SPITS array in relation to Novaya Zemlya, the former Soviet nuclear test sites, and seismicity in the European Arctic. *The blue symbols indicate location estimates of seismic events from the NORSAR-reviewed regional event bulletin between 1998 and 2008. The red/white squares indicate the approximate locations of north and south Soviet nuclear test sites on Novaya Zemlya. The orange symbols refer to events on or close to Novaya Zemlya in the time period 1992-2010. Symbol size relates to event magnitude.*

However, the area around the SPITS array is characterized by permafrost and several neighboring glaciers which regularly generate seismic disturbances. During certain time intervals, e.g. in melting or freezing periods, these transients can dominate the seismic data and significantly reduce the detection capability of the seismic array. An example is shown in Figure 2 in the form of a helicorder type plot for 1 June 2011. The time interval between 15:00 and 18:00 UTC contains a large number of such transients.

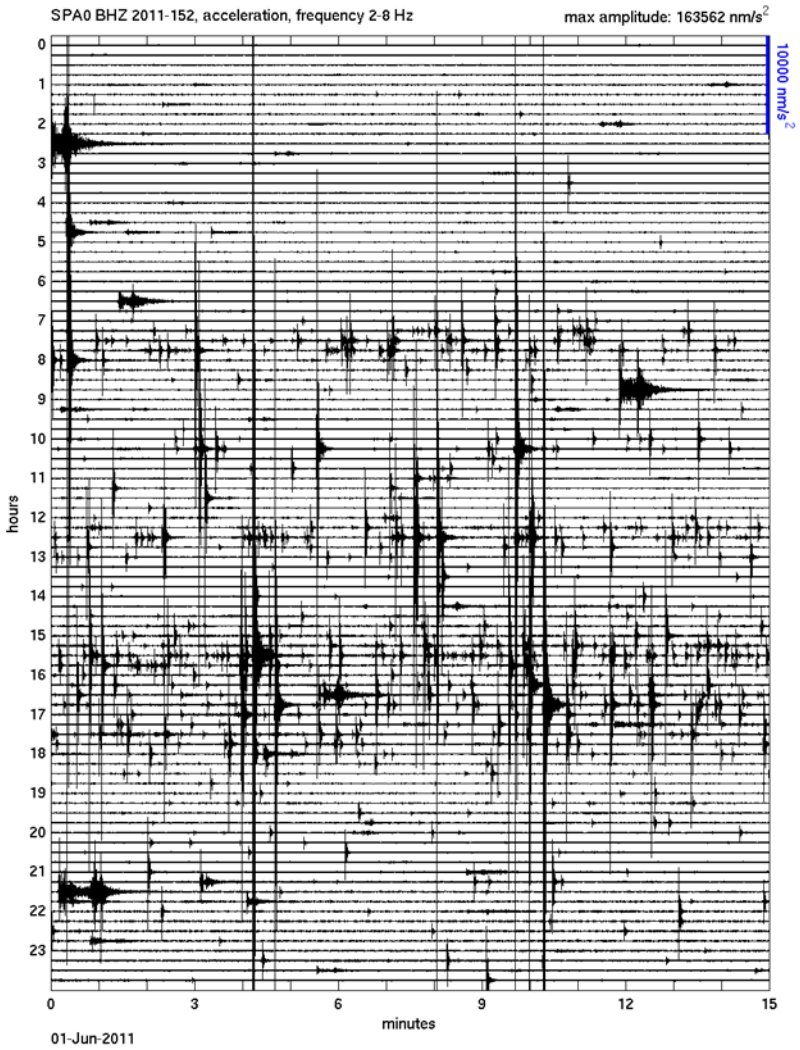

Figure 2. Mock helicorder plot for the central element vertical sensor of the SPITS array for 1 June 2011 (see http://www.norsardata.no/cgi-bin/spdatashow.cgi?year=2011&doy=152&sta=SPI). Signals from events at regional distances have wavetrains up to several minutes in duration. The vertical scale is dominated by very short duration signals (duration 5-10 seconds) of which many thousand are detected in this 24 hour period.

Frequently undesirable local transient waveforms are repetitive, being produced by a relatively small number of sources of recurrent events. When this situation obtains, a pre-processing strategy for "tagging" or removal of nuisance transients based on waveform correlation techniques is feasible. This strategy could take the form of detection and marking for transients, much like glitch or dropout detection and marking is implemented with a masking channel in some systems. A more ambitious approach

might remove nuisance transients by waveform estimation and subtraction. This paper examines this latter possibility.

Waveform estimation and subtraction closely resembles noise cancellation approaches used to remove continuous interference (e.g. noise from rotating machines); we choose to frame the problem in a mathematical structure close this formulation (Figure 3).

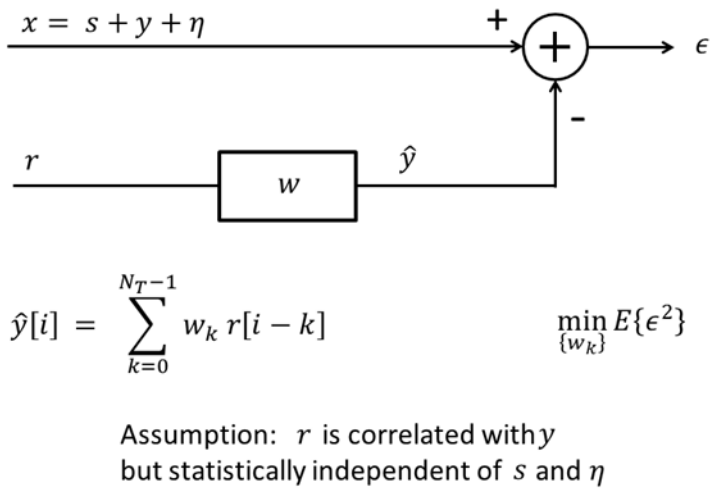

Figure 3. Block diagram of the classical cancellation algorithm.

Noise cancellers, such as the adaptive Widrow-Hoff [1960] algorithm, model the observed waveform as a sum of a desired signal, s, undesired interference, y, and possibly uncorrelated background noise, η. The methods require an independent observation, r, of the interference process which is free of the signal one wishes to observe. This independent reference must be correlated with the interference in the primary signal channel. Cancellers, then, estimate a finite-duration transfer function, w, to "shape" the reference interference measurement to match (as $\hat{y}$) the interference in the primary channel. The estimate $\hat{y}$ is subtracted from the primary waveform to suppress y. This kind of interference suppression can be highly effective and simple to apply, provided it is possible to identify and instrument the sources of interference with separate, dedicated sensors [for a geophysical application, see e.g. Harris et al., 1991]. It is possible to apply cancellation against multiple sources of interference simultaneously with a suite of reference observations targeting all of the sources. This extension of the process leads to a multichannel variant of the simple canceller shown in Figure 3.

In the application we contemplate in this paper, an independent observation of the noise source is not available, but an estimate of the transfer function is known, being measured from past observations. Our application is the more difficult dual of the common cancellation problem, in that we propose to swap the roles of r and w, estimating an interference source of effectively unbounded duration.

We consider two approaches to the problem of estimating the interference source. In the first we treat the source as a completely general waveform, solving a least-square problem for estimating the discretized samples of the interference source. This approach, though completely general, has two principal disadvantages: the first is that its computational cost is huge, and the second is that it biases estimates of the desired signal. The second approach makes a simplifying assumption that the interference consists of highly repetitive transient waveforms, so that the source consists of a series of delta functions and the transfer function is identical to the repetitive waveform that we seek to remove from the data. This approach first performs a detection step, using a generalized correlation detector to identify the number and times of the impulsive events that constitute the interference source. A cancellation step follows detection, scaling the impulses to match amplitudes of predicted and observed occurrences of the interfering transients

In fact, our model for the interference transfer function is somewhat more nuanced. We generalize the transfer function with a subspace representation to allow a range of variation in signals being targeted for cancellation. In principle, this generalization allows some variation in source location, mechanism or time history [Harris, 1989]. Although we have described the source time history as largely quiescent but punctuated by bursts of impulsive activity, in fact the source may be characterized by a near continuous range of activity with occasional very large excitations subsiding to a range of smaller excitations grading to an indistinct rumble. Because the detection step requires that we choose a threshold for detection, the smaller occurrences of interfering transients will not be removed from the data. The advantages of this approach are that it is comparatively much faster than the first method and, for desired signals, largely distortion free. Its disadvantage is that it removes less interference, particularly low-level interference. We observe that as the threshold is reduced the second method asymptotically approaches the first.

We also describe a post-processing technique for reducing bias in desired signals in the first cancellation method, based on a heuristic for identifying when the method is attempting to cancel a transient that is not a copy of the interference waveform.

This application and its solutions have much in common with finite fault decomposition (source slip estimation) using empirical Green's functions. The difference is that we consider problems with just a few source locations and continuous streams, rather than finite duration waveforms decomposed to hundreds of source locations gridding a fault surface. In addition our objective is not directly estimation of the source, which is responsible for unwanted interference in our application, but removal of the consequent transient interference in the continuous stream data.

The data we consider are multichannel: array data. In fact, the number of available channels controls a tradeoff between desired removal of transient interference and undesired distortion of waveforms from events of interest. The larger the number of data channels, the more reliably interference can suppressed while simultaneously minimizing distortion of desired waveforms.

The paper is organized as follows: in the next section (II) we present our mathematical model of transient interference intermixed with waveforms from events of interest. In section III, we derive the two principal solutions to the source estimation problem, discussing practical algorithms and the heuristic for suppressing bias in the continuous-source cancellation method. Following that, in section IV we apply the methods to data from the Spitsbergen array, comparing results.

II. Mathematical Formulation

We assume the seismic observations are multichannel, with N_c channels, and represent the sampled continuous signals as a vector sequence:

$$\mathbf{x}[n] = \begin{bmatrix} x_1[n] \\ x_2[n] \\ \vdots \\ x_{N_c}[n] \end{bmatrix} = \mathbf{s}[n] + \mathbf{y}[n] + \boldsymbol{\eta}[n] \qquad (1)$$

Streams from the individual channels are designated by $x_i[n]$; $i = 1, \cdots, N_c$ with time index n. We use italic characters to indicate scalar quantities and bold characters to indicate vectors or matrices. The observations consist of three components: unmodeled waveforms from events of interest, $\mathbf{s}[n]$**,** undesired interference transients, $\mathbf{y}[n]$**,** and background noise, $\boldsymbol{\eta}[n]$. We assume each of possibly several interference components to have the form:

$$\sum_k \mathbf{g}[n-k]\mathbf{f}[k] \tag{2}$$

The $\mathbf{g}[n]$ are finite duration transfer functions (Green's functions or sums of Green's functions) relating forcing functions $\mathbf{f}[n]$ at the interference sources to the resulting transients observed in the data. N_T is the duration of the transfer functions in seconds. In equation (2) $\mathbf{g}[\cdot]$ is assumed to be $\mathbf{0}$ for values of its index outside of the interval $[0, \cdots, N_T - 1]$. Taking into account the possibility of N_S independent interference sources, the complete representation for the data sequence is:

$$\mathbf{x}[n] = \mathbf{s}[n] + \sum_{i=1}^{N_S} \sum_k \mathbf{g}_i[n-k]\mathbf{f}_i[k] + \boldsymbol{\eta}[n] \tag{3}$$

As mentioned in the introduction, interfering transients often are repetitive. We use this fact to form empirical estimates of the transfer functions for groups of related interference events. Following [Harris, 2006], we use a 5 step process:

1.) Construct a pool of events by running a power (STA/LTA) detector over a period of data when interference sources are active. Extract event data segments of duration N_T.

2.) Cross-correlate waveforms from all events in the pool.

3.) Cluster the events using a hierarchical agglomerative algorithm using waveform correlation as a distance measure. The clustering threshold determines the number of clusters. Select clusters which have more than some minimum number of events. The number of such clusters is an estimate of the number of discrete sources producing interfering transients. We expect such sources to be prolific; in fact, it is this characteristic which distinguishes nuisance transients from events of interest.

4.) For each significant cluster, align the corresponding waveforms using delays estimated from peaks in the cross-correlation functions, then construct a data matrix of aligned waveforms with one column for each event.

5.) For each significant cluster, construct an orthonormal representation for the column space by performing a singular value decomposition (SVD) of the data matrix and retaining the left singular vectors corresponding to the top singular values of the decomposition. We determine the dimension d (rank) of the representation by finding the d largest singular values with cumulative energy greater than some specified

fraction of the total energy in all singular values. The matrix of left singular vectors is orthonormal by construction and constitutes the representation of the transfer function for a particular source.

Step 4 requires some elaboration. Since the data are multichannel, we multiplex the waveforms for a single event in channel-sequential form to form a column of the data matrix. This produces a consistent packed representation for a vector waveform (referring to equation 1):

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}[0] \\ \mathbf{x}[1] \\ \vdots \\ \mathbf{x}[N_T - 1] \end{bmatrix} \tag{4}$$

that preserves moveout among the waveforms across an array. The dimension of $\mathbf{x}$ is $N_c \cdot N_T \times 1$. With N_E events, the complete data matrix has the form:

$$\mathbf{X} = [\mathbf{x}_1 \quad \cdots \quad \mathbf{x}_{N_E}] \tag{5}$$

It is common to normalize the event data so that $\mathbf{x}_i^T \mathbf{x}_i = 1$ for all events, so that no single event or subset of events comes to dominate the data matrix and the corresponding SVD.

The SVD of step 5 produces the decomposition:

$$\mathbf{X} = \boldsymbol{Z}\,\boldsymbol{\Sigma}\,\mathbf{V}^T \tag{6}$$

with diagonal singular value matrix and orthonormal left $\mathbf{Z}$ and right $\mathbf{V}$ matrices of singular vectors:

$$\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_{N_E} \end{bmatrix} \tag{7}$$

We select the dimension of the representation d as the smallest integer that satisfies:

$$\frac{\sum_{i=1}^{d} \sigma_i^2}{\sum_{i=1}^{N_E} \sigma_i^2} \geq \gamma \tag{8}$$

where γ is called the fractional energy capture threshold. Here the singular values are assumed to be sorted into descending order. The partition of $\mathbf{W}$ that we retain as a representation for the collection of related waveforms is:

$$\mathbf{U} \;=\; [\mathbf{z}_1 \quad \mathbf{z}_2 \quad \cdots \quad \mathbf{z}_d] \tag{9}$$

The dimension of $\mathbf{U}$ is $N_c \cdot N_T \;\times\; d$ and it is an orthonormal matrix:

$$\mathbf{U}^T\mathbf{U} \;=\; \mathbf{I}_{d\times d} \tag{10}$$

Referring to (2), our working hypothesis is that $\mathbf{U}$ spans the matrix of Green's functions:

$$\mathbf{G} \;=\; \begin{bmatrix} \mathbf{g}[0] \\ \mathbf{g}[1] \\ \vdots \\ \mathbf{g}[N_T-1] \end{bmatrix} \tag{11}$$

so that $(\mathbf{I} \;-\; \mathbf{U}\,\mathbf{U}^T)\,\mathbf{G} \;\approx\; \mathbf{0}$.

The representation matrix $\mathbf{U}$ can be partitioned into separate time steps:

$$\mathbf{U} \;=\; \begin{bmatrix} \mathbf{U}[0] \\ \mathbf{U}[1] \\ \vdots \\ \mathbf{U}[N_T-1] \end{bmatrix} \tag{12}$$

which aids in the construction of a model for the continuous stream with embedded transients. Using this sequential form, our model for the interference component of equation (3) becomes:

$$\hat{\mathbf{y}}[n] \;=\; \sum_{i=1}^{N_S}\sum_{k} \mathbf{U}_i[n-k]\mathbf{a}_i[k] \tag{13}$$

For values of the indices of $\mathbf{U}_i[\cdot]$ outside of the interval $[0,\cdots,N_T-1]$**,** we consider $\mathbf{U}_i = \mathbf{0}$**.** Our objective in cancelling the interference is to estimate the weight sequences $\mathbf{a}_i[k]$ so as to minimize the energy in the residual sequence:

$$\sum_{n\,\in\, I} (\mathbf{x}[n] - \hat{\mathbf{y}}[n])^2 \tag{14}$$

over some suitably chosen interval I.

III. Algorithms

Our algorithms operate on finite time intervals. In this section, we define as $\mathbf{x}_n$ a segment of continuous data consisting of $N > N_T$ time steps and beginning at time n:

$$\mathbf{x}_n = \begin{bmatrix} \mathbf{x}[n] \\ \mathbf{x}[n+1] \\ \vdots \\ \mathbf{x}[n+N-1] \end{bmatrix} \tag{15}$$

The convolution operation defining our model of interference (from one source) then can be described by a matrix product:

$$\hat{\mathbf{y}}_n = \begin{bmatrix} \mathbf{U}[0] & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \mathbf{U}[0] & \mathbf{0} & \cdots & \vdots \\ \mathbf{U}[N_T-1] & \vdots & \mathbf{U}[0] & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{U}[N_T-1] & \vdots & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{U}[N_T-1] & \cdots & \mathbf{U}[0] \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{U}[N_T-1] \end{bmatrix} \begin{bmatrix} \mathbf{a}[n] \\ \mathbf{a}[n+1] \\ \vdots \\ \mathbf{a}[n+N-N_T] \end{bmatrix} \tag{16}$$

which simplifies the description of our first algorithm.

For a succinct description, we define the weight vector:

$$\mathbf{a}_n = \begin{bmatrix} \mathbf{a}[n] \\ \mathbf{a}[n+1] \\ \vdots \\ \mathbf{a}[n+N-N_T] \end{bmatrix} \tag{17}$$

of dimension $(N - N_T + 1) \cdot d \times 1$ and the model matrix:

$$\mathbf{M} = \begin{bmatrix} \mathbf{U}[0] & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \mathbf{U}[0] & \mathbf{0} & \cdots & \vdots \\ \mathbf{U}[N_T-1] & \vdots & \mathbf{U}[0] & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{U}[N_T-1] & \vdots & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{U}[N_T-1] & \cdots & \mathbf{U}[0] \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{U}[N_T-1] \end{bmatrix} \tag{18}$$

of dimension $N_c \cdot N \times (N - N_T + 1) \cdot d$. With these definitions our estimate of the interference is:

$$\hat{\mathbf{y}}_n = \mathbf{M}\,\mathbf{a}_n \tag{19}$$

Under the simplifying assumption that $\boldsymbol{s}_n$ and $\boldsymbol{\eta}_n$ can be modeled as a uniform, independent, white Gaussian processes, a solution to the estimation problem can be found by choosing $\mathbf{a}_n$ to minimize the squared residual norm:

$$\|\mathbf{x}_n - \mathbf{M}\,\mathbf{a}_2\|_2^2 \tag{20}$$

Continuous-source algorithm

Our first algorithm minimizes (20) directly, with the solution:

$$\mathbf{a}_n = (\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^T\,\mathbf{x}_n \tag{21}$$

This least-squares problem minimization is well-posed, provided the matrix:

$$\mathbf{R} = \mathbf{M}^T\mathbf{M} \tag{22}$$

is full-rank. Note that $\mathbf{R}$ is banded, and, in fact, block Toeplitz:

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_{N_T-1} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{R}_{-1} & \mathbf{R}_0 & \mathbf{R}_1 & \vdots & \mathbf{R}_{N_T-1} & \ddots & \vdots \\ \vdots & \mathbf{R}_{-1} & \ddots & \mathbf{R}_1 & \vdots & \ddots & \mathbf{0} \\ \mathbf{R}_{-N_T+1} & \cdots & \mathbf{R}_{-1} & \mathbf{R}_0 & \mathbf{R}_1 & \ddots & \mathbf{R}_{N_T-1} \\ \mathbf{0} & \mathbf{R}_{-N_T+1} & \cdots & \mathbf{R}_{-1} & \mathbf{R}_0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \mathbf{R}_0 & \mathbf{R}_1 \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{R}_{-N_T+1} & \cdots & \mathbf{R}_{-1} & \mathbf{R}_0 \end{bmatrix} \tag{23}$$

Where, for $0 \leq p \leq N_T - 1$:

$$\mathbf{R}_p = \sum_{j=p}^{N_T-1} \mathbf{U}^T[j]\;\mathbf{U}[j-p] \tag{24}$$

Note that the $\mathbf{R}_p$ are dimensioned $d \times d$ and, by construction, $\mathbf{R}_0 = \mathbf{I}_{d\times d}$. Also, $\mathbf{R}$ is symmetric, with:

$$\mathbf{R}_{-p} = \mathbf{R}_p^T \tag{25}$$

The core calculation is the evaluation of equation (21), which we break into two parts. The first is evaluation of $\mathbf{C} = \mathbf{M}^T\mathbf{y}$, which, by inspection, is a matrix convolution operation, and can be accomplished efficiently with overlap-add algorithms [Oppenheim and Shafer (1975)]. Given the result of this operation, the second part is evaluation of $\mathbf{a} = \mathbf{R}^{-1}\mathbf{C}$. Since $\mathbf{R}$ is block Toeplitz, there is an efficient Levinson-Durbin recursion for the solution of this problem, which does not require construction of the full $\mathbf{R}$ matrix or its inverse. For the reader who is not familiar with the Levinson-Durbin recursion, a derivation is provided for convenience in the appendix.

