Evaluating and Improving the ElarmS Earthquake Early Warning Algorithm

By

Holly Moore Brown

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Earth and Planetary Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Richard M. Allen, Chair

Professor Douglas S. Dreger

Professor David R. Brillinger

Fall 2012

Abstract

Evaluating and Improving the ElarmS Earthquake Early Warning Algorithm

by

Holly Moore Brown

Doctor of Philosophy in Earth and Planetary Science

University of California, Berkeley

Professor Richard M. Allen, Thesis Committee Chair

Earthquake Alarm Systems, or ElarmS, is a network-based earthquake early warning (EEW) system developed in California. This dissertation concerns the transition of ElarmS from research prototype to production-grade code. We test ElarmS' performance for large earthquakes with an extensive dataset of large events from Japan. Using the Japanese dataset and results, we develop a statistical error model for the magnitude, location, and ground motion prediction algorithms. We adapt ElarmS to run continuously throughout the state and analyze system latencies due to various seismic networks and instruments. We then rewrite ElarmS completely in newer, more efficient code, while redesigning the association and alert algorithms for improved performance. Finally, we apply an Artificial Neural Network (ANN) at the end of system processing, to detect and block false alerts. The ANN allows ElarmS to send earlier, faster alert messages, when fewer stations have contributed data to an event estimate. ElarmS now sends realtime alert messages to the California Integrated Seismic Network's ShakeAlert EEW system.

Table of Contents

List of Figures

List of Tables

Acknowledgements

I owe sincere thanks to my advisor, Richard Allen. He was a ton of fun to work with, and possesses more patience and humor than perhaps any other professor of all time. I greatly enjoyed his company and mentorship while I worked on campus, and was humbled by his generous flexibility when I went on maternity leave and then resumed my studies while working from home. Truly I could not have finished without his guidance. Best of all, his enthusiasm and vision for earthquake early warning made the project inspiring and fun, from start to finish.

I am grateful to Professor Doug Dreger for his excellent seismology classes and moral support, and to Professor David Brillinger for his statistical expertise and encouragement.

My family, especially my mother and my mother-in-law, made the last year of schoolwork possible by setting aside their own affairs to play with my son every time I had a meeting on campus. Their help was priceless, and transformed each trip to Berkeley into a delightful outing.

Finally my deepest thanks are for my husband, Tom. He cheered me on as I worked late in the night, proofread papers, listened to practice talks, brainstormed research ideas, and made me laugh when I might otherwise have given up. I thank my lucky stars every day for his friendship and companionship.

As for my son, Roger, he made all the work worthwhile.

Chapter 1

Introduction

Earthquake Early Warning and ElarmS

Earthquake Early Warning (EEW) is a method of recognizing earthquakes as they occur, and providing warning before severe shaking reaches major metropolitan areas. It depends on dense networks of seismic detection instruments, rapid and reliable connections to central computers, and efficient software that recognizes and characterizes the earthquake in a matter of moments. The resulting alert messages can provide a few seconds to tens of seconds of warning before damaging shaking arrives.

Simple EEW devices have been in use since the 1980s and have gradually increased in both complexity and capability over the ensuing decades. Currently there are operational EEW systems in Mexico, Japan, Taiwan, Turkey, and Romania, sending realtime warnings to government agencies and private corporations by direct computer link, and in some cases to the general public via cell phone.

In California, EEW has been in research development for more than ten years, and is just in the last few years making the transition to production-grade systems, available for public and private use.

This dissertation concerns a particular California EEW algorithm, called Earthquake Alarm Systems, or ElarmS, and its transition from research prototype to robust production code. ElarmS is part of the new ShakeAlert EEW system run by the California Integrated Seismic Network (CISN), currently delivering prototype alert messages to government and industrial users throughout California.

Evaluating Uncertainties and Performance for Large Events

EEW systems are designed specifically for large earthquakes, but there have not been enough recent large earthquakes in California to substantially test ElarmS' performance for these events. In 2007 and 2008, we applied the ElarmS algorithm to a dataset of large earthquakes in Japan. This gave us the opportunity to observe ElarmS' performance with many earthquakes much larger than those typically seen in California. The expansive new dataset was also useful for evaluating uncertainties in the algorithm's estimates of events. We used the uncertainties to develop a statistical error model for ElarmS. The results of this study are described in Chapter 2, which was published in its entirety in *Seismological Research Letters* (Brown, *et al.*, 2009).

Realtime Implementation

Chapter 3 describes the transition of ElarmS from processing events offline, long after an event has ended, to realtime processing as the event unfolds. This transition involved a tremendous group effort. New hardware connections and software capability had to be established to handle hundreds of simultaneous incoming data streams. The ElarmS algorithm had to be modified to run continuously, to tolerate triggers arriving out of sync with each other, and to send realtime alert messages to the development group. Of new relevance were latencies in the system; where and by how much were data transmissions delayed? How did these latencies affect event processing? Chapter 3 investigates these questions and other aspects of realtime processing, and was published in its entirety in *Soil Dynamics and Earthquake Engineering* (Brown, *et al.*, 2011).

Developing the Second-Generation ElarmS Code

Once ElarmS was running successfully in realtime, we were theoretically ready to connect it to the CISN ShakeAlert EEW system. But the ElarmS code was still a research prototype, albeit heavily upgraded and modified. In 2010 and 2011 we rewrote ElarmS from scratch into a new, production-grade code called E2. Many parts of E2, such as the thoroughly vetted magnitude scaling relations, were functionally identical to the original ElarmS, just written more cleanly in a more modern coding language. Other parts, such as the association algorithm and the waveform processing module, were significantly rewritten. Chapter 4 details the author's contribution to the new E2 code.

Improving ElarmS Performance with Assistance from an Artificial Neural Network

Finally, we took the realtime, production-grade E2 code, and applied an Artificial Neural Network (ANN) at the end of the processing stream to filter out false alerts. The operational version of ElarmS currently requires four seismic stations to register a P-wave arrival before it will send out an alert message. Waiting for four stations greatly decreases the chance of false alarms, but delays the alert message by about two seconds. The ANN improves the false alarm rate for events with only two or three stations, allowing ElarmS to send faster alerts. Chapter 5 summarizes the ANN study and results, and has been submitted for publication in the *Bulletin of the Seismological Society of America*.

Between 2007 and 2012 we greatly improved and updated ElarmS, resulting in a true metamorphosis from research prototype to real-world product. Today ElarmS operates continuously, processing seismic waveforms from every EEW-compatible seismic station in the state and sending realtime alert messages to the CISN ShakeAlert system. The research and updates described here have the potential to reduce loss of life due to earthquakes in California for years to come.

Chapter 2

Testing ElarmS in Japan

This chapter has been published in *Seismological Research Letters* as "Testing ElarmS in Japan", by Brown, Allen and Grasso, 2009. Issue 80(5), pages 727-739.

2.1 Introduction

Earthquake early warning systems use seismic networks, rapid telemetry, and software algorithms to detect an earthquake immediately after its inception, estimate its damage potential, and disseminate a warning to surrounding communities before peak ground shaking occurs. Earthquake Alarm Systems, or ElarmS, is an earthquake early warning system developed in California. The ElarmS algorithm recognizes earthquakes from the initial P-wave arrivals at seismometers near the epicenter. The characteristics of the P-wave, including amplitude and frequency, are used to estimate a final magnitude for the event. P-wave arrival times from several stations are combined to estimate the hypocenter of the event. Finally, the estimated magnitude and hypocenter are applied to attenuation relations to produce a prediction of ground shaking levels in the region.

ElarmS has been tested extensively with datasets of earthquakes from northern and southern California (Allen and Kanamori, 2003; Allen, 2007; Wurman, *et al*., 2007, Tsang, *et al*., 2007; Allen*, et al*., 2009). While the test datasets from California included a large range of locations and source types, there are a limited number of recent, well-recorded, large earthquakes available for testing the early warning system. In this study we take ElarmS to another geographic and seismic setting, and test the algorithms with a dataset of 84 large-magnitude earthquakes, including 43 of magnitude 6.0 or greater, in Japan. The Japanese test dataset is valuable both for the insight into ElarmS' processing of large events, and for the chance to process events in a completely different geologic setting. The offshore and deep nature of many of the events presents new challenges to the methodology. The Japanese earthquakes offer an opportunity to improve the robustness of ElarmS, extend its abilities to other settings, and confirm its relevance for large-magnitude events.

2.2 Earthquakes

The dataset contains 84 earthquakes recorded by Japan's Kyoshin Net (K-NET) strong-motion seismic network. K-NET consists of 1,000 digital strong motion seismometers, distributed across Japan with approximately 25km spacing. Each station is capable of recording accelerations up to 2,000 $^{cm}/_{s^2}$, with a sampling frequency of 100 samples per second and a dynamic range of 100dB.

The 84 test earthquakes occurred between September 1996 and June 2008 (Figure 2.1). The waveforms were downloaded directly from the K-NET website (www.k-net.bosai.go.jp). All events occurred within 100km hypocentral distance of at least three stations. The JMA magnitude estimates range from 4.0 to 8.0. Forty-three events have magnitudes of 6.0 or greater. The largest event in the dataset is the Tokachi-Oki earthquake of 26 September, 2003, which had a magnitude of 8.0.

2.3 Warning Methodology

2.3.1 Location

When only one station has triggered on a P-wave arrival, ElarmS sets the estimated hypocenter directly beneath the station, at a depth of 8km. When a second station triggers, or if the first two stations trigger simultaneously, ElarmS locates the event on the great circle between the two stations, at a location dependent on their arrival times, again at a fixed depth of 8km. When three stations have triggered, ElarmS triangulates an event epicenter using a two-dimensional grid search, the arrival times at each of the stations, and a typical P-wave velocity. The depth is still fixed at 8km. In California, this third method continues to be used when four, five, or more stations have triggered. Note that the use of a two-dimensional grid search and fixed depth are acceptable for the shallow seismogenic zone in most of California.

In the Japanese subduction zone, however, this is inappropriate. For the Japanese dataset a three-dimensional location algorithm is needed. The first three scenarios, for one, two, or three triggers, remain the same. When four or more stations have triggered, ElarmS creates a three-dimensional grid of possible event hypocenters, with depths ranging from 0km to 100km in steps of 10km. Observed P-wave travel times are compared to those predicted by travel time curves for seismic waves originating at each point of the grid. The best match of arrival times is deemed to be the event hypocenter. As additional stations trigger, the grid search is repeated and the estimated hypocenter is updated.

2.3.2 Magnitude

ElarmS uses two P-wave parameters, peak amplitude and maximum predominant period, to create two independent estimates of final event magnitude. The estimates from each parameter are then averaged together to form the ElarmS magnitude for the event.

Maximum predominant period, τ_p^{max}, was developed first and provided the original framework for ElarmS (Allen and Kanamori, 2003). For a given region, observed τ_p^{max} values from the first few seconds of the P-wave are plotted as a function of final event magnitude. A least-squares fit to the data results in τ_p^{max} vs. magnitude scaling relations for the region (Allen and Kanamori, 2003; Lockman and Allen, 2007; Wurman*, et al*., 2007; Tsang*, et al*., 2007). These empirically determined scaling relations are then used to estimate magnitude.

For a given event, the first station to recognize a P-wave arrival reports a τ_p^{max} value after one second of observation. As additional seconds pass, the station may update that value if a larger τ_p^{max} is observed. When more stations trigger, they too initially report one-

second τ_p^{max} values, which may increase with additional seconds of data. Each second, ElarmS applies the scaling relations to the most current τ_p^{max} observation from each triggering station to determine an estimated magnitude. All station magnitudes are then averaged together to create a single τ_p^{max}-based magnitude. This averaged magnitude estimate is adjusted every second, as current stations update their τ_p^{max} observations and new stations trigger.

Peak displacement amplitude, P_d, of the P-wave was added as an ElarmS parameter by Wurman, et al., in 2007. For a given region, observed P_d values are recorded at each station, scaled to a common epicentral distance, and plotted against final event magnitude. Again, a least squares fit is used to determine regional P_d vs. magnitude scaling relations. During event processing, P_d observations are combined in the same way as τ_p^{max} observations. Each triggered station reports a P_d observation every second. The P_d value may increase if a large displacement is observed later in the P-wave. ElarmS converts each P_d observation to an estimated magnitude, and then averages all the estimated magnitudes together to create a single P_d-based magnitude. As more P_d observations become available, they are incorporated into the estimate.

ElarmS performs each of these scaling calculations independently, resulting in one τ_p^{max}-based magnitude estimate and one P_d-based magnitude estimate each second. The two magnitudes are averaged together to create the "ElarmS" magnitude estimate, which is used for predicting ground shaking. ElarmS uses a simple linear average of the two methods, as we have yet to observe an improvement with the use of a weighted average (Wurman*, et al.*, 2007).

2.3.3 AlertMaps

Once a hypocenter and magnitude are estimated, ElarmS predicts regional ground accelerations from attenuation relations. For Japan, ElarmS uses the same attenuation relations that the NEIC uses to create ShakeMaps for Japanese events. For events shallower than 20km or with magnitude less than 7.7, the attenuation relations are those defined by Boore, et al. (1997). For all other events in Japan, the attenuation relations are those from Youngs, et al. (1997).

The initial AlertMap is generated using only the estimated magnitude and event location. As stations begin observing peak ground acceleration (PGA), the observations are incorporated into the AlertMap and the attenuation function is adjusted to best fit the available data. The intent of including PGA observations is to correct for any errors in the ElarmS estimate of magnitude. As each PGA observation is added, the AlertMap predictions are adjusted closer to the true observed ground motions. If the catalog magnitude and location are used for an event, and all PGA observations are included, then by definition the AlertMap exactly equals the USGS ShakeMap for that event. Using the estimated magnitude and location, any error in the final AlertMap, after all PGA observations are included, is thus due to errors in the magnitude and location estimates.

2.3.4 Error Calculations

Errors are calculated by comparing ElarmS' output to published or observed values. For magnitude and location, the ElarmS estimate is compared to the K-NET published magnitude and location. For ground motions, the ElarmS prediction for any station which has not yet observed peak ground shaking is compared to the final observation at that station. Only predictions that are made before peak ground shaking occurs are considered for the error analysis.

2.4 System Performance

2.4.1 Magnitude Estimation

We begin by determining the scaling relationship between peak displacement (P_d) and magnitude in Japan. Acceleration data from each station is double integrated to displacement, and scaled to a common epicentral distance of 10km. These scaled displacements are then plotted as a function of the K-NET catalog magnitude (M_L). Applying a least-squares fit to the observed $\log_{10}(P_d)$ values gives a scaling relation of $\log_{10}(P_d) = 0.66*M_L - 4.02$, shown in Figure 2.2a. The dashed line in the figure is the scaling relation for Northern California, $\log_{10}(P_d) = 0.73*M_w - 3.77$ (Wurman, *et al.*, 2007). The scaling relation for Japan is of a similar slope to that of Northern California, but with lower displacements at all magnitudes. This implies that Japan has higher attenuation than Northern California.

P_d observations from all events are weighted equally in the determination of the P_d scaling relations. We note that the largest event, the M8.0 Tokachi-Oki earthquake, does not fall on the scaling relations. Using the observed displacements for Tokachi-Oki and the P_d scaling relations determined from all events, ElarmS underestimates the event magnitude by more than one magnitude unit. Other studies have shown that peak displacements may saturate at near-source stations during large magnitude events (Wurman*, et al.*, 2007; Zollo*, et al.*, 2006). This is the effect we observe for the largest earthquake in our dataset. The relatively low amplitudes near the source can lead to underestimation of magnitude. This suggests that P_d should not be used alone in regions that are prone to very large earthquakes.

The second magnitude estimation method is maximum predominant period, or τ_p^{max}. The observed τ_p^{max} values at each station are plotted in Figure 2.2b against the final catalog magnitude of each event. A least-squares fit to the data produces a scaling relation of $\log_{10}(\tau_p^{max}) = 0.21*M_L - 1.22$, shown as the solid line in Figure 2.2b. Wurman et al, 2007, found a scaling relation for Northern California of $\log_{10}(\tau_p^{max}) = 0.15*M_w - 0.78$, shown as the dashed line. The observed predominant periods in Japan are of similar values to those of Northern California, but the best-fit slope is steeper in Japan. τ_p^{max} does not appear to display the saturation effects that P_d does. For the M8.0 Tokachi-Oki event, using the best fit scaling relation and τ_p^{max} observations leads to an estimated magnitude of 8.2, using τ_p^{max} alone. However, τ_p^{max} shows more scatter than P_d, particularly for the lowest magnitude events. This agrees with similar results found by previous studies (Olson and Allen, 2005; Wurman*, et al.,* 2007).

ElarmS produces a single event magnitude by averaging together the magnitudes from P_d and τ_p^{max}. Figure 2.2c shows the three magnitude estimates for each event. The green points are the magnitude estimates for each event using only P_d observations at each station. The blue points are the magnitude estimates using only τ_p^{max}, and the red points are the final magnitude estimate for each event when τ_p^{max} and P_d magnitudes have been averaged together. The red line is the linear best fit to the averaged magnitudes. The dashed black line is the desired 1-1 fit, for which the ElarmS estimated magnitude would be exactly equal to the K-NET catalog magnitude. The averaged magnitudes fall close to the desired 1-1 fit, improving on both the saturation effects of P_d at high magnitudes and the scatter of τ_p^{max} at low magnitudes. For the M8.0 event, the averaged ElarmS magnitude is 7.5, one half unit lower than the published magnitude.

Figure 2.2d shows a histogram of the error in the average magnitude estimates from Figure 2.2c. The red curve in the figure is the bestfit Gaussian distribution. The mean error is 0.0, with a standard deviation of 0.4 magnitude units. All but one event is within one magnitude unit. The one deviant event is a magnitude 4.1. The τ_p^{max} observations for that event are high, leading to an average magnitude of 5.18.

2.4.2 Magnitude Estimation of Largest Events

Of primary concern is the accuracy of ElarmS magnitude estimates for large magnitude events. We consider subsets of the dataset, analyzing events within specific magnitude ranges. For events with magnitude 6 or greater, the mean error is 0.04 magnitude units, with a standard deviation of 0.46. This is only slightly higher than the mean and standard deviation for all events. For events of magnitude 7 or greater (of which there are eight in this dataset, including the M8.0 Tokachi-Oki event), the ElarmS magnitude has a mean error of -0.21, with a standard deviation of 0.53. This underestimation for the largest events is partly due to the saturation of P-wave amplitudes for large events, but is also related to the offshore location of the largest events. Poor azimuthal coverage causes large errors in the estimated epicentral distance, which in turn contaminates the magnitude determined by peak displacement.

2.4.3 Magnitude Dependence on Time

Each station reports to ElarmS once every second, updating its observed values of displacement and period. The peak displacement may occur in the first second after a trigger, or it may occur later. If a larger displacement is observed in later seconds, the peak value for that station will be increased accordingly. If the one-second value continues to be the largest as additional seconds of data are recorded, that one-second peak is kept throughout. Thus the peak displacement recorded in the first second is the minimum possible value for that station. The same is true for maximum predominant period, τ_p^{max}. The one-second observation of τ_p^{max} is the minimum value for that station. It may be increased with additional seconds of data, but not decreased.

Magnitude error is thus directly dependent on the number of seconds of data available at each station. In Figure 2.3 we calculate separate scaling relations for each time window. Figure 2.3a shows the resulting scaling relations for peak displacement, P_d, given a 1-,2-,3-,4-, or 5-second window. Using one second of data at each station results in a shallow

slope to the scaling relation, leading to systematic underestimation of magnitude for larger events. Each additional second of data increases the slope and improves the magnitude estimate for large events. The change from four to five seconds is minimal. We therefore determine that four seconds is the optimal time window, providing the most data while still allowing rapid response to the earthquake.

Figure 2.3b shows the effect of time window on maximum predominant period. τ_p^{max} is less sensitive to time window, but still shows a slight increase in slope and improvement in magnitude estimate with additional seconds of data. The fit to large events in particular improves with additional seconds. This verifies the finding of Allen and Kanamori (2003), that the initial magnitude using one second of data is a minimum estimate, and additional seconds of data increase the magnitude estimate for large events. Again the four-second window maximizes data availability and timeliness.

2.4.4 Magnitude Dependence on Stations

Any earthquake early warning system is dependent on data that arrives one station at a time. We test the sensitivity of the scaling relations to the number of stations reporting triggers. The stations are sorted by distance to the epicenter, and only the closest 1, 2, 3, etc., stations are used to estimate magnitude. Figure 2.4 shows the absolute value of error in ElarmS magnitude estimates as additional station data is incorporated. The green and blue dashed lines are the error in the independent magnitude estimates made using only P_d or τ_p^{max}, respectively. The solid line is the error in the averaged magnitude estimate, which ElarmS uses to predict ground motions in the region. This combined P_d and τ_p^{max} estimate has an average absolute error of less that 0.42 magnitude units using only the single closest station to the epicenter, and the error decreases with the addition of more stations. P_d by itself produces an average error of less than half a magnitude unit for all numbers of stations. τ_p^{max} by itself has higher error values than P_d, but still less than 0.6 magnitude units on average. Using only the closest one or two stations to the epicenter, the average error (solid line) is lower than both individual errors. This is because one method may overestimate the magnitude and the other may underestimate it; their average is closest to the true magnitude. Once three or more stations are providing data the average error in the magnitude using P_d alone is slightly lower than the average estimate. However, given the uncertainties (error bars in Figure 2.4) this difference is insignificant and we prefer to use two independent magnitude estimates rather than relying on just one.

2.4.5 Magnitude Error Distributions

Every ElarmS magnitude estimate is dependent on both the number of seconds of data at each station, and the number of stations reporting. If we wish to know the uncertainty in a given magnitude estimate, we must consider the quantity of data included in the estimate. Figure 2.5 shows magnitude error histograms and best-fit Gaussian probability distributions, determined for specific quantities of data. The mean and standard deviation for all the error distributions are shown in Table 2.1. Figure 2.5a and b use exactly one second of P-wave arrival at each station. Figure 2.5a shows the error histogram when the closest 5 stations each contribute 1 second of data. The curves in Figure 2.5b represent the best-fit Gaussian error distributions when 1, 2, 3, 4 or 5 stations each contribute 1

second of data to the magnitude estimate. Figures 2.5c and d are for 2 seconds of data, Figure 2.5e and f for 3 seconds, up to 5 seconds of data in Figure 2.5i and j. The probability distributions get more peaked in each row, indicating that additional seconds of data improve the accuracy of the magnitude estimate.

More data for a magnitude estimate is always desired, but an early warning system must be prompt to be useful. ElarmS creates its initial magnitude estimate using the first second of data at the first triggered station, and then updates the estimate as more data becomes available. The error distributions shown in Figure 2.5 can be used to assign an uncertainty to the ElarmS magnitude estimate from the first estimate when 1 sec of data is available.

2.4.6 Location

The Japanese events presented a challenge in that many occurred offshore (Figure 2.1). The ElarmS location algorithm depends on triangulating between several stations, but this can be hindered by poor azimuthal coverage. For many events all stations are to the west, and the process of locating the hypocenter accurately requires more stations than it does for onshore events. Many events are also very deep, and a minimum of four stations must be available before hypocentral depth can be estimated. The mean error in the hypocentral location for all events and any number of stations providing arrival times, i.e. one station up to the total number of triggering stations for each event, is eighteen kilometers. The greatest errors are for those occurring furthest offshore or deepest. Figure 2.6a shows a histogram of hypocentral location errors, using P-wave arrival times from 5 stations. The errors using 1 to 5 stations are listed in Table 2.1. We determine the best fitting log-normal distributions (Figure 2.6b) for location error. These distributions can be used to estimate the error in any given ElarmS location, as a function of the number of stations reporting triggers.

2.4.7 Attenuation Relations

The final source of error is the attenuation relation used to translate magnitude and location into a prediction of local ground acceleration. To isolate this error we give ElarmS the catalog magnitude and location for each event. ElarmS applies those parameters to the ShakeMap attenuation relations for Japan to create an initial AlertMap prediction of ground motions. As stations report peak ground shaking observations, ElarmS incorporates them into the model and adjusts the predictions accordingly. Thus the only sources of error in the final PGA predictions are the attenuation calculations themselves, and any scatter in the local PGA observations.

The predictions are then compared to the actual observed ground motions recorded at any station whose data is not yet incorporated into the model. Figure 2.7a shows a histogram of the resulting errors when five PGA observations are included. Table 2.1 lists the mean errors using 0 to 5 PGA observations. Figure 2.7b shows the probability distributions derived from the error data. The yellow curve is the initial ground motion estimates, using only magnitude and location with no peak shaking observations. As stations report peak ground shaking, their observations are used to adjust the AlertMap PGA predictions up or down. The remaining curves show the error in the adjusted PGA predictions, when

actual PGA observations are included. Using one or two PGA observations results in the most error, more than the zero-observations case. The increased error using only one or two PGA observations occurs because each peak ground acceleration observation is affected by unpredictable path effects, and may have significant variability compared to the average regional ground shaking. Only after several individual station observations are included do their individual errors cancel each other out. The errors using three, four, or five PGA observations are better than those of the zero-observation case. This suggests ElarmS should use the initial, zero-observation model until at least three PGA observations are available.