A note about implementation: the matrices of (18) and (23) are enormous in any realistic problem and cannot be constructed explicitly. Consequently, the block Toeplitz structure must be exploited with the Levinson-Durbin recursion or some similar algorithm. Even exploiting that structure, the solution is computationally intensive on today's servers. For example, for a rank-3 subspace representation and a data block size of 3 minutes at 40 samples per second, a template size of 10 seconds, the system to be solved is $(3 \cdot 170 \cdot 40 = 20{,}400)$ $20{,}400 \times 20{,}400$. For continuous stream processing we found it necessary to break the data into overlapping blocks, process each block separately and piece the results back together. We use three minute blocks, with one minute overlaps on each end, retaining the residual only for the central minute of each block. With this much overlap we do not notice any discontinuities in the processed residual segments when the continuous result is assembled. Although this method of processing contains a significant number of redundant calculations, processing larger blocks to reduce redundancy (say 7 minute blocks with the same overlap) is not more efficient, because the cost of solution by Levinson-Durbin recursion is $O(N^2)$.

Another point regarding the continuous source algorithm is that when more than one source of interference is present, with model matrices $\mathbf{U}_i$, we simply concatenate the model matrices to form a combined model $\mathbf{U} = [\mathbf{U}_1 \quad \mathbf{U}_2 \quad \cdots]$.

Intermittent-source algorithm

As noted in the introduction, a much faster solution is possible by approximating the excitation at each interfering source as a collection of discrete events. Instead of estimating the entire source time history as a continuous stream, we treat it as an intermittent, sparse sequence of excitations. We then perform a waveform fit and subtraction for distinct excitation events. Referring to the data model of equation (3), we represent each of the continuous sources $\mathbf{f}_i[n]$ as a superposition of M_i impulsive

forcing functions (corresponding to discrete interfering events) at distinct time indices $\{n_{ij}\}$:

$$\mathbf{f}_i[n] \;=\; \sum_{j=1}^{M_i} \mathbf{f}_{ij}\delta\big[n - n_{ij}\big] \tag{26}$$

In our empirical representation for the interference (equation 13) we treat each weight sequence as a similar superposition of delta functions:

$$\mathbf{a}_i[n] \;=\; \sum_{j=1}^{M_i} \mathbf{a}_{ij}\delta\big[n - n_{ij}\big] \tag{27}$$

The consequent interference model due to one source (i) is:

$$\sum_k \mathbf{U}_i[n-k] \sum_{j=1}^{M_i} \mathbf{a}_{ij}\delta\big[k - n_{ij}\big] \;=\; \sum_{j=1}^{M_i} \mathbf{U}_i\big[n - n_{ij}\big]\,\mathbf{a}_{ij} \tag{28}$$

The interference over all N_S sources is:

$$\hat{\mathbf{y}}[n] \;=\; \sum_{i=1}^{N_S}\sum_{j=1}^{M_i} \mathbf{U}_i\big[n - n_{ij}\big]\,\mathbf{a}_{ij} \tag{29}$$

Under this formulation, our cancellation objective is to find the collections of indices $\{n_{ij}\}$ and weights $\{\mathbf{a}_{ij}\}$ that minimize the functional:

$$\sum_n \left(\mathbf{x}[n] \;-\; \sum_{i=1}^{N_S}\sum_{j=1}^{M_i} \mathbf{U}_i\big[n - n_{ij}\big]\,\mathbf{a}_{ij} \right)^2 \tag{30}$$

We solve this problem with a type of alternation approach. We first estimate the indices, then the weights, then iterate the process. To estimate the indices, we perform a detection step, running a subspace detector [Harris, 2006; Harris and Dodge, 2011] over the continuous multichannel stream for each distinct interference source. Subspace detectors generate continuous detection statistics that range between 0 and 1, much like correlation coefficients. Each detector operates by running a sliding

window of length N_T over the continuous stream sample-by-sample, projecting the data in each successive window position into the subspace defined by the interference model representation $\mathbf{U}_i$. The detection statistic for a given time step is the energy of the projection normalized by the energy of the data in the window. The maximum possible projected energy is the energy in the data itself, so the maximum value of the ratio is one and equals one only when the data fall completely within the subspace (i.e. the trace in the window is a target interference waveform). We note, that if the model representation has rank one (i.e., if the representation consists of a single waveform), then the subspace detection statistic is exactly the square of the sample correlation coefficient between the data and the interference waveform.

The indices are estimated as the points at which the detection statistics exceed some threshold value. When several interference sources are in play, the detection process may be complicated by near simultaneous detections of an interference event by two or more detectors. The solution to this problem is to collect near-simultaneous triggers from all detectors and promote the one with the maximum statistic as the single detection. This reconciliation process determines the value of i in n_{ij}.

Once the indices of interference events are determined, the weights $\{\mathbf{a}_{ij}\}$ are estimated by minimizing the functional of equation (30). This process is simple if interfering events are sparse and do not overlap. In this case, interference events can be treated independently; the weights are obtained by a simple projection of the waveform data in the trigger window onto the subspace:

$$\mathbf{a}_{ij} = \begin{bmatrix} \mathbf{U}_i[0] \\ \mathbf{U}_i[1] \\ \vdots \\ \mathbf{U}_i[N_T - 1] \end{bmatrix}^T \begin{bmatrix} \mathbf{x}[n - n_{ij}] \\ \mathbf{x}[n - n_{ij} + 1] \\ \vdots \\ \mathbf{x}[n - n_{ij} + N_T - 1] \end{bmatrix} \tag{31}$$

For very active interference sources, such as the ones shown in Figure 1, interference waveforms frequently are superimposed. Then the estimation procedure is complicated by the need to estimate several weight vectors simultaneously over the union of time intervals of the interference events. In addition, with overlapped events, the subspace detection statistics may fall below the detection threshold for at least one of the events. Interference transients, even those well-characterized by one of the interference models, can be missed in the detection step. It is this effect which forces an iterative approach.

Our complete algorithm consists of estimating the indices $\{n_{ij}\}$ and weights $\{\mathbf{a}_{ij}\}$ in each iteration incrementally from the cancellation residual of the previous iteration. It often happens that cancellation of an interference transient in one pass of the algorithm exposes a new interference waveform to be detected in the next pass. In each pass, we add the indices of new detections to the accumulating set of indices for all events and re-estimate the weights for all detections. This process continues until no new detections are obtained. In the examples we consider, as many as four iterations produce usable detections.

Post-Processing Correction to the Continuous-Source Canceller

In the above, we have described two procedures for obtaining an estimate of the interference, $\mathbf{y}[n]$, resulting from a set of repeating transient noise sources: a continuous source algorithm and an intermittent source algorithm. Figure 4 displays a schematic picture of the behavior typically observed by the two algorithms in our case study of the SPITS icequakes. In the top trace (a) we see an original waveform containing a signal of interest in addition to background noise and high amplitude transients that we would like to remove from the data stream. In trace (b) we see the model for $\mathbf{y}[n]$ generated by the intermittent source estimator. This estimate is zero for all time samples except for those immediately following subspace detections; it is non-zero wherever there are sufficiently good matches with the interference source templates. Only two of the three unwanted transient signals are modelled in this example. While the subspace detector allows some variability in the source-time function, the sequence of multiple events which led to the superposition seen here results in a detection statistic lower than our predetermined threshold. As a consequence, the composite signal is ignored. Trace (c) is the residual, $\mathbf{s}[n] + \boldsymbol{\eta}[n]$, obtained from our observation $\mathbf{x}[n]$ by subtracting the interference estimate, i.e.:

$$\mathbf{s}[n] \; + \; \boldsymbol{\eta}[n] = \; \mathbf{x}[n] - \; \mathbf{y}[n] \tag{32}$$

The two transients present in (b) are cancelled with the rest of the waveform unchanged. As noted in the previous section, sometimes composite signals can be partially cancelled in one pass of the intermittent source algorithm, with further cancellation in subsequent iterations.

Trace (d) shows the estimate of $\mathbf{y}[n]$ from the continuous source model. The algorithm assumes that the interference source is always present. As a consequence, this estimate generally is non-zero. All three of the interference transients are well represented in the model, even the one resulting from the composite event. However,

the amplitude of $\mathbf{y}[n]$ is a significant fraction of the amplitude of the desired signal $\mathbf{s}[n]$, such that when the subtraction is performed (trace *e*) that signal has been corrupted. Observation of long durations of data indicate that the amplitude of $\mathbf{y}[n]$ is a relatively constant fraction of the amplitude of $\mathbf{x}[n]$ for a given interference template (with characteristics controlled by, for example, time-bandwidth product of the interference waveform, number of channels observing). A greater number of channels (e.g. with a 3-component array rather than a vertical-only array) or a more complex transient signal will match windows of unrelated data less well, reducing this fraction. Distortion of the original waveform is particularly prominent when a sharp contrast in amplitudes occurs over a time-window that is short in comparison with the length of the interference template. This is particularly problematic since it is prone to occur exactly at the onset of high SNR seismic phase arrivals.

Fortunately, it is relatively easy to identify the data segments for which it is beneficial to subtract $\mathbf{y}[n]$ from $\mathbf{x}[n]$. If the transient model $\mathbf{y}[n]$ represents the observation $\mathbf{x}[n]$ well at a given time, the match is good both in form and in amplitude. As a post-processing correction to the continuous source cancellation algorithm, we seek a scalar windowing function, $\mathrm{W}[n]$, such that

$$\mathbf{s}[n] \; + \; \boldsymbol{\eta}[n] = \; \mathbf{x}[n] - \; \mathrm{W}[n]\mathbf{y}[n] \tag{33}$$

with $\mathrm{W}[n] = 1$ where cancellation should be performed and $\mathrm{W}[n] = 0$ where the original waveform should be preserved. We begin by setting $\mathrm{W}[n] = 1$ for all samples *n* and we evaluate the instantaneous envelope functions $|\mathbf{x}[n]|$ and $|\mathbf{y}[n]|$ using the data traces and the corresponding Hilbert transforms. For any sample *n* where either ratio $|\mathbf{x}[n]|:|\mathbf{y}[n]|$or $|\mathbf{y}[n]|:|\mathbf{x}[n]|$ falls below a nominal threshold (e.g. 0.9) we set $\mathrm{W}[n] = 0$ which results in a binary trace which is predominantly unity over data windows containing a nuisance signal and predominantly zero elsewhere. There are small clusters of exceptions due to waveform variability and coincidental similarity over very short windows. Applying a median filter with a length of a few seconds eliminates these, resulting in a rectangular function that is unity over durations comparable to those of the interference transients and zero elsewhere. Adding a short-duration cosine taper to each of the zero-to-unity and unity-to-zero jumps in this function provides a windowing function $\mathrm{W}[n]$ that does not result in discontinuities (Figure 4, trace *f*). The lower two traces in Figure 4 display the windowed continuous source model and the resulting residual, in which the interference transients are removed with no distortion of the signal at other times.

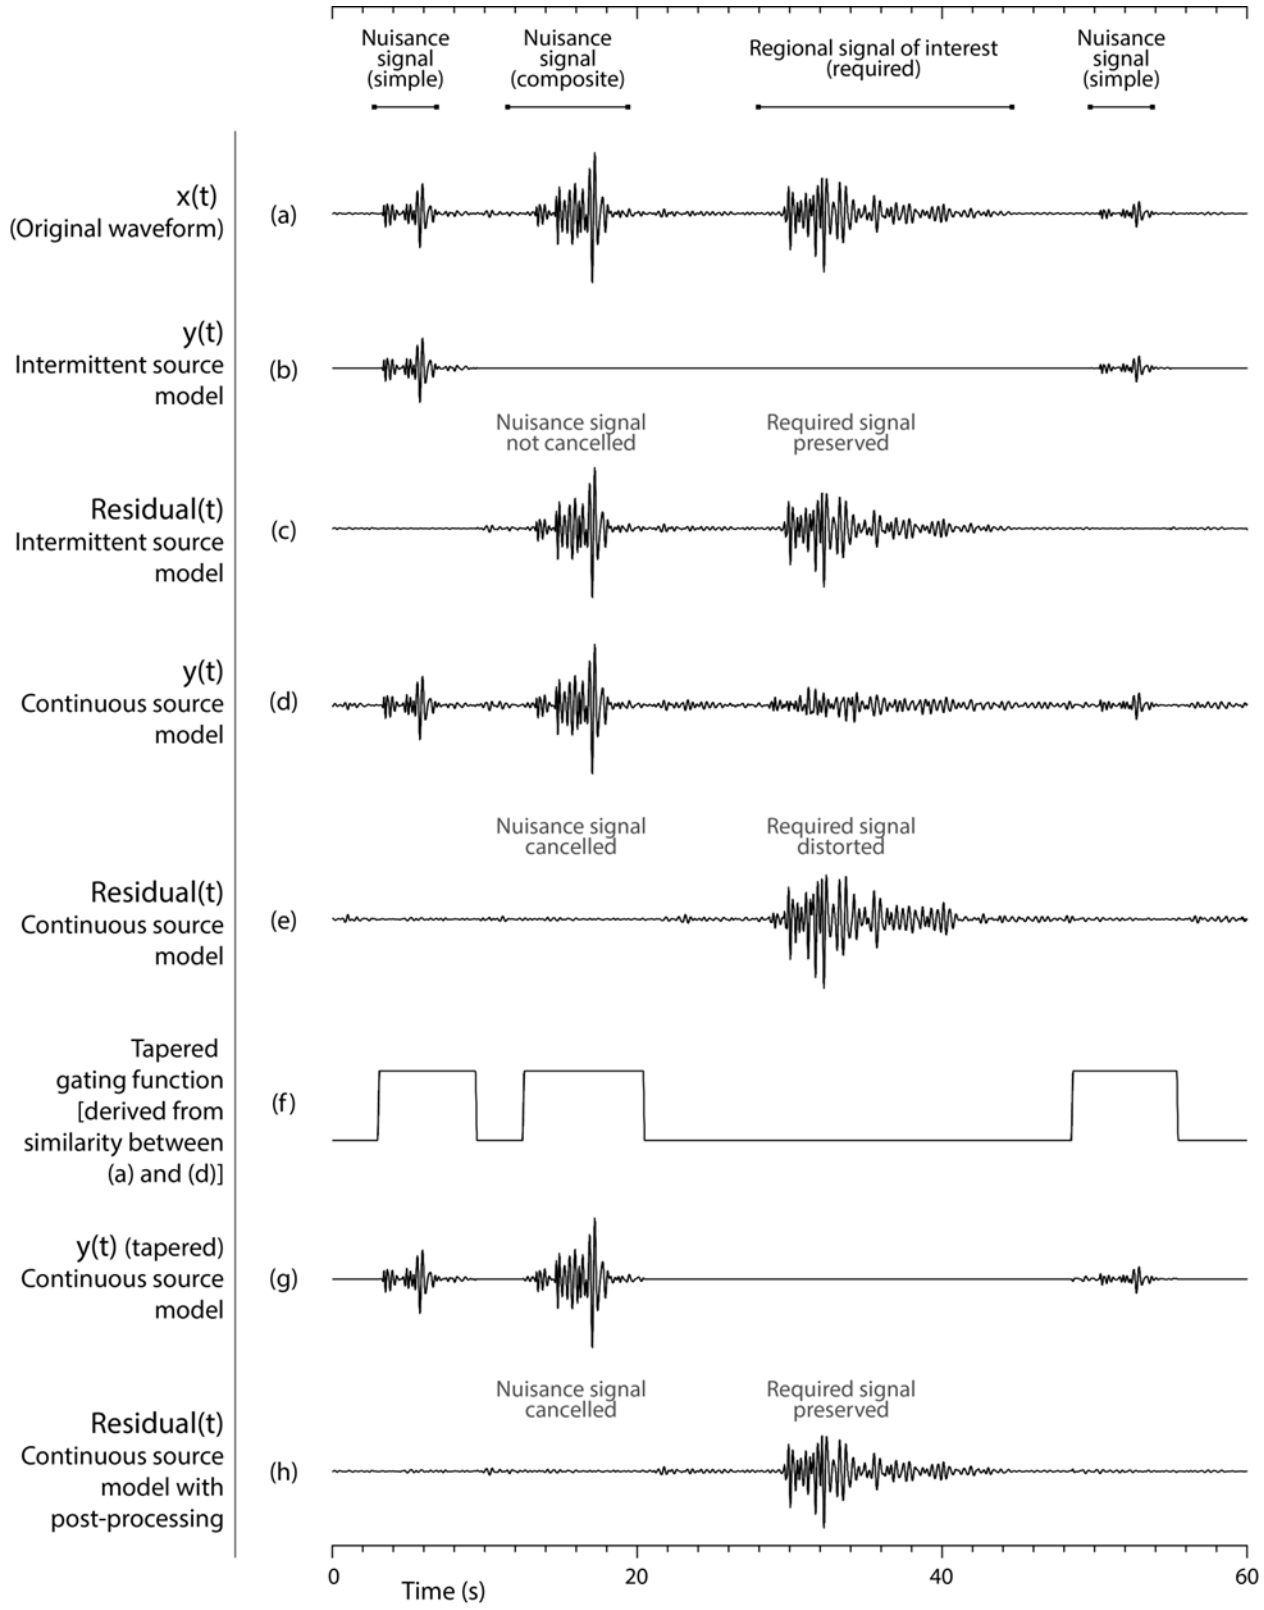

Figure 4. Schematic depiction of transient signal cancellation using intermittent and continuous source models of the interference.

IV. Spitsbergen Array Example

As an example of the application of the algorithms described in the previous section, we apply both to the continuous stream from the SPITS array during a period of prolific icequake occurrence. As described in the introduction, icequakes in the region surrounding the SPITS array dominate observations during periods of melting and freezing. P, S and Rg phases from these very local events are observed at the SPITS array and are, in many cases, reported by the automatic detection processor. Low propagation velocities, often in the range 1.3-2.0 km/s, are quite diagnostic for the Rg phases. However, for such nuisance P and S phases, the apparent velocities observed are indistinguishable from those estimated for P and S arrivals from regional events of monitoring interest. These nuisance detections cannot be eliminated from the detection stream without applying rigorous diagnostics to other properties of the signals (e.g. duration). Nuisance phases can constitute a very large fraction of all SPITS detections; failure to remove them from the processing pipeline can lead to spurious event hypotheses that slow the generation of seismic bulletins. For the year 2011 about 25% of SPITS detections appear to be slowly-propagating Rg phases. Figure 5 shows the azimuthal distribution of the 195,000 Rg-type phases reported in 2011, plotted at the geographical location of the SPITS array. We observe that the peaks of the azimuthal distribution coincide with the main directions towards neighboring glaciers.

During periods with high local event activity we often observe that “interesting” signals from regional or teleseismic distances are missed by the automatic detector because of data contamination by local signals. Even if nuisance transients do not overlap signals of interest, they may interfere with detection statistics, for example by reducing the Signal-to-Noise Ratio (SNR) for STA/LTA type detectors.

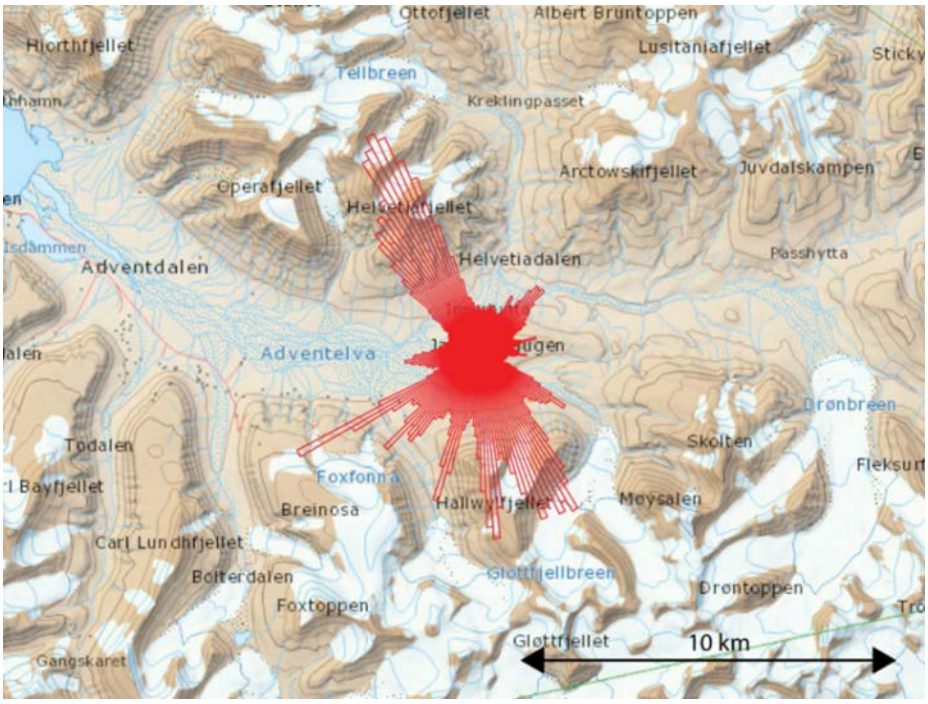

Figure 5. The rose diagram (red) shows the azimuthal distribution of 195,000 Rg-type phases reported in the SPITS detection lists for 2011. The rose diagram is centered at the location of the SPITS array.

As an example of cancellation performance, we apply both algorithms to continuous data from day 152 (June 1) of 2011. The SPITS array consists of 9 elements, a central element surrounded by a ring of 3 elements, with both surrounded by a larger concentric ring of 5 elements. The central element and the elements of the outer ring are three-component sensors. Elements of the inner ring are vertical sensors only. The data at this array are sampled at a rate of 80 samples per second. We apply the cancellers to two array configurations: the first consists of just the 9 vertical components, and the second consists of all 21 components.

The 5-step prescription described in section II was followed to identify the number of interference sources and to construct subspace representations for the waveforms they produce. In order to make a consistent comparison between the vertical and three-component data sets, the clustering and energy capture thresholds were manipulated in the design process to produce two clusters in each case, with a dimension of 3 for the subspace representing waveforms from the first cluster and dimension of 2 for the second subspace.

Results for the 21-channel cancellation example are shown in Figure 6 in a helicorder-style plotting format. At left in the figure is the original record from channel SPA1 BHZ for day 152 filtered into the 2 to 8 Hz band. It is clear from this panel that the amplitudes in the filtered waveform data are dominated by very large interference transients, such that events of interest are difficult to identify visually. The middle panel of the figure shows the cancellation residual obtained with the continuous algorithm and appears to show a very substantial reduction in the total energy of interference events in the recording. Regional events of interest, unrelated to the interference transients now are the dominant signals in the display. The data were decimated to 40 samples per second prior to applying the continuous canceller. The continuous canceller is quite slow, completing cancellation of 24 hours of data in perhaps 6 hours on a 12-core Xeon server. This slow speed occurs despite the fact that the Java implementation is concurrent, making use of all 12 cores of the server.

The panel at right displays the residual obtained with the intermittent-source canceller after four iterations. The detection threshold was set to 0.6 in this case. Cancellation of the interference transients is not quite as thorough as for the continuous canceller, particularly for two rather large events and numerous small events. The small interference events probably were not detected, due to low SNR. The two large interference events may have been superimposed on regional waveforms. While cancellation may not have been quite as complete with this algorithm, it nonetheless has two distinct advantages: it is orders of magnitude faster than the continuous canceller and it has significantly less distortion of the remaining waveforms from regional events of interest. However, this latter advantage is eliminated by the post-processing correction to the continuous-source canceller.

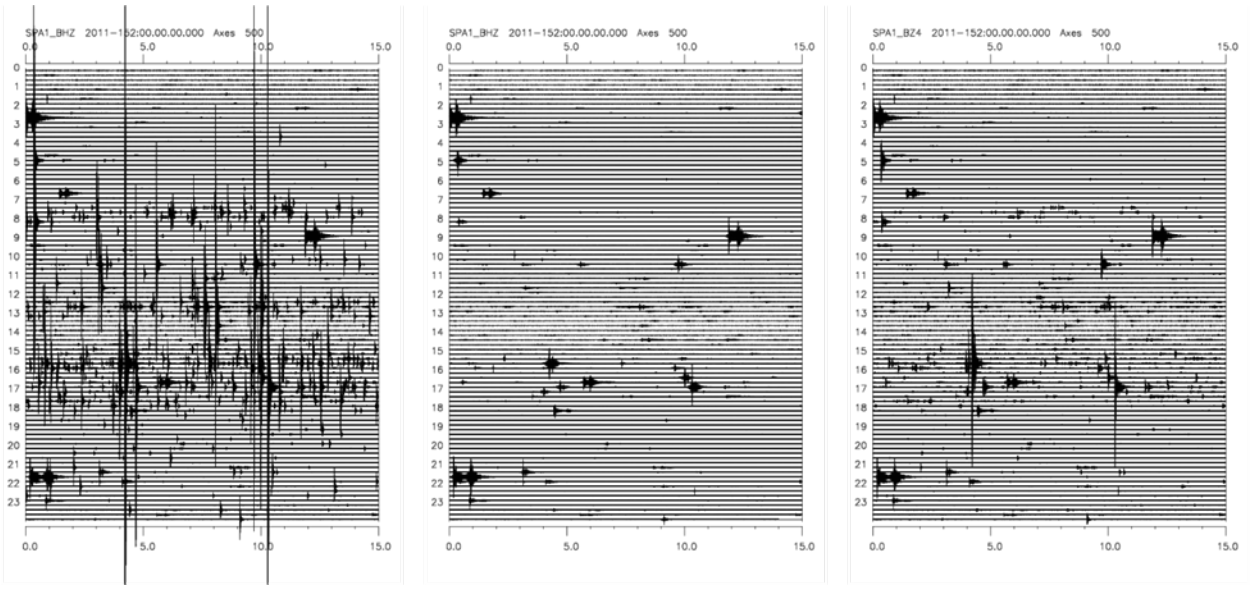

Figure 6. Helicorder type record of one channel (SPA1 BHZ) of the SPITS array (left), and cancellation residuals for the continuous canceller algorithm (middle) and the detecting canceller algorithm (right). *Icequake removal is more thorough with the continuous algorithm, but at the price of greater computational cost.*

For the intermittent-source canceller, it is illustrative to examine what is actually removed from the data stream in each iteration (Figure 7, bottom panels). It is clear that each iteration removes progressively fewer transients, and with lower amplitudes, than the previous iteration. However, since the computational cost of executing a single iteration of the detector is insignificant compared to the cost of the complete process, it is worth iterating the procedure until no subsequent change occurs in the waveform. In this example, this termination condition occurred after 4 iterations – the 5th iteration produced no further detections or cancellations.

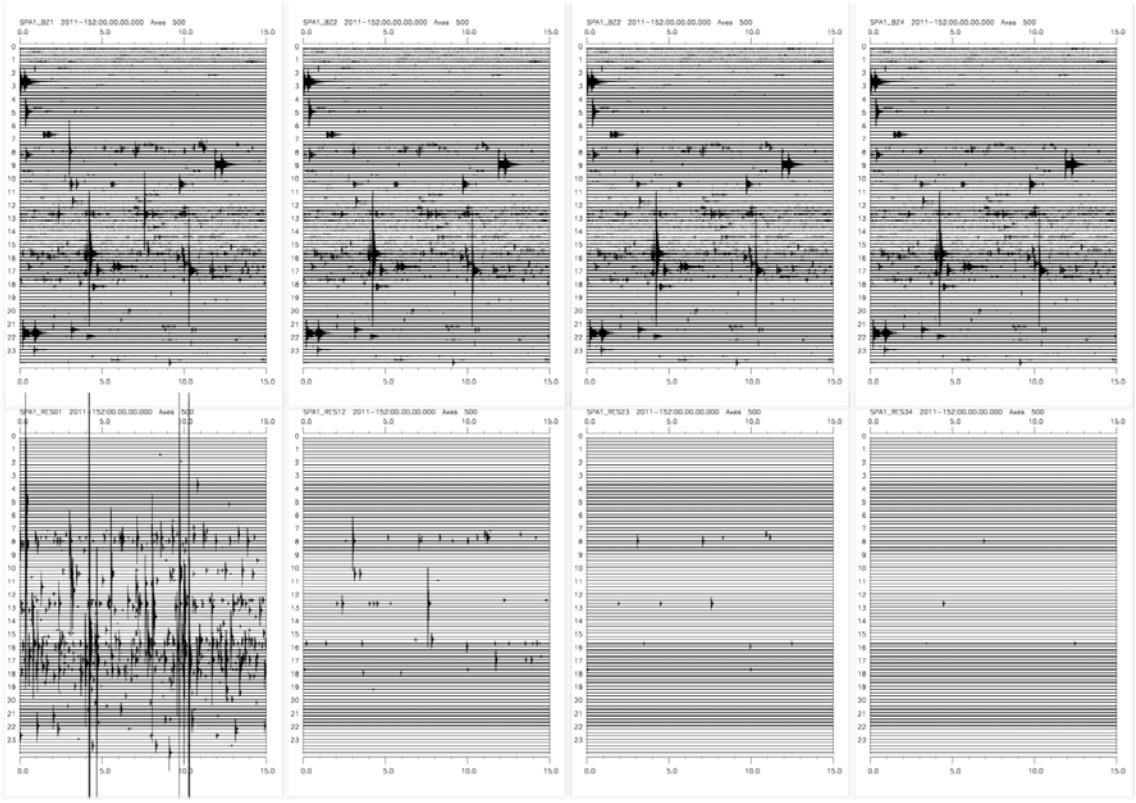

Figure 7. Composite of cancellation residuals for four iterations of the detection algorithm (top) and differences between the residuals from successive iterations (bottom). As anticipated, cancellation is subject to diminishing returns with iteration. The first two iterations produce significant numbers of detections and cancellations. There are no further cancellations after four iterations.

The advantages of the continuous source canceller are apparent in Figure 8, which compares details of the intermittent and continuous cancellation residuals over a 3-minute data segment. There are two clear icequake events in this time window. The intermittent source canceller reduces the amplitude of the first signal by an order of magnitude but fails to eliminate it to the point where it would not result in a detection subsequently with an STA/LTA detector. The second of the icequake signals is simply missed by the detector in the intermittent-source algorithm. By contrast, both of the signals are reduced to the level of the background noise by the continuous-source canceller. The continuous-source cancellation residual has been modified by the post-processor designed to preserve desired signals. With the tapered window function, no changes are made to the waveform outside of the time intervals covering the icequake signals.

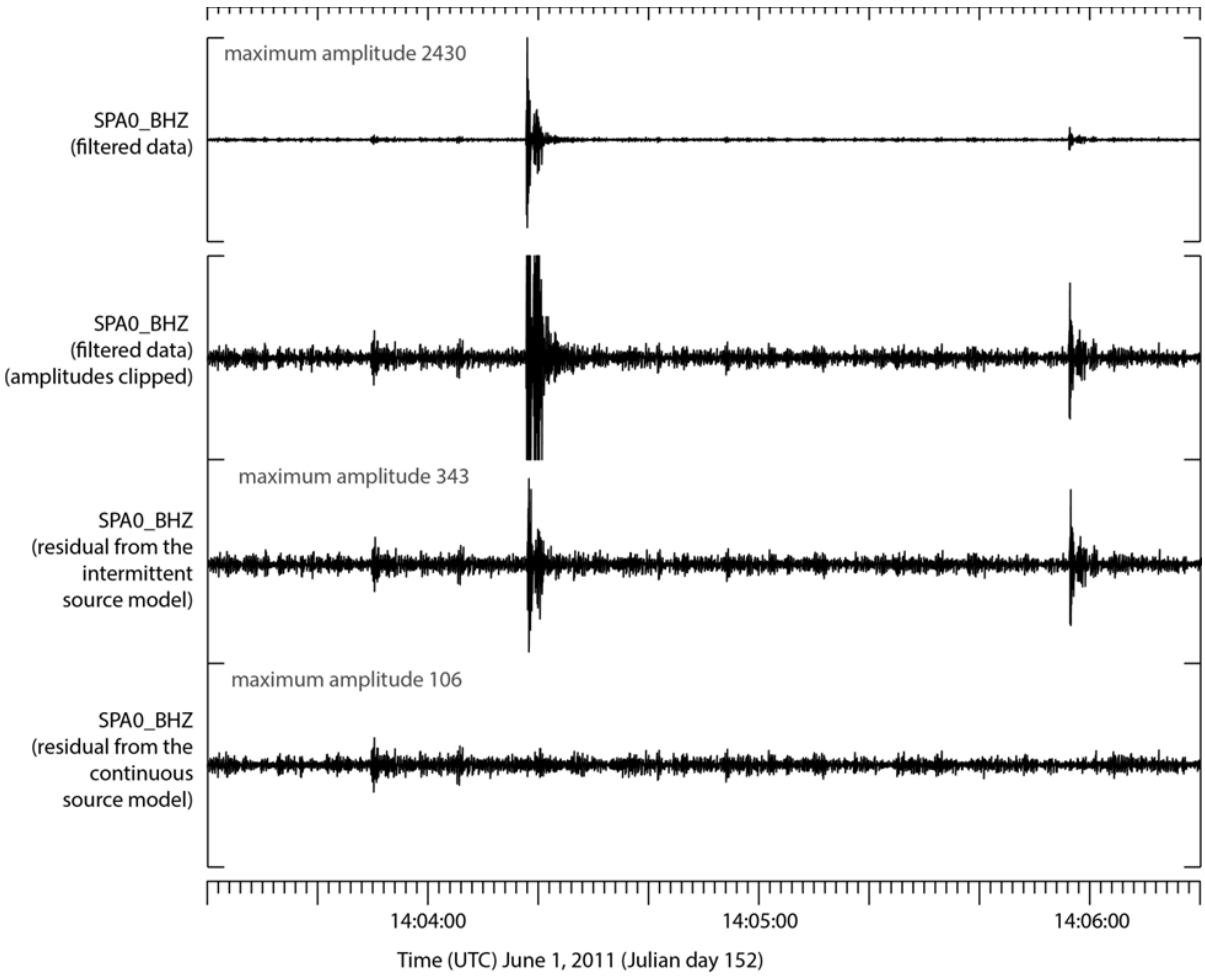

Figure 8. A 3-minute long data segment from the SPA0_BHZ channel. The top trace shows the full dynamic range of the filtered data and the lower 3 traces are all displayed to the same vertical scale. Traces 2, 3, and 4 from the top show respectively the original filtered data at higher gain, the cancellation residual from the intermittent-source algorithm, and the residual from the continuous-source algorithm.

The next 3 figures provide additional information on the necessity of the post-processing correction to the continuous-source canceller. These figures show a detailed section (1200 seconds) of the stream and various cancellation residuals. Figure 9 displays the results of cancellation using the 9 vertical elements of the SPITS array. The design process for the cancellation algorithms resulted in two clusters of interference waveforms; one was represented with a subspace template of dimension 2 and the other of dimension 3. The top trace of the figure displays the original data from channel SPA0 BHZ. It contains a number of icequakes and a large regional event waveform.

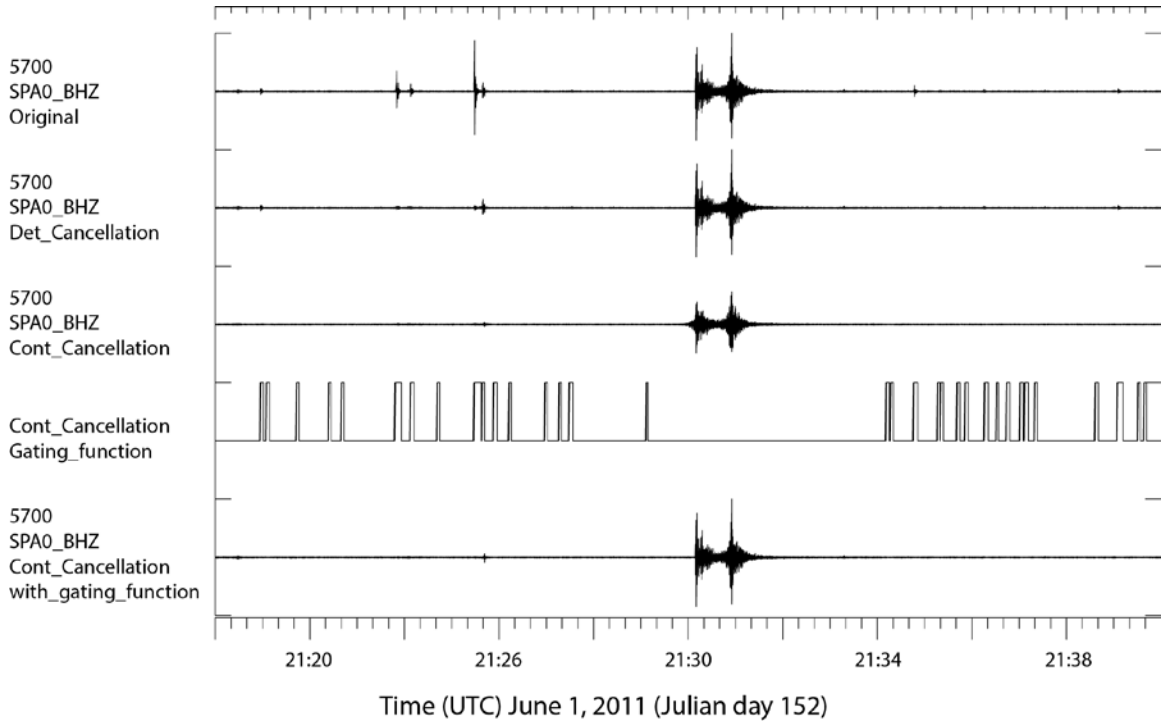

Figure 9. Cancellation options for the 9-element array of SPITS vertical sensors show that it is possible to suppress icequakes without distorting desired regional signals. The top trace is an original filtered signal. The second and third traces are cancellation residuals for the intermittent-source and continuous-source algorithms respectively. The fourth trace is the gating function generated by the post-processor, and the last trace is the result of correcting the continuous-source residual with the gating function.

The second trace is the residual from the intermittent-source canceller (also called detection canceller) using the two interference source representations of dimensions 2 and 3. It largely eliminates or suppresses the icequakes without distorting the regional waveform. The third trace is the residual from the continuous-source canceller, which is better at suppressing the icequakes, but at the cost of distorting the regional waveform. At this scale, the distortion is most obvious in the loss of amplitude of the P and S phases. The fourth trace is the gating function generated by the post-processor indicating where the cancellation residual should be stitched into the original stream. The final trace is the corrected continuous-source residual, showing that the gating function preserves the regional waveform while allowing cancellation of the icequakes. This residual is the cleanest of the cancellation options.

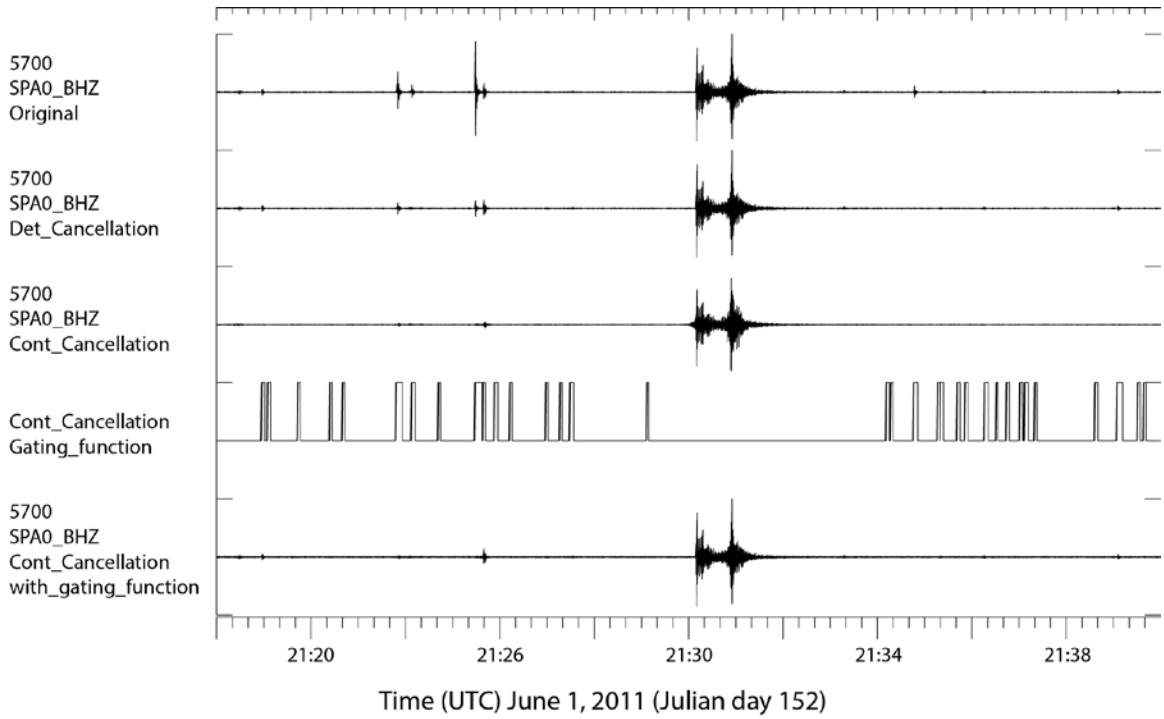

Figure 10. Cancellation residuals for the <u>full array of 21 vertical and horizontal elements of the SPITS array</u> show that larger numbers of channels mitigate distortion. The traces are the equivalents of those in Figure 9 for cancellation performed on 21 channels instead of 9.

Figure 10 provides cancellation residuals from the full 21-element SPITS array for comparison with Figure 9. In this case the cancellers also used two sources with 2- and 3-dimension subspace representations. Since the total number of degrees of freedom in the cancellers is that same as the prior example, but the number of channels is greater, it may be anticipated that constraints on distortion of the regional signal are improved. This is the case, as can be seen in the third trace of the Figure 10. Amplitude loss in the regional event is less pronounced in this figure. It also is apparent that the icequakes are suppressed less effectively by the intermittent-source algorithm (trace 2) when 21 channels are used instead of 9. This effect is a consequence of the larger number of constraints that have to be met in fitting 21 channels of data with 5 degrees of freedom instead of 9 channels. As before, the continuous-source cancellation residual is the best among the three cancellation options.