2.4.8 Example Earthquake

Figure 2.8 shows a step-by-step progression of ElarmS processing for an example event in Japan. The event is a magnitude 6.4, occurring 26 July, 2003, at a depth of 11.9km. The first two triggers are recognized at two stations simultaneously (Figure 2.8a). ElarmS places the event location between the two stations, at a fixed depth of 8km. One second later (Figure 2.8b), ElarmS combines the P_d and τ_p^{max} observations from these two stations to create a magnitude estimate of 6.6 and predicts the distribution of ground shaking. A third station triggers, and the epicenter is estimated, with depth still fixed at 8km. One second later (Figure 2.8c), which is two seconds after the first trigger, the third station's data is incorporated into the estimate, and magnitude is decreased to 6.2. Two more stations trigger, and the five total triggers so far are used to estimate a hypocenter, at a depth of 10km. Finally, Figure 2.8d shows the ElarmS AlertMap when all data for the event is available, 12 seconds after the first trigger. The final ElarmS magnitude estimate is M6.4. The magnitudes include P_d and τ_p^{max} observations from all 26 available stations, and the AlertMap is adjusted for peak ground shaking observations from all 26 stations.

2.5 ElarmS Errors

2.5.1 The Error Model

To determine the total error in the ElarmS prediction of ground shaking, we combine the errors determined for magnitude, location, and attenuation relations. The mean and standard deviation from each probability distribution above are used to run a Monte Carlo simulation, generating random errors within each distribution. These errors are then factored into the ground motion predictions using the NEIC's preferred attenuation relations for Japan. The NEIC global attenuation relations use either the Boore, et al., (1997) or the Youngs, et al., (1997) model, depending on depth and magnitude of the event. ElarmS follows the same criteria, choosing between the two models accordingly. For the example below, we show the error calculations from Boore, et al., (1997). The coefficients (B1, B2, B3, B5) are those recommended by Boore for a reverse mechanism event.

The "expected" ground motions are determined by applying the catalog magnitude and location to the attenuation relations with no errors. This represents the ideal output from ElarmS, if all estimates were perfect.

Ideal output:
$\ln(PGA)_{ideal} = B_1 + B_2*(M-6) + B3*(M-7) + B5*\ln(R)$

M is the magnitude, R is the distance from the event epicenter to the location whose PGA is being predicted, and B1,B2,B3, and B5 are coefficients.

To calculate the ground motion estimates provided by ElarmS, the catalog magnitude and location are again applied to the attenuation relations, but now with the addition of the errors generated by the Monte Carlo simulation.

Estimated output:
$\ln(\hat{PGA}) = B_1 + B_2*(M+\varepsilon_M - 6) + B3*(M+\varepsilon_M - 7) - B5*\ln(R \pm \varepsilon_R) + \varepsilon_{Att}$

The total error in the ElarmS estimated ground motions is the difference between the ideal and estimated values.

Error:
$\varepsilon_{PGA} = \ln(PGA)_{ideal} - \ln(\hat{PGA})$

These error estimates have no units and are the natural logarithm of the ratio of the ideal PGA estimate to the ElarmS estimate. A positive error means the predicted PGA was lower that the ideal. A factor of two difference between the ideal and predicted PGA results in an error of 0.7, and a factor of 10 in an error of 2.3.

Monte Carlo simulations can now be used to determine the error distribution, ε_{PGA}, for any warning scenario. These errors include contributions from the magnitude estimation, the location estimation, and the attenuation relations. Each of these factors has multiple error distributions, depending on the number of reporting stations. Thus the total error, ε_{PGA}, is a function of the number of stations contributing to the location estimate, the number of stations providing P_d and τ_p^{max} values, the number of seconds of data available at each of those stations, and the number of stations reporting peak ground motion observations. We only consider situations where the total number of triggered stations is greater than or equal to the number of stations providing magnitude estimates, which is in turn greater than or equal to the number of stations providing peak ground shaking estimates. We consider up to 5 stations reporting each of the observational parameters and generate a 1000-run Monte Carlo simulation for each of the 1086 combinations of station contributions.

The parameters for each of the resultant distributions are kept in an internal ElarmS library, to choose from for any given scenario. For example, when ElarmS makes a prediction of impending ground motions, based on a location derived from 2 station trigger times, 2 magnitude estimates (one with 2 seconds of data and one with 1 second of data), and 0 peak ground shaking observations, it reports that the expected error in the predicted shaking 100 km from the epicenter is -0.13 ± 0.45, determined from the appropriate (pre-calculated) Monte Carlo simulation. The error distribution for this scenario is shown as the red line in Figure 2.9a, along with the distributions for two other

example scenarios. The green line is the PGA error distribution given 3 stations contributing to location, 2 magnitude estimates (one with 3 seconds of data and one with 2 seconds of data), and 1 peak ground shaking observation. The blue line is the distribution given 5 stations contributing to location, 5 magnitude estimates (4 with four seconds and 1 with three seconds), and 3 peak ground shaking observations. Table 2.2 lists the errors predicted by the model for the example event in section 4.8, at various distances from the hypocenter. Figure 2.9b shows the distributions from all 1086 combinations of station contributions. The mean errors for these distributions range from -0.17 to 0.20, with a median of 0.04. Standard deviations range from 0.32 to 0.56, with a median of 0.39.

2.5.2 Sensitivity Analysis

We perform a sensitivity analysis on the Monte Carlo error distributions, to explore the source of errors in the final ground motion predictions. Each error distribution is generated with the input of error terms for magnitude, location, and attenuation relations. One by one we set each of these error terms equal to zero, leaving the others at their observed values and recalculating all 1086 Monte Carlo simulations.

Figure 2.9c shows the resulting distributions if the magnitude error is zero, meaning the simulated ElarmS magnitude estimate is exactly equal to the catalog magnitude. All error is due to the location estimate and attenuation relations. The mean errors of these distributions range from -0.04 to 0.21, with a median of 0.07. The standard deviations range from 0.31 to 0.49, with a median of 0.37. The mean errors are now positive and no longer centered on zero compared to the complete error model (shown in Figure 2.9b). The standard deviations are 5% lower than those of the complete error model, on average, resulting from the fact that there is reduced uncertainty in the PGA estimates.

Figure 2.9d shows the distributions if the location error is zero. That is, the simulated ElarmS location estimate is exactly equal to the catalog location, and any error is due to the magnitude estimate algorithm and attenuation relations. The mean errors of these distributions range from -0.11 to 0.22, with a median of 0.04. The standard deviations range from 0.28 to 0.50, with median of 0.34. The mean errors are similar to those of the complete error model (Figure 2.9b), while the standard deviations are 13% lower.

Figure 2.9e shows the distributions if the error from the attenuation relation is set to zero. Here the attenuation relations are assumed to be perfect, and the only error comes from the magnitude and location estimates. The mean errors of these distributions range from -0.25 to 0.25, with a median of -0.06. The standard deviations range from 0.15 to 0.44, with a median of 0.20. This is a substantial improvement over the complete error model shown in Figure 2.9b, with a mean still centered around zero, but standard deviations that are 49% lower. From this we conclude that the inherent variability in ground motion at a point with respect to even the best fitting attenuation relation is the largest source of error in the ElarmS prediction of ground shaking.

2.6 Conclusions

a) The scaling relations between P_d and magnitude and between τ_p^{max} and magnitude are clearly evident for this Japanese dataset. This is an important result, given the large number of large (M>6) earthquakes. It implies that the basic ElarmS magnitude algorithms remain robust and useful for large magnitude events. For the entire dataset the average magnitude error and standard deviation is 0.0±0.4. For events with JMA magnitude of 6.0 or greater the error is 0.0±0.5, for M≥7 events it is -0.2±0.5. This indicates a "saturation" effect for the M≥7 earthquake, which is partly due to a saturation in P-wave amplitude and partly due to difficulty in rapid and accurate locations of the large events which are all offshore.

b) Both of the scaling algorithms, P_d and τ_p^{max}, are independently effective at estimating final magnitude from the first few seconds of the P-wave. Combining the two working methods reduces error in the final magnitude estimate when only a small number of stations are reporting and increases the overall robustness of the system. Peak displacement is vulnerable to saturation at the highest magnitudes and initial uncertainty in event location. Maximum predominant period shows more scatter for the low-magnitude events, but is less sensitive to saturation at high magnitudes and uncertainty in location. A system based on both P_d and τ_p^{max} is therefore more robust for all events.

c) While ElarmS estimated the location of the majority of events within 18km, it struggled to accurately locate events that were far offshore. The addition of a new algorithm for determining hypocentral depth improved the location estimates of deep events, but accurate location of deep, distant events remains a challenge. The errors from these poor location estimates affect the final ground motion predictions in this region. They also contribute to errors in the magnitude estimate made from peak displacement observations, since these are scaled by epicentral distance. We note that for current realtime processing in California, these location difficulties are not pertinent, as most events are on- or near-shore, and nearly all are shallower than 20km (Hill*, et al.*, 1990).

d) A new error model for ElarmS provides a library of error distributions generated by Monte Carlo simulations. Every ElarmS prediction of magnitude, location, and ground motions can now be provided with an associated uncertainty, based on the quantity of data contributing to the prediction. Uncertainty estimates are essential for both internal study of the system, and for potential end users, who must decide whether to act on a given prediction. Using Monte Carlo simulations we explore the full range of errors in PGA predictions from the first estimate using 1 sec of P-wave data at the first station to trigger to using P-wave data and PGA observations at 5 stations. The error distributions have mean errors of -0.2 to 0.2 (median 0.0) and standard deviations of 0.3 to 0.6 (median 0.4). A factor of 2 error in the predicted PGA relative to the observed corresponds to an error of 0.7.

e) We find that the most significant contribution of the error in ElarmS' final ground motion prediction comes from the inherent variability in peak ground motion at a given location with respect to even the best fitting attenuation relations. Calculating regional

ground motions from attenuation relations using only the estimated magnitude and location of the earthquake resulted in less error than did the same calculation with the addition of the first 1 or 2 observations of peak ground motion. Only when three or more station observations are combined does their inclusion in ground shaking estimation improve the accuracy of the predictions.

The continued improvement of the ElarmS methodology increases its utility in California and in other regions. An accurate, prompt, and reliable early warning system in any seismic setting has the potential to reduce loss of life and money during a damaging earthquake. The developments from this study bring ElarmS one step closer to providing reliable realtime warnings to the public.

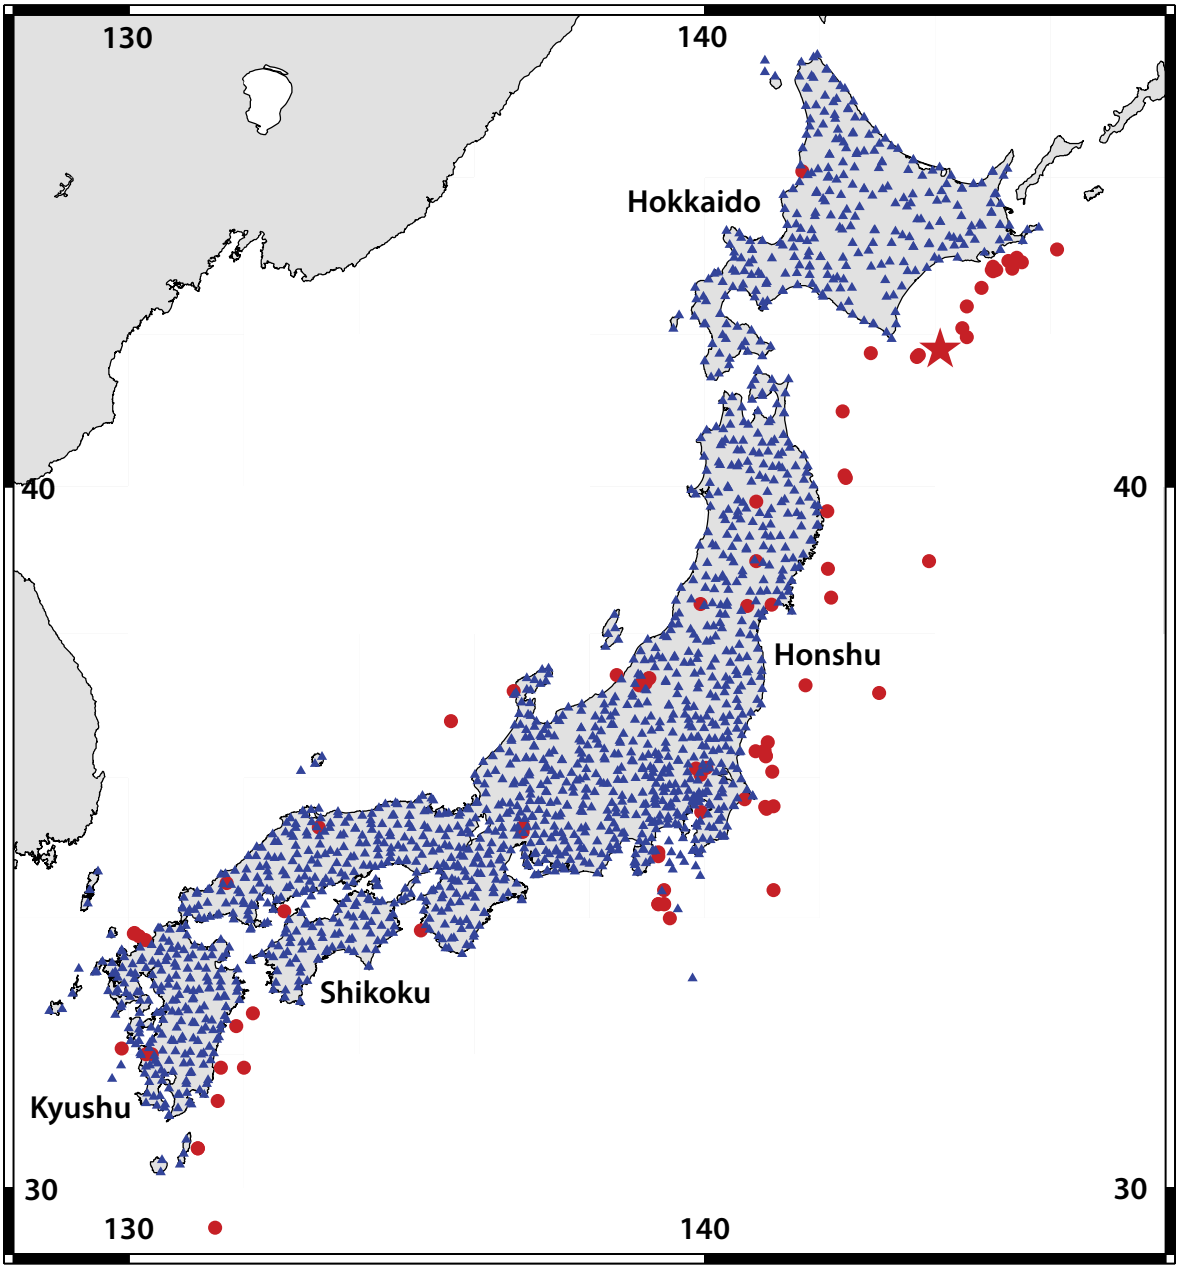

Figure 2.1: Location map. Red circles are events used in this study, blue triangles are K-NET stations. The red star is the largest event in our dataset, the M8.0 Tokachi-Oki earthquake of 2003.

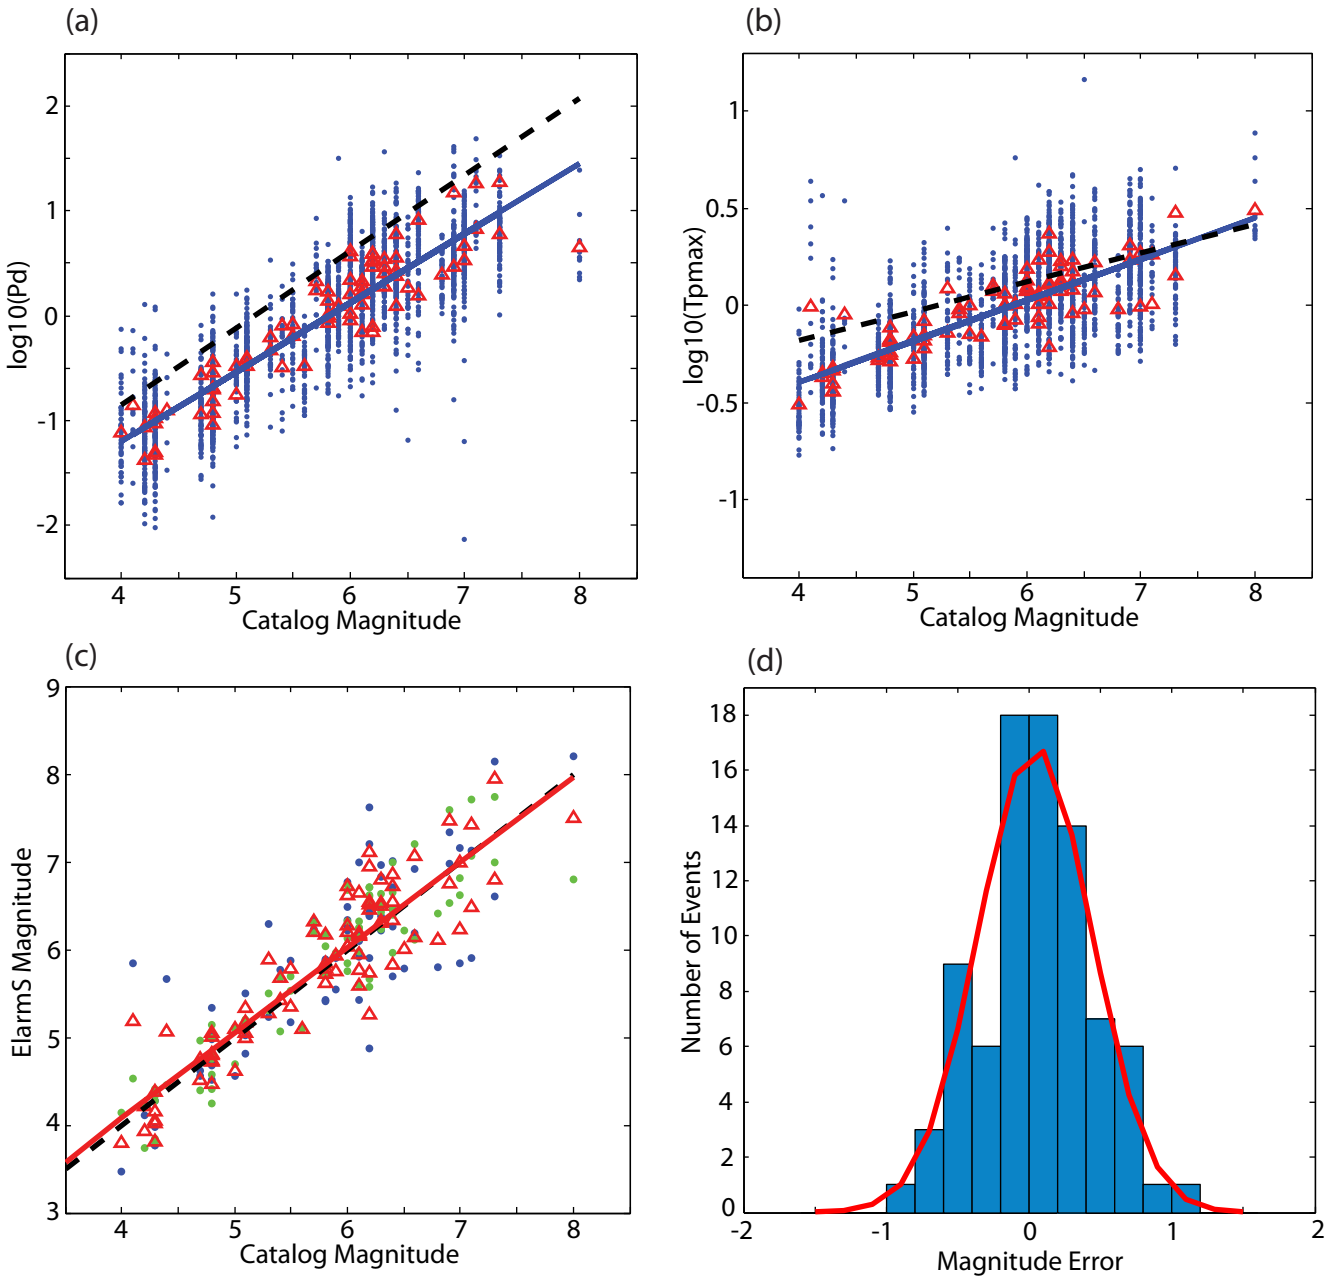

Figure 2.2: Scaling Relations. (a) Scaling relations for peak displacement. Blue circles are log10(Pd) values observed at individual stations and corrected for distance. Red triangles are average peak displacements for each event. Solid blue line is the linear best fit to this data, log10(Pd) = 0.66*ML – 4.02. Dashed black line is the linear best fit for Northern California, log10(Pd) = 0.73*Mw – 3.77(Wurman, et al., 2007). (b) Scaling relations for maximum predominant period. Blue circles are log10(τpmax) values observed at individual stations. Red triangles are average τpmax for each event. Solid blue line is linear best fit to this data, log10(τpmax) = 0.21*ML – 1.22. Dashed black line is linear best fit for Northern California, log10(τpmax) = 0.15*Mw – 0.78 (Wurman, et al., 2007). (c) ElarmS magnitude for each event. The green circles are magnitudes using only Pd, the blue circles are magnitudes using only τpmax , and each red triangle is the average of the Pd and τpmax magnitudes for that event. The solid red line is the linear best-fit to the average magnitudes (red triangles). The black dashed line is the ideal 1-1 fit, for which every ElarmS magnitude exactly equals the catalog magnitude for that event. (d) The histogram of errors in the average ElarmS magnitude estimates from (c). The red line is the best-fit Gaussian distribution for the magnitude errors and has a mean of 0.0 and standard deviation of 0.4.

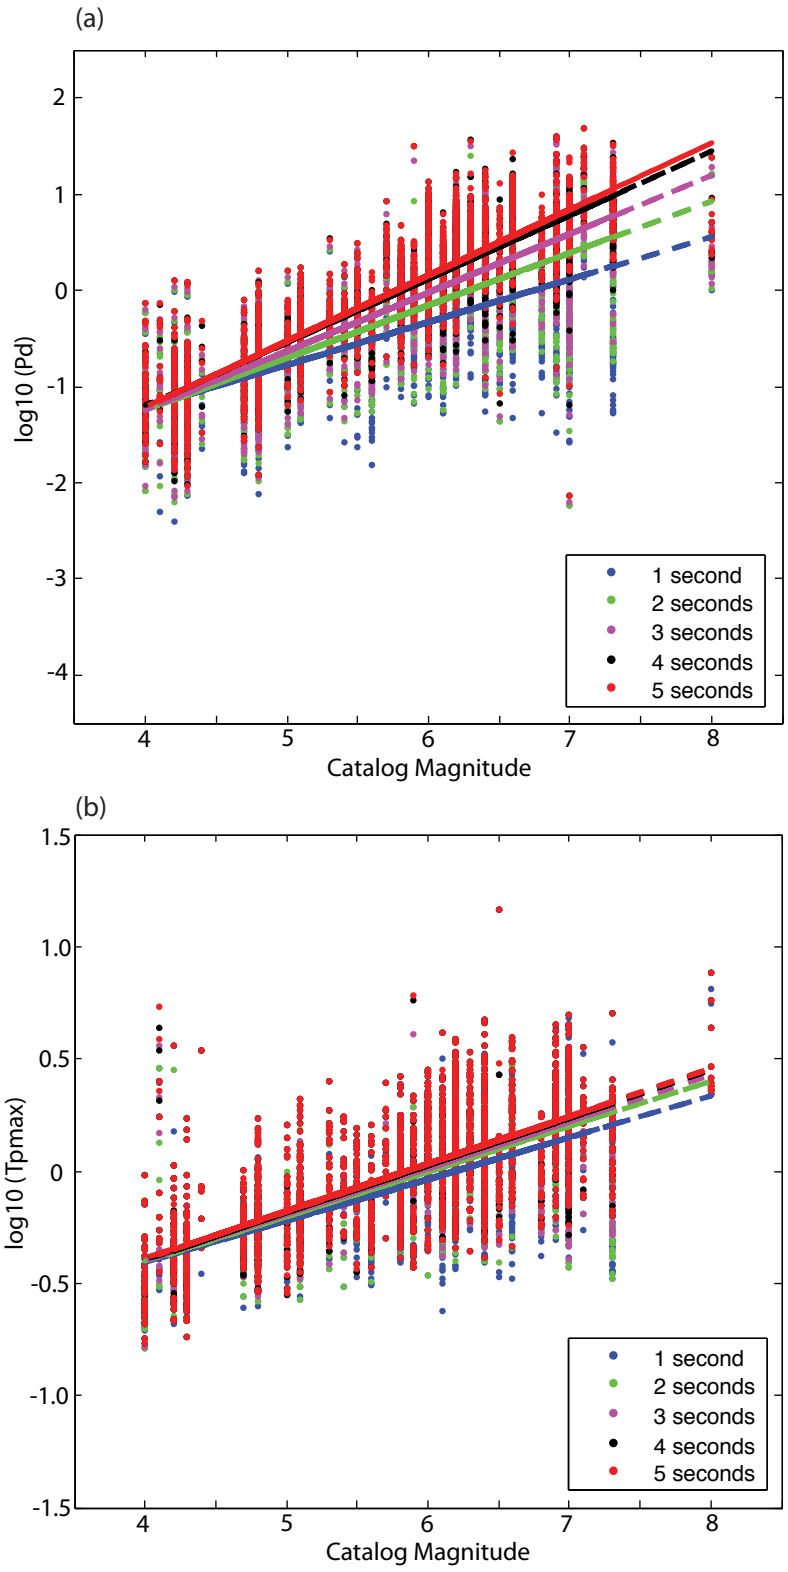

Figure 2.3: Effect of number of seconds of P-wave data on scaling relations. Circles are observations at individual stations, and lines are linear best-fit scaling relations to circles of the same color. Blue circles are observations made using only one second of P-wave data at each station; blue line is linear best fit scaling relation using only one second of data. Green is two seconds of data, purple is three seconds, black is four seconds, and red is five seconds. (a)Effect of number of seconds of P-wave data on peak displacement (Pd) scaling relations. (b)Effect of number of seconds of P-wave data on maximum predominant period (τpmax) scaling relations.