Detail of the distortion created by the continuous canceller and corrected by the gating function of the post-processor is displayed in Figure 11. The top trace is again the original filtered data from station SPA0 BHZ. The next two traces are the residuals from the continuous-source cancellers acting on 9 channels and 21 channels of data respectively. This distortion manifests itself in amplitude reduction in the waveforms,

but more disturbingly in P-wave precursor artifacts. Distortion definitely is reduced with the addition of more channels to the processed data stream (comparing trace 3 to trace 2). However, the post-processing algorithm restores the regional signal (bottom trace) without adding noticeable artifacts to the rest of the stream. With this observation, the continuous-source canceller, corrected by the post processor appears would be the algorithm of choice were it not for the fact that the continuous canceller is several (~3) orders of magnitude slower than the intermittent-source algorithm.

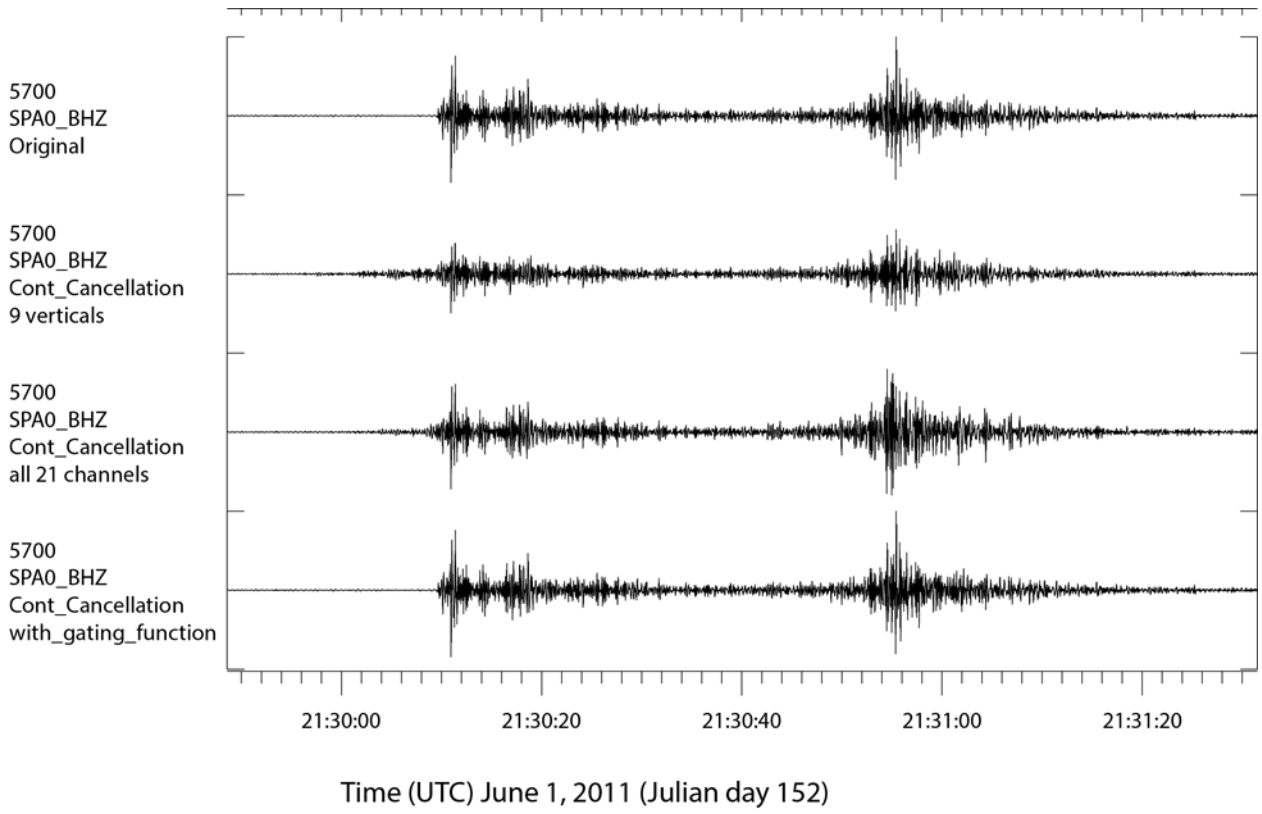

Figure 11. Details of cancellation residuals for the three algorithms around the regional waveform shown in Figures 9 and 10. The top trace is the original SPA0 BHZ data, filtered. The next two traces are the residuals from the continuous-source cancellers using 9 and 21 channels of the SPITS array, respectively. The bottom trace is the residual from the continuous-source canceller corrected by the post-processor. Waveform distortion is avoided in this case.

V. Conclusions

In this chapter we have proposed a collection of algorithms for removing repetitive nuisance transients from array data streams. The algorithms are applicable whenever an array is surrounded by sources producing highly correlated, short-duration transient waveforms. At the moment, we have identified one practical algorithm, the intermittent-source canceller, which can be applied to a fairly large number of sources surrounding an array and is sufficiently fast to be embedded in a monitoring pipeline. We have, in addition, examined a continuous-source canceller, which appears to be more effective in removing nuisance transients, but is computationally prohibitive in its present implementation and requires a post-processing step to prevent it from distorting signals from events of interest. The intermittent-source canceller is asymptotically equivalent to the continuous-source algorithm in the sense that it tends to the latter algorithm as the threshold in the detection step is reduced to zero.

We will report separately (next chapter) on a study quantifying the performance of the intermittent-source canceller, but state here that it appears to be highly effective in reducing the number of interference detections that currently create bottlenecks in pipeline associators. Our expectation is that next-generation pipelines could include adaptive cancellers for each array that would clean the array data streams of locally-generated transient interference. We envision hierarchical adaptive pipelines that would contain, at their base levels, adaptive mini-pipelines designed to detect repetitive nuisance events at each array and generate, then apply, cancellers automatically. The research that remains to be done in order to achieve this vision is for the autonomous identification of repetitive transients in a continuously adapting pipeline preprocessor or lower-level component of the hierarchy. What we have done here is to demonstrate that cancellation is technically feasible given a set of templates that represent a majority of interfering transients.

Appendix A Levinson-Durbin Algorithm for Solution of Block Toeplitz Matrix Equation

The system of equations requiring solution involves a block Toeplitz matrix:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{a}[0] \\ \mathbf{a}[1] \\ \vdots \\ \mathbf{a}[k] \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \mathbf{C}[k] \end{bmatrix} \tag{A.1}$$

The solution proceeds by recursion. Suppose we have the solution for the $(k-1)^{st}$ order problem:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_{k-1} \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_{k-1}^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{a}^{k-1}[0] \\ \mathbf{a}^{k-1}[1] \\ \vdots \\ \mathbf{a}^{k-1}[k-1] \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \mathbf{C}[k-1] \end{bmatrix} \tag{A.2}$$

We construct the solution to the k^{th} order problem given that of the $(k-1)^{st}$ order problem. Note that:

$$\begin{bmatrix} & [\mathbf{R}]_{k-1} & & \mathbf{R}_k \\ & & & \mathbf{R}_{k-1} \\ & & & \vdots \\ \mathbf{R}_k^T & \mathbf{R}_{k-1}^T & \cdots & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{a}^k[0] \\ \mathbf{a}^k[1] \\ \vdots \\ \mathbf{a}^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \mathbf{C}[k] \end{bmatrix} \tag{A.3}$$

and:

$$\begin{bmatrix} \mathbf{R}_0 & \cdots & \mathbf{R}_{k-1} & \mathbf{R}_k \\ \vdots & & & \\ \mathbf{R}_{k-1}^T & & [\mathbf{R}]_{k-1} & \\ \mathbf{R}_k^T & & & \end{bmatrix} \begin{bmatrix} \mathbf{a}^k[0] \\ \mathbf{a}^k[1] \\ \vdots \\ \mathbf{a}^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{C}[0] \\ \mathbf{C}[1] \\ \vdots \\ \mathbf{C}[k] \end{bmatrix} \tag{A.4}$$

where:

$$[\mathbf{R}]_{k-1} = \begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_{k-1} \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_{k-1}^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \tag{A.5}$$

The solution proceeds by carrying along recursions for two special cases in addition to the main problem of (A.1):

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha}^k[0] \\ \boldsymbol{\alpha}^k[1] \\ \vdots \\ \boldsymbol{\alpha}^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \tag{A.6}$$

and:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\beta}^k[0] \\ \boldsymbol{\beta}^k[1] \\ \vdots \\ \boldsymbol{\beta}^k[k] \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} \tag{A.7}$$

By inspection, from (A.3) and (A.6) we have:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha}^{k-1}[0] \\ \boldsymbol{\alpha}^{k-1}[1] \\ \vdots \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \boldsymbol{\gamma}_k \end{bmatrix} \tag{A.8}$$

with:

$$\boldsymbol{\gamma}_k = \sum_{i=0}^{k-1} \mathbf{R}_{k-i}^T \; \boldsymbol{\alpha}^{k-1}[i] \tag{A.9}$$

From (A.4) and (A.7), we have a similar relation:

$$\begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_1 & \cdots & \mathbf{R}_k \\ \mathbf{R}_1^T & \mathbf{R}_0 & \mathbf{R}_1 & \vdots \\ \vdots & \mathbf{R}_1^T & \ddots & \mathbf{R}_1 \\ \mathbf{R}_k^T & \cdots & \mathbf{R}_1^T & \mathbf{R}_0 \end{bmatrix} \begin{bmatrix} \mathbf{0} \\ \boldsymbol{\beta}^{k-1}[0] \\ \vdots \\ \boldsymbol{\beta}^{k-1}[k-1] \end{bmatrix} = \begin{bmatrix} \boldsymbol{\mu}_k \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} \tag{A.10}$$

with:

$$\boldsymbol{\mu}_k = \sum_{i=0}^{k-1} \mathbf{R}_{i+1} \; \boldsymbol{\beta}^{k-1}[i] \tag{A.11}$$

We need to find matrices $\boldsymbol{\theta}_k$ and $\boldsymbol{\varphi}_k$ s.t.

$$\begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \boldsymbol{\gamma}_k \end{bmatrix} \boldsymbol{\theta}_k + \begin{bmatrix} \boldsymbol{\mu}_k \\ \mathbf{0} \\ \vdots \\ \mathbf{I} \end{bmatrix} \boldsymbol{\varphi}_k = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \tag{A.12}$$

There are two non-trivial constraints in (A.12):

$$\begin{aligned} \boldsymbol{\theta}_k + \boldsymbol{\mu}_k \boldsymbol{\varphi}_k &= \mathbf{I} \\ \boldsymbol{\gamma}_k \boldsymbol{\theta}_k + \boldsymbol{\varphi}_k &= \mathbf{0} \end{aligned} \tag{A.13}$$

Solving for $\boldsymbol{\theta}_k$:

$$\boldsymbol{\theta}_k - \boldsymbol{\mu}_k\boldsymbol{\gamma}_k\boldsymbol{\theta}_k = \mathbf{I} \tag{A.14}$$

$$\boldsymbol{\theta}_k = (\mathbf{I} - \boldsymbol{\mu}_k\boldsymbol{\gamma}_k)^{-1} \tag{A.15}$$

Solving for $\boldsymbol{\varphi}_k$:

$$\boldsymbol{\varphi}_k = -\boldsymbol{\gamma}_k(\mathbf{I} - \boldsymbol{\mu}_k\boldsymbol{\gamma}_k)^{-1} \tag{A.16}$$

Thus:

$$\begin{bmatrix}\boldsymbol{\alpha}^k[0]\\ \boldsymbol{\alpha}^k[1]\\ \vdots\\ \boldsymbol{\alpha}^k[k]\end{bmatrix} = \begin{bmatrix}\boldsymbol{\alpha}^{k-1}[0]\\ \boldsymbol{\alpha}^{k-1}[1]\\ \vdots\\ \mathbf{0}\end{bmatrix}(\mathbf{I} - \boldsymbol{\mu}_k\boldsymbol{\gamma}_k)^{-1} - \begin{bmatrix}\mathbf{0}\\ \boldsymbol{\beta}^{k-1}[0]\\ \vdots\\ \boldsymbol{\beta}^{k-1}[k-1]\end{bmatrix}\boldsymbol{\gamma}_k(\mathbf{I} - \boldsymbol{\mu}_k\boldsymbol{\gamma}_k)^{-1} \tag{A.17}$$

Similarly, for the $\boldsymbol{\beta}^k[i]$, we need to find matrices $\acute{\boldsymbol{\theta}}_k$ and $\acute{\boldsymbol{\varphi}}_k$ s.t.

$$\begin{bmatrix}\mathbf{I}\\ \mathbf{0}\\ \vdots\\ \boldsymbol{\gamma}_k\end{bmatrix}\acute{\boldsymbol{\theta}}_k + \begin{bmatrix}\boldsymbol{\mu}_k\\ \mathbf{0}\\ \vdots\\ \mathbf{I}\end{bmatrix}\acute{\boldsymbol{\varphi}}_k = \begin{bmatrix}\mathbf{0}\\ \mathbf{0}\\ \vdots\\ \mathbf{I}\end{bmatrix} \tag{A.18}$$

Again there are two non-trivial constraints:

$$\begin{aligned}\acute{\boldsymbol{\theta}}_k + \boldsymbol{\mu}_k\acute{\boldsymbol{\varphi}}_k &= \mathbf{0}\\ \boldsymbol{\gamma}_k\acute{\boldsymbol{\theta}}_k + \acute{\boldsymbol{\varphi}}_k &= \mathbf{I}\end{aligned} \tag{A.19}$$

With the result:

$$\acute{\boldsymbol{\varphi}}_k = (\mathbf{I} - \boldsymbol{\gamma}_k\boldsymbol{\mu}_k)^{-1} \tag{A.20}$$

$$\acute{\boldsymbol{\theta}}_k = -\boldsymbol{\mu}_k(\mathbf{I} - \boldsymbol{\gamma}_k\boldsymbol{\mu}_k)^{-1} \tag{A.21}$$

Thus:

$$\begin{bmatrix}\boldsymbol{\beta}^k[0]\\ \boldsymbol{\beta}^k[1]\\ \vdots\\ \boldsymbol{\beta}^k[k]\end{bmatrix} = -\begin{bmatrix}\boldsymbol{\alpha}^{k-1}[0]\\ \boldsymbol{\alpha}^{k-1}[1]\\ \vdots\\ \mathbf{0}\end{bmatrix}\boldsymbol{\mu}_k(\mathbf{I}-\boldsymbol{\gamma}_k\boldsymbol{\mu}_k)^{-1} + \begin{bmatrix}\mathbf{0}\\ \boldsymbol{\beta}^{k-1}[0]\\ \vdots\\ \boldsymbol{\beta}^{k-1}[k-1]\end{bmatrix}(\mathbf{I}-\boldsymbol{\gamma}_k\boldsymbol{\mu}_k)^{-1} \quad \text{(A.22)}$$

Returning to the solution of the original problem, we try:

$$\begin{bmatrix}[\mathbf{R}]_{k-1} & \mathbf{R}_k\\ & \mathbf{R}_{k-1}\\ & \vdots\\ \mathbf{R}_k^T \quad \mathbf{R}_{k-1}^T \quad \cdots & \mathbf{R}_0\end{bmatrix}\begin{bmatrix}\mathbf{a}^{k-1}[0]\\ \mathbf{a}^{k-1}[1]\\ \vdots\\ \mathbf{0}\end{bmatrix} = \begin{bmatrix}\mathbf{C}[0]\\ \mathbf{C}[1]\\ \vdots\\ \boldsymbol{\varepsilon}_k\end{bmatrix} \quad \text{(A.23)}$$

where:

$$\boldsymbol{\varepsilon}_k = \sum_{i=0}^{k-1}\mathbf{R}_{k-i}^T\ \mathbf{a}^{k-1}[i] \quad \text{(A.24)}$$

To update the $\mathbf{a}^k[i]$, we need:

$$\begin{bmatrix}\mathbf{C}[0]\\ \mathbf{C}[1]\\ \vdots\\ \boldsymbol{\varepsilon}_k\end{bmatrix} + \begin{bmatrix}\mathbf{0}\\ \mathbf{0}\\ \vdots\\ \mathbf{I}\end{bmatrix}(\mathbf{C}[k]-\boldsymbol{\varepsilon}_k) = \begin{bmatrix}\mathbf{C}[0]\\ \mathbf{C}[1]\\ \vdots\\ \mathbf{C}[k]\end{bmatrix} \quad \text{(A.25)}$$

By inspection:

$$\begin{bmatrix}\mathbf{a}^k[0]\\ \mathbf{a}^k[1]\\ \vdots\\ \mathbf{a}^k[k]\end{bmatrix} = \begin{bmatrix}\mathbf{a}^{k-1}[0]\\ \mathbf{a}^{k-1}[1]\\ \vdots\\ \mathbf{0}\end{bmatrix} + \begin{bmatrix}\boldsymbol{\beta}^k[0]\\ \boldsymbol{\beta}^k[1]\\ \vdots\\ \boldsymbol{\beta}^k[k]\end{bmatrix}(\mathbf{C}[k]-\boldsymbol{\varepsilon}_k) \quad \text{(A.26)}$$

3.2 A Revised Detection Framework Architecture

We significantly revised the data handling architecture in adapting the detection framework [Harris and Dodge, 2011; Kværna et al., 2014] to perform signal cancellation. In part this was because the previous strategy of sharing data buffers among multiple objects made it more difficult to implement a process that would be modifying those data buffers. More importantly, as we move towards development of a detection framework that can scale out a cluster of computers, we can no longer take advantage of shared memory. In the new architecture, the various computational units no longer have any shared state. Instead, they pass data messages to one another. The framework is still implemented as a conventional Java application that runs within a single JVM, but we plan to begin an implementation based on Apache Spark this summer.

Figure 12 shows the flow of data through the current architecture. Data is produced by the StreamServer which is responsible for assembling data into blocks of the required size. The StreamServer produces objects of type StreamSegment which contain the time series data, and the necessary metadata to allow correct routing. In the current architecture, the StreamSegment objects are passed through a BlockingQueue to an object called the StreamProcessor. Within the StreamProcessor 0-N StreamModifier objects may consume a StreamSegment and emit a different StreamSegment. Currently, there are two possible StreamModifiers (DownSampler and SegmentedCancellor). The final modified StreamSegment is passed to a StreamTransformer which emits a TransformedStreamSegment. This is a StreamSegment that contains in addition a DFT of itself.

At this point, a parallel section is entered. Within this section, for each Detector:

1. The TransformedStreamSegment is used to create a DetectionStatistic
2. The DetectionStatistic is consumed by a DetectionStatisticScanner which emits
3. A Collection of TriggerData.

At this point, the threads join and all the TriggerData is ingested by a TriggerProcessor. The TriggerProcessor archives the triggers and produces a collection of Trigger objects which contain additional state created by the TriggerProcessor.

The Triggers are then passed to the DetectionProcessor which applies the necessary rules to promote some triggers to Detections. Any such Detections are then handed off to a DetectorCreator and to a DetectionArchiver.

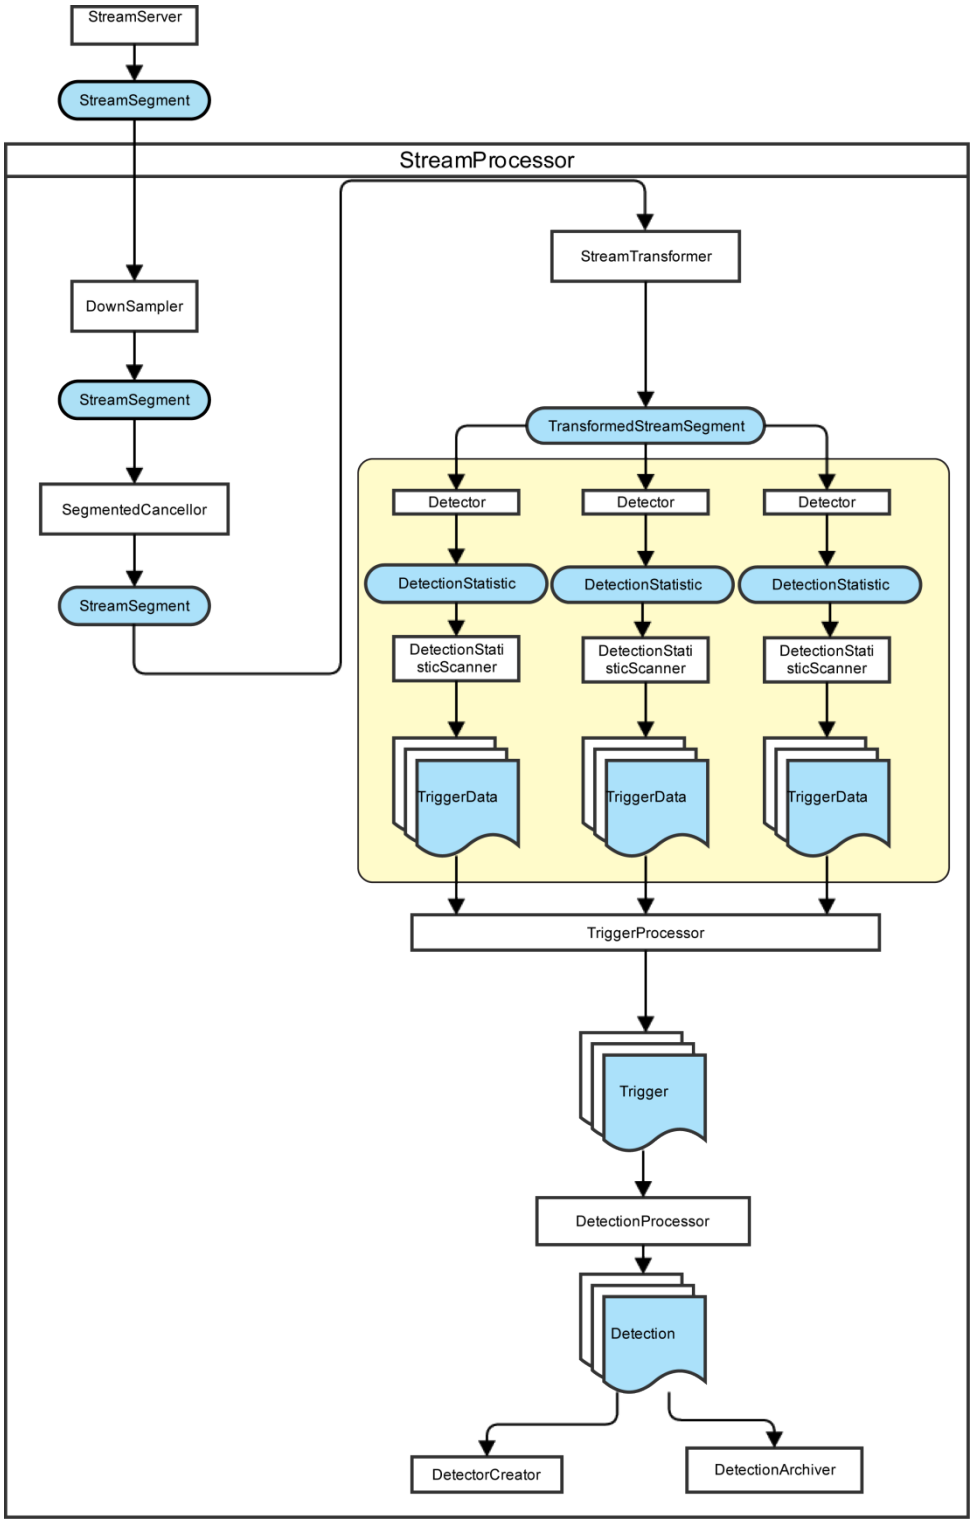

Figure 12. Flow of data through the revised detection framework architecture.

Test of the Signal Cancellation Component in the Detection Framework

To test the cancellation option, the framework was run against a 3-day segment (day 2011151 through day 2011153) of vertical-component data from the SPITS array. During this time interval intense icequake activity had occurred, with hundreds of such signals visible at SPITS. Figure 13 shows the SPA0-BHZ channel for the time period with an inset showing the form of a typical signal.

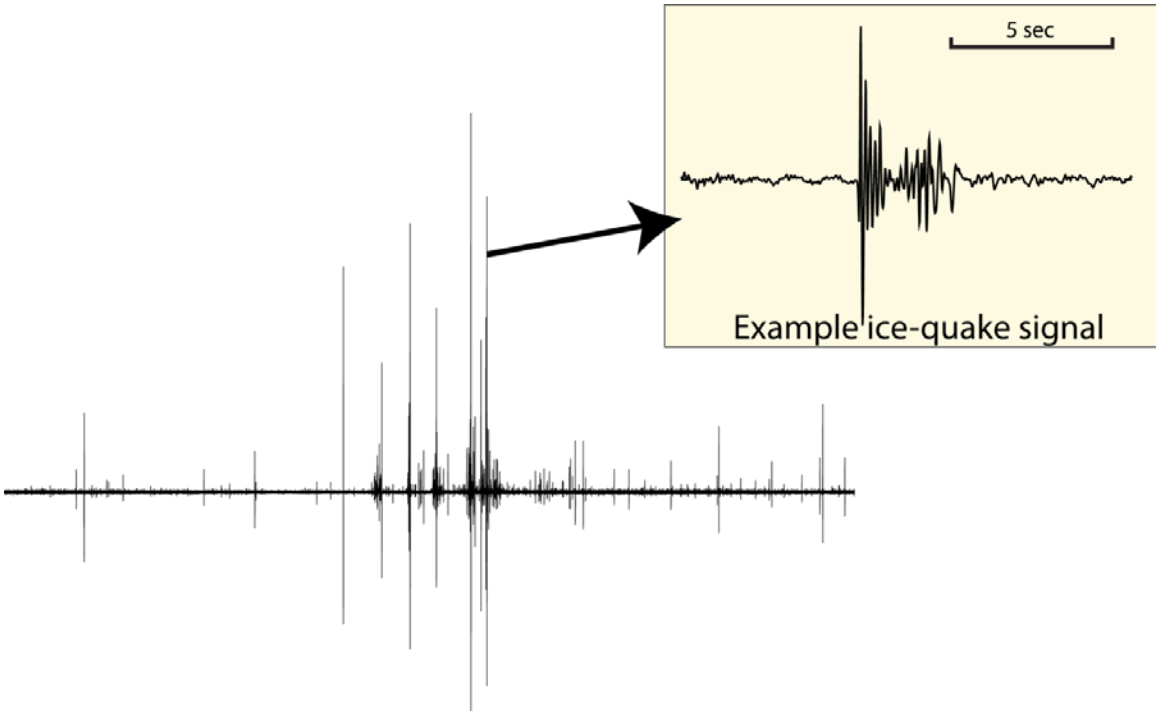

Figure 13. 3 days of SPAO-BHZ data which is dominated by signals from nearby icequakes.

For reference, we first ran the detection framework without cancellation enabled. The framework was configured to use all the vertical-component channels filtered into the band 2-8 Hz and with a single STA/LTA spawning detector. New templates created by the framework were constrained to be between 10 and 30 seconds in length, and the detection threshold of the correlation detectors was 0.6.

This run produced 49 correlation detectors and a total of 291 detections. Two of the correlators accounted for 242 of the detections, and all of those signals appear to be due to icequakes. These detections are shown in Figure 14.

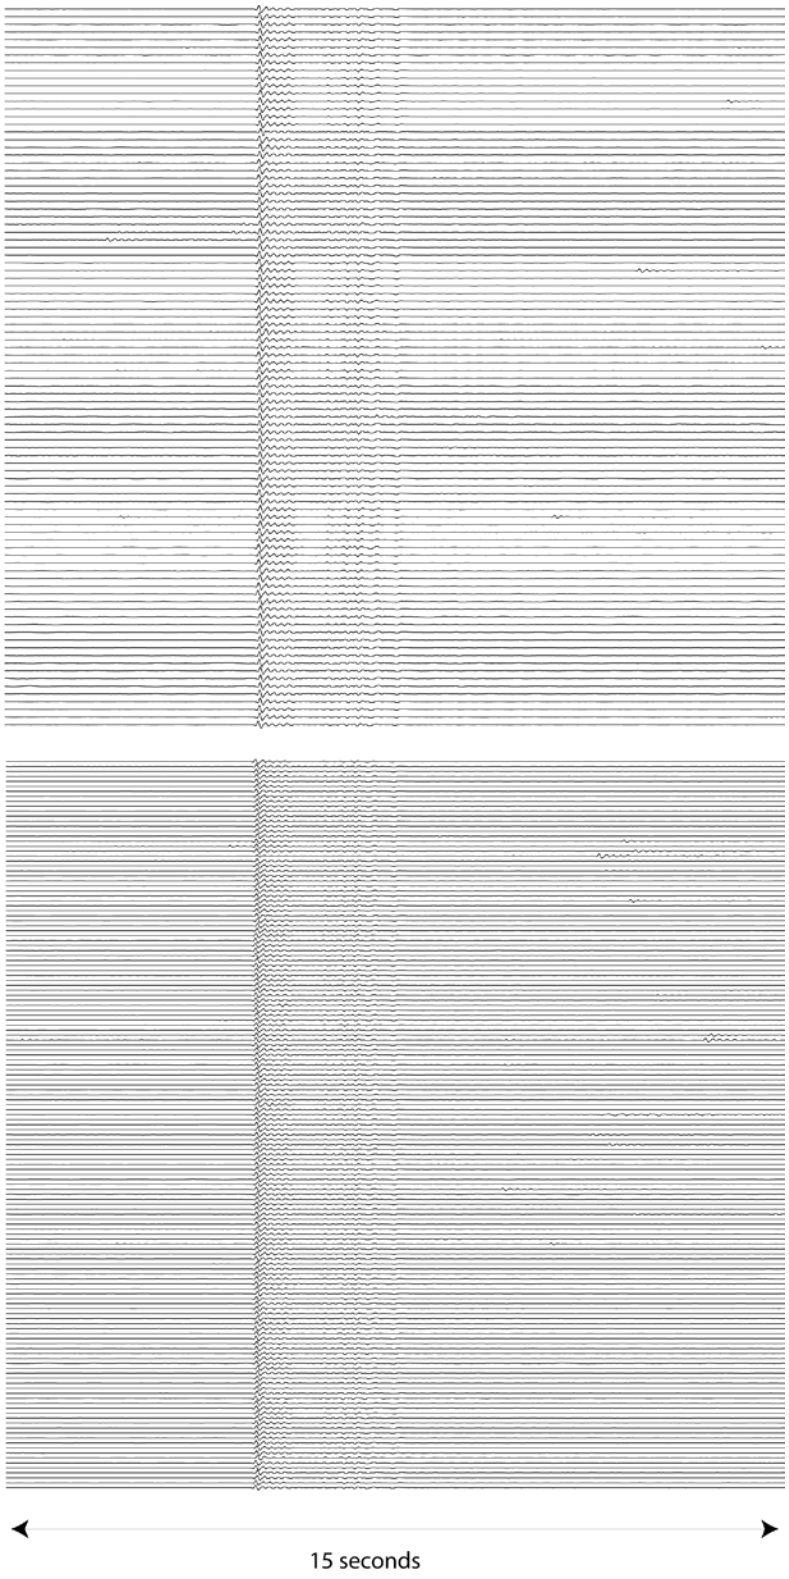

Figure 14. (Top) 94 detections produced by detector 92532 and (bottom) 148 detections from detector 92541 produced during the first run of the framework.

The remaining detections are distributed among the other detectors, with most having only a single detection. These appear to be mostly local earthquakes although a few apparent icequake signals are present as well. Figure 15 shows all these detections.

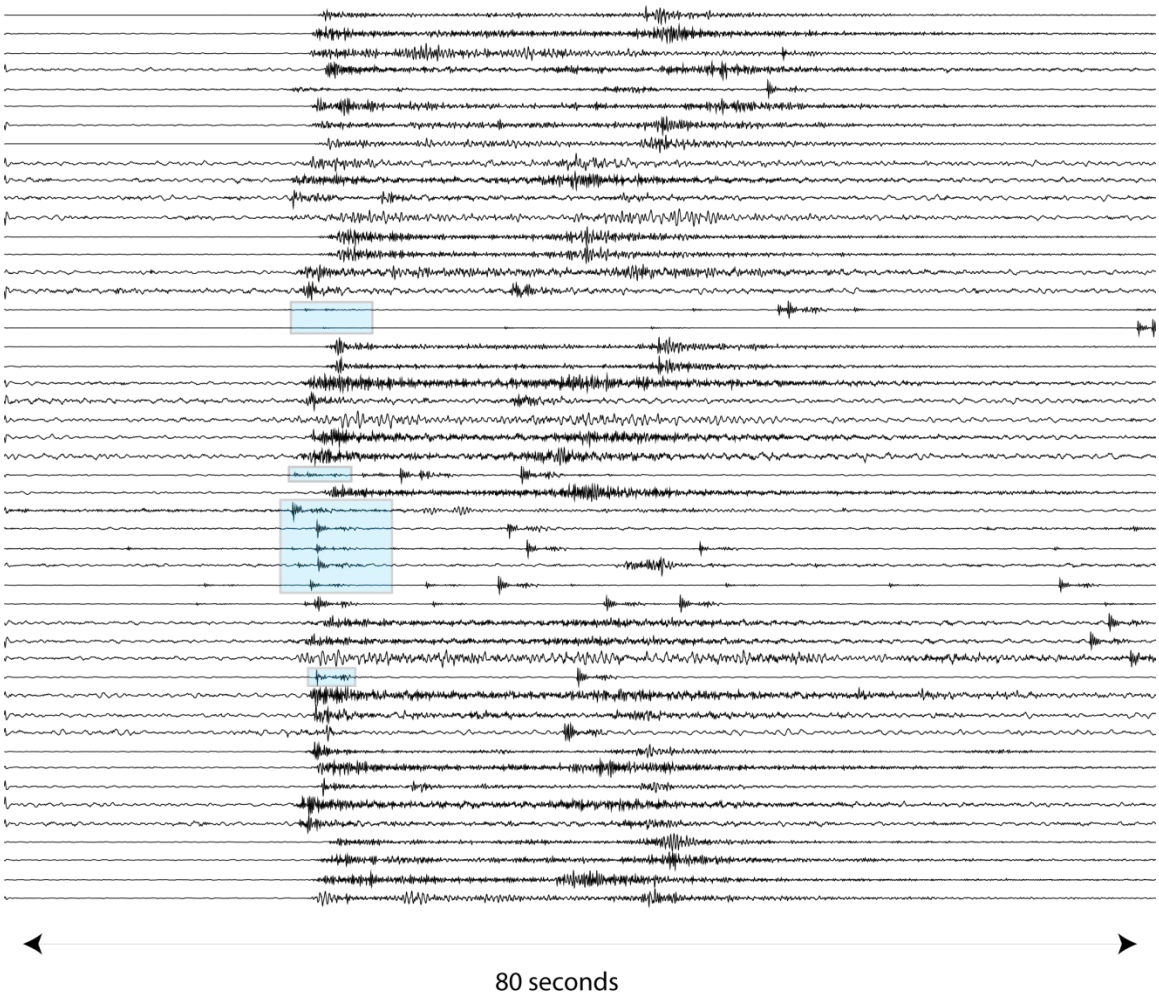

Figure 15. The 49 detections produced by the remaining detectors. The blue-shaded signals also appear to be icequake detections, but, perhaps because of interfering signals, were assigned to unique detectors.

Cancellation run (RUNID 365)

The discrete segmented cancellation processor provided by Deschutes Signal processing uses a two-step algorithm. First the signals to be cancelled are detected using subspace detectors. At each detection point, a best-fit to the templates is constructed and then subtracted from the signal. This process is iterated until no more detections are found.

The detectors used by the canceller are produced by the cancellation software given an input set of nuisance signals. This processing is controlled by a number of parameters, some of which are shown below.

Table 1. Parameters controlling detector creation and cancellation behavior as they were set for this experiment.

templateLength	6.0
maxOffset	0.5
clusteringThreshold	0.9
minNumEvents	10
energyCapture	0.9
detectionThreshold	0.4
peakHalfWidth	0.2
simultaneityThreshold	2.0
numberOfIterations	5

For the execution of RUNID 365 the framework was configured identically to that of its previous run, except that the cancellation flag was set, and a group of 272 icequake signals was provided to serve as input for the signal canceller. From these, the software created two cancellation detectors. The effect of the cancellation can be seen by visual inspection of the signals pre- and post-cancellation. This is shown in Figure 16 below.

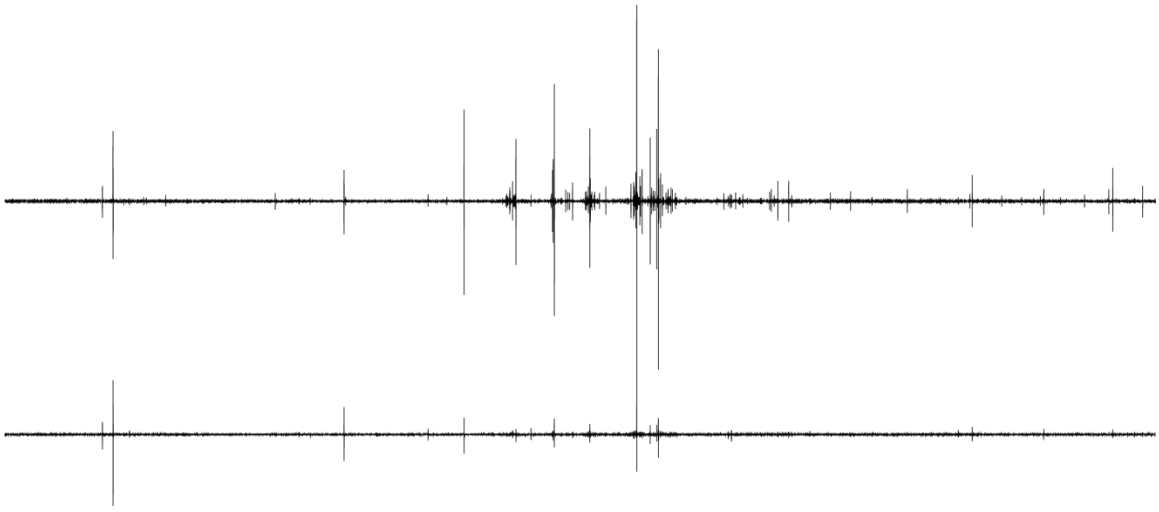

Figure 16. 3 days of SPA0-BHZ data prior to cancellation (top) and after cancellation (bottom). *The traces are plotted at the same scale.*

Comparison of Results

In total, the canceller performed 2047 signal cancellation operations. However, because the cancellation process is iterative, the number of independent icequake signals is smaller. To the nearest second, there were 952 distinct signals cancelled, and to the nearest 10-seconds there were 817 distinct cancellations.

In contrast to the first run of the framework, RUNID 365 produced only 71 detections. Of these 17 were detected by two detectors (92584 and 92594) and are icequake signals. Table 2 shows these detections grouped by detector and associated by time with detections from the first run. All 17 are associated with detections by (RUNID 364) detectors 92541 and 92532, both of which were formed on icequake signals. Visual inspection of the detection segments confirms the signal type.

Table 2. The 17 detections from RUNID 365 that match (nearly) in time with icequake detections from RUNID 364.

364 detectorid	364 detectionid	365 detectorid	365 detectionid	Time difference
92541	195546	92584	195804	-6.6625
92541	195537	92584	195802	-6.6615
92532	195649	92584	195822	5.775
92541	195611	92584	195819	-6.675
92541	195654	92584	195824	-6.6625
92532	195514	92584	195795	5.769445
92541	195550	92584	195806	-6.6615
92541	195567	92584	195811	-6.6625
92541	195763	92594	195842	-4.9
92541	195685	92594	195829	-4.9
92541	195661	92594	195825	-4.8865
92541	195598	92594	195815	-4.9125
92541	195606	92594	195818	-4.9
92541	195604	92594	195816	-4.9
92532	195564	92594	195810	7.526389
92541	195651	92594	195823	-4.9125
92541	195582	92594	195813	-4.9

The cyan-shaded group is by detector 92584 and the green-shaded group is by detector 92594. The first two columns are the detector number and detection-id from the first run. The third and fourth columns show the corresponding information from the second run, and the last column is the difference in detection times.

Given the success of the signal canceller in removing so many other nuisance signals it was curious that these detections still remained. Also, why are the detection time differences so large? Both of these groups of detections were formed by detectors whose template included small icequake signals followed by a very much larger icequake signal near the end of the template. The residual from the canceller was still a detectable signal, but the detailed structure was sufficiently different to change the point at which the STA/LTA detector could exceed its threshold. Figure 17 shows an example from the template event for detector 92584. In this example, two icequake signals are present. In run 1, the detection is declared at the onset of the first signal. But, after cancellation, the first signal is almost entirely removed, so the detection point shifts by several seconds.