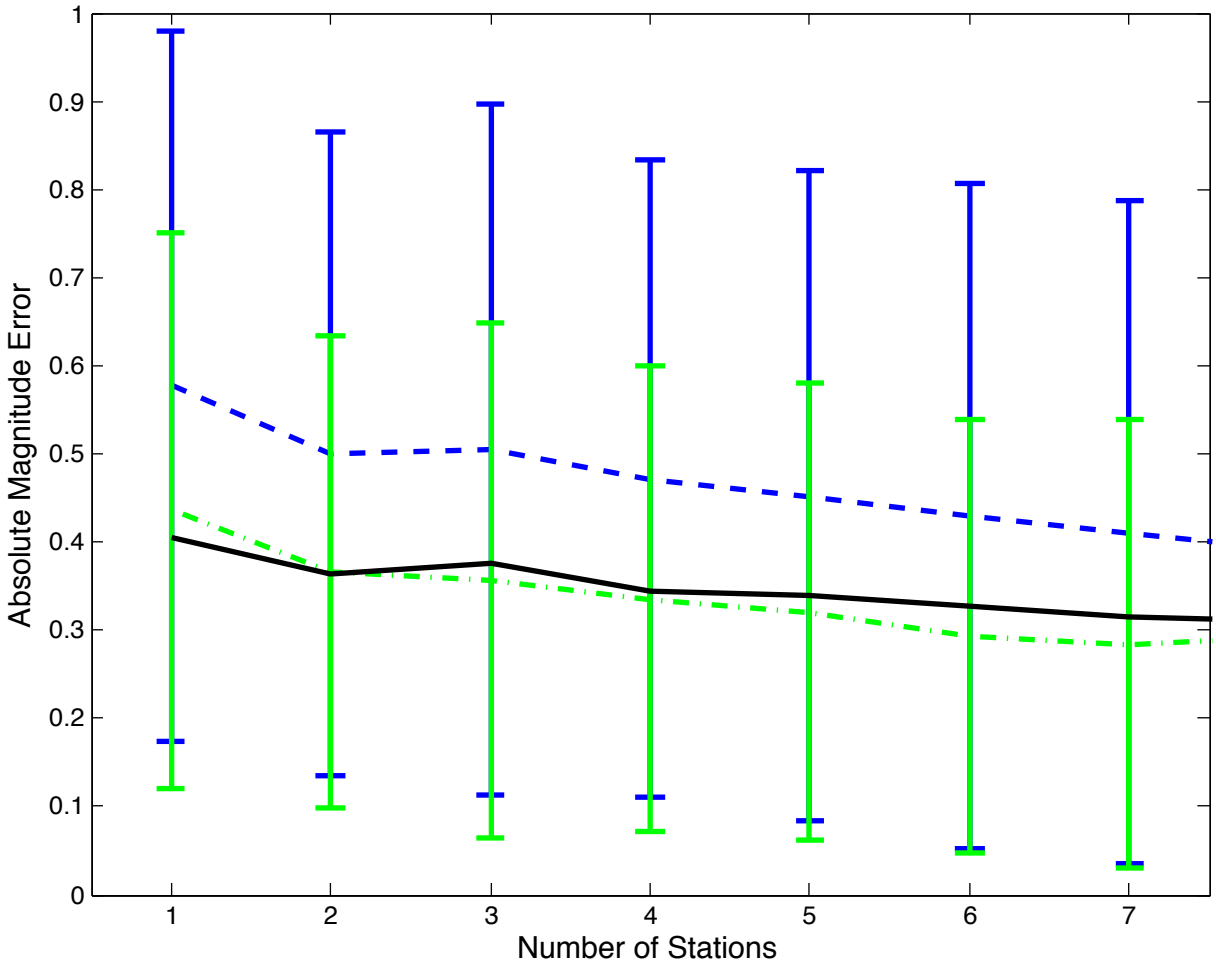

Figure 2.4: Effect of number of stations on the accuracy of magnitude estimates. The blue dashed line is the average error in magnitude estimates made using only maximum predominant period, τpmax. The green dashed line is the average magnitude error using only peak displacement, Pd. The solid black line is the error when the two estimates are averaged.

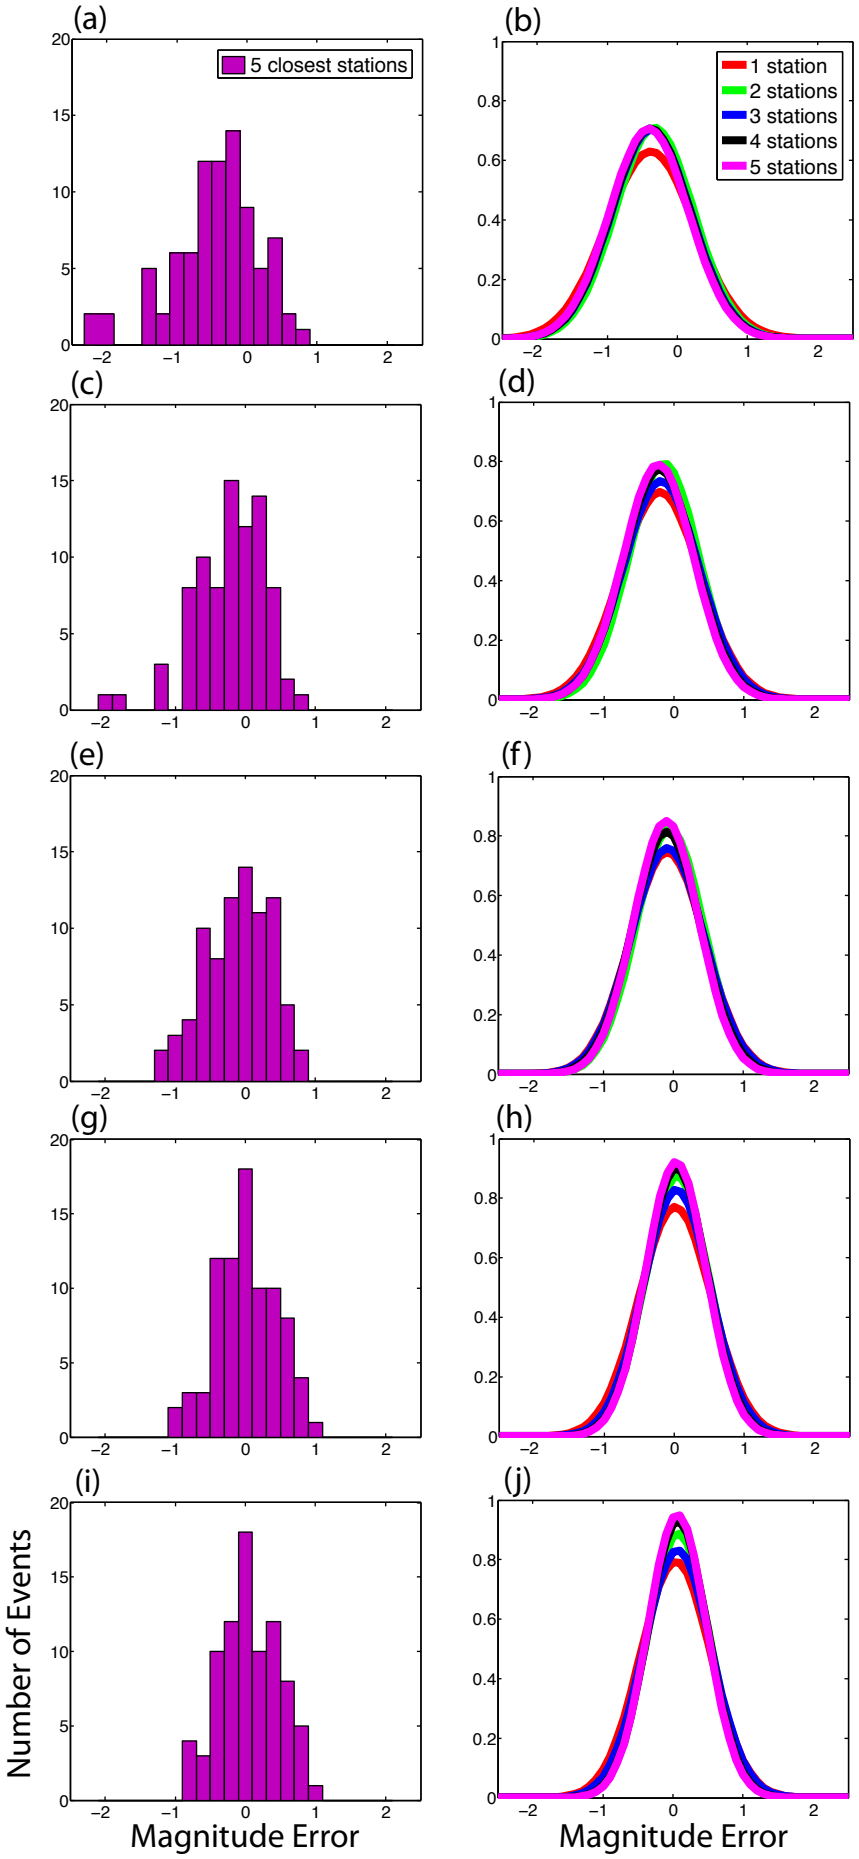

Figure 2.5: Error histograms for magnitude, and corresponding probability distributions showing the effects of the number of stations and number of seconds of data. Each row contains one sample histogram and five probability distributions for 1, 2, 3, 4 and 5 stations providing a specific number of seconds of P-wave data. The colors within each row indicate the number of stations used. Red is one station, green is two stations, blue is three stations, black is four stations, and purple is five stations. (a) Magnitude error when the first one second of P-wave arrival is used from the closest five stations. (b) Error distributions using 1 to 5 stations, all with one second of P-wave each. (c, d) Error using the first two seconds of P-wave at each station. (e,f) Error using the first three seconds. (g,h) Error using the first four seconds. (i,j) Error using the first five seconds.

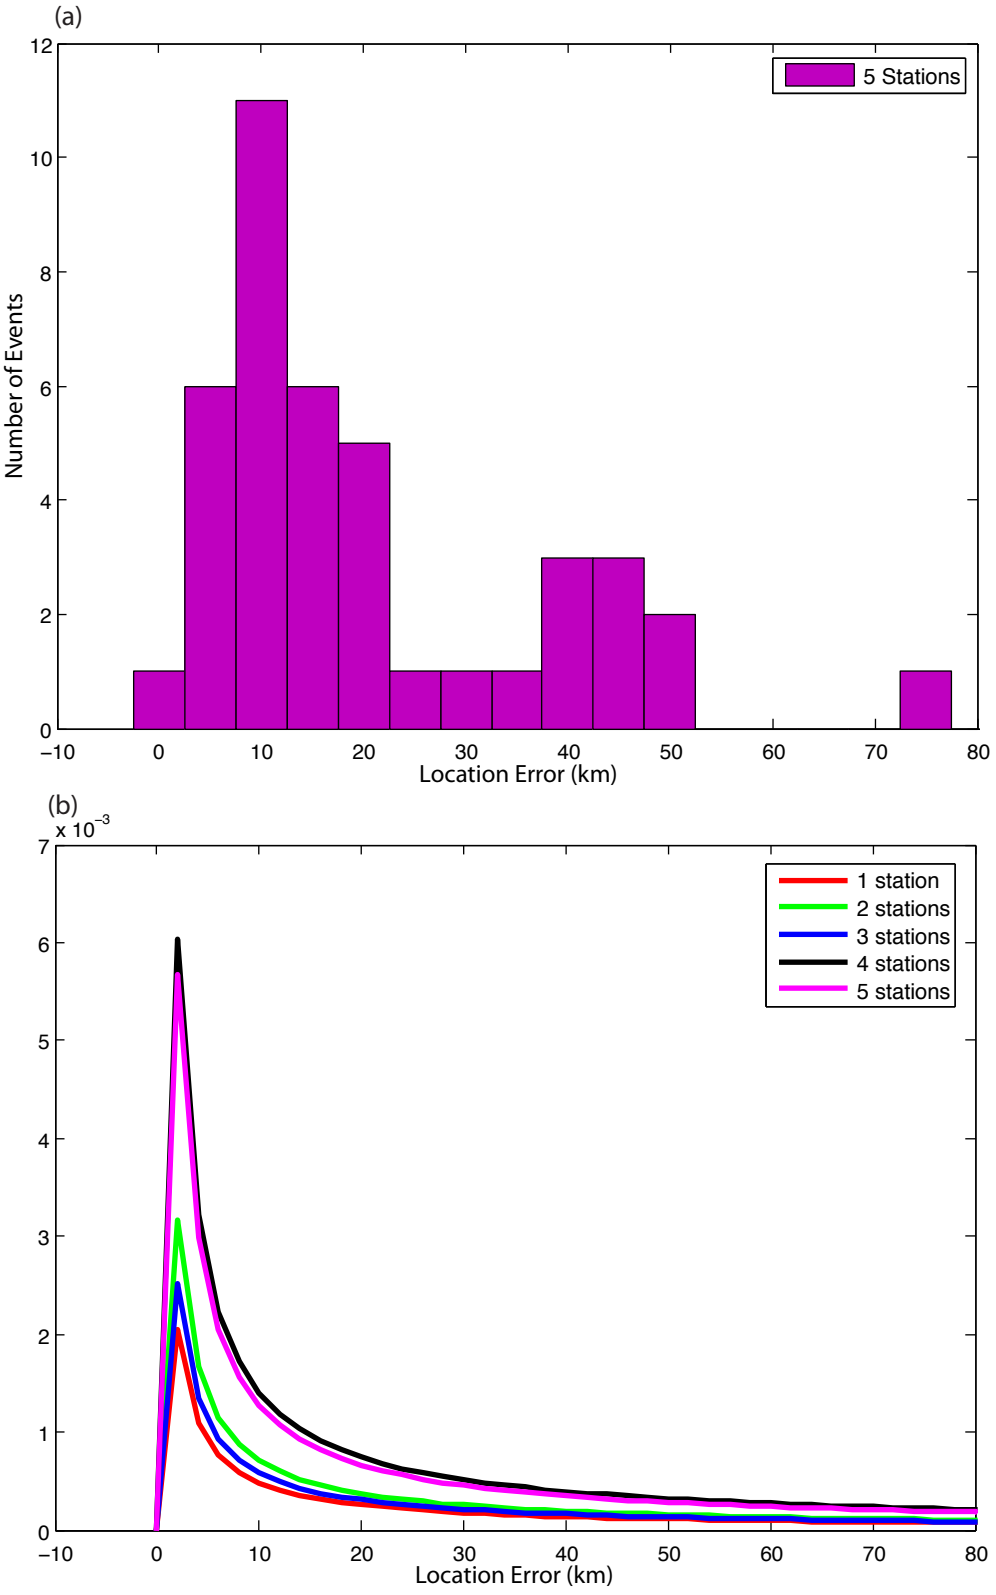

Figure 2.6: (a) Histogram of error in location estimate, when five stations provide P-wave arrival times. (b) Best-fit log-normal distributions to location errors as a function of the numbers of stations providing trigger times. Red is error in location estimates when only one station trigger is used to estimate event location. Green is error when two stations are used. Blue uses three stations, purple uses four stations, and black uses five stations.

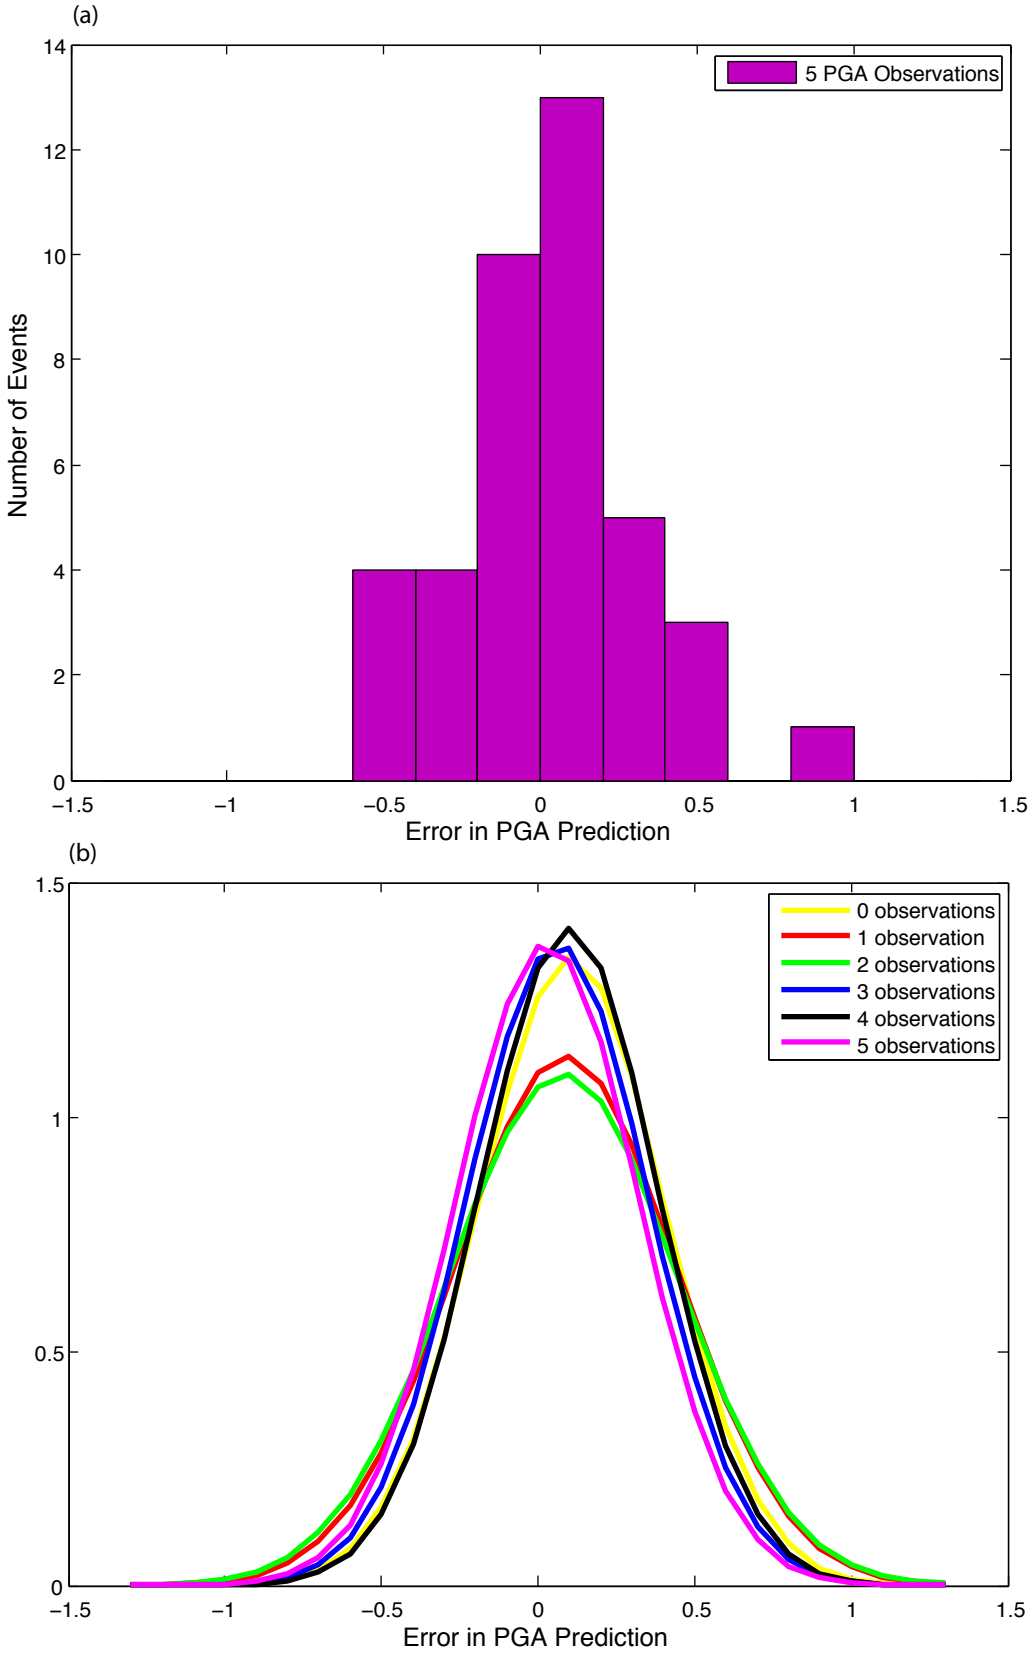

Figure 2.7: (a) Sample histogram of error in ground acceleration estimate, when five stations provide PGA observations. (b) Best-fit Gaussian distributions to the PGA errors when various numbers of stations provide PGA observations. Yellow is error in estimates made with zero observations of peak ground shaking. Red is error using one station observation of peak ground shaking. Green uses two observations, blue uses three observations, black four observations, and purple five observations.

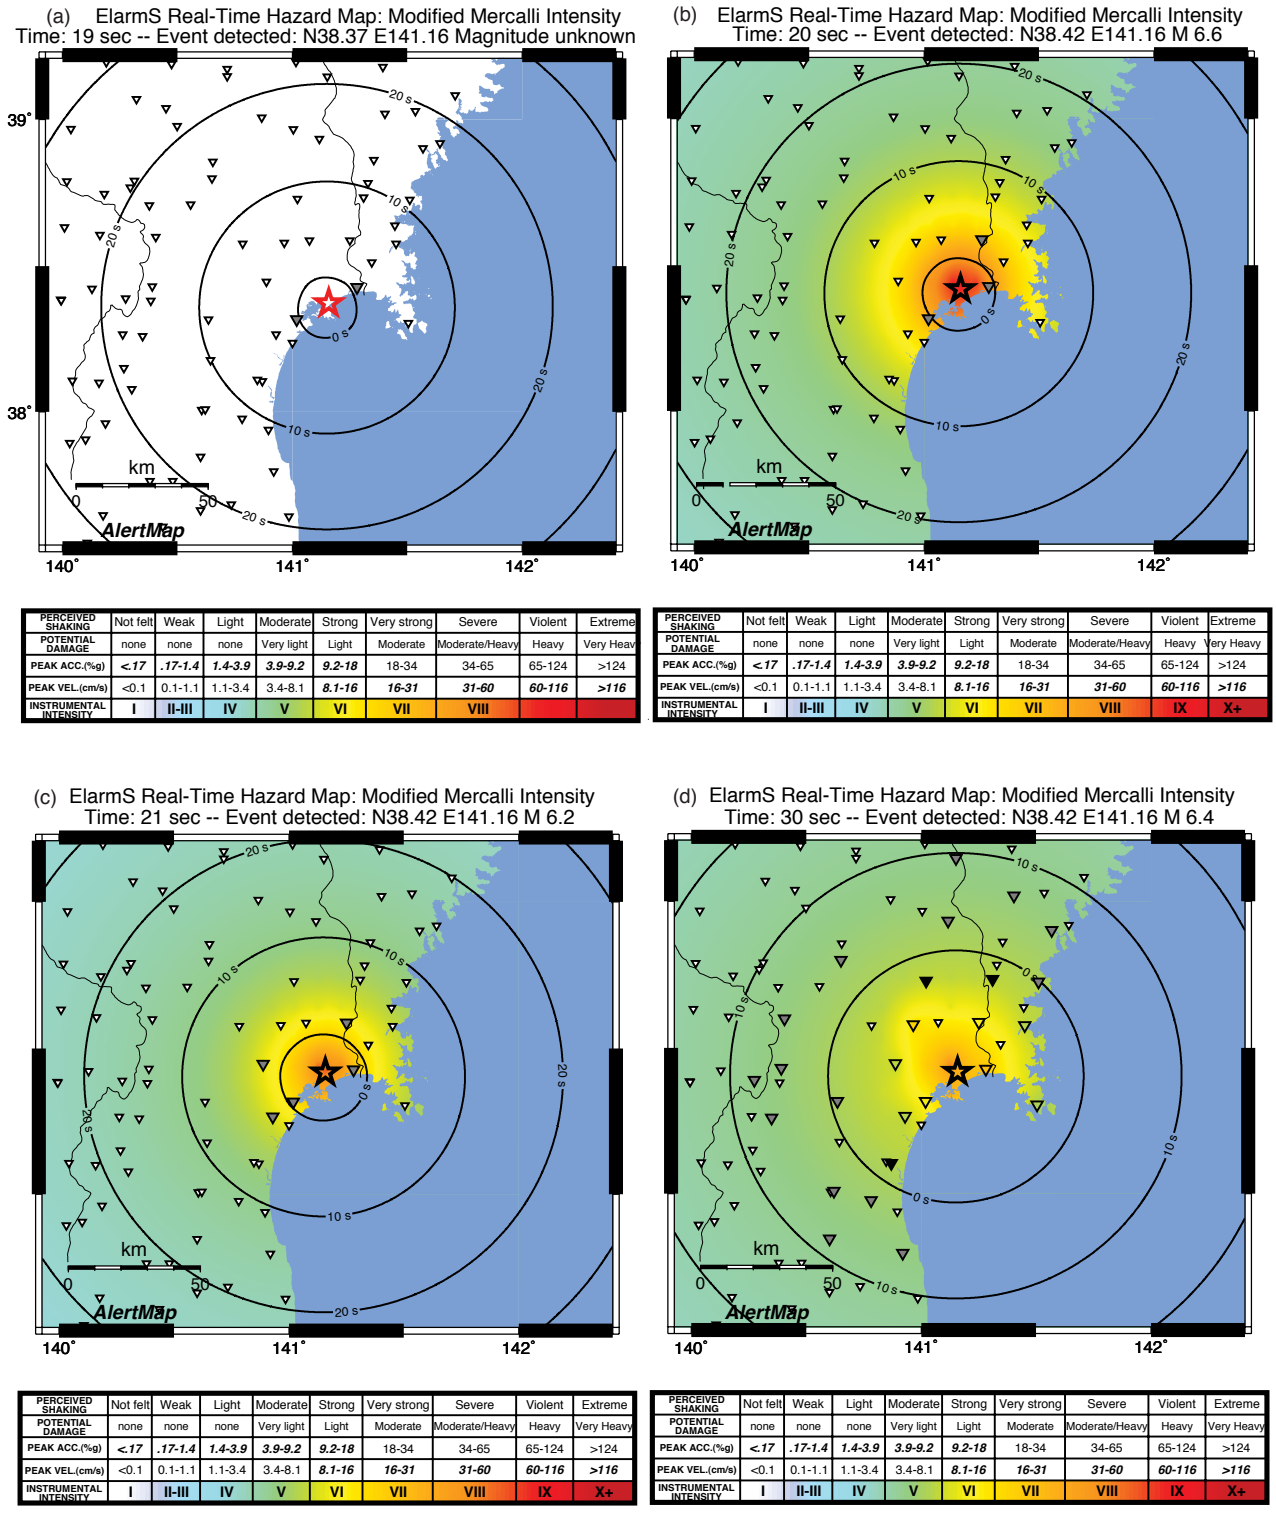

Figure 2.8: AlertMap for an example earthquake, Magnitude 6.4, 26 July 2003, at a depth of 11.9km. (a)The first two triggers occur simultaneously. Event location (star) is set between the triggering stations, at a depth of 8km. Circular contours show the warning times as a function of location. (b) One second later, the first magnitude estimate, M6.6, is available and translated into ground shaking intensity across the region. A third station triggers, and the event epicenter is located by triangulating between the three triggering stations. Depth remains set at 8km. (c) One second later, all three stations are now contributing to the magnitude estimate, which decreases to M6.2. Two more stations trigger, and the five stations total are used to estimate an event hypocenter, at a depth of 10km. (d) The final ElarmS AlertMap, twelve seconds after the first trigger. Magnitude is M6.4. 26 stations are contributing to the magnitude, location, and ground motion estimates.