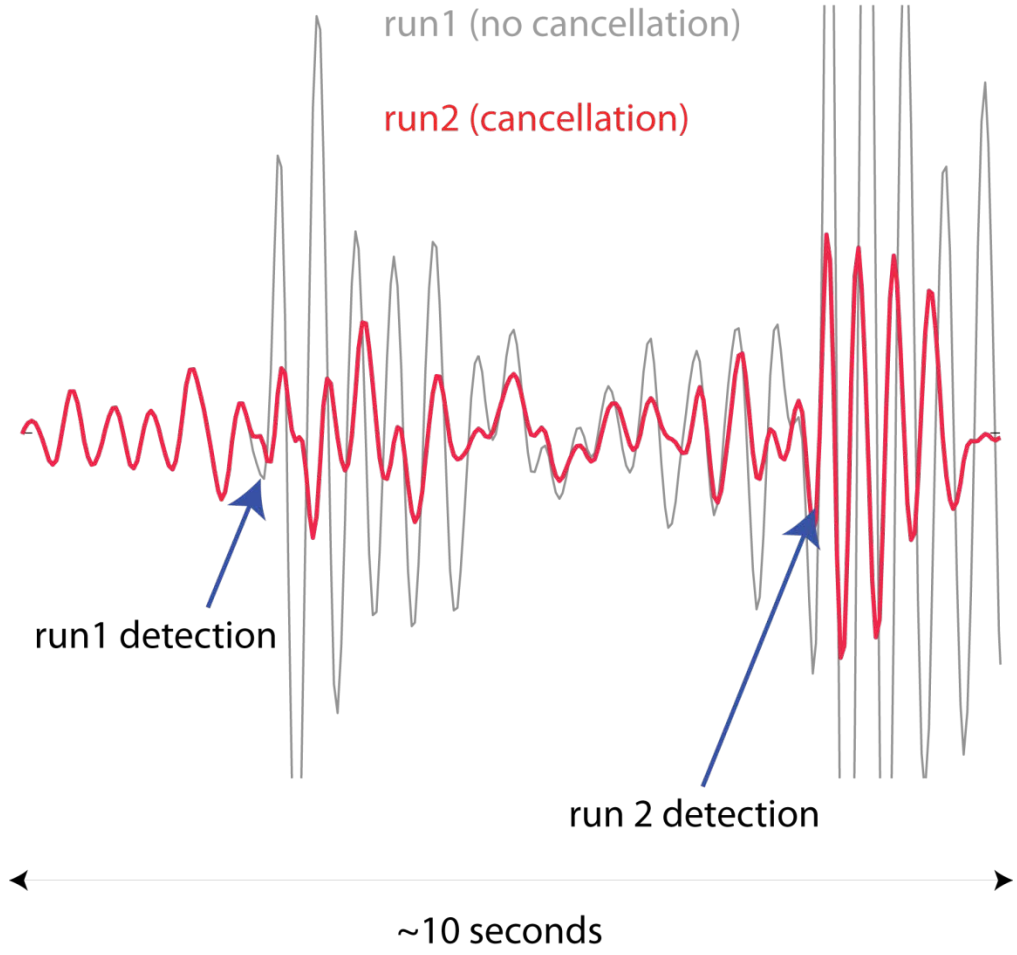

Figure 17. This figure shows how partial cancellation can cause changes in detection time for cases when multiple instances of the nuisance signal are present in close proximity.

Of the remaining 54 detections in RUNID 365, 39 are exact matches in time to non-icequake detections from RUNID 364. There are 11 detections from RUNID 364 that were not seen in RUNID 365. Another 8 of the detections in RUNID 365 are on partially-cancelled signals. There are also 8 detections in RUNID 365 that were not seen in RUNID 364. The complete comparison

of the detection results between runs 364 and 365 is presented in Table 3 below. The unshaded rows summarize the exact matches showing the detectorid and detectionid from each run as well as the detection time.

Table 3. Accounting of all detections in RUNID 364 and RUNID 365 excepting 242 known icequake detections from RUNID 364 detectors 92532 and 92541 and 17 icequake detections from RUNID 365 detectors 92584 and 92594.

92517	195486	92567	195777	2011/05/31 (151) 06:06:14.000
92518	195487	92568	195778	2011/05/31 (151) 07:26:22.000
92519	195488	92569	195779	2011/05/31 (151) 17:36:25.000
92520	195489	92570	195780	2011/05/31 (151) 17:46:47.000
92521	195490	92571	195781	2011/05/31 (151) 17:57:37.000
92522	195491	92572	195782	2011/05/31 (151) 18:25:17.000
92523	195492	92573	195783	2011/05/31 (151) 19:30:49.000
92524	195493	92574	195784	2011/05/31 (151) 21:14:12.000
92525	195494	92575	195785	2011/05/31 (151) 21:47:19.000
92526	195495	92576	195786	2011/06/01 (152) 01:13:47.000
92527	195496	92577	195787	2011/06/01 (152) 01:32:20.000
92528	195497	92578	195788	2011/06/01 (152) 02:11:27.000
92529	195498	92579	195789	2011/06/01 (152) 02:29:58.000
92530	195499	92580	195790	2011/06/01 (152) 04:30:49.000
92531	195500	92581	195791	2011/06/01 (152) 05:54:13.000
92529	195501	92579	195792	2011/06/01 (152) 06:31:23.000
92533	195505	CANCELLED		2011/06/01 (152) 07:34:35.000
Detection on partial cancellation		92582	195793	2011/06/01 (152) 07:38:15.000
92534	195508	CANCELLED		2011/06/01 (152) 07:47:02.000
Detection on partial cancellation		92583	195794	2011/06/01 (152) 07:50:42.000
92535	195522	92585	195796	2011/06/01 (152) 08:56:50.000

Table 3 (continued). Accounting of all detections in RUNID 364 and RUNID 365 excepting 242 known icequake detections from RUNID 364 detectors 92532 and 92541 and 17 icequake detections from RUNID 365 detectors 92584 and 92594.				
92536	195523	92586	195797	2011/06/01 (152) 09:15:11.000
92537	195524	92587	195798	2011/06/01 (152) 09:16:42.000
92538	195525	92588	195799	2011/06/01 (152) 09:25:31.000
92539	195527	92589	195800	2011/06/01 (152) 09:34:06.000
92540	195528	92590	195801	2011/06/01 (152) 10:00:19.000
Correlation detection 2 of 3 (Large icequake signal masked it in run 1.)		92585	195803	2011/06/01 (152) 11:07:52.000
92542	195549	92591	195805	2011/06/01 (152) 12:09:55.000
Detection on partial cancellation		92592	195807	2011/06/01 (152) 12:28:16.000
Detection on partial cancellation		92593	195809	2011/06/01 (152) 12:37:06.000
Correlation detection based on partially-cancelled signal		92592	195808	2011/06/01 (152) 12:37:52.000
Detection of signal masked by icequake signal?		92595	195812	2011/06/01 (152) 13:00:06.000
92543	195586	92596	195814	2011/06/01 (152) 14:23:09.000
92544	195587	CANCELLED		2011/06/01 (152) 14:26:39.000
92545	195594	CANCELLED		2011/06/01 (152) 15:01:28.000
92546	195596	CANCELLED		2011/06/01 (152) 15:04:19.000
92547	195599	CANCELLED		2011/06/01 (152) 15:16:56.000
Detection on partial cancellation		92597	195817	2011/06/01 (152) 15:28:38.000
Detection on partial cancellation		92598	195820	2011/06/01 (152) 15:33:58.000
92548	195624	CANCELLED		2011/06/01 (152) 15:44:04.000
Signal onset masked by icequake signal		92599	195821	2011/06/01 (152) 16:35:34.000
92549	195659	CANCELLED		2011/06/01 (152) 16:58:31.000
92550	195668	92600	195826	2011/06/01 (152) 17:22:20.000
92551	195670	This is a 2nd detection on the same signal that only occurred in run 1		2011/06/01 (152) 17:22:21.000
92552	195671	92601	195827	2011/06/01 (152) 17:26:45.000

Table 3 (continued). Accounting of all detections in RUNID 364 and RUNID 365 excepting 242 known icequake detections from RUNID 364 detectors 92532 and 92541 and 17 icequake detections from RUNID 365 detectors 92584 and 92594.

Detection on partial cancellation		92602	195828	2011/06/01 (152) 17:35:34.000
Signal preceded by partially-cancelled icequake signal		92603	195830	2011/06/01 (152) 18:04:15.000
92553	195704	CANCELLED		2011/06/01 (152) 19:57:35.000
92554	195708	92604	195831	2011/06/01 (152) 21:08:40.000
Correlation detection 3 of 3 (Large icequake signal in coda caused correlation to fail in run 1)		92579	195832	2011/06/01 (152) 21:45:08.000
92555	195719	92605	195833	2011/06/01 (152) 23:35:32.000
Detection of signal masked by icequake signal?		92606	195834	2011/06/02 (153) 01:05:27.000
Signal preceded by partially-cancelled icequake signal		92607	195835	2011/06/02 (153) 01:17:51.000
92556	195740	CANCELLED		2011/06/02 (153) 07:44:20.000
92557	195746	92608	195836	2011/06/02 (153) 10:00:17.000
Detection of signal masked by icequake signal?		92609	195837	2011/06/02 (153) 12:35:06.000
92558	195752	92611	195839	2011/06/02 (153) 12:39:49.000
92559	195754	92612	195840	2011/06/02 (153) 13:08:28.000
92560	195759	92613	195841	2011/06/02 (153) 16:39:04.000
92561	195765	92614	195843	2011/06/02 (153) 20:30:03.000
92535	195769	92585	195844	2011/06/02 (153) 21:25:06.000
92562	195771	92615	195845	2011/06/02 (153) 22:18:27.000
92563	195772	92616	195846	2011/06/02 (153) 22:25:35.000
92564	195773	92617	195847	2011/06/02 (153) 22:44:32.000
92565	195776	In Last Buffer. Not seen for detection		2011/06/03 (154) 00:02:06.000

The yellow shaded rows present information for cases when a detection occurred in run 364 but there was no corresponding detection in run 365. The green-shaded rows show information for cases when there was a detection in run 365 but no corresponding detection in run 364 and which are probably valid non-icequake signals. The pink-shaded rows are for cases when a detection was produced on a partially-cancelled icequake signal. The cyan-shaded rows are new detections of questionable validity.

Missing Detections in RUNID 365

There were 11 detections produced in RUNID 364 that were not seen in RUNID 365 and that were not in the two groups of known icequake signals. These are shown in yellow in Table 3 and the (un-cancelled) signals are shown below in Figure 18. All but the last of these signals are visually similar to the icequake signal shown in Figure 13, but they appear to have higher noise or other changes in structure sufficient that they did not correlate well enough with the templates of detectors 92532 or 92541 to be included in their detections. However, the multi-rank subspace detectors used by the cancellation processor were able to detect these signals, so they were cancelled sufficiently to not be detected in run 365.

The last signal is a special case. The time stamp of this detection is 2011/06/03 (154) 00:02:06.000. This is 2 minutes past the latest data the framework was requested to process, so it is almost certainly the last buffer of data pushed into the canceller. The framework simply shut down before this buffer was retrieved from the canceller. Although this is a missed detection in the context of this experimental run, in a continuously-operating system there is no "last buffer", so this situation could not arise in practice.

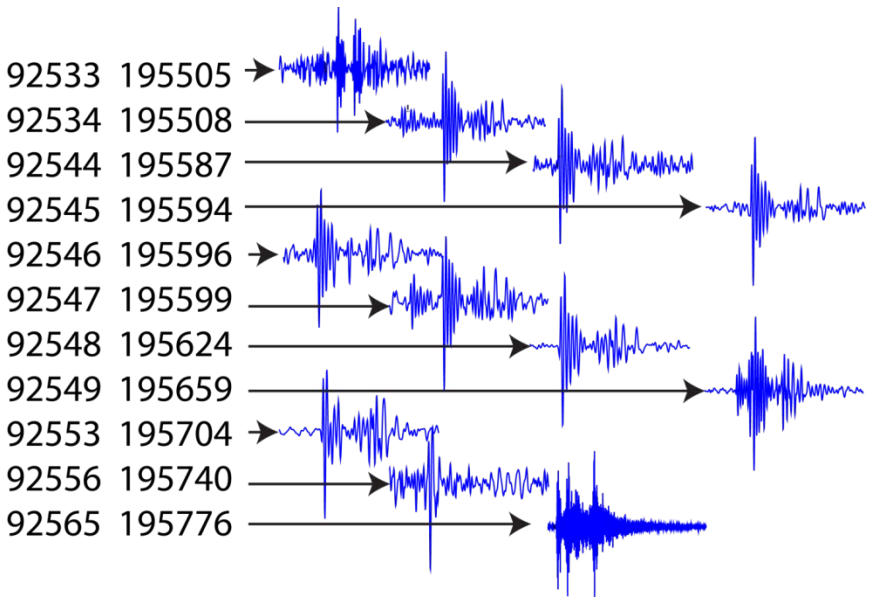

Figure 18. Signals that were detected in the first run of the framework, but which were not detected in the run with cancelling turned on. The numbers are respectively the detector number and detection-id from run1. In all but the last case, the signals appear to be from icequakes, and were removed by the canceller. The last signal was not detected because it occurred in the very last data buffer processed by the system.

Detections on Partial Signal Cancellations

There are 8 cases of unwanted detections in run 365.

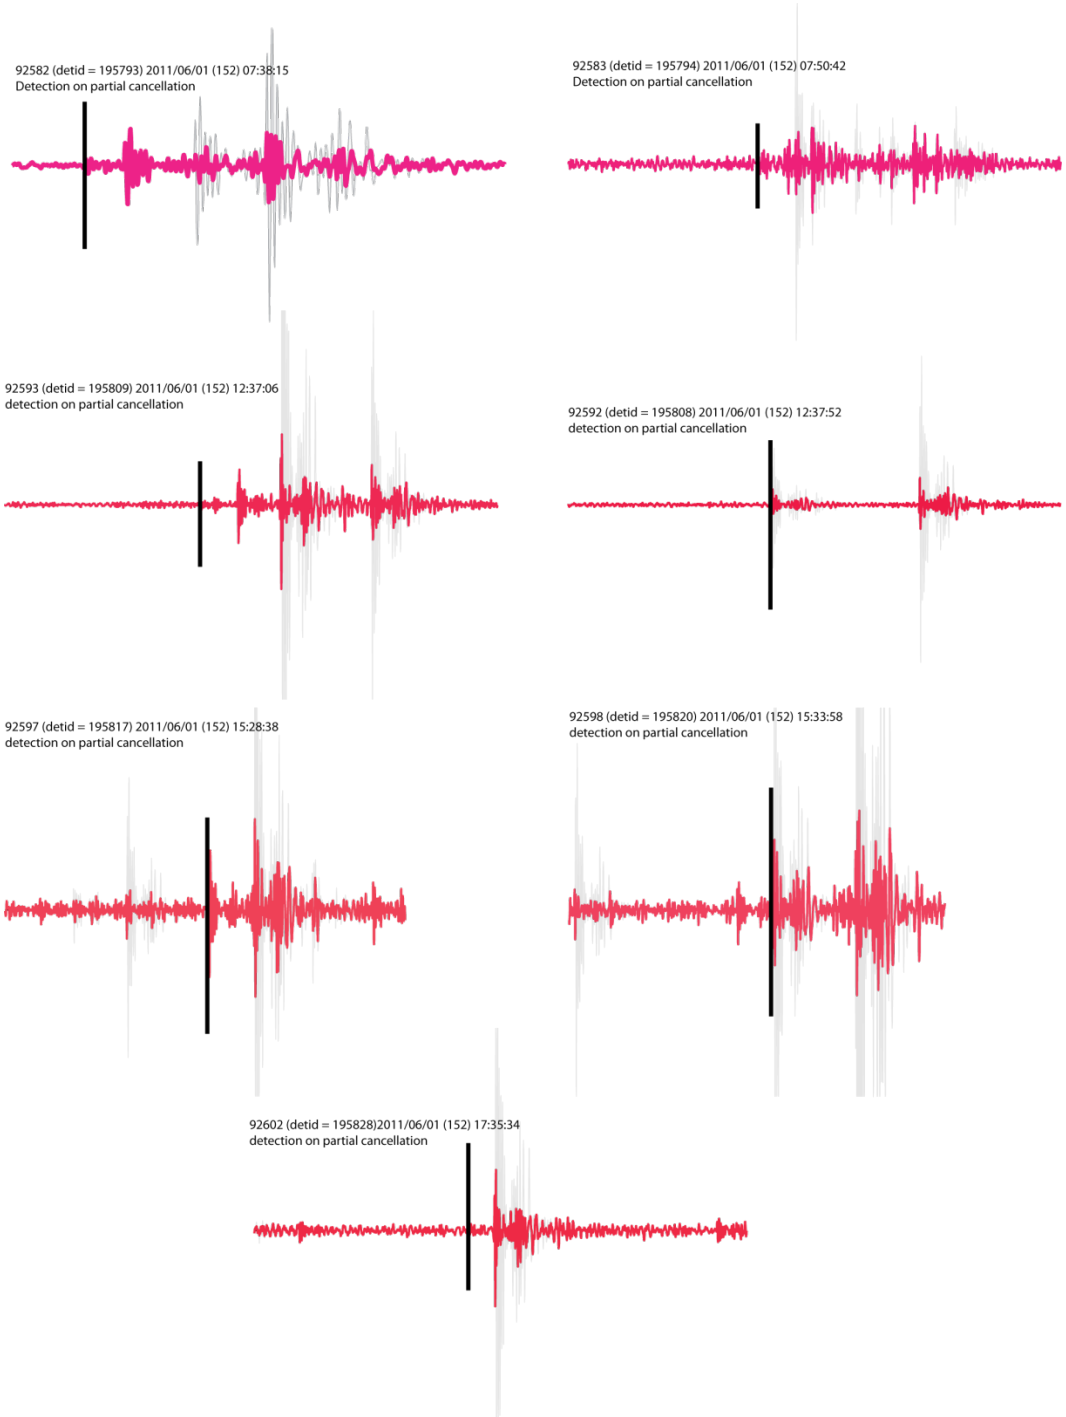

Figure 19. Instances in which the framework produced detections on incompletely cancelled icequake signals. The magenta traces are the post-cancellation signals, and the light gray shows the original signal. The detection time is indicated by the vertical black line.

In these cases, a detection occurred in run 365 with no corresponding detection in run 364, and the detection was on a partially-cancelled icequake signal. These cases are shown in Table 3 as the pink-shaded rows. Figure 19 shows the signals involved. (Only 7 signals are shown. The 8th was a correlation detection on the template of detector 92592.) Figure 19 shows the original signals in light gray and the cancelled signals on which the detections were made in magenta. The reason I think the residual signals are from icequake sources is that they maintain the same phase structure as the un-cancelled signals and have very little energy outside the bounds of the un-cancelled signals.

What makes these detections a little problematic is that the cancellation process and mixing of signals can produce residuals that have some resemblance to local earthquake seismograms. Clearly, an operational system that included cancellation would have to make it very easy for analysts to be able to observe the un-cancelled signal in context to help sort out these (hopefully rare) occurrences.

New detections

There are also 5 cases where a detection was made in run 365 with no corresponding detection in run 364 and in which the detection is likely to be for a valid non-icequake signal. These are shown in Table 3 in green, and the seismograms are shown in Figures 8-10.

In the cases shown in Figure 20 (detections 195821, 195830, 195835) a signal failed to be detected in run 364 because it fell within the blackout interval following an icequake signal detection.

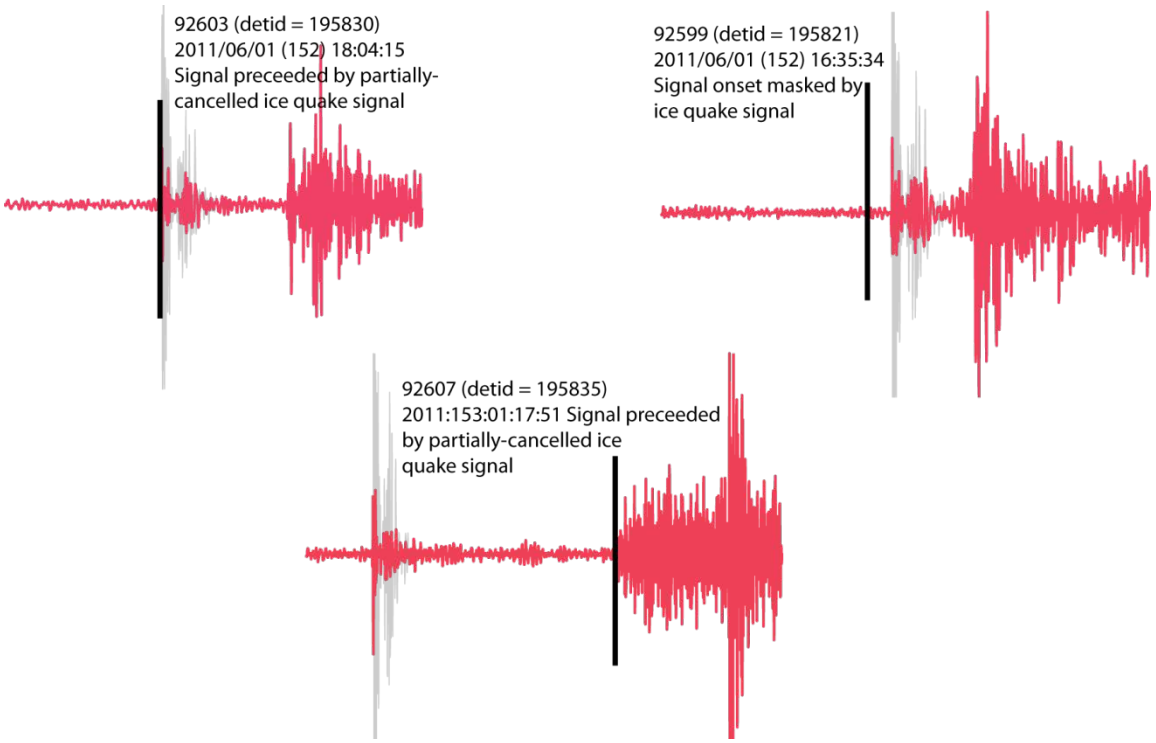

Figure 20. The three cases where a valid signal was not detected in run 365 because the signal onset was in the blackout interval for an icequake detection.

After cancellation, in two cases the detection is still at the onset of the icequake signal, but the detection is by the STA/LTA detector instead of by one of the correlators. In the third case, the cancelled icequake signal trigger was discarded in favor of the trigger on the new signal.

In two cases (detections 195803 and 195832) a potential correlation detection failed because of the presence of a much larger icequake signal. These are shown in Figure 21.

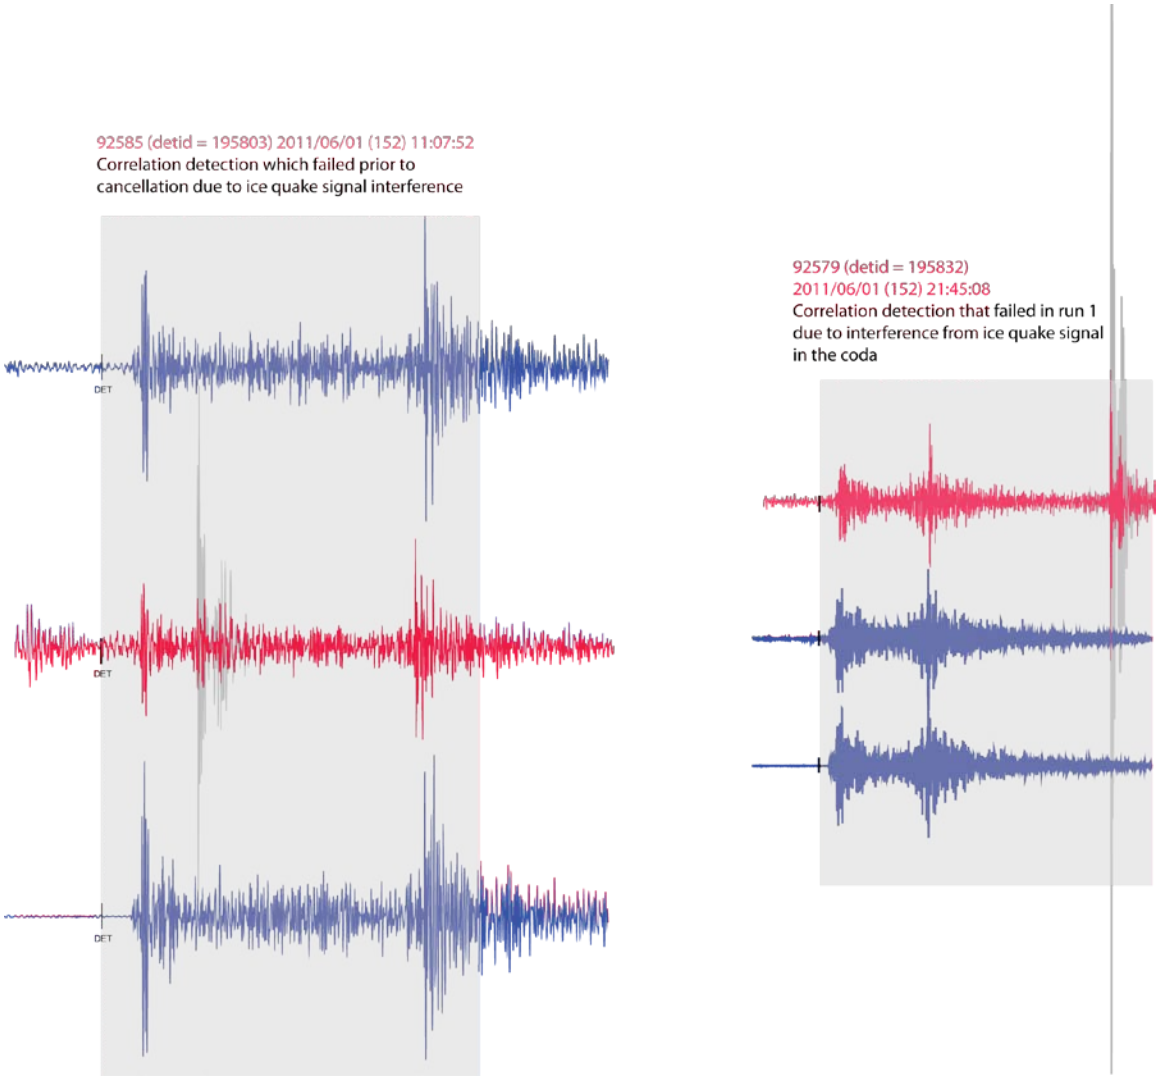

Figure 21. The two cases where a signal failed to be detected by correlation detectors because of the presence of a strong icequake signal.

In both cases, the signals shown in blue were detected by correlators in run 364, but the third signal was missed. In run 365, the cancelled icequake signal was sufficiently attenuated that the third signal was detected.

Finally, there are three possible "new" detections (see Figure 22). The residual signals may be the remnants of the cancelled icequake signal, but they do have a shape very consistent with that of a local earthquake signal. Also, the first (detection 195812) begins shortly before the start of the cancelled icequake signal.

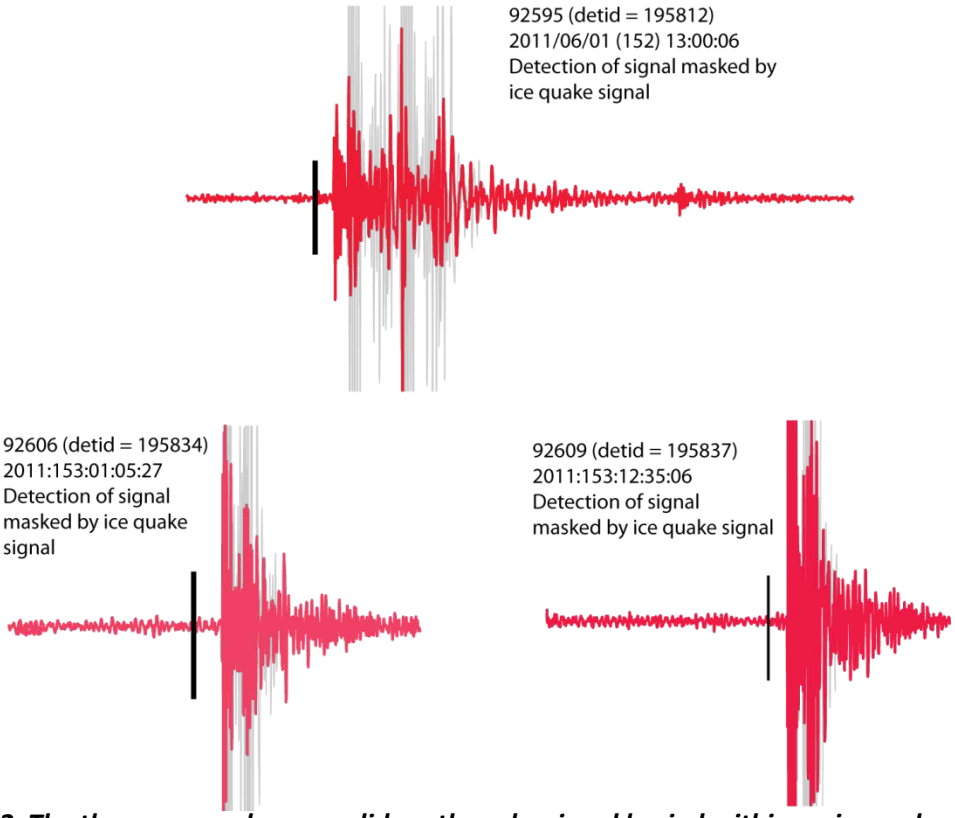

Figure 22. The three cases where a valid earthquake signal buried within an icequake signal may have been detected.

3.3 A Context-Based Procedure for Identification of Aftershocks

Under the existing processing pipeline at the International Data Center (IDC), the difficulty in progressing from the most advanced of the fully automatic event bulletins (SEL3) to the final, analyst-reviewed, event bulletin (REB) increases substantially when an extensive aftershock sequence is underway. This is illustrated in Figure 23. The events in the automatic SEL3 bulletin for October 7, 2005, are displayed in panel (a) together with the corresponding set of analyst-confirmed events displayed in panel (b). The patterns of symbols in these two panels are not too dissimilar. The size of some symbols (event magnitude) changes as the analyst reassesses the arrival times, phase identifications and amplitude measurements. Other symbols are removed completely as the analyst rejects spurious event hypotheses, and others are added as the human eye identifies valid combinations of phase arrivals that were missed, or incorrectly grouped, in the initial phase association. Otherwise, on this fairly typical day, the SEL3 event bulletin provides a reasonably good indicator of how the reviewed REB event bulletin will look.

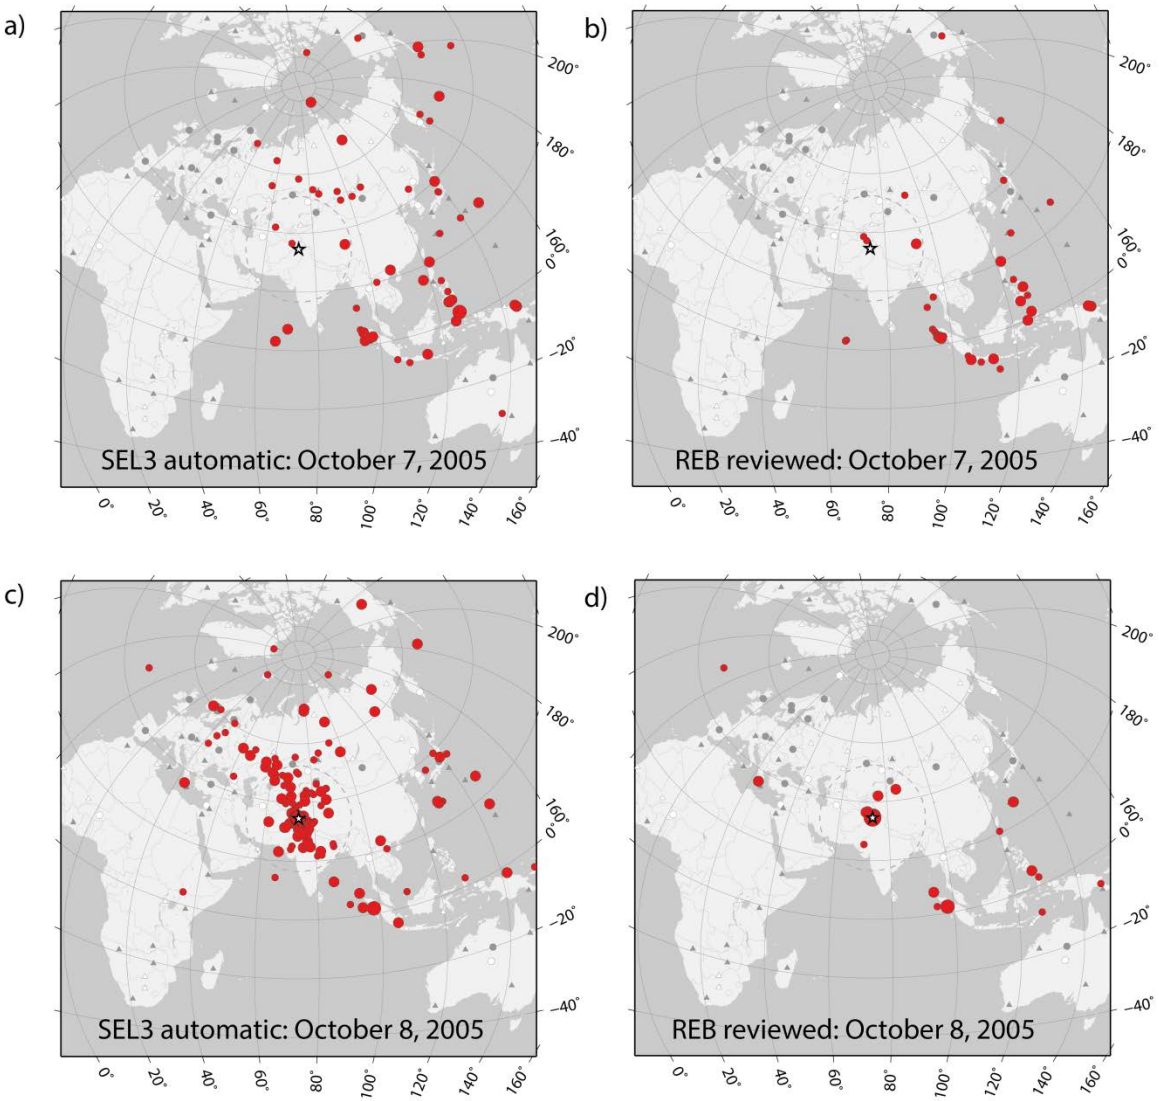

Figure 23. Locations of seismic events (red symbols) in the automatic SEL3 event bulletin at the IDC (a and c) and in the analyst reviewed REB (b and d) on the dates indicated. October 8, 2005, was the date of the magnitude 7.6 Kashmir earthquake (location given by the star) which was followed by many thousands of aftershocks. The locations of IMS array stations are given by circles and the locations of IMS 3-component stations are given by triangles. The stations that were operational at the time are shaded dark and the stations not working at the time are indicated as white symbols.

On October 8, 2005, a magnitude 7.6 earthquake struck the Pakistan-administered region of Kashmir and was followed by several thousands of aftershocks detectable using regional and global seismic networks. Panel (c) of Figure 23 displays the locations of the seismic event hypotheses of the automatic SEL3 bulletin for this day, with the analyst-reviewed event locations are displayed in panel (d). The automatic event location estimates essentially cover the entire Eurasian continent and the human effort

required to progress to the pattern seen in panel (d), with several hundreds of events occurring within a very confined geographical region, is enormous. It is exactly these demands on human resources that we seek to ameliorate in this project.

Alternative procedures for the association of seismic phases and creation of event hypotheses (e.g. Arora et al., 2013) may provide a solution (or partial solution). The other strategy is to be able to identify and classify as accurately as possible the set of events occurring within the aftershock region, using an independent procedure, and to then find all the detections resulting from this set of identified events and then to repeat the event building stage in a new iteration in the absence of these detections.

Pattern detectors (primarily subspace detectors) have been used to great effect for classifying aftershock sequences from somewhat smaller events, e.g. the m_b = 6.5 San Simeon sequence in 2003 (Harris and Dodge, 2011) or the m_b=6.0 Svalbard sequence in 2008 (Junek et al., 2015). However, the dimensions of the source region and subsequent diversity of the waveforms make this problematic for increasingly large events and the 2005 Kashmir sequence was shown by Slinkard et al. (2013) to contain only a very small number of truly repeating events. We instead form specially targeted detectors focused at the source region of interest and propose to optimize detection for phases generated by events in the sequence.

Figure 24 displays traces at 5 stations which are demonstrably sensitive to signals from the site of the 2005, October 8, Kashmir earthquake. One of these traces is simply an optimally bandpass-filtered vertical component channel from a 3-component station. All of the other traces are specialized beams deployed to detect new signals generated in this source region. Given the huge dynamic range of the signals, and the vast number of closely-spaced seismic arrivals, it is quite challenging to identify each of the individual signal arrivals. Triggers on STA/LTA traces are an obvious starting point, although we have in addition experimented with an AR-AIC procedure (AutoRegressive-Akaike Information Criterion: see, e.g. Leonard and Kennett, 1999, and references therein). The ratio of the short-term average to the long-term average for a given signal arrival is often down-weighted due to very high amplitude signals in the preceding data which increase greatly the value of the LTA. In contrast, the AR-AIC function as presented by Leonard and Kennett (1999) provides, for any given time-sample, a measure of the contrast in amplitude and frequency content of the signal immediately before and immediately after.

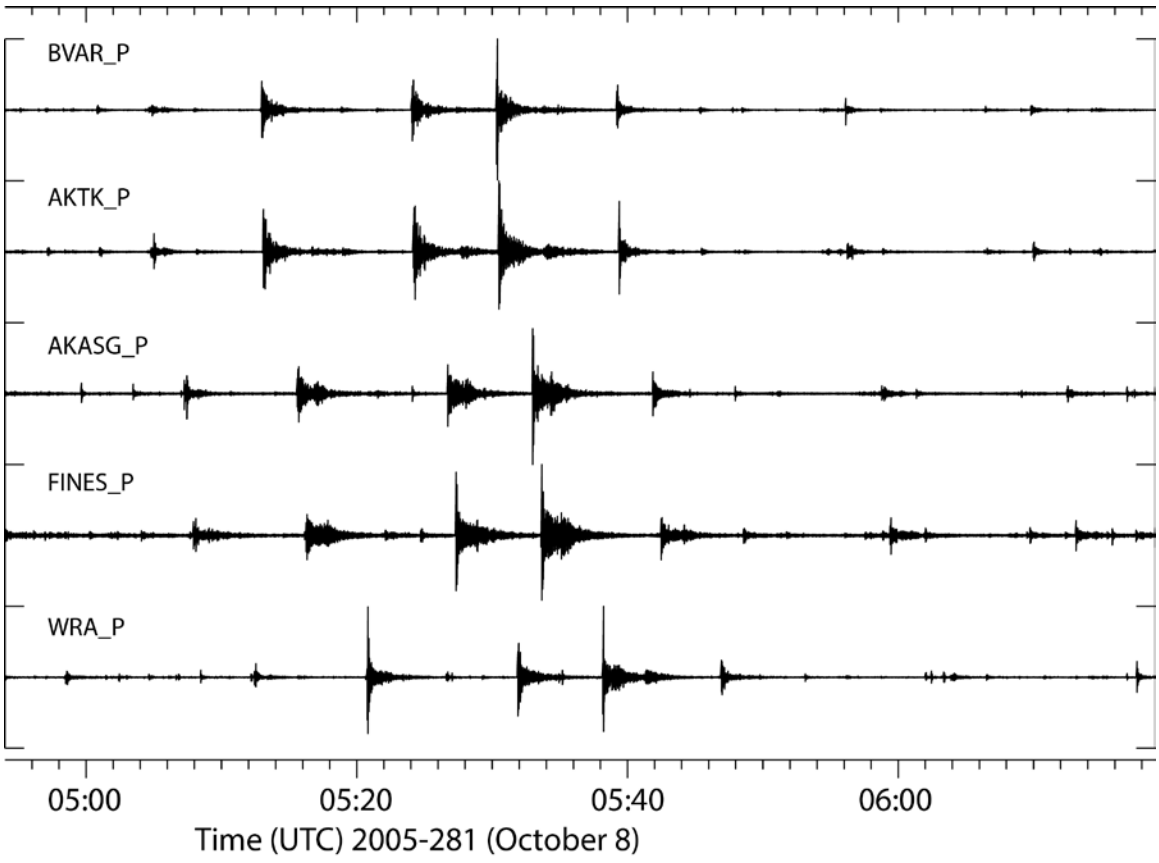

Figure 24. Traces that are deemed to be optimal on the stations indicated for the phases sought from this aftershock sequence, given the location of the main shock. All of these stations are at teleseismic distances from the source. For one of the stations, AKTK, this is simply the vertical component trace. For all the other stations, this is an optimally steered beam – not part of the standard beam deployment used for general detection purposes. Each trace is, in addition, bandpass filtered in a frequency band deemed to be optimal.

In the examples provided by Leonard and Kennett (1999), the autoregressive model is calculated in a short time-window prior to a known seismic arrival and the AR-AIC function is evaluated over a longer data window including the arrival. There is of course no reason why this procedure cannot be applied continuously, irrespective of whether or not there is a signal in the interval of interest. A continuous AR-AIC trace can be synthesized as a summation of the AR-AIC traces in each of the overlapping windows, and the resulting traces are displayed in Figure 25. Signal onsets are readily identified as significant local maxima.