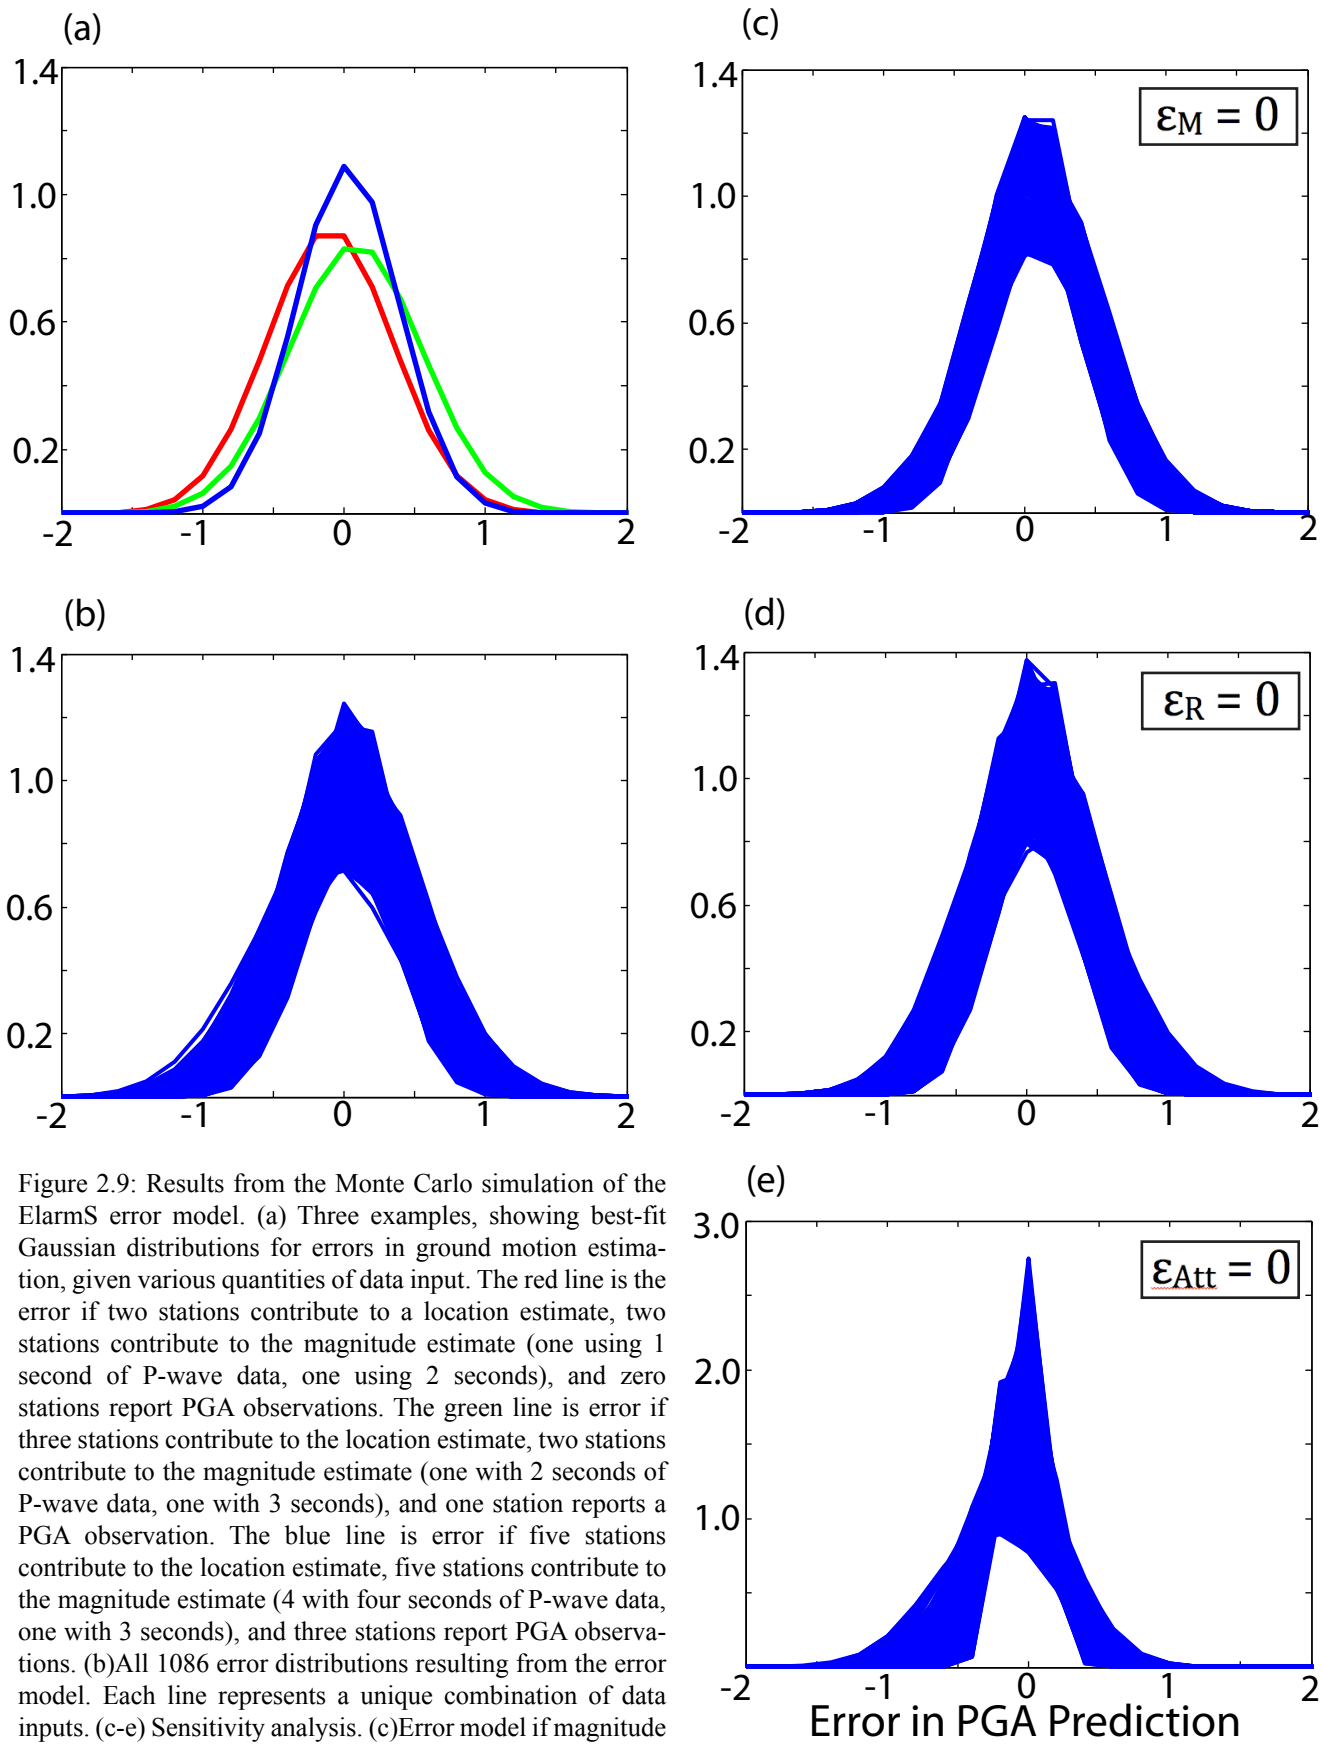

Figure 2.9: Results from the Monte Carlo simulation of the ElarmS error model. (a) Three examples, showing best-fit Gaussian distributions for errors in ground motion estimation, given various quantities of data input. The red line is the error if two stations contribute to a location estimate, two stations contribute to the magnitude estimate (one using 1 second of P-wave data, one using 2 seconds), and zero stations report PGA observations. The green line is error if three stations contribute to the location estimate, two stations contribute to the magnitude estimate (one with 2 seconds of P-wave data, one with 3 seconds), and one station reports a PGA observation. The blue line is error if five stations contribute to the location estimate, five stations contribute to the magnitude estimate (4 with four seconds of P-wave data, one with 3 seconds), and three stations report PGA observations. (b)All 1086 error distributions resulting from the error model. Each line represents a unique combination of data inputs. (c-e) Sensitivity analysis. (c)Error model if magnitude estimate contains no error. (d)Error model if location estimate contains no error. (e)Error model if ground motion estimate contains no error.

Table 2.1: Mean ± standard deviation of error distributions used by error model.						
	0 stations	1 station	2 stations	3 stations	4 stations	5 stations
Mag, 1 sec	-	-0.38 ± 0.63	-0.33 ± 0.56	-0.37 ± 0.57	-0.39 ± 0.56	-0.41 ± 0.56
Mag, 2 sec	-	-0.2 ± 0.57	-0.15 ± 0.5	-0.18 ± 0.54	-0.21 ± 0.52	-0.22 ± 0.50
Mag, 3 sec	-	-0.09 ± 0.53	-0.05 ± 0.48	-0.08 ± 0.52	-0.10 ± 0.49	-0.10 ± 0.47
Mag, 4 sec	-	0.01 ± 0.52	0.04 ± 0.46	0.03 ± 0.48	0.03 ± 0.44	0.02 ± 0.43
Mag, 5 sec	-	0.04 ± 0.50	0.07 ± 0.45	0.07 ± 0.48	0.07 ± 0.43	0.06 ± 0.42
Location	-	33.6 ± 17.9	32.1 ± 21.4	32.5 ± 18.7	18.8 ± 13.6	21.1 ± 16.8
PGA	0.11 ± 0.30	0.09 ± 0.35	0.08 ± 0.37	0.06 ± 0.29	0.10 ± 0.28	0.03 ± 0.30

Table 2.2: PGA Errors for Example Event, 26 July, 2003			
	(b) 1 second	(c) 2 seconds	(d) 12 seconds
Mean Observed Error	0.05 ± 0.25	-0.05 ± 0.25	0.13 ± 0.26
Predicted Error, 20km	-0.09 ± 1.38	-0.20 ± 1.36	-0.18 ± 1.37
Predicted Error, 50km	-0.13 ± 0.77	-0.11 ± 0.56	-0.11 ± 0.64
Predicted Error, 100km	-0.01 ± 0.39	-0.02 ± 0.38	0.01 ± 0.37

Chapter 3

Development of the ElarmS Methodology

This chapter has been published in *Soil Dynamics and Earthquake Engineering* as "Development of the ElarmS methodology for earthquake early warning: Realtime application in California and offline testing in Japan", by Brown, Allen, Hellweg, Khainovski, Neuhauser and Souf, 2011.

Preface

The ElarmS development described in this chapter required considerable effort by all members of the Berkeley Seismological Laboratory EEW group. Consequently this chapter, and the associated *Soil Dynamics and Earthquakes Engineering* paper, represents work done by the entire group. The dissertation author, Holly Brown, wrote the chapter and compiled the figures, and was individually responsible for the research pertaining to system latencies and testing ElarmS in Japan, including the subsequent upgrade of the location algorithm from two dimensions to three.

3.1 Introduction

Earthquake early warning (EEW) systems are algorithms that detect the initial P-waves from an earthquake, rapidly estimate the location and magnitude of the event, and then predict subsequent ground shaking in the surrounding region. EEW systems offer the potential for a few seconds to a few tens of seconds warning prior to hazardous ground shaking: enough time for individuals to get to a safe location, perhaps under a sturdy table, for shutdown of utilities, slowing of trains, and other automated steps to reduce hazards from ground shaking.

In July 2009, the California Integrated Seismic Network (CISN) completed a three-year investigation into the viability of an EEW system in California. Three algorithms were expanded, tested, and compared during the study: Onsite, a single-station method that uses τ_c and P_d (Böse, *et al.*, 2009), Virtual Seismologist, a network-based method that uses peak amplitudes and Bayesian statistics (Cua, *et al.*, 2009), and ElarmS, a network-based method that uses τ_p^{max} and $P_{d/v}$ (Allen, *et al.*, 2009).

The goal of the three-year project was to determine whether EEW is feasible in California. Results from each algorithm were continuously reported to a central database run by the Southern California Earthquake Center (SCEC) for analysis. By the end of the three years, all three algorithms had successfully predicted ground shaking before it was

felt for many earthquakes in the state. At the end of the study the CISN determined that EEW is feasible, potentially desirable, and within reach for California. In August 2009 a second three-year study was initiated, to integrate the three test algorithms into a single prototype EEW system and provide realtime warning to a small group of test users by the end of the study in summer 2012.

Here we delineate the methodology, progress, and results of the ElarmS algorithm, which is now an integral part of the forthcoming prototype CISN EEW system. The ElarmS algorithms for magnitude and location estimation were developed offline with two datasets of events from Northern and Southern California. Those algorithms are now used in realtime, continuously processing waveforms from throughout the state of California and producing predictions of ground shaking within seconds of event detection. A separate dataset of events from Japan was processed offline to test ElarmS' performance for large events. From the Japan results we developed an error model which can be used in realtime to estimate the uncertainty in any ElarmS prediction.

3.2 Development and Methodology

3.2.1 Overview

Earthquake Alarm Systems, or ElarmS, is a network-based EEW system. The algorithm detects P-wave arrivals at several stations around an event epicenter and uses the amplitude and frequency content of the P-wave to rapidly estimate the magnitude and hypocenter of the event. Estimates from several stations are combined to improve accuracy and minimize the chance of a false alarm. ElarmS then applies the estimated magnitude and location to CISN ShakeMap regional ground motion prediction equations (GMPEs) to produce a realtime prediction of impending ground shaking. Predictions above a certain threshold prompt an automatic alert message that can be sent to users.

The ElarmS algorithm is divided into a waveform processing module and an event monitoring module. The waveform processing module analyzes raw waveforms from all contributing stations, detects P-wave arrivals, and calculates the necessary ElarmS parameters: predominant period, peak amplitudes, signal-to-noise ratio (SNR), peak ground acceleration and velocity (PGA and PGV), and trigger times. These parameters are then passed to the event monitor, which associates the triggers into an event, estimates the event location, estimates the magnitude, and predicts ground shaking. As additional stations record P-wave arrivals, the waveform processing module passes their parameters to the event monitor, which includes them into the event analysis (Allen, 2007; Allen, *et al.*, 2009).

3.2.2 Location

Event location is estimated by a four-stage algorithm, defined by the number of station triggers. When a single station triggers, the event is located directly beneath the station, at a depth of 8km. When two stations have triggered, the event is located between them based on arrival times, again at a depth of 8km. When three stations have triggered, ElarmS uses a two-dimensional grid search at a depth of 8km to determine the hypocenter and origin time that minimizes arrival time residuals. Finally, once four or more stations have triggered, ElarmS performs a three-dimensional grid search, with depth intervals

every 10km, to estimate the hypocenter and origin time that minimizes arrival time residuals. In California, most events occur at depths of 5-15 km and the average depth is 8km (Hill, *et al.*, 1990). Rather than determining depth, ElarmS sets the depth of all California earthquakes to 8km. When processing events in Japan, all four stages are used including the depth determination.

3.2.3 Magnitude

ElarmS was originally developed from an empirically observed relationship between maximum predominant period, τ_p^{max}, and final event magnitude (Allen, 2004; Allen and Kanamori, 2003; Lockman and Allen, 2005; Olson and Allen, 2005). For any vertical channel (broadband HHZ, or strong motion HLZ, HNZ), the predominant period time series is defined recursively by:

$$\tau_{p,i} = 2\pi (X_i/D_i)^{½}$$

where $X = \alpha X_{i-1} + x_i^2$ and $D_i = \alpha D_{i-1} + (^{dx}/_{dt})_i^2$. The constant α is a smoothing factor equal to *1 - dt*, where *dt* is the sample interval, and x_i is the ground velocity of the last sample. Acceleration waveforms are integrated to velocity first, and all waveforms are filtered with a causal 2-pole, 3-Hz, low-pass Butterworth filter. τ_p^{max} is then the maximum observed τ_p value during the first four seconds of P-wave arrival.

To determine the empirical scaling relations, all τ_p^{max} values for a given region are plotted against the final magnitude of each event. A least squares fit to the data produces the scaling relation, which is then used in realtime to estimate magnitude (see section 3.1).

In 2007 ElarmS was updated to utilize a second P-wave parameter, the peak amplitude (Wurman, *et al.*, 2007). As before, vertical-component waveforms are filtered with a 3Hz low-pass Butterworth filter. Peak amplitudes observed during the first four seconds of P-wave arrival are scaled to an epicentral distance of 10km and compared to the final catalog magnitude for the event. A least squares fit to the data provides a scaling relation for the region. Note that the peak amplitude scaling relations are dependent on the epicentral distance of the amplitude observation. In Northern California, peak displacement is used for broadband (HH) instruments and peak velocity is used for strong motion (HL and HN) instruments. Peak displacement has a theoretically longer period signal and thus less high frequency noise than peak velocity, but numerically integrating the acceleration signal twice (from acceleration to velocity, and again from velocity to displacement) introduces errors. We found that for acceleration instruments in Northern California, peak velocity provides a more robust scaling relation than does peak displacement. In Southern California and Japan, peak displacement produced the strongest scaling relation for all instruments, despite the double integration from acceleration. In general, we refer to the peak amplitude scaling relations as $P_{d/v}$ with the understanding that we may use P_d or P_v for any given site.

Although the scaling relations for τ_p^{max} and $P_{d/v}$ are determined using four seconds of P-wave arrival, waiting for a full four seconds of P-wave to be available during realtime processing wastes valuable seconds of potential warning time. Instead ElarmS begins to apply the scaling relations and estimate magnitude as soon as a single station has observed a single full second of P-wave arrival (the first half-second is discarded). As

additional seconds of P-wave become available, ElarmS recalculates τ_p^{max} and $P_{d/v}$ accordingly. Since both τ_p^{max} and $P_{d/v}$ are the maximum or peak values, they can only increase with additional seconds of data. The initial one-second magnitude estimate is therefore always a minimum estimate.

To ensure that early arriving S-waves at near-field stations do not interfere with the magnitude estimate which is P-wave based, ElarmS also utilizes a simple P/S filter, based on an S-P moveout of 8km/s (with a minimum S-P time of 1 second, assuming most events are 8km deep). The S-P time is estimated at each station given the event location and the P-waveform is only used up to the S-wave arrival. One potential drawback of this filter is that location errors may cause valid P-wave data to be discarded as misidentified S-waves.

For each triggering station, τ_p^{max} and $P_{d/v}$ are scaled separately to create two independent estimates of magnitude. The estimates are then averaged to form a single event magnitude for that station. As additional stations report P-wave triggers, their magnitude estimates are averaged into the event magnitude, to provide an increasingly accurate description of the event as time passes.

3.2.4 Ground Motions

Once location and magnitude have been estimated for an event, ground motion is predicted at each triggered station by applying the location and magnitude to CISN-defined ShakeMap GMPEs for the region (Wald, *et al.*, 1999a). The resulting "AlertMap" displays predicted ground shaking in the familiar ShakeMap format, i.e. a map of predicted shaking intensity. As peak ground shaking is observed at individual stations, the observations are integrated into the shaking intensity map. ElarmS incorporates a bias correction by scaling the GMPE up or down to best-fit the available observations. Eventually, when all stations have reported peak ground shaking, the AlertMap looks much the same as the post-event ShakeMap.

The ElarmS algorithm has been tested with datasets from Northern California, Southern California, and Japan (Allen, 2006; Allen, 2007; Allen and Kanamori, 2003; Allen, *et al.*, 2009; Brown, *et al.*, 2009; Lockman and Allen, 2005; Olson and Allen, 2005; Tsang, *et al.*, 2007, Wurman, *et al.*, 2007). Each test dataset provided regional scaling relations for τ_p^{max} and $P_{d/v,}$ and utilized GMPEs specific to that location. Most recently ElarmS has been adapted to run in realtime throughout the state of California.

3.3 Application of ElarmS to California

3.3.1 Scaling and GMPEs

Offline tests of California earthquake datasets have produced separate scaling relations for Northern and Southern California events (Tsang, *et al.*, 2007; Wurman, *et al.*, 2007). The magnitude scaling relations are determined empirically by comparing observed τ_p^{max} and $P_{d/v}$ values to final catalog magnitude for a dataset of test events, with as wide a range of magnitudes as possible. Once determined, the scaling relations are used in realtime to estimate event magnitude, based on realtime observations of P-wave frequency and amplitude.

For northern California, Wurman, *et al.,* analyzed a dataset of 43 events recorded by Berkeley Digital Seismic Network (BK) and Northern California Seismic Network (NC) seismometers (Figure 3.1) between 2001 and 2007, with magnitudes ranging from 3.0 to 7.1. The analysis resulted in the following scaling relations:

$M_w = 5.22 + 6.66*\log_{10}(\tau_p^{max})$ for τ_p^{max} on HHZ, HLZ, HNZ channels
$M_w = 1.04*\log_{10}(P_d) + 1.27*\log_{10}(R) + 5.16$ for P_d on HH channels
$M_w = 1.37*\log_{10}(P_v) + 1.57*\log_{10}(R) + 4.25$ for P_v on HL channels
$M_w = 1.63*\log_{10}(P_v) + 1.65*\log_{10}(R) + 4.40$ for P_v on HN channels

where R is the epicentral distance to the station. The τ_p^{max} and P_d relations are shown in Figure 3.2ab. These scaling relations are now used by ElarmS for all events north of the Gutenberg-Byerly line (shown on Figure 3.1 as the line between regions mSA/eCAn and BB/eCAs).

For southern California, Tsang, *et al.*, analyzed a dataset of 59 earthquakes recorded by the Southern California Seismic Network (CI) between 1992 and 2003, with magnitudes ranging from 3.0 to 7.3. The analysis resulted in the following scaling relations (Figure 3.2cd):

$M_w = 6.36 + 6.83*\log_{10}(\tau_p^{max})$ for τ_p^{max} on HHZ, HLZ, HNZ channels
$M_w = 1.24*\log_{10}(P_d) + 1.65*\log_{10}(R) + 5.07$ for P_d on HH, HL, HN channels

These scaling relations are used by ElarmS for all events south of the Gutenberg-Byerly line.

Ground motions in Northern and Southern California are predicted using the Boatwright, *et. al.*, GMPE, as preferred by CISN ShakeMap version 3.2:

$\log_{10}(PGA, PGV) = A + B*(M-M_s) - \log_{10}(R_g) + k*R + B_v*\log_{10}(V_s/V_a)$

where M is event magnitude, V_s is a site correction, $R = \sqrt{(R_e^2 + d^2)}$, R_e is epicentral distance, d is depth, and $R_g = R$, if $R \leq R_0$, or $R_g = R_o*(R/R_0)^g$, if $R > R_0$. Remaining coefficients are specific for large events (M>5.5) or small events (M≤5.4), and are shown in Table 3.1.

3.3.2 Realtime Processing

ElarmS was adapted to run in realtime in Northern California in October 2007, and expanded statewide in November 2008. The system now processes waveforms from all realtime-capable stations in the state: a total of 603 velocity and accelerations sensors at 383 sites (Figure 3.1). The ElarmS waveform processing module is distributed among three regional processing centers, which receive the continuously streamed waveforms. Data from the Berkeley Digital Seismic Network (BK) are streamed to UC Berkeley, data from the Northern California Seismic Network (NC) and from some stations in the USGS Strong Motion Network (NP) are streamed to USGS Menlo Park, and data from the Southern California Seismic Network (CI), the Anza Network (AZ), and the remaining NP stations are streamed to Caltech/USGS Pasadena. At these regional processing centers, the waveform processing module distills the waveforms to their essential parameters: trigger times, peak predominant period, peak amplitudes (acceleration, velocity, and displacement), peak ground shaking observations, and signal-to-noise ratio. These parameters are then forwarded to UC Berkeley, where a single event monitor

integrates data from all of California to identify and analyze earthquakes in realtime. When an event is determined to be above a certain magnitude threshold, an alert message can be sent to users notifying them of the event location, origin time, estimated magnitude, and number of triggers. Currently alerts are sent to the authors and the SCEC database for CISN EEW analysis.

3.3.3 System Latency

The total ElarmS processing time, from when a P-wave arrives at a station until ElarmS outputs event information, can be divided into two types: telemetry of data and computer analysis time. Data telemetry includes the time while a station collects data into a packet for transmission, transit from individual stations to the regional processing centers where the waveforms are processed, and transit time from the processing centers to UC Berkeley where the single event monitor is located. Stations transmit data to the processing centers by frame-relay, internet, private intranet, radio, or microwave, depending on the station. The processing centers transmit data to Berkeley by internet or private intranet.

The primary source of telemetry latencies is the packetization of data by station data loggers. A data logger will not send its data to the waveform processing module until the data packet is full. Packet sizes are usually of a configurable byte size, but many station data loggers are currently set for packet sizes equivalent to 4-6 seconds of data. Manually reconfiguring these data loggers to require packets equivalent to 1-2 seconds of data would greatly decrease the delays. In addition, all BK data loggers and most CI data loggers will be upgraded to data loggers with short 1 second packets in the next two years with recently provided US Federal stimulus funding as part of the American Recovery Reinvestment Act (ARRA).

Figure 3.3a shows the data latencies for transmission to the waveform processing site by each seismic network. These delays are the difference in seconds between when a P-wave arrives at a station and when the waveform packet is received by the regional processing center. They are thus composed of the time for a packet to fill and the time in transit to the regional processing center. The median latencies for each network are shown in Table 3.2. The median latency across all networks is 5.23 seconds. Each histogram is characterized by an extended tail at the high latencies (the figure is truncated at 20 seconds for clarity, but the distributions continue to higher latencies, up to several hundred seconds, for a small number of stations). The tail indicates stations that are drastically delayed, due to poor telemetry availability, temporary telemetry failure or station disruption.

NC has the fastest median of 2.5 sec due to a large number of NC station data loggers configured for a packet size equivalent to 1-2 seconds of data. However there is a substantial tail to the distribution, indicating that the remaining stations are significantly slower. The Gaussian-like distribution for BK, with a median of 6.2 seconds, illustrates the nearly uniform hardware, software and telemetry configuration for all stations in the network, with few excessively delayed stations. CI uses much the same equipment as BK and shows a similar distribution with a slightly faster median of 5.2 sec. NP is a little

slower with a median of 7.4 sec. The NP distribution shows a peak around 2 or 3 seconds, similar to NC, but a multitude of slower stations add a significant tail to the distribution, increasing the median. AZ has the highest median latency, 9.3 seconds, which is due to an extra telemetry step as the data is forwarded through the Scripps Oceanographic Institute before arriving at the Caltech regional processing center.

Figure 3.3b shows the delays by data logger type, independent of network. Again, the delays are the difference between when a P-wave arrives at a station and when the waveform packet is received by the regional processing center. The distribution statistics are shown in Table 3.2. The fastest data logger is the K2 used at many of the USGS sites and designed to send 1 sec data packets. The Quanterra Q330 comes second, again due to the fact that it sends out 1 sec data packets, although there is a wider range of the total telemetry latencies which is likely due to software discrepancies between the different networks. The Berkeley processing software was designed for the older model data loggers and has not yet been updated to accommodate the Q330. This software will be upgraded by Spring 2010. The older Quanterra data loggers (the Q730, Q680, Q980 and Q4120) are slower. In the network upgrade that is now underway the majority of these older and slower data loggers are being upgraded to Q330s. The combined effect of new dataloggers and revised software will reduce the latencies at these stations by 3 to 5 sec.

3.3.4 Alert Criteria

The station distribution in California is not uniform (Figure 3.1). Not surprisingly, the performance of a network-based system is directly related to the density of the network. Accuracy improves when more stations contribute to an event estimate, but potential warning time is lost while waiting for those stations to trigger, especially when the stations are far apart. ElarmS performs best in the heavily instrumented regions around Los Angeles, San Diego, and San Francisco (LA, sSA, SFBA in Figure 3.1). In these regions the mean station separation is only 20km, and the system often receives two or three triggers in the first second after an earthquake begins. In regions with lower station density the system must wait, as valuable seconds pass, until enough stations have reported P-wave arrivals. Regions with less dense instrumentation also suffer from higher false alarm rates, as there are fewer stations to contradict a false trigger. We therefore tailor the alert requirements to each region.

In regions SFBA, LA and sSA, where inter-station spacing is approximately 20km, the system requires at least 4 triggers within 30 km of the epicenter before an alert can be sent for an event. In southeastern California (eCAs), the Big Bend region (BB), the middle San Andreas (mSA), and the northern San Andreas (nSA), where stations are separated by 20-100km, we require 5 or more stations within 100km to trigger before an alert is generated. And in the Mendocino Triple Junction (MTJ), northeastern California (eCAn), and the Channel Islands (cIS), where stations are more than 100km apart, we require 10 or more stations (at any epicentral distance) to trigger. These regional boundaries and requirements continue to be refined as we monitor the realtime system.

3.3.5 False and Missed Alerts

Figure 3.4 shows all detected, false, and missed alerts with magnitude 3 or greater which occurred in Northern California during a ten-week test period from 8 August 2009 and 20 October 2009. A false alert is defined as an ElarmS event that meets the alert criteria for its region but does not correspond to a an event in the Advanced National Seismic System (ANSS) catalog. A missed alert is an ANSS M>3 event for which no ElarmS alert message was issued; ElarmS may have not detected the event, or it may have detected the event but not satisfied the criteria required to issue an alert. For this ten-week test period there were 63 real events M>3. ElarmS detected 45 of them and missed 18. Eleven of the missed events were part of an aftershock sequence described below. ElarmS also sent four false alert messages for nonexistent events.

The false and missed alarm rates are related to two factors: the station density, and whether an earthquake is occurring during a swarm such as during an aftershock sequence. In the SFBA region, where inter-station spacing is approximately 20km, there were 8 detected events and no false or missed alerts for this time period (Figure 3.4). In mSA there were 3 detected events and 1 false alert. In nSA there were 8 detected events, 1 false alert and 2 missed alerts. Performance is moderate in the mSA and nSA regions as the station spacing is 20-100km.

In the eCAn and eCAs regions performance is much poorer due to the much lower station density. In eCAn there were two missed alerts and one false alert. In the eCAs region in the lower right of the map there is a cluster of green (detected) and red (missed) squares. These represent two M5 events on October 1st and 3rd, and their aftershock sequences. ElarmS successfully detected the M5.1 event on October 1st, but missed the M5.2 event two days later. It caught 20 out of 31 total aftershocks of magnitude 3 or greater. ElarmS missed the second large event due to increased background noise and concurrent aftershock activity from the first event.

This illustrates the challenge of defining optimal alert criteria for each region. Criteria which are too strict (requiring too many stations to trigger) may fail to be met by a moderate size event, resulting in no alert message even though the event is real, or will slow down the time until an alert is issued. Criteria which are too loose (requiring too few stations) may be met by unrelated, erroneous triggers, resulting in an alert message when there is no real event. As with all associators the performance is also reduced during swarms of seismicity or aftershock sequences. Improvements to the associator scheme specifically for early warning applications would be beneficial.

3.4 Sample Events

We illustrate ElarmS performance in California with three sample events from different regions of the state, all processed by the realtime system.

3.4.1 M_w5.4 Alum Rock, SFBA Region

Figure 3.5 shows the M_w5.4 Alum Rock event, which occurred on 30 October, 2007. This was the largest event in the San Francisco Bay Area since the 1989 Loma Prieta M_w6.9 event. At the time of the Alum Rock earthquake, ElarmS had been running in realtime for

less than a month and used only stations from the BK network. The event begins in Figure 3.5a when two stations trigger simultaneously. The location is estimated between the stations, at a depth of 8km. One second later (Figure 3.5b), the magnitude is estimated at 5.2, using the observed τ_p^{max} and $P_{d/v}$ values from the two triggered stations. A third station triggers and the location is triangulated based on the arrival times at the three stations. The estimated location and magnitude are applied to local GMPEs to produce a prediction of ground shaking around the epicenter. The mean errors in the PGA and PGV predictions are -0.2 and -0.3, respectively, at this time. PGA and PGV errors are the difference of the logarithm of the observed minus the predicted ground motions; a factor of two difference between the predicted and observed PGA corresponds to an error of 0.7, and a factor of 10 to an error of 2.3. One second later (Figure 3.5c), the τ_p^{max} and $P_{d/v}$ values from the third station are incorporated, and the magnitude estimate rises to M5.8. The errors in PGA and PGV change to 0.0 and -0.4. One second later (Figure 3.5d), a fourth station triggers, the location is adjusted, and the magnitude estimate rises to M5.9. The predictions of peak ground shaking are adjusted to account for the new location and magnitude estimates and a second peak ground motion observation. The mean PGA and PGV errors change to 0.1 and -0.2. As additional seconds pass, more stations trigger and their P-wave parameters are incorporated into the evolving estimates of location and magnitude, and the predictions of ground shaking. Figures 3.8a and 3.8b show the errors in the magnitude and location estimates as time progresses.

ElarmS uses a bias correction to shift the GMPEs up or down to match available ground motion observations. In Figure 3.5c the AlertMap shows a decrease in expected ground motions, despite the increase in magnitude. In this case there is only one observation available (represented by the light blue diamond just southeast of the epicenter), and it lowers the predictions for the whole region until more observations are available in the next second. Iervolino, *et al.*, found that GMPEs contribute significantly more error to EEW predictions of ground shaking than do magnitude or location estimates. This is due to the inherent variability in peak ground motion at a given location with respect to even the best fitting attenuation relations. However, Iervolino, *et al*., also found that ground motion predictions stabilize as more information is incorporated. While the inclusion of a single ground motion observation may increase the error in the AlertMap (Brown, *et al.*, 2009), as more observations are included their individual errors cancel each other out. Future versions of ElarmS will wait until three or more ground motion observations are available before including them in the prediction, to avoid the increased uncertainty associated with using just one or two observations.

Figure 3.5e shows the CISN ShakeMap published after the Alum Rock event. From the time of the first magnitude estimate, one second after the first P-wave detection, the predictive AlertMap (Figure 3.5b) is a close match to the ShakeMap. Figure 3.5f shows a seismogram recorded in San Francisco during the Alum Rock earthquake. The timeline denotes the times at which the data used in (a),(b),(c), and (d) was available. At the time ElarmS applied a 15 second buffer to the incoming waveforms, to reduce latency differences between stations. Despite the 15 second buffer, the data used to create (b-d) was available four to two seconds before the S-waves reached San Francisco and peak ground shaking began. This event represented the first "proof-of-concept" event for the

realtime ElarmS system as it illustrates that hazard information is available before shaking is felt.

3.4.2 M_w5.4 Chino Hills, LA Region

Figure 3.6 shows the M_w5.4 Chino Hills event, which occurred on 29 July, 2008. At the time ElarmS was midway through the conversion to statewide coverage, and was receiving data from only 15 southern California stations. ElarmS was still able to estimate magnitude, location and ground shaking using only the three stations within 100km of the epicenter. When the first station triggered, the event was located directly beneath the station at a depth of 8km. The observed τ_p^{max} and $P_{d/v}$ values were used to estimate a magnitude of 5.4. From that location and magnitude, local GMPEs were used to predict peak ground shaking in the region (Figure 3.6a). After a second station triggered, the location was adjusted between the stations based on arrival times, at a depth of 8km. The τ_p^{max} and $P_{d/v}$ magnitudes for the second station were averaged together with those from the first station, producing a new event magnitude of M5.8. The new location and magnitude were used to update the predictions of ground shaking (Figure 3.6b). Figure 3.6c shows the CISN ShakeMap for comparison. The ShakeMap is published after the event, using observations from all available stations. The ElarmS predictive AlertMap is reasonably similar to the ShakeMap, considering ElarmS used data from only two stations (the third and final available station triggered six seconds later and did not significantly change the AlertMap). Figures 3.8c and 3.8d show the progression of magnitude and location errors with time.

3.4.3 M_w4.4 Lone Pine, eCAs Region

The Lone Pine M_w4.4 occurred on October 3, 2009, in the eCAs region. In this region the stations are separated by 20-100km, so ElarmS requires at least 5 stations to trigger before issuing an alert. In Figure 3.7a the event is detected when two stations trigger simultaneously, four seconds after the event origin time. One second later (3.7b) the event magnitude is estimated at 4.0. Four more seconds pass before a third station triggers, at which point the location is adjusted and the magnitude estimate is raised to 4.1 (3.7c). The thin station coverage necessitates waiting longer in this region than in the previous examples. The five station requirement for alert issuance is not met until two seconds later (3.7d), eleven seconds after the event begins. The fourth and fifth stations did not appreciably change the magnitude, location, or ground motion predictions in this case, but they ensured that the event was real (Figure 3.8ef).

3.5 Application of ElarmS to Japan

3.5.1 Scaling and GMPEs

While ElarmS has been tested with many datasets in California, there are few recent, well-recorded, large earthquakes in California. Since an early warning system is designed specifically to warn people of large events, we are especially interested in its performance for these events. Thus we tested the system with a dataset of large events from Japan (Brown, *et al.*, 2009). The Japanese events also provided insight into ElarmS' performance in a subduction zone environment.

The dataset included 84 Japanese events that occurred between September 1996 and June 2008 (Figure 3.9). The magnitudes ranged from 4.0 to 8.0, with 43 events of magnitude 6.0 or greater. The largest event was the M8.0 Tokachi-Oki earthquake of 26 September, 2003. The events were recorded by Japan's Kyoshin Net (K-NET) strong-motion seismic network. K-NET consists of 1,000 digital strong motion seismometers, distributed across Japan with approximately 25km spacing. Each station is capable of recording accelerations up to 2,000 $^{cm}/_{s}^{2}$, with a sampling frequency of 100Hz and a dynamic range of 108dB.

The events were processed offline, using all available data, using the same methodology as described above. The first step is to determine scaling relations between the predominant period and peak amplitudes of the P-waves and the magnitude for the event dataset. The observed scaling relations for Japan are shown in Figure 3.2ef and are:

$$M_{JMA} = 4.76*\log_{10}(\tau_p^{max}) + 5.81$$

$$M_{JMA} = 5.82 + 1.52*\log_{10}(P_d) + 1.39*\log_{10}(R)$$

where M_{JMA} is the JMA catalog magnitude and R is the epicentral distance. The predominant periods observed in Japan are of similar values to those of Northern and Southern California, but the best-fit slope is steeper in Japan. The peak amplitude values are higher than those in Northern California and lower than those in Southern California, with a slightly shallower slope in Japan.

For the prediction of peak ground shaking, we used the GMPEs that the global ShakeMap system uses for Japanese events. The global ShakeMap GMPEs use either the Boore, *et al.*, or the Youngs, *et al.*, model, depending on depth and magnitude of the event. For events shallower than 20km or smaller than magnitude 7.7, the relations are defined by Boore, *et al.*, with numerical coefficients specified for reverse faulting:

$$\ln(PGA) = -0.117 + 0.527*(M-6) + 0.778*\ln(R)$$

where R is defined by

$$R = (R_{jb}^2 + h^2)^{½}$$

R_{jb} is the closest distance in km to the surface projection of the fault and h is a model coefficient representing depth. We substitute the epicentral distance for R_{jb}.

For events deeper than 20km or greater than magnitude 7.7, global ShakeMap and ElarmS use the GMPEs defined by Youngs, *et al.*:

$$\ln(PGA) = 0.2418 + 1.414*M - 2.552*\ln(R_{jb} + 1.7818 \exp(0.554*M)) + 0.00607*h$$

where again we substitute the epicentral distance for $R_{jb.}$

3.5.2 Performance for Large Magnitudes

Once the necessary scaling relations had been developed all 84 events were processed in a simulated realtime environment to provide ElarmS predictions of ground shaking. We assumed zero data latency and processed data sequentially according to the time-stamp on the waveform data. After the events were processed we analyzed ElarmS performance for different magnitude ranges. Figure 3.10 shows the resulting ElarmS magnitude error histograms. The blue histogram is the magnitude error for all events in the Japanese dataset, with magnitudes from 4.0 to 8.0. The mean error for all events was 0.0 magnitude units, with a standard deviation of 0.4. The green histogram is the magnitude

error for all events magnitude 6.0 or greater (of which there are 43). The mean error for this distribution is again 0.0, with a standard deviation of 0.5. This is a similar distribution statistically to that for all events. The red histogram is the magnitude error for events magnitude 7.0 or greater (of which there are seven in this dataset). Of the seven events M>7, four of the magnitudes are underestimated, two are overestimated, and one is accurately estimated. The mean error for this distribution is -0.2 magnitude units, with a standard deviation of 0.5. This lower mean error means that ElarmS underestimates the magnitude of the largest events by 0.2 magnitude units on average. An underestimation of 0.2 magnitude units is within our tolerance for ElarmS magnitude estimates, but we recognize that the magnitude algorithm may need to be adjusted to prevent underestimation in the future. A first step may be to weight the average of τ_p^{max} and $P_{d/v}$ in favor of τ_p^{max} for high magnitude events, since τ_p^{max} is less prone to saturation effects at the highest magnitudes (Brown, *et al.*, 2009).

3.5.3 Methodological Improvements

The Japanese dataset provided some methodological challenges. The majority of the events were offshore. The resulting limited azimuthal coverage (all stations are onshore) slowed down our location algorithm, requiring more station trigger times and therefore more seconds to produce a reasonable epicentral estimate. Many of the events were also deep. The original California location algorithm assumed a depth of 8km for all events, and found the hypocenter on a 2D grid at that depth. For the subduction zone events we expanded the algorithm into a 3D grid search, finding hypocenters at depths down to 80km, in 10km increments. Figure 3.11 shows a histogram of location estimate errors using the new 3D grid search. The histogram includes all hypocentral location estimates for each event, from the initial 1-trigger estimate to the final estimate using all available stations. The median location error, across all events and all number of triggers, is 11km.

3.5.4 Error Model

As part of the Japan dataset testing, we developed an error model similar to that of Iervolino, *et al.*, to analyze the errors in ElarmS' output (Brown, *et al.*, 2009). We separated the algorithm into its location, magnitude, and ground motion steps, and isolated the errors produced during each step. Errors were calculated by comparing the estimated location or magnitude to the catalog location or magnitude, and the predicted ground shaking at all stations and times prior to recording ground shaking to the eventual observation of peak ground shaking at that station. Predictions of peak ground shaking at stations after the peak shaking had occurred were not included in the error analysis. The errors of each component of the system are shown in Table 3.3.

The accuracy of any given step is dependent on the amount of data available. The error in the location estimate, for example, is dependent on the number of stations reporting P-wave arrivals. The error in the magnitude estimate is dependent on both the number of stations providing information and the number of seconds of P-wave that have arrived at each station. The error in the prediction of peak ground shaking is dependent on the number of stations whose observations of peak ground shaking have been used to adjust the prediction. The "0 stations" error is when no stations have yet recorded peak ground shaking, and the prediction of ground shaking is based on the GMPEs alone.

The errors calculated (Table 3.3) were then used to produce an error model for ElarmS' final prediction of ground shaking, given any combination of inputs. If there were no errors at all in the system, then the ElarmS prediction of ground shaking would be based on the same magnitude and location that the catalog uses. Since ElarmS uses the global ShakeMap GMPEs, an error-free ElarmS AlertMap should look much like the global ShakeMap. Therefore, the error contributed by ElarmS is the difference between the ShakeMap calculation of ground shaking and the AlertMap prediction of ground shaking. The ideal, error-free output is defined by the GMPEs for an event. For example, for an event shallower than 20km depth with a magnitude less than 7.7, the error-free output would simply be the Boore, *et al.*, GMPE. For peak ground acceleration (PGA):

$$\ln(PGA)_{ideal} = -0.117 + 0.527 * (M-6) + 0.778 * \ln(R) \qquad \text{Ideal, error-free output}$$

where M is magnitude and R is the distance from the event epicenter to the location where PGA is being predicted.

We then introduce errors into the calculation, using the error distributions we observed for our Japan dataset.

$$\ln(\hat{PGA}) = -0.117 + 0.527*(M+\varepsilon_M - 6) + 0.778*\ln(R \pm \varepsilon_R) + \varepsilon_{Att} \qquad \text{ElarmS output}$$

where M is the catalog magnitude, R is the epicentral distance, and ε_M, ε_R, and ε_{Att} are the errors in magnitude, location, and GMPEs, respectively.

The difference between PGA_{ideal} and $\hat{PGA}$ is the error in our final prediction of ground shaking.

$$\varepsilon_{PGA} = \ln(PGA)_{ideal} - \ln(\hat{PGA}) \qquad \text{Error}$$

This represents the total error in the entire algorithm. ε_{PGA} is a unitless value; a factor of two difference between the ideal and estimated PGA corresponds to an error of 0.7, and a factor of 10 to an error of 2.3.

The errors for each step (ε_M, ε_R, ε_{Att}) are dependent on the quantity of data included (the number of trigger times, the number of τ_p^{max} and $P_{d/v}$ values, etc.) and vary within the probability distributions defined in Table 3.3. Thus the error model is similarly dependent. We calculated ε_{PGA} 1000 times for every combination of data inputs, 1086 combinations, each time choosing the error values by a Monte Carlo simulation based on the mean and standard deviation of the error distributions (Table 3.3). The resulting 1000 values for ε_{PGA} are used to create a probability distribution for ε_{PGA} given that specific combination of data inputs. Figure 3.12a shows three sample ε_{PGA} distributions, and Figure 3.12b shows all 1086 ε_{PGA} distributions, corresponding to 1086 unique combinations of data inputs (number of stations contributing to location estimate, number stations contributing to magnitude estimate, number of seconds of P-wave for each station, and number of observations of peak ground shaking). The mean errors for all ε_{PGA} distributions range from -0.2 to 0.2, with a median of 0.0. Standard deviations range from 0.3 to 0.6, with a median of 0.4. The standard deviations of all error distributions are less than 0.7 meaning less than a factor of 2 error in the PGA prediction. These calculated error distributions are stored in an internal library, accessible during processing for realtime estimates of uncertainty in PGA predictions.

Finally we analyzed the error contributions of each step of the algorithm separately. By assuming zero error in the magnitude estimate, for example, we could remove that error contribution from the system and observe how much the 1086 error distributions change. In all cases the error distributions remained centered about zero, but the median standard deviation decreased, indicating a decrease in the range of errors. Removing the GMPE error from the system (by setting $\varepsilon_{Att}=0$) decreased the median standard deviation by 49%, compared with a decrease of 13% when location error was removed ($\varepsilon_R=0$) and 5% when magnitude error was removed ($\varepsilon_M=0$) (Brown, *et al.*, 2009).This result confirmed that of Iervolino, *et al*., who demonstrated conclusively that GMPEs contribute the most uncertainty to an EEW prediction of ground shaking. Iervolino, *et al*., also showed that predictions of peak ground motions only stabilize when data is included from multiple stations. Since an EEW must use GMPEs, the safest recourse is to wait until multiple stations are providing data before issuing a prediction of ground motions. As always, there is a trade-off of speed versus accuracy in any EEW prediction (Iervolino, *et al.*, 2009).

3.6 Conclusions

The three-year CISN project gave us the opportunity to combine the offline development of ElarmS in California with the error analysis performed in Japan and produce a statewide realtime system. Already we have integrated data from five disparate networks, adapted our algorithms to run in realtime using data that is unevenly delayed by telemetry, and added the ability to send alert messages within seconds of event detection. While improvements to the seismic networks in California would improve ElarmS performance, ElarmS has successfully predicted ground shaking for many events even with the current network of stations.