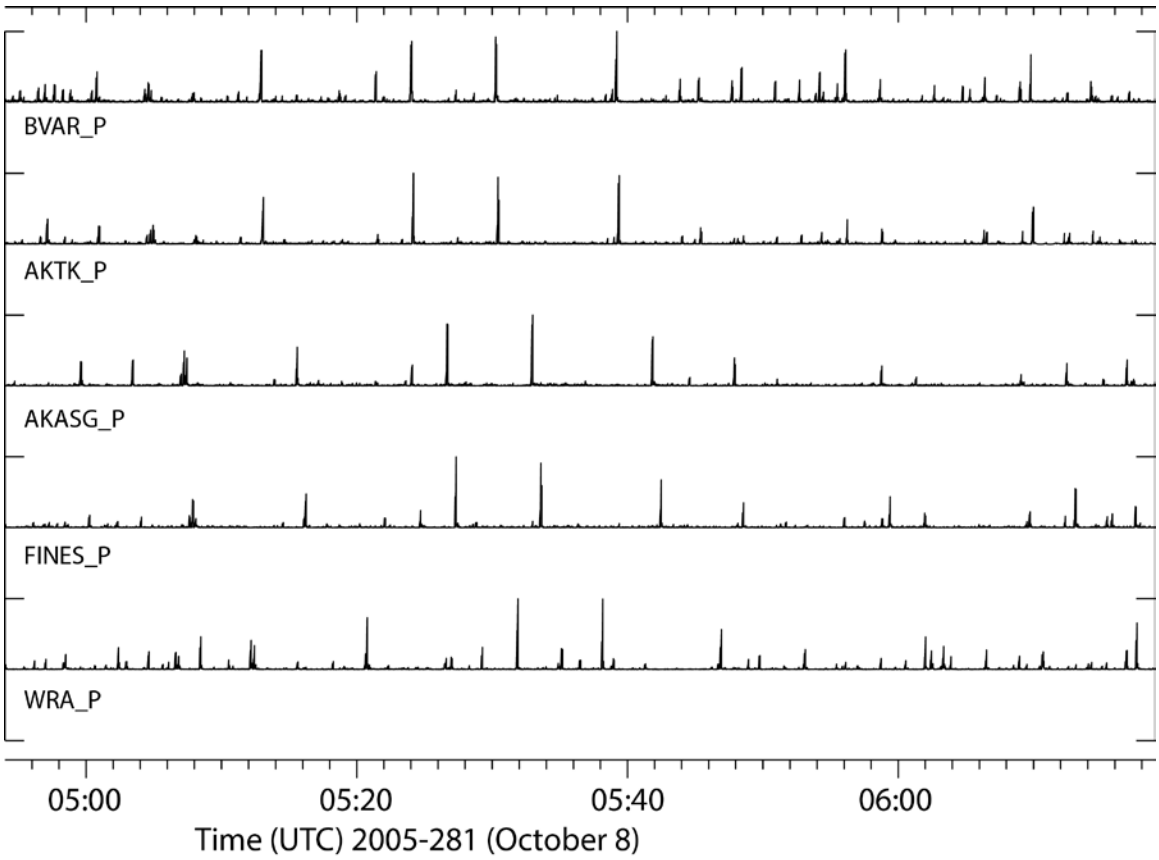

Figure 25. Continuous AR-AIC functions for the optimal traces displayed in Figure 24. Each of these traces attains a maximum value at times at which the contrast between the frequency and amplitude content of the waveform before and after attains a local maximum.

Figure 26 displays 6 more traces, optimized for detecting arrivals from this particular aftershock sequence: this time for stations at regional distances. This provides an additional complication. Whereas the stations at teleseismic distances (displayed in Figure 24) only usually show a P-arrival with a significant signal to noise ratio, the regional stations usually show both P and S phases. As indicated in Figure 27, for each station, most signal arrivals leave a significant signature both on the trace optimized for P-phases and on the trace optimized for S-phases. However, it appears that for P-phases the maximum AIC-value on the P-AIC trace almost always exceeds the AIC-value on the S-AIC trace and, similarly, for S-phases, the S-AIC maximum almost always exceeds the P-AIC maximum. Figure 28 demonstrates how, for each local maximum, a simple binary operation whereby the AIC trace with the smallest value is set to zero provides a clear indication of P- and S-phases for each event. A preliminary algorithm for associating each of the AIC maxima using a geographical grid search over the source region is currently undergoing testing.

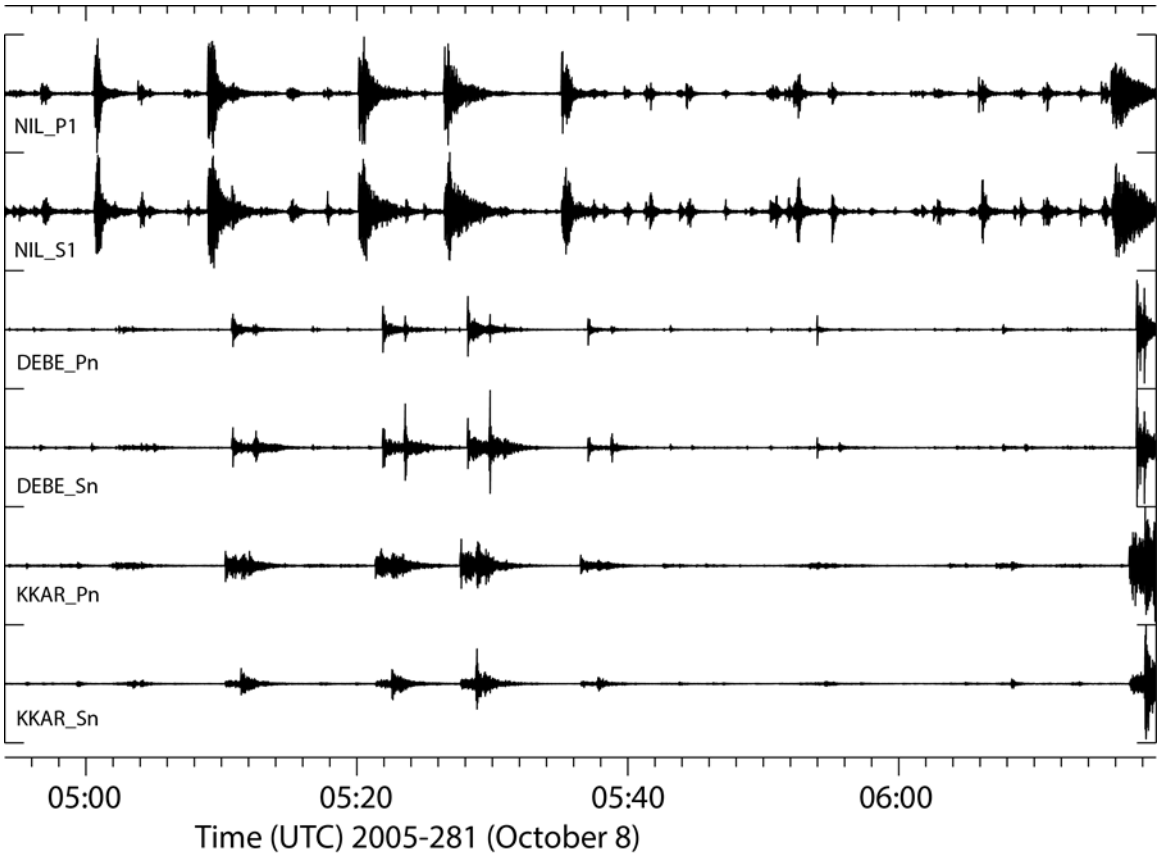

Figure 26. Traces optimized for the detection of regional P and regional S at 3 stations within 10 degrees of the main shock. For the KKAR array, both of these traces are beams steered with parameters chosen to optimize the SNR for the sought-after phases. For the 3-component stations NIL and DEBE, we use the vertical component only for the P-phases and the transverse horizontal rotation for the S-phases. The frequency bands used are generally higher for the P-phases than for the S-phases. Note that the NIL waveforms are clipped for the largest events meaning that only the P-arrival is usually seen for these events.

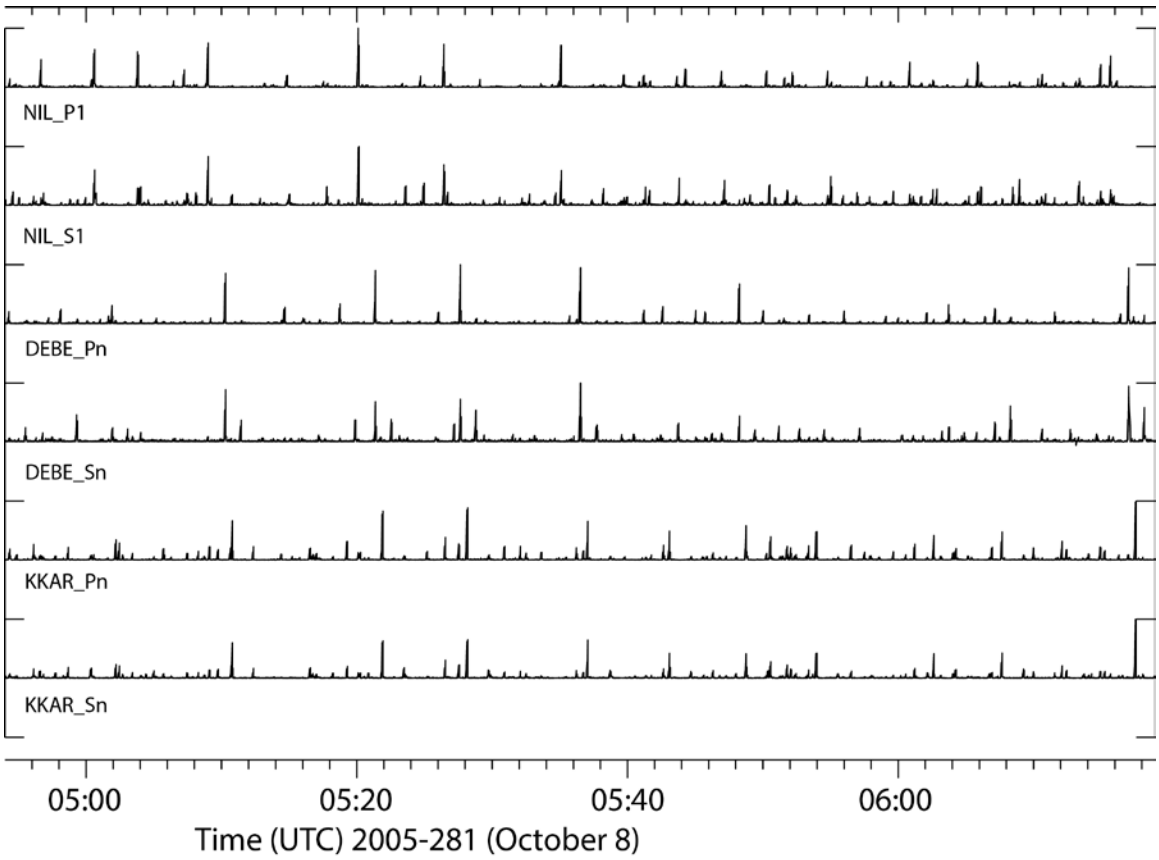

Figure 27. Continuous AR-AIC functions for the traces displayed in Figure 26.

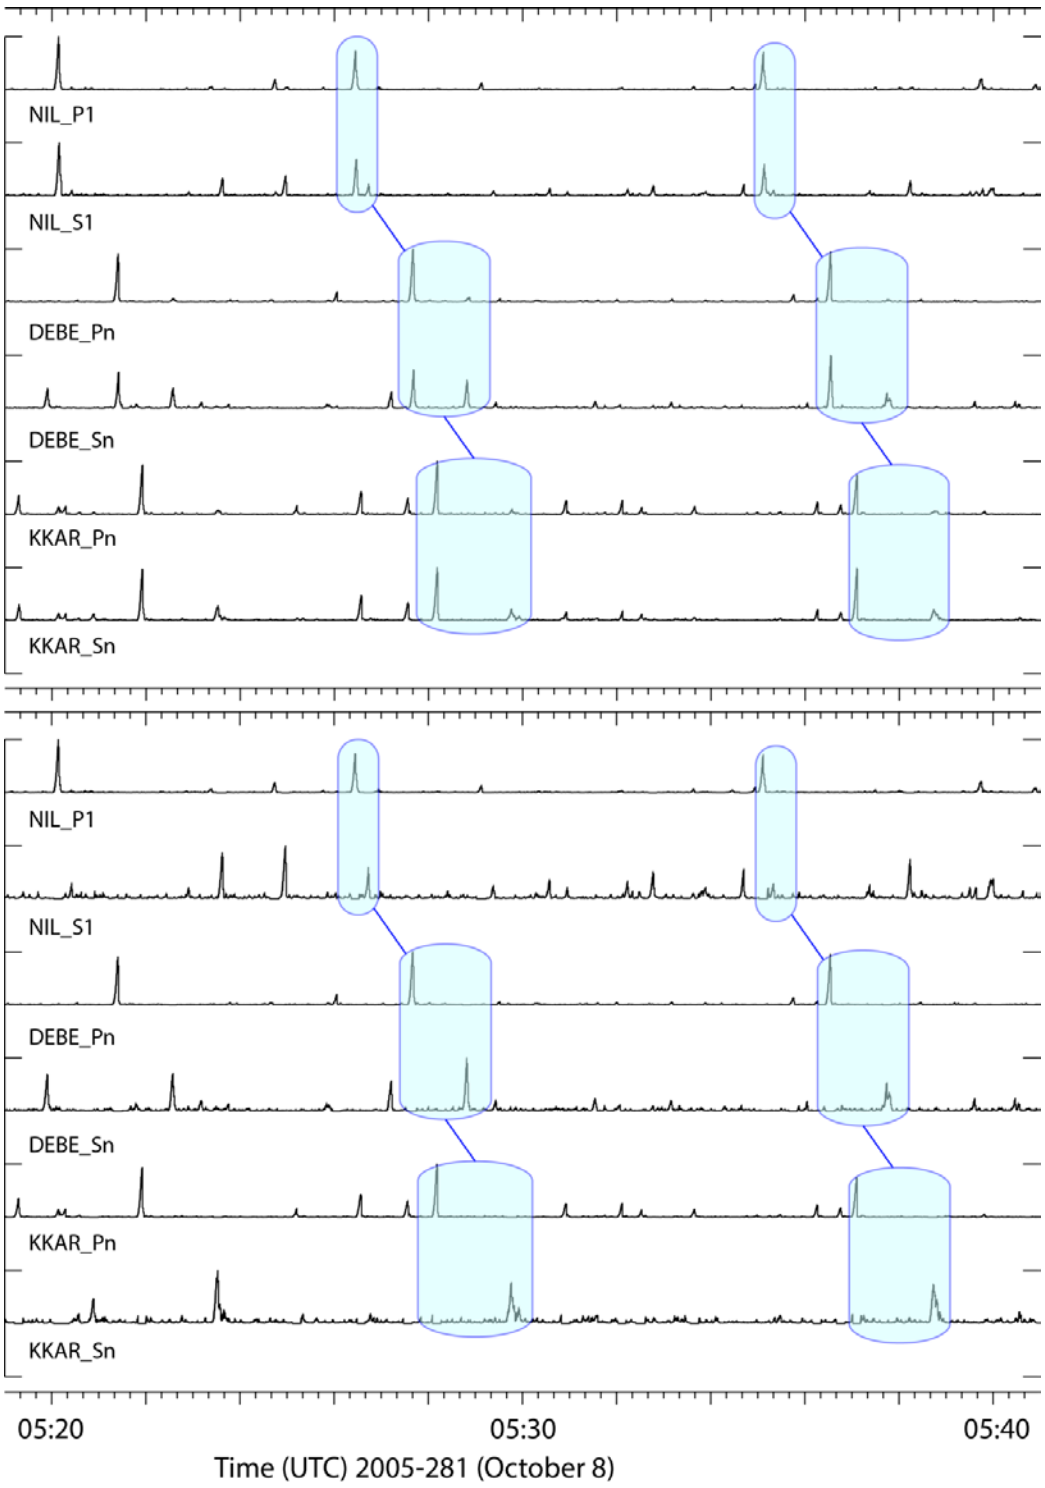

Figure 28. Close-up of the traces displayed in Figure 27 in original form (upper panel) and in reduced form (lower panel).

4 RESULTS AND DISCUSSION

Each of the following three subsections considers in turn the progress made during the second 12 months of this contract on the material described in subsections (3.1), (3.2), and (3.3).

4.1 Subspace Signal Cancellation

The substantial modification to the detection framework architecture has included the subspace signal cancellation. The conclusions from the implementation of this procedure are as follows:

1. Cancellation results in far fewer nuisance signals being detected.
2. No non-nuisance signals are missed as a result of cancellation.
3. A few additional valid detections of non-nuisance signals occur after cancellation.
4. A few nuisance signals are only partially cancelled by the subspace signal cancellation procedure. This is usually the result of multiple event sources and does have a tendency to result in detections. The number of partially cancelled signals is likely to be reduced significantly by implementation of the continuous cancellation scheme. The primary disadvantage of this scheme is the significantly increased computational effort.

A draft manuscript detailing the cancellation process has been prepared and is included as section 3.1 of this annual report.

4.2 Redesign of the detection framework architecture

The data handling architecture of the detection framework has been revised significantly. In part this was because the previous strategy of sharing data buffers among multiple objects made it more difficult to implement a process that would modify those data buffers. Also, the use of shared memory in the old implementation prohibited the distribution of processes over a cluster. In the new architecture, the various computational units no longer have any shared state. Performance of the new framework has been evaluated and deemed to be satisfactory.

4.3 A Context-Based Procedure for Identification of Aftershocks

Extensive aftershock sequences overwhelm current event detection pipelines. A promising means of mitigating this problem in situations where pattern detectors struggle to completely characterize the source region in question is described. The procedure consists of three main stages:

1. Determination of the stations that have the best detection capability for signals from the source region in question, and designing optimal detectors (e.g. specialized beams) for each phase in question.
2. Calculation of continuous AR-AIC traces which result in significant local maxima for each signal arrival of interest. These AR-AIC traces can then be reduced (i.e. local maxima that can be demonstrated not to correspond to the phase of interest can be eliminated).
3. A grid search algorithm for identifying events from the source region of interest based upon the reduced AR-AIC traces. All phase arrivals from the initial lists of detections that correspond to these events are then removed from the parametric data stream and a new iteration is performed.

This procedure has been used for detection and initial location of aftershocks from the October 8, 2005, M = 7.6 Kashmir earthquake.

5 CONCLUSIONS

In the first 12 months of this contract, significant progress was made on three important components of a context-based processing pipeline: cancellation of transient signals, empirical matched field processing and adaptive beamforming, and the evaluation of event hypotheses. In the second 12 months of the contract, the signal cancellation procedure was revised significantly and tested extensively. This procedure has reduced significantly the number of detections resulting from nuisance signals without resulting in any signals of interest being missed. The detection framework architecture has been revised significantly to replace shared memory with message passing, allowing ease of modification to the architecture and to facilitate the use of distributed architectures. The revised architecture has been implemented and tested. Finally, a context-based procedure for optimal characterization of events from a limited source region has been devised and tested. This procedure can comprise a component in an iterative pipeline which classifies extensive aftershock sequences completely, allowing a new iteration of the phase association procedure to construct event hypotheses not related to the aftershock sequence on a cleaned, sparser list of detections.

REFERENCES

Arora, N. S., Russell, S., and Sudderth, E., "NET-VISA: Network Processing Vertically Integrated Seismic Analysis," *Bulletin of the Seismological Society of America*, **103**, pp. 709-729, doi:10.1785/0120120107, 2013.

Gibbons, S. J., Schweitzer, J., Ringdal, F., Kværna, T., Mykkeltveit, S., and Paulsen, B., "Improvements to Seismic Monitoring of the European Arctic Using Three-Component Array Processing at SPITS," *Bulletin of the Seismological Society of America*, **101**, pp. 2737-2754, doi:10.1785/0120110109, 2011.

Harris, D. B., "Characterizing source regions with signal subspace methods: theory and computational methods," UCID-21848, Lawrence Livermore National Laboratory, 1989.

Harris, D. B., "Subspace Detectors: Theory," *UCRL-TR-222758*, Lawrence Livermore National Laboratory, 2006

Harris, D. B. and Dodge, D. A., "An autonomous system for grouping events in a developing aftershock sequence," *Bulletin of the Seismological Society of America*, **101**, pp. 763-774, 2011.

Harris, D. B., Jarpe, S. P. and Harben, P. E., "Seismic noise cancellation in a geothermal field," *Geophysics*, 56(10), pp. 1677-1680, 1991.

Junek, W. N., Kværna, T., Pirli, M., Schweitzer, J., Harris, D. B., Dodge, D. A., and Woods, M. T., "Inferring Aftershock Sequence Properties and Tectonic Structure Using Empirical Signal Detectors," *Pure and Applied Geophysics*, **172**, pp. 359-373, doi:10.1007/s00024-014-0938-0, 2015.

Kværna, T., Harris, D. B., Gibbons, S. J. and Dodge, D. A., "Adapting Pipeline Architectures to Track Developing Aftershock Sequences and Recurrent Explosions," *AFRL-RV-PS-TR-2014-0075*, NORSAR, Kjeller, Norway, 2014.

Leonard, M. and Kennett, B. L. N., "Multi-component autoregressive techniques for the analysis of seismograms," *Physics of the Earth and Planetary Interiors*, **113**, pp. 247-263, doi:10.1016/s0031-9201(99)00054-0, 1999.

Oppenheim, A. V. and Schafer, R. W., **Digital Signal Processing**, Englewood Cliffs, N.J., Prentice-Hall. ISBN 0-13-214635-5, 1975.

Slinkard, M. E., Carr, D. B., and Young, C. J., "Applying Waveform Correlation to Three Aftershock Sequences," *Bulletin of the Seismological Society of America*, **103**, pp. 675-693, doi:10.1785/0120120058, 2013.

Widrow, B. and Hoff, M. E, "*Adaptive switching circuits*," Proc. Of WESCON Conv. Rec., part 4, pp. 96-140, 1960.

LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

CTBTO	Comprehensive Nuclear-Test-Ban Treaty Organization
IDC	International Data Center
REB	Reviewed Event Bulletin (an analyst reviewed event bulletin generated at the IDC)
SEL3	Standard Event List 3 (a fully automatic seismic event bulletin generated at the IDC)

DISTRIBUTION LIST

DTIC/OCP
8725 John J. Kingman Rd, Suite 0944
Ft Belvoir, VA 22060-6218 1 cy

AFRL/RVIL
Kirtland AFB, NM 87117-5776 2 cys

Official Record Copy
AFRL/RVBYE/Dr. Robert Raistrick 1 cy