There are opportunities for improvement in the next three-year phase of the project. Our algorithm continues to struggle with false alarms, especially in the regions with low station density. Honing our regional trigger requirements may be the primary step needed to reduce the false alarm rate. In addition, the event associator needs to be improved to better tolerate aftershock sequences, so that we don't risk missing a large mainshock due to its foreshock or other nearby events.

Data latencies are also a significant problem, claiming much of the potential warning time. Some latencies may be reduced by more efficient code design, such as updating the BK network software to accommodate the faster Q330 data loggers. Others require reformatting individual station data loggers, or upgrading data logger hardware. In the next two years ARRA stimulus funding will be used to upgrade many data loggers throughout the CISN, reducing latency by 3-5 seconds at these sites. The current statistical median latency is 5.2 seconds. With these upgrades we anticipate this will be reduced to 2-3 seconds.

We expect many improvements to the ElarmS code and the CISN networks during the coming three years. Learning from the realtime experience of the last three years, the ElarmS, VS, and Onsite methodologies will be integrated into a single prototype system.

New code is being written to reduce processing delays and hardware upgrades will reduce data transmission latencies. The CISN is currently identifying a small group of about 10 test users who will soon start to receive alerts from the new prototype system, called the CISN ShakeAlert System.

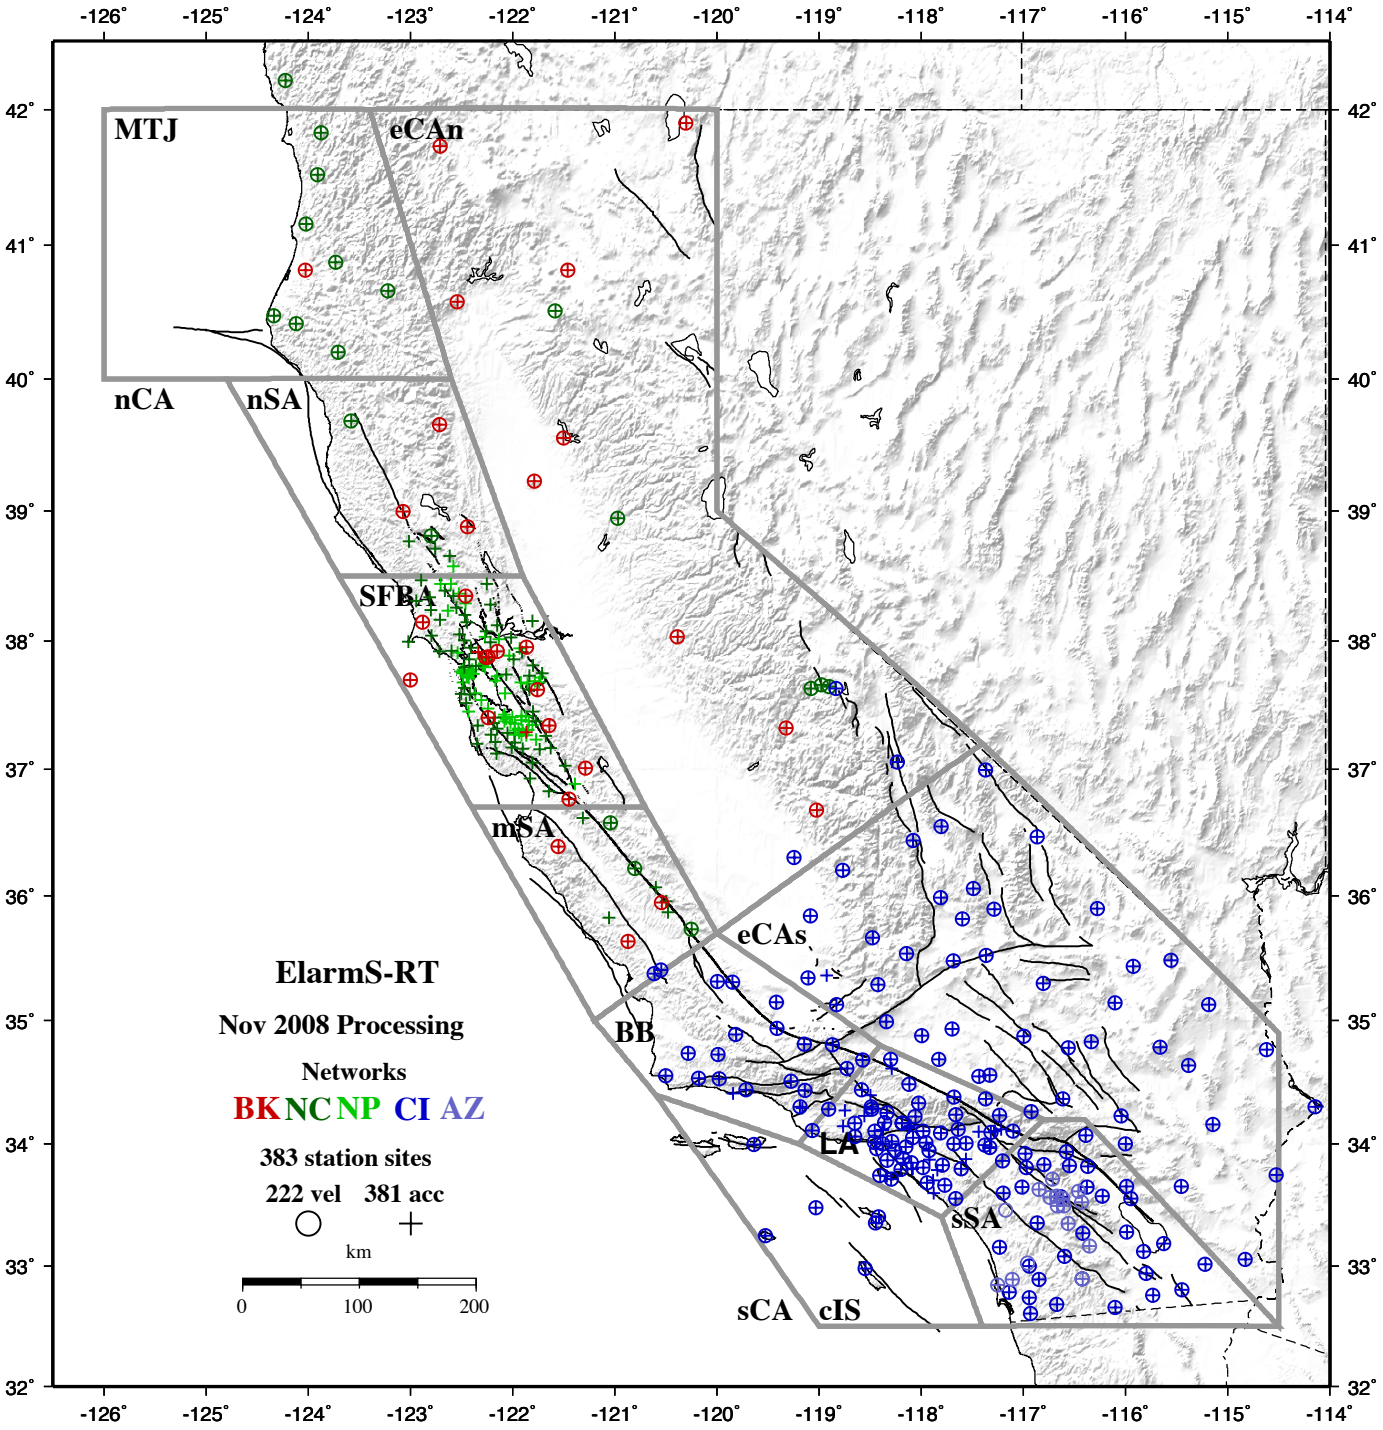

Figure 3.1. Realtime seismic stations used by ElarmS in California. Circles are velocity instruments, and crosses are accelerometers. Many stations have co-located velocity and acceleration sensors. The grey boxes indicate regions used for alert requirements: Mendocino Triple Junction (MTJ), north San Andreas (nSA), San Francisco Bay Area (SFBA), middle San Andreas (mSA), Big Bend (BB), Los Angeles (LA), south San Andreas (sSA), Channel Islands (cIS), east California south (eCAs), and east California north (eCAn). The straight line between regions mSA/eCAn and BB/eCAs is the Gutenberg-Byerly line dividing northern and southern California.

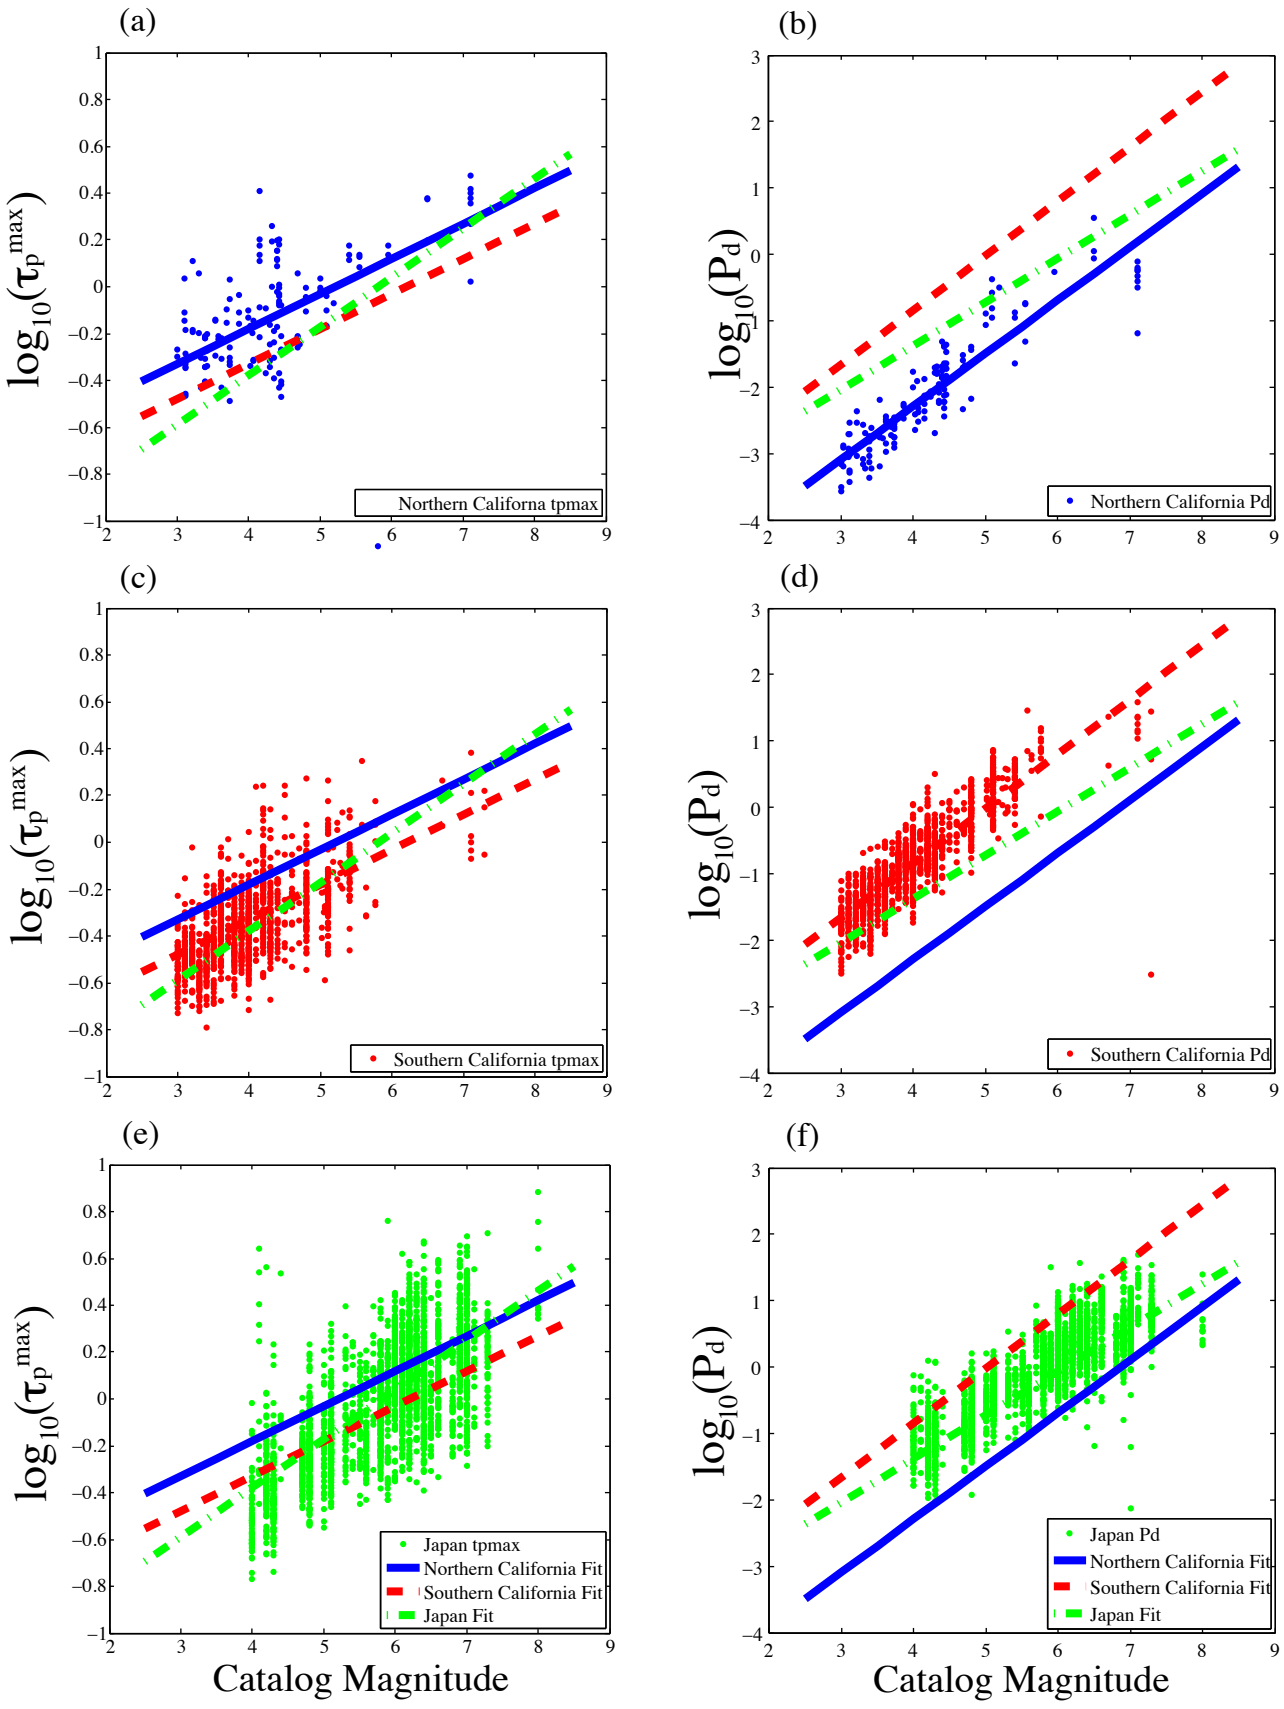

Figure 3.2. Scaling relations. (a) τpmax, northern California; (b) Pd, northern California; (c) τpmax, southern California; (d) Pd, southern California; (e) τpmax, Japan; (f) Pd, Japan. Circles are individual station observations of τpmax or Pd. Lines are regional scaling relations defined by the linear best fit to the data. The best-fit linear relations for all three regions are shown on all plots.

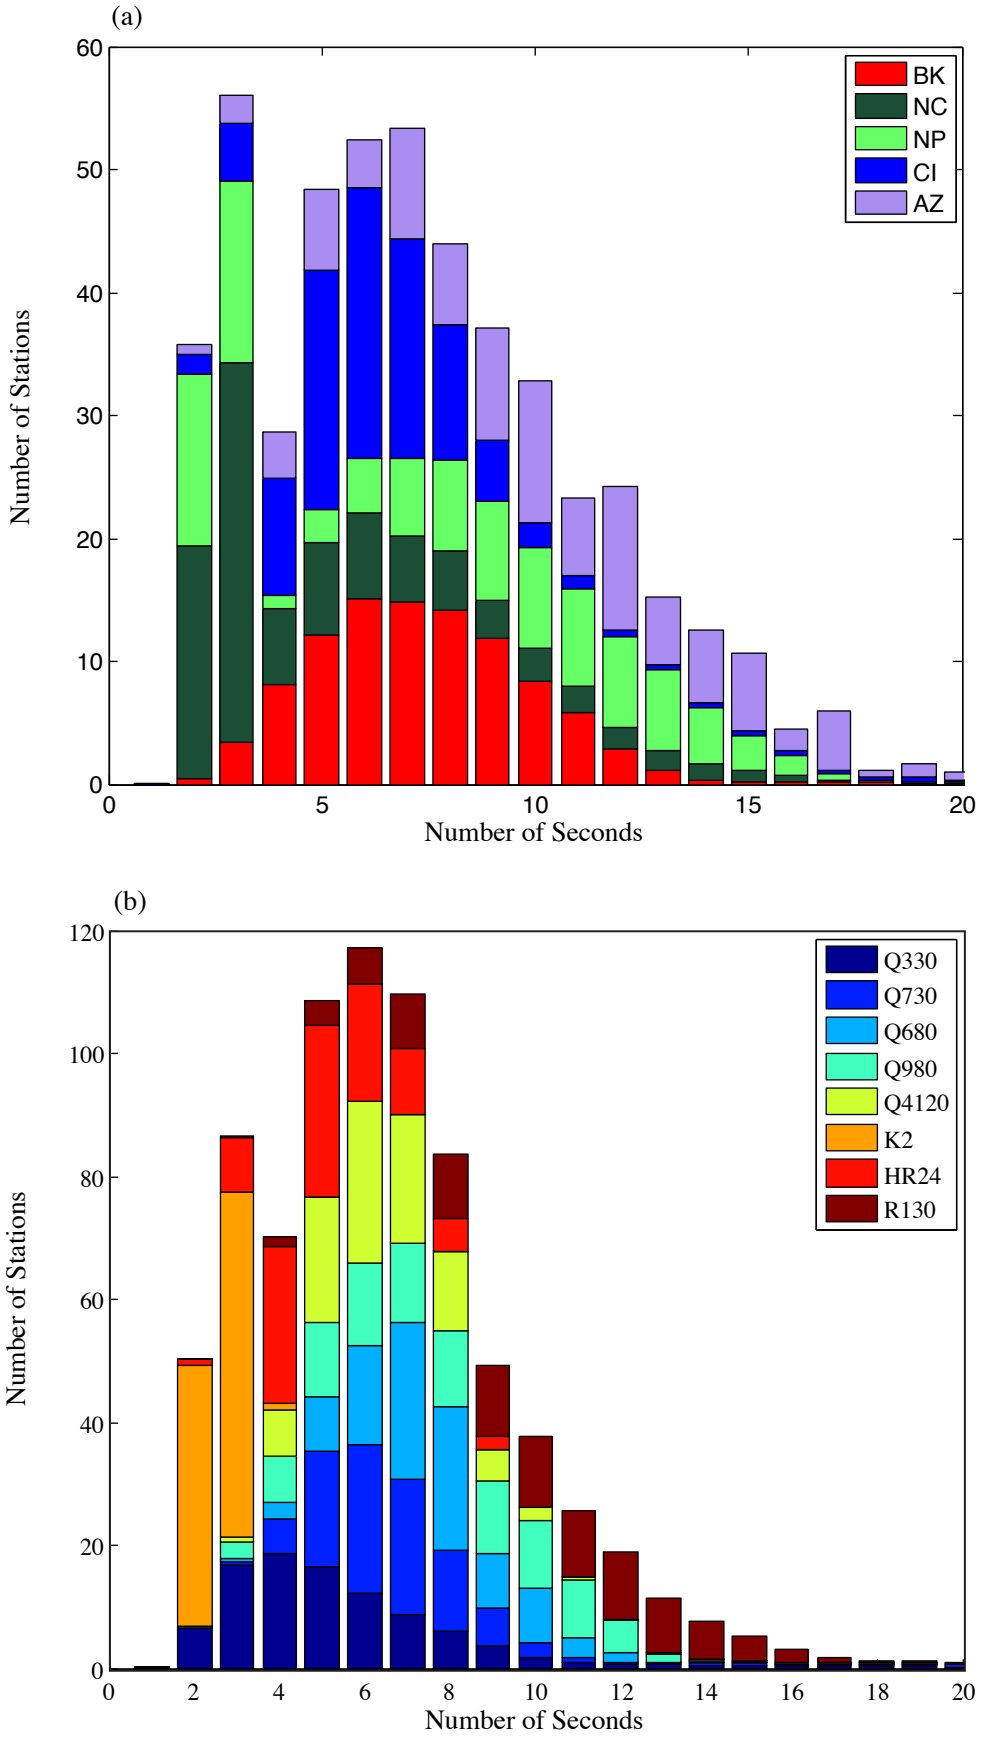

Figure 3.3. A stacked histogram of latencies by (a) network, and (b) data logger type. Both histograms are truncated at 20 seconds for clarity, but the long tail to the histogram continues, with columns of 0-2 data points, up to as much as 200 seconds.

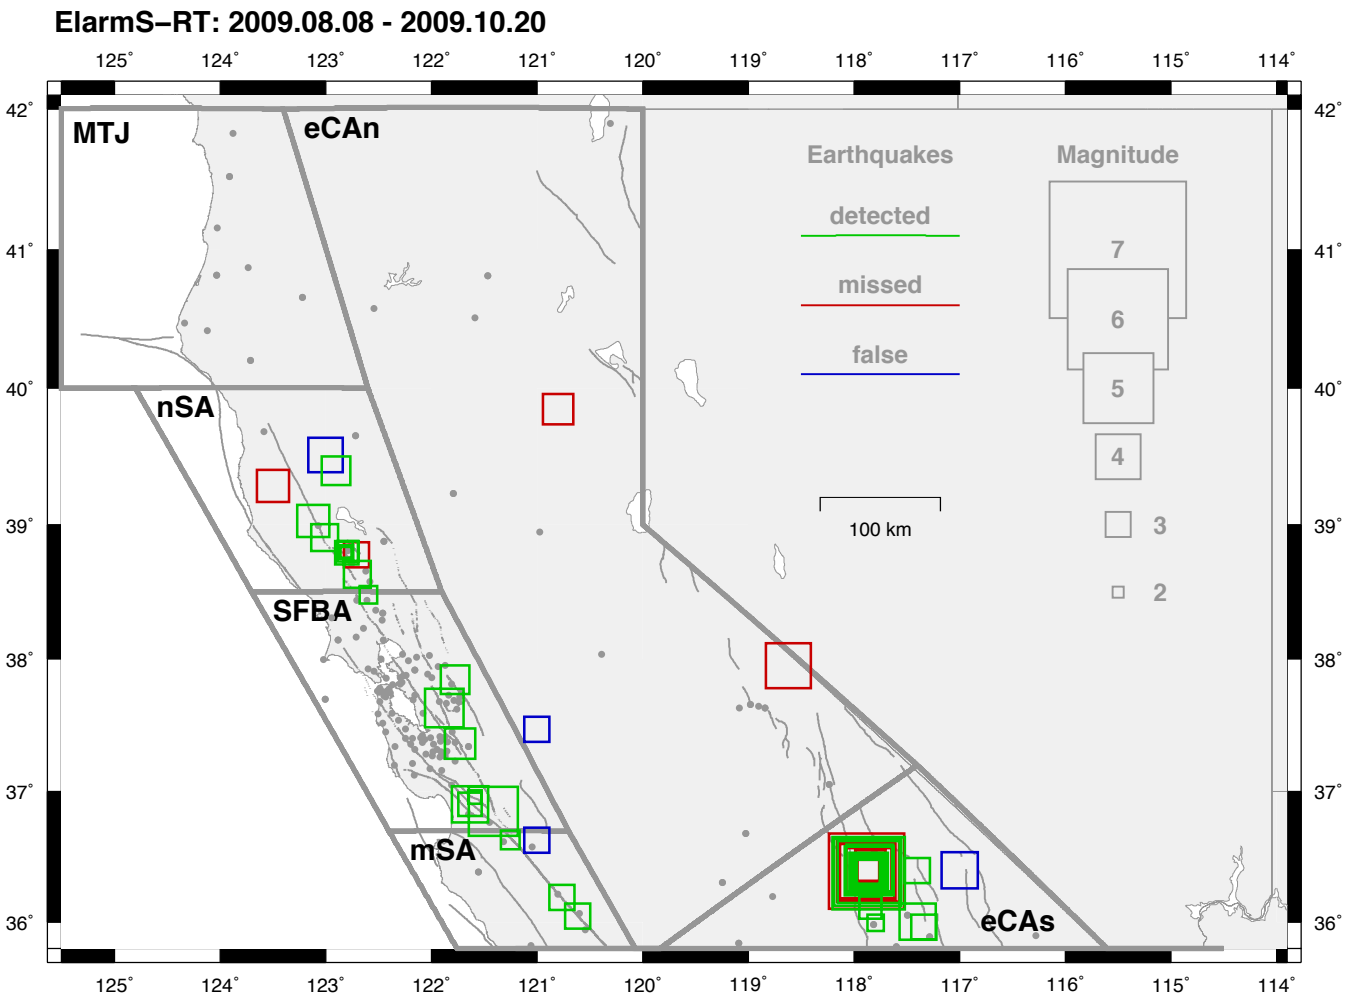

Figure 3.4. Map showing all ElarmS-detected earthquakes with M>3, and all false and missed alerts in Northern California, during a ten-week test period from 8 August 2009 until 20 October 2009. Green, red, and blue boxes are detected, missed, and false alerts, respectively. Grey dots are seismic stations.

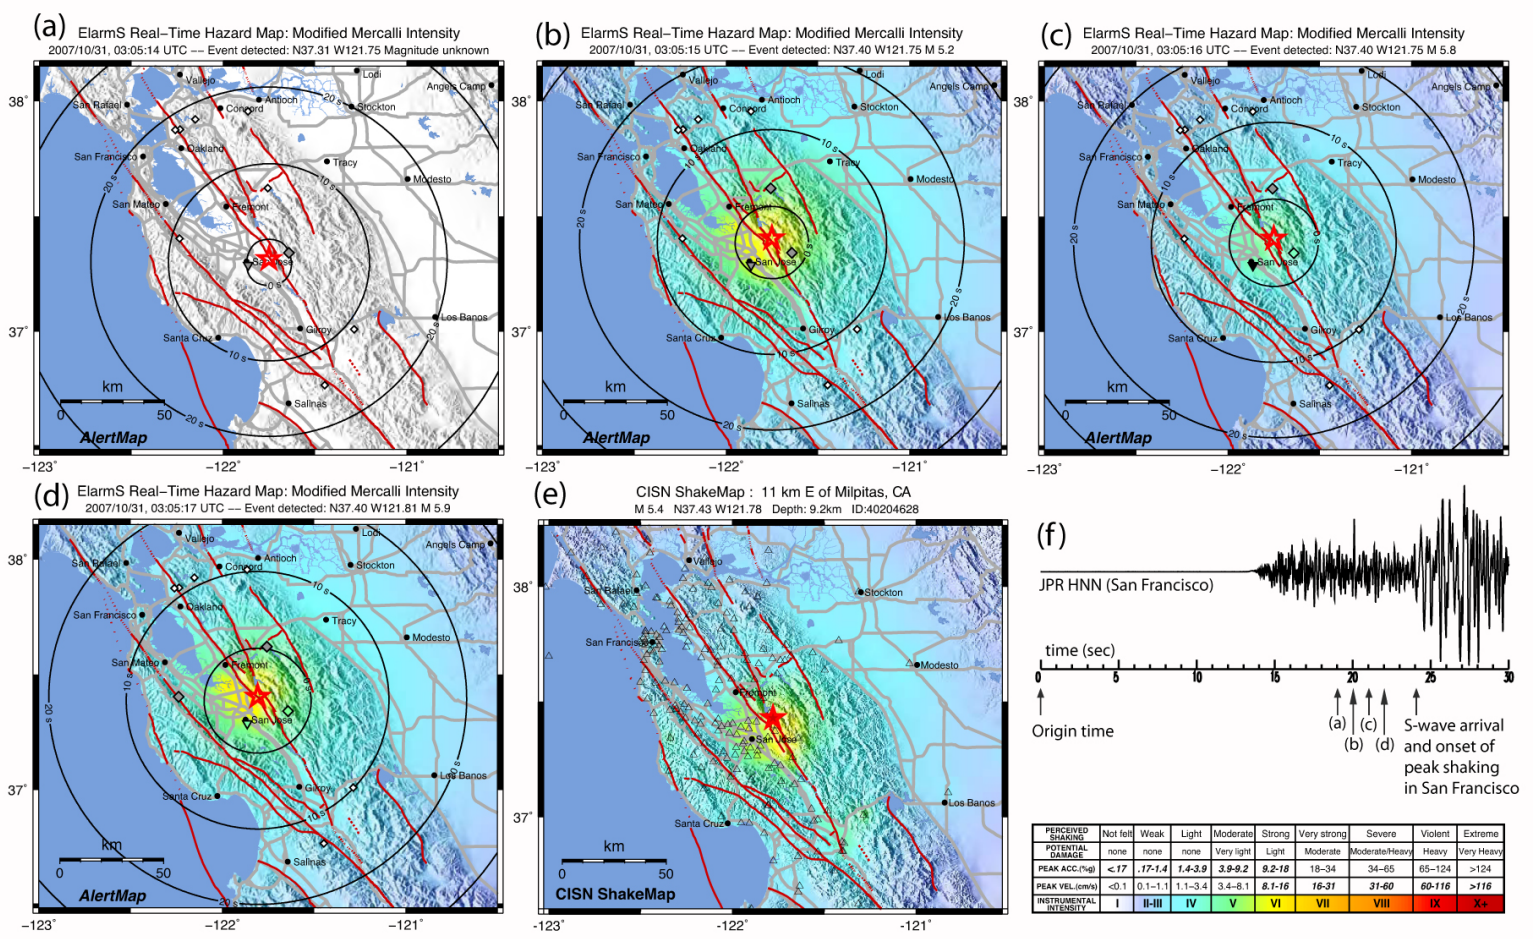

Figure 3.5. Example of ElarmS event processing for the 30 October 2007 Alum Rock Mw5.4 earthquake. (a-d) Progressive AlertMaps as stations trigger and the event is analyzed in realtime. The AlertMaps themselves were produced after the event, but the data used to create them was available at the time indicated on the map. (e) CISN ShakeMap published after the event. (f) Timeline comparing when the data used to create the AlertMaps was available with respect to the arrival of peak ground shaking in San Francisco.

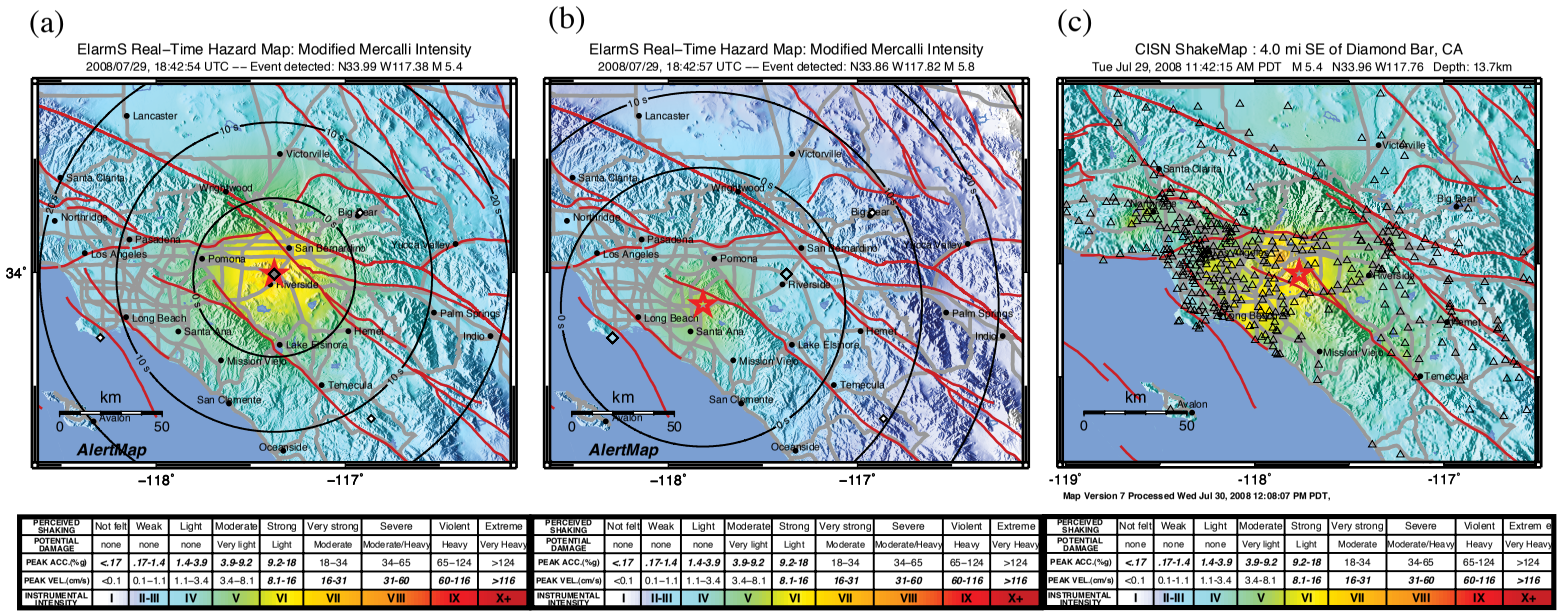

PERCEIVED SHAKING	Not felt	Weak	Light	Moderate	Strong	Very strong	Severe	Violent	Extreme
POTENTIAL DAMAGE	none	none	none	Very light	Light	Moderate	Moderate/Heavy	Heavy	Very Heavy
PEAK ACC.(%g)	*<.17*	*.17-1.4*	*1.4-3.9*	*3.9-9.2*	*9.2-18*	18–34	34–65	65–124	>124
PEAK VEL.(cm/s)	<0.1	0.1–1.1	1.1–3.4	3.4–8.1	*8.1-16*	*16-31*	*31-60*	*60-116*	*>116*
INSTRUMENTAL INTENSITY	I	II-III	IV	V	VI	VII	VIII	IX	X+

PERCEIVED SHAKING	Not felt	Weak	Light	Moderate	Strong	Very strong	Severe	Violent	Extreme
POTENTIAL DAMAGE	none	none	none	Very light	Light	Moderate	Moderate/Heavy	Heavy	Very Heavy
PEAK ACC.(%g)	*<.17*	*.17-1.4*	*1.4-3.9*	*3.9-9.2*	*9.2-18*	18–34	34–65	65–124	>124
PEAK VEL.(cm/s)	<0.1	0.1–1.1	1.1–3.4	3.4–8.1	*8.1-16*	*16-31*	*31-60*	*60-116*	*>116*
INSTRUMENTAL INTENSITY	I	II-III	IV	V	VI	VII	VIII	IX	X+

PERCEIVED SHAKING	Not felt	Weak	Light	Moderate	Strong	Very strong	Severe	Violent	Extrem e
POTENTIAL DAMAGE	none	none	none	Very light	Light	Moderate	Moderate/Heavy	Heavy	Very Heavy
PEAK ACC.(%g)	*<.17*	*.17-1.4*	*1.4-3.9*	*3.9-9.2*	*9.2-18*	18-34	34-65	65-124	>124
PEAK VEL.(cm/s)	<0.1	0.1-1.1	1.1-3.4	3.4-8.1	*8.1-16*	*16-31*	*31-60*	*60-116*	*>116*
INSTRUMENTAL INTENSITY	I	II-III	IV	V	VI	VII	VIII	IX	X+

Figure 3.6. Example of ElarmS processing for the 29 July 2008 Chino Hills Mw5.4 earthquake. (a) AlertMap showing predictions of ground shaking after one station had triggered. (b) AlertMap showing adjusted predictions after two stations had triggered. (c) CISN ShakeMap published after the event.

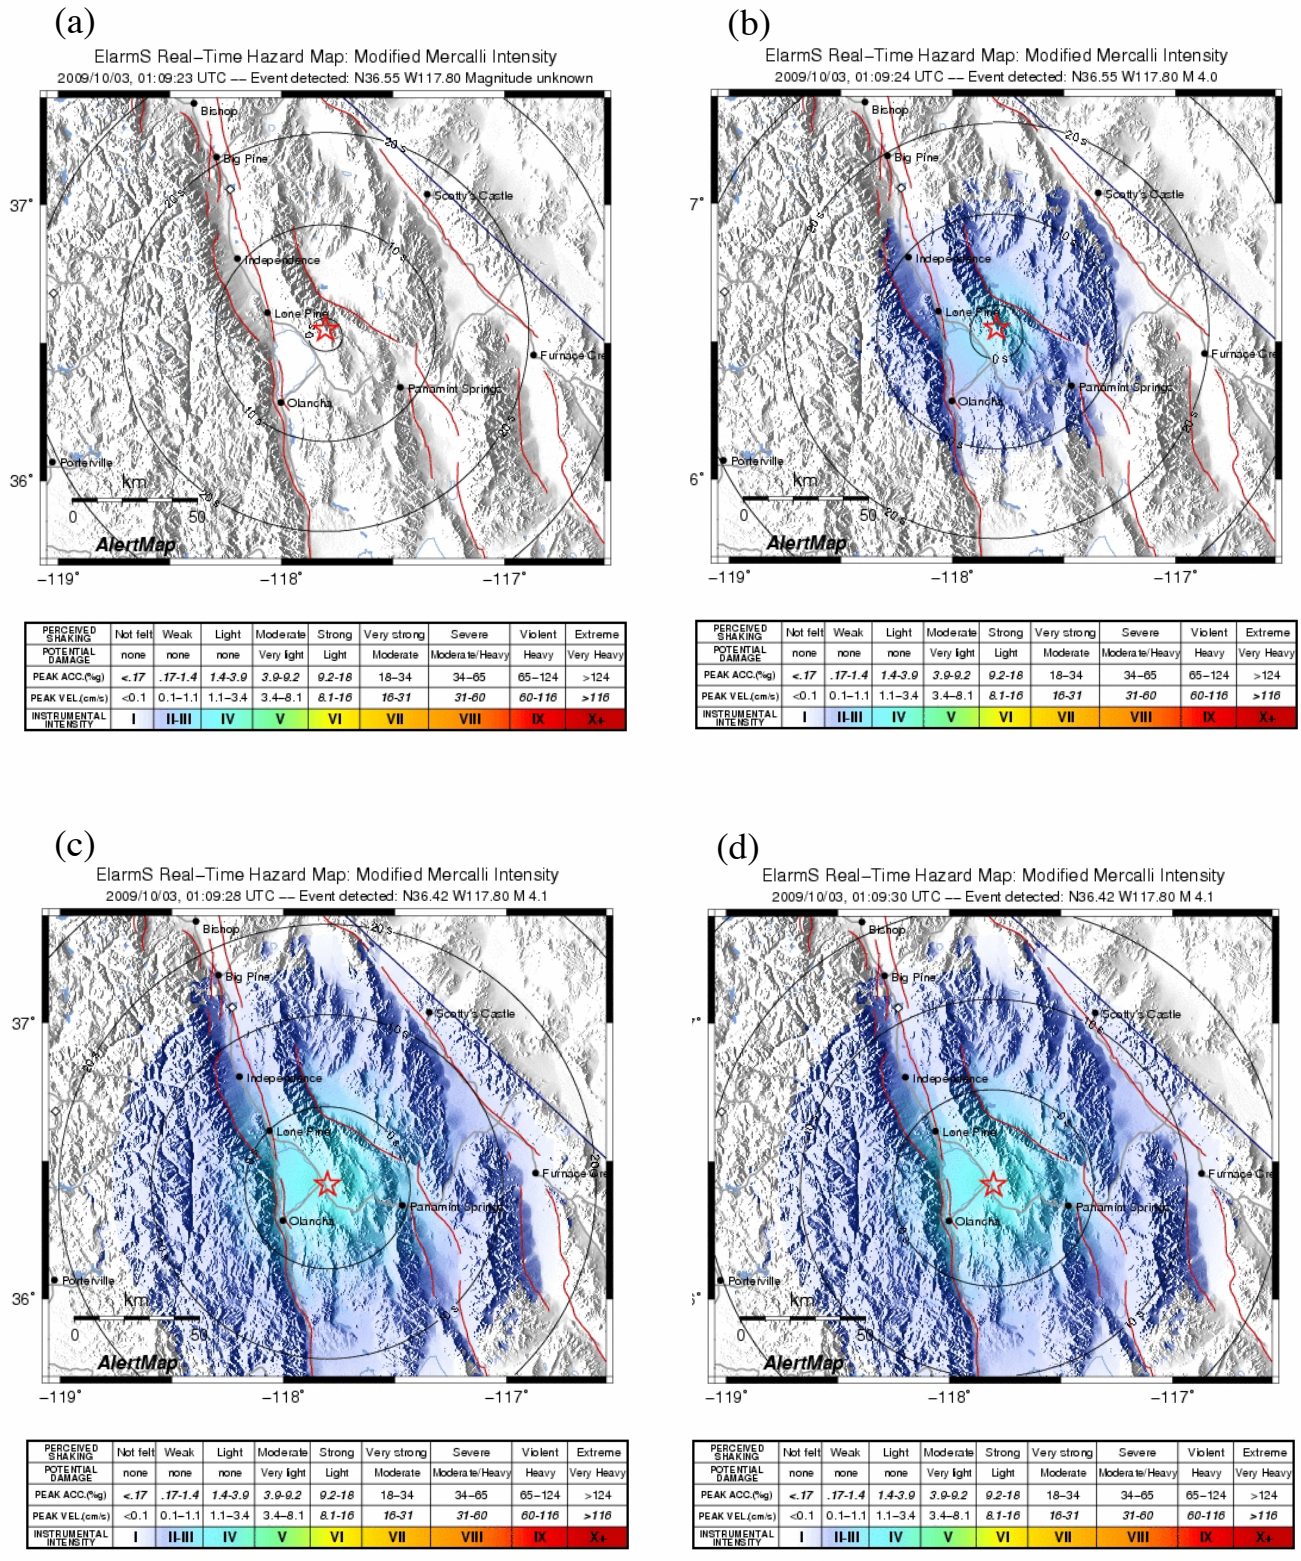

Figure 3.7. Example of ElarmS processing for the 3 October 2009 Lone Pine Mw4.4 earthquake. (a) Hypocenter was estimated when two stations triggered, 4 seconds after the event began. (b) One second later (OT + 5 seconds) magnitude was estimated at 4.0, using P-wave parameters from the two triggering stations. (c) Four seconds later (OT + 9) a third station triggered. Location, magnitude, and ground shaking predictions were adjusted. (d) One second later (OT + 11), the five station requirement was met and an alert was issued (to the authors) for this event.

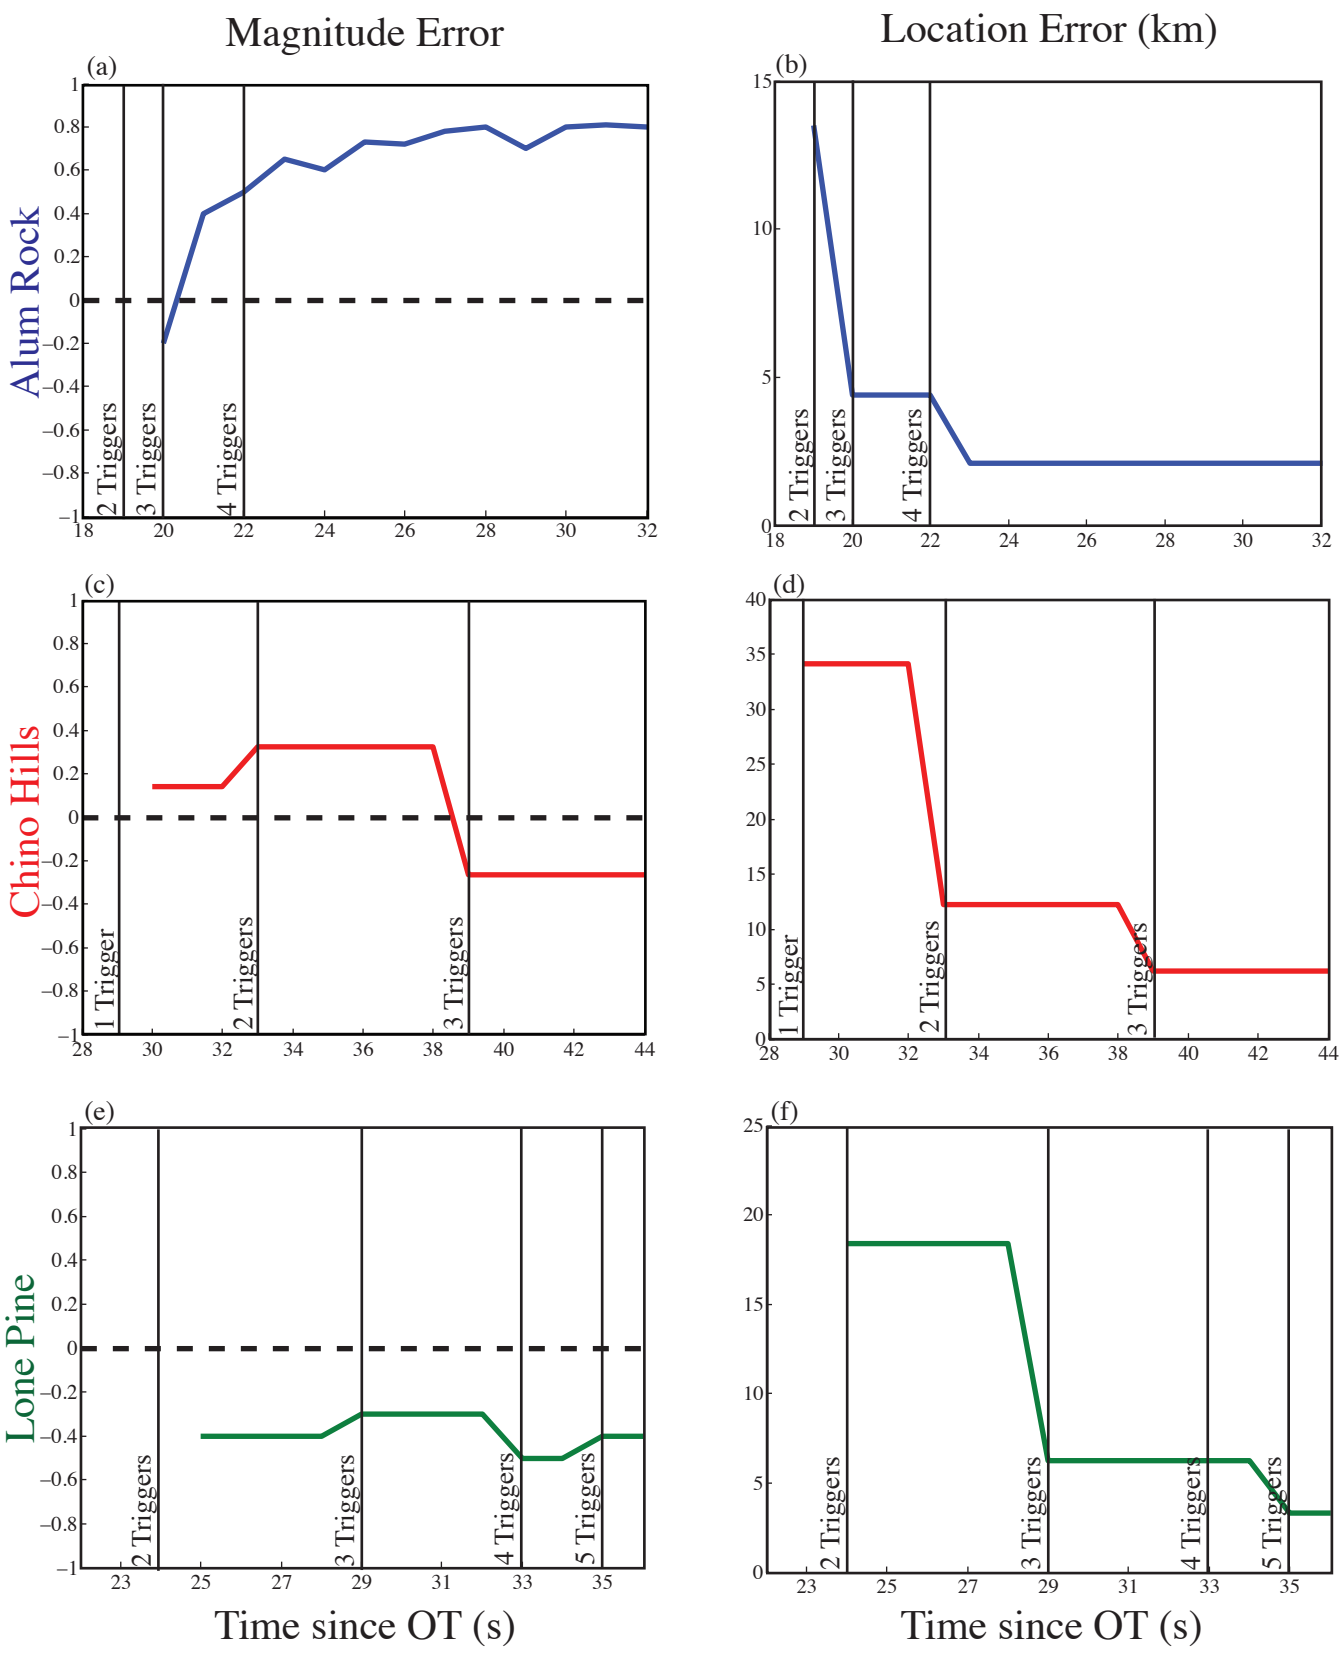

Figure 3.8. Magnitude and location error with time for the three California sample events: Alum Rock (a, b), Chino Hills (c, d), and Lone Pine (e, f). Horizontal axis is time in seconds since origin time of the earthquake. These times include a 15-second buffer for Alum Rock and 20-second buffers for Chino Hills and Lone Pine. Vertical axis is error in magnitude estimate (magnitude units) or epicentral location estimate (km).

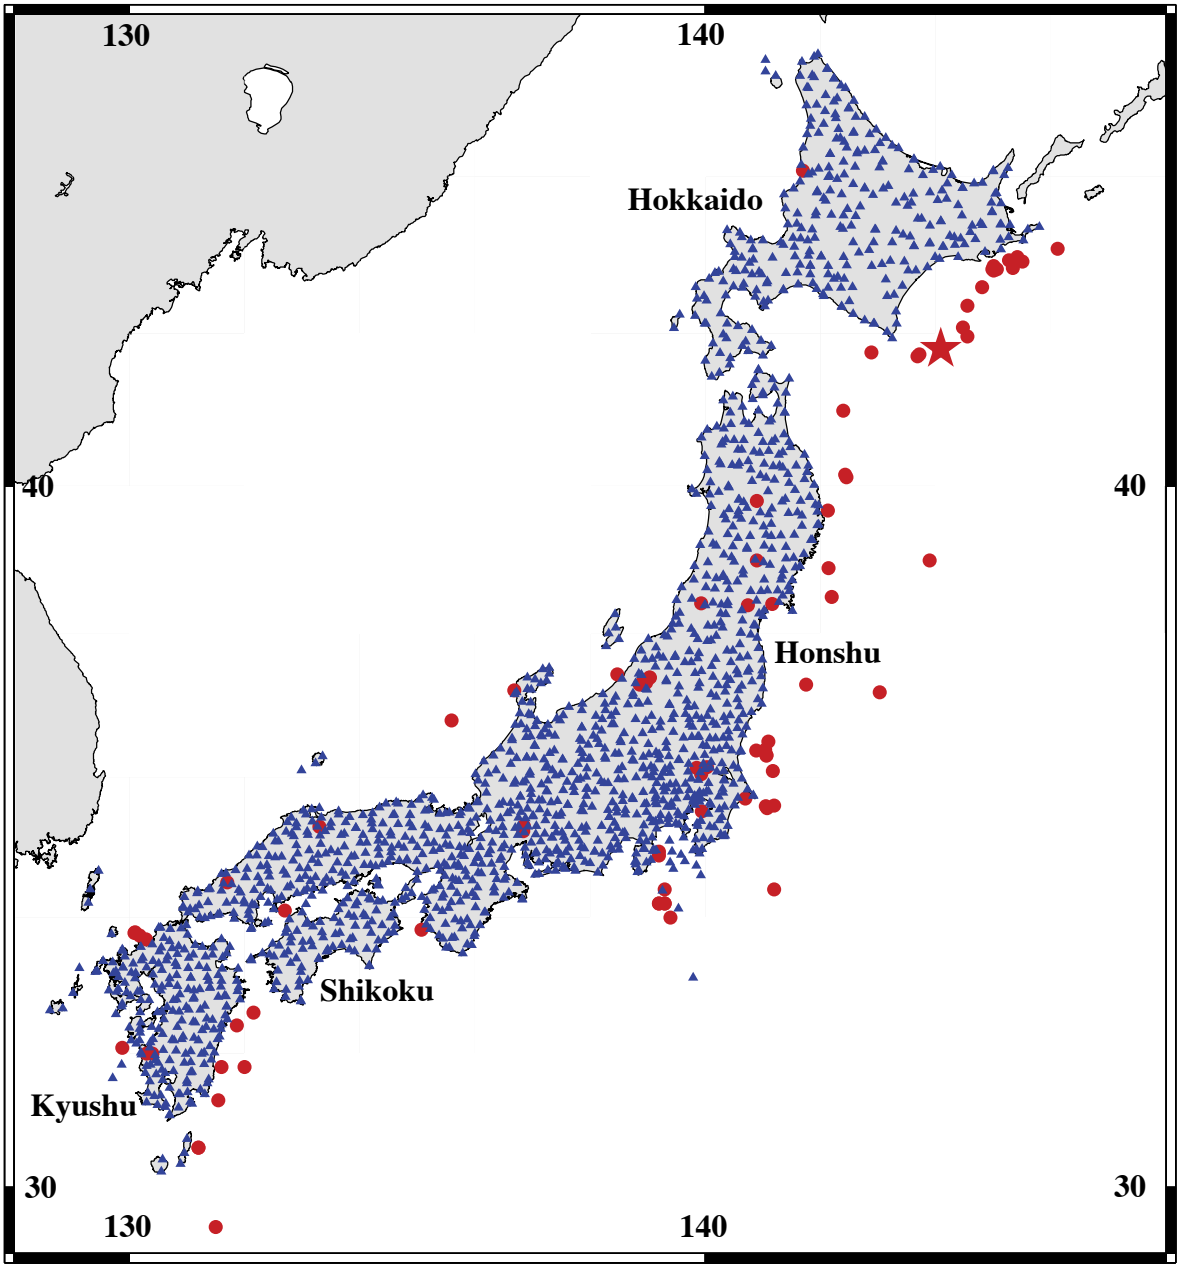

Figure 3.9. Events and stations used in the Japan test dataset. Red circles are events, blue triangles are K-NET stations. The red star is the largest event in the study, the M8.0 Tokachi-Oki earthquake of Sept 26, 2003.

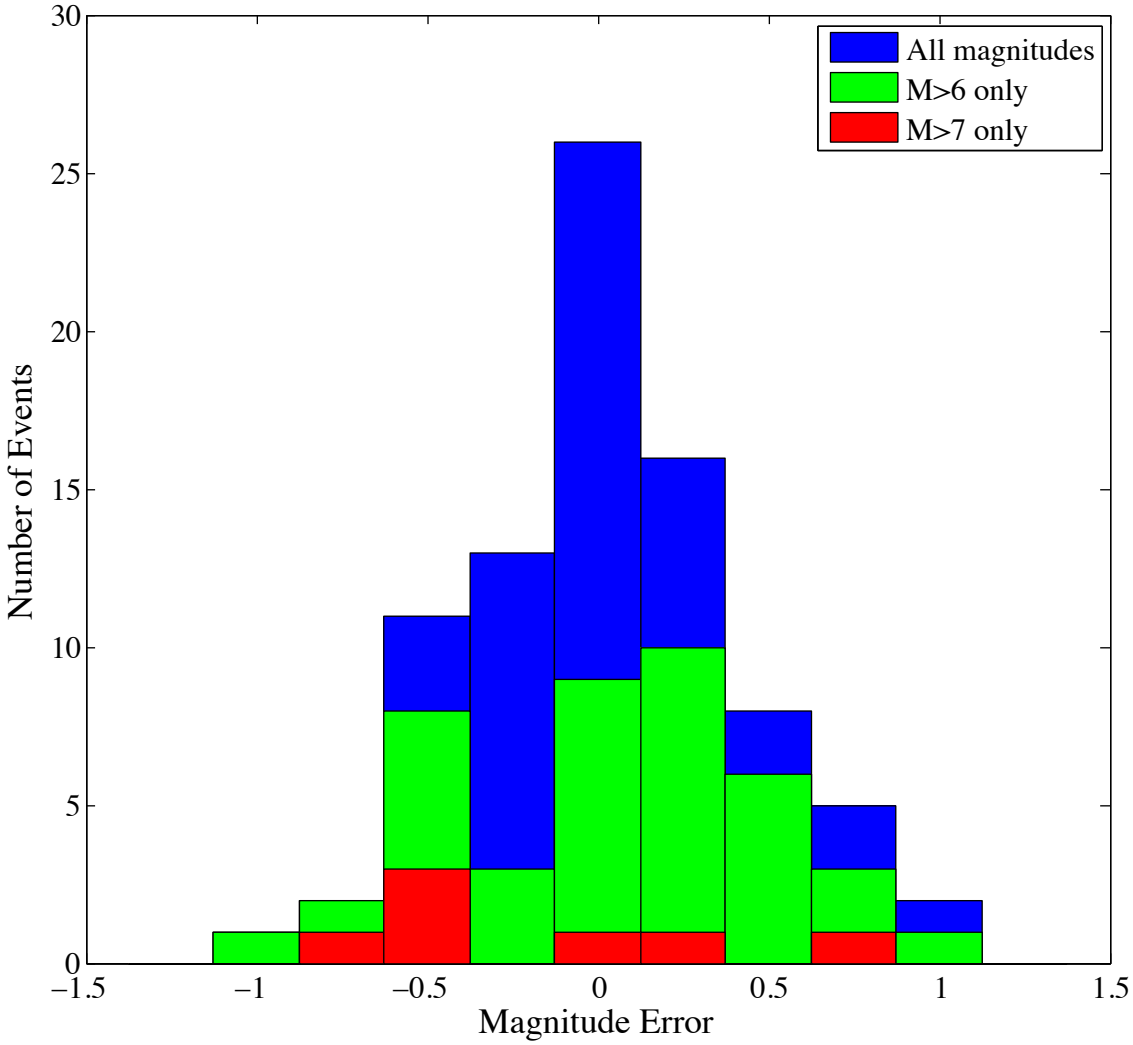

Figure 3.10. Histogram of magnitude errors for Japan dataset. The blue histogram is the distribution of magnitude error for all 84 events in the Japan dataset, M4.0 to M8.0. The green histogram is the distribution for the subset of 43 events with magnitude 6.0 or greater (up to and including magnitude 8.0) and overlays the blue histogram. The red histogram, again overlaying the green histogram, is for the subset of 7 events with magnitude 7.0 or greater (up to and including magnitude 8.0).

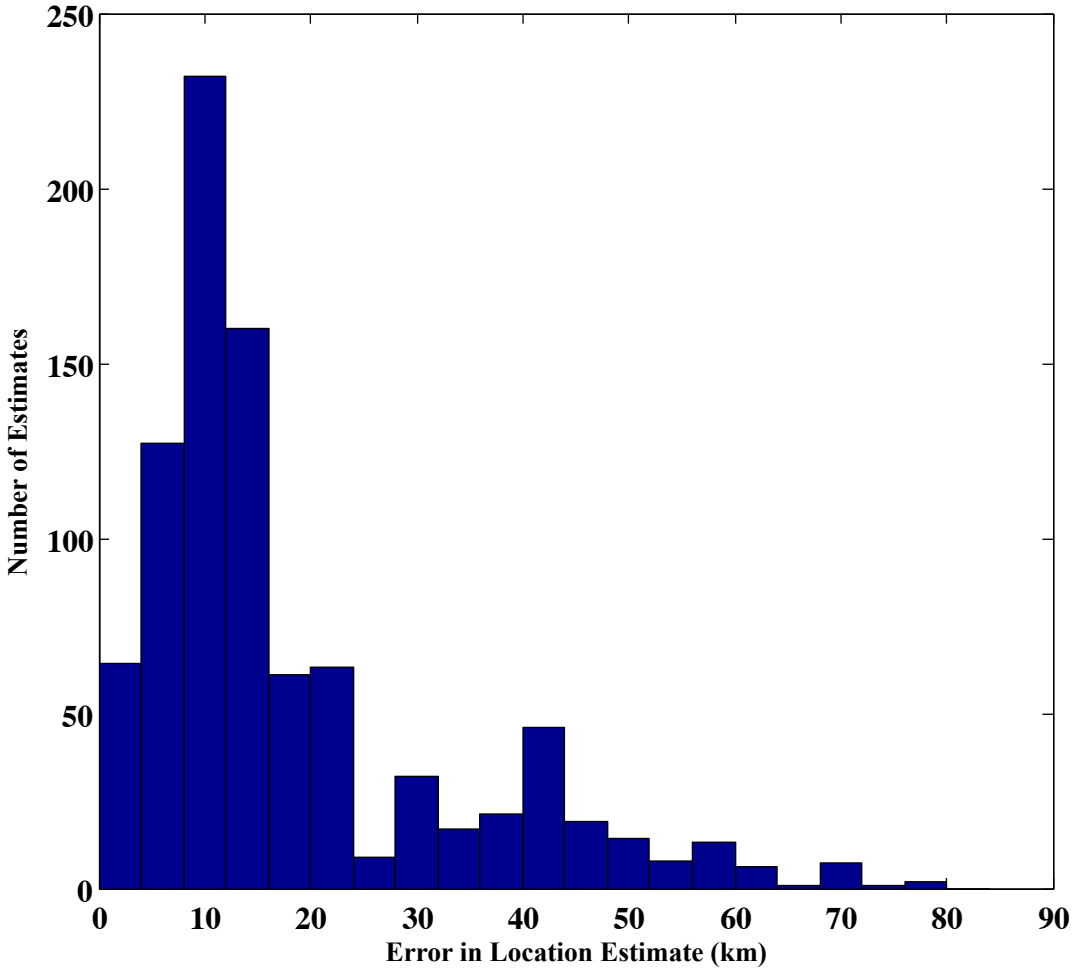

Figure 3.11. Histogram of errors in location estimate. Each event contributes an initial 1-trigger estimate, a 2-trigger estimate, and so on until all available stations are included. The median error across all estimates, with any number of triggers, is 11km.

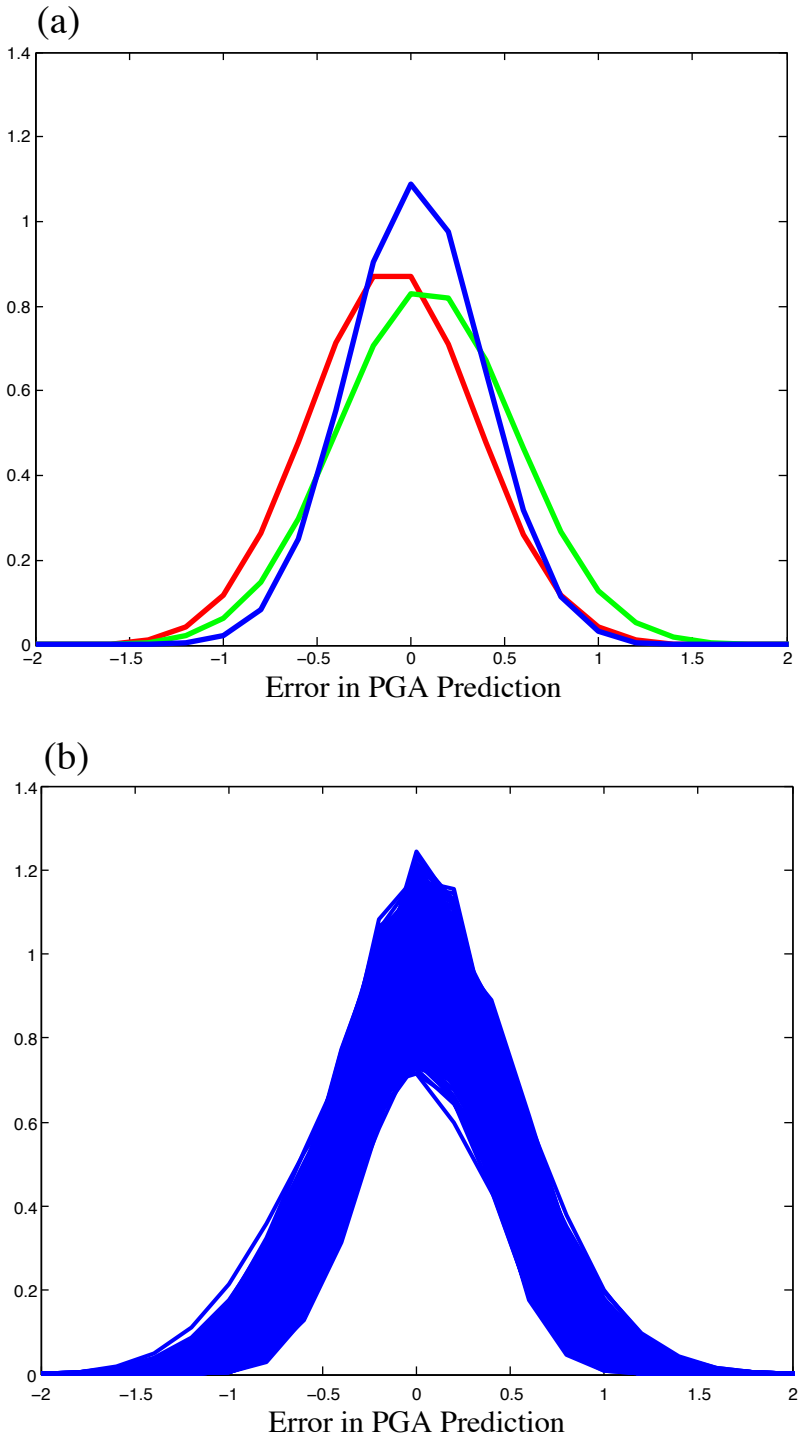

Figure 3.12. Error model distributions. (a) Three examples, showing best-fit Gaussian distributions for errors in ground motion estimation, given various quantities of data input. The red line is the error if two stations contribute to a location estimate, two stations contribute to the magnitude estimate (one using 1 second of P-wave data, one using 2 seconds), and zero stations report PGA observations. The green line is error if three stations contribute to the location estimate, two stations contribute to the magnitude estimate (one with 2 seconds of P-wave data, one with 3 seconds), and one station reports a PGA observation. The blue line is error if five stations contribute to the location estimate, five stations contribute to the magnitude estimate (4 with four seconds of P-wave data, one with 3 seconds), and three stations report PGA observations. (b)All 1086 error distributions resulting from the error model. Each line represents a unique combination of data inputs.

Table 3.1. Coefficients for the Boatwright, et al., ground motion prediction equation used in California.

	A	B	k_0	R_0	g	M_s	p	B_v	V_a	k
PGA, M>5.5	2.52	0.31	-0.0073	27.5	0.7	5.5	0.3	-0.371	560	$k_0 * 10^{(p*(Ms-M))}$
PGV, M>5.5	2.243	0.58	-0.0063	27.5	0.7	5.5	0.3	-0.371	560	$k_0 * 10^{(p*(Ms-M))}$
PGA, M≤5.4	2.52	1	-0.0073	27.5	0.7	5.5	0	-0.371	560	k_0
PGV, M≤5.4	2.243	1.06	-0.0063	27.5	0.7	5.5	0	-0.371	560	k_0

Table 3.2. Median values for the telemetry latencies shown in Figure 3.3.

		Median Delay (s)
Network	BK	6.2
	NC	2.5
	NP	7.4
	CI	5.2
	AZ	9.3
Data Logger	Q330	4.0
	Q730	5.5
	Q680	6.3
	Q980	6.6
	Q4120	5.3
	K2	1.6
	HR24	4.0
	R130	9.1

Table 3.3. Parameters of error distributions for magnitude, location, and ground motion.

	0 stations	1 station	2 stations	3 stations	4 stations	5 stations
Mag, 1 sec	-	-0.4 ± 0.6	-0.3 ± 0.6	-0.4 ± 0.6	-0.4 ± 0.6	-0.4 ± 0.6
Mag, 2 sec	-	-0.2 ± 0.6	-0.2 ± 0.5	-0.2 ± 0.5	-0.2 ± 0.5	-0.2 ± 0.5
Mag, 3 sec	-	-0.1 ± 0.5	-0.1 ± 0.5	-0.1 ± 0.5	-0.1 ± 0.5	-0.1 ± 0.5
Mag, 4 sec	-	0.0 ± 0.5	0.0 ± 0.5	0.0 ± 0.5	0.0 ± 0.4	0.0 ± 0.4
Mag, 5 sec	-	0.0 ± 0.5	0.1 ± 0.4	0.1 ± 0.5	0.1 ± 0.4	0.1 ± 0.4
Location	-	34 ± 18	32 ± 21	32 ± 19	19 ± 14	21 ± 17
PGA	0.1 ± 0.3	0.1 ± 0.4	0.1 ± 0.4	0.1 ± 0.3	0.1 ± 0.3	0.0 ± 0.3

Chapter 4

Developing the Second Generation ElarmS Code

4.1 Introduction

Earthquake Early Warning (EEW) is a method of recognizing earthquakes in progress and sending immediate alerts to surrounding population centers seconds before damaging ground shaking begins. Magnitude, location, and origin time are estimated from P-wave arrivals at nearby seismic stations. Earthquake Alarm Systems, or ElarmS, is an EEW methodology developed at the University of California, Berkeley. ElarmS currently operates as part of the larger ShakeAlert EEW project run by the California Integrated Seismic Network (CISN). Here we detail the development of the second generation ElarmS code, designed to improve speed and accuracy for successful integration with the production-grade ShakeAlert system.

4.2 ShakeAlert

ShakeAlert combines three research early warning systems in California into a single, production-grade system to provide warnings to industrial, government, and corporate recipients, and eventually to the general public. ShakeAlert is based on three research EEW systems: ElarmS at the University of California Berkeley, OnSite at the California Institute of Technology, and Virtual Seismologist at ETH Zurich (Böse, *et al*., 2012). Each system has a unique method of trigger association, magnitude estimation, and false alarm filtering (Cua, *et al.,* 2009, Böse, *et al.,* 2009, Brown, *et al.,* 2009). By combining output from all three systems, ShakeAlert benefits from the strengths of each algorithm and minimizes the weaknesses. The ShakeAlert Decision Module receives alerts from each algorithm, identifies when algorithms are describing the same event, and combines algorithm output into a single summary of each earthquake. The combined event information is then sent as a single alert message to users' computer screens. A popup message warns of impending shaking, with an estimated magnitude and an audible countdown to the predicted shaking. Since January 2012, ShakeAlert is providing prototype warnings to several government and industrial entities in Northern and Southern California (Böse, *et al*., 2012).

4.3 ElarmS Method

ElarmS is a network-based EEW, developed in California over the last ten years. A scaling relation between P-wave frequency ($\tau_p$max) and magnitude was empirically determined from a calibration dataset of earthquakes, first in southern California (Allen and Kanamori, 2003) and then in northern California (Tsang*, et al*., 2007). Then a second set of scaling relations, between P-wave amplitude (P_d) and magnitude, was empirically determined, again for northern and southern California (Wurman*, et al*., 2007). In 2007 the system was tested offline with a dataset of large earthquakes in Japan (Brown*, et al.,* 2009), and also expanded to run in realtime throughout California (Allen*, et al*., 2009). In

2009 ElarmS was connected to the ShakeAlert Decision Module for events in the greater San Francisco Bay Area (Brown*, et al.,* 2011).

ElarmS consists of two primary parts, a waveform processing module and an event monitor. The waveform processing module, WP, is distributed among three seismic processing centers, at UC Berkeley, USGS Menlo Park, and Caltech/USGS Pasadena. Waveforms are telemetered from an individual seismic station to a processing center, where they are processed by the ElarmS WP. The WP detects P-wave arrivals, and stores key waveform parameters immediately following the P-wave: the peak amplitude (displacement, velocity, acceleration) and peak predominant period values every tenth of a second for four seconds. These values, along with the P-wave trigger time, are sent in one-second packets to the ElarmS Event Monitor, running in a single installation at UC Berkeley.

The Event Monitor receives the one-second packets from all triggered stations in California, and associates related triggers together to identify earthquakes in progress. The event monitor then estimates the earthquake's location and origin time using a grid search algorithm, and the earthquake's magnitude using the τ_p^{max} and P_d/P_v magnitude scaling relations. When it determines that an earthquake has a large magnitude, it sends an alert message to the ShakeAlert Decision Module.

The original ElarmS code has been running successfully in realtime for several years, but it was written as a research prototype, not production code. In order to successfully integrate with the realtime ShakeAlert system, we developed a second generation ElarmS code, designed specifically to maximize the current network, hardware, and software performance capabilities.
The new code, referred to as ElarmS-2, or E2, consists of both a new waveform processing module and a new event monitor module. Here we describe the new event monitor.

4.4 Second Generation Code, ElarmS-2

The effectiveness of any early warning system depends on the timeliness of its warnings. E2 is written entirely in C++ for improved processing speed. The new event monitor features a modular code design, facilitating easy replacement of individual algorithm pieces (association, location, magnitude, and alert filter) at any time, without disrupting the processing stream (Figure 4.1). The new location and magnitude modules are functionally identical to their ElarmS-1 counterparts, just translated into C++ and rewritten more concisely. (See Brown, *et al.*, 2009, for more details about the location and magnitude algorithms.) The association and alert filter modules are significantly redesigned, to improve accuracy and reduce false alerts.

4.4.1 New Associator

The purpose of the associator is to recognize that triggers occurring near each other in space and time represent P-wave arrivals from a single earthquake, and to determine which triggers belong to that earthquake and which don't. Later blocks of code will

specify the earthquake's exact location and magnitude, but first the associator needs to recognize that the earthquake is occurring.

Space/Time Box
The foundation of the associator is a simple space/time box (Figure 4.2). When a new trigger enters the system, the associator measures its proximity, in time and distance, from each existing event in the system. If the distance is within the acceptable bounds defined by the space/time box, the trigger is associated with the event, and the event is passed to the location and magnitude modules for recalculation.

The new E2 associator has a wider space/time box, allowing more P-wave triggers to be considered for association. This is important because P-wave times are compared to the estimated event origin. In the early association stages, the estimated event origin may be off by several seconds and up to tens of kilometers. By widening the space/time box and allowing "negative time" for very near triggers, we ensure that slightly inaccurate event origin times and locations do not prevent association of valid triggers.

Multistage Event Creation
The original associator created a new single-station event from any new trigger which could not be associated with an existing event. The new E2 associator has additional levels of event creation (Figure 4.1). If a new trigger cannot be associated with an existing event, it is added to a "hopper" of unassociated triggers. The system then scans through the hopper, looking for any set of 3 or more triggers which can be associated together into a new event. Table 4.1 shows the specific requirements to create a new event with three or more stations. This multi-trigger event step is ideal for large earthquakes, for which many triggers may be detected simultaneously. If the system cannot generate a multi-trigger event, it scans through the hopper again, looking for any two triggers which are near each other in space and time and have significantly large P-wave amplitudes and frequencies. This would indicate the first detections of a large earthquake. Finally, if the system cannot create a two-station event, it scans through the hopper one last time, looking for any one trigger with a P-wave amplitude or frequency large enough to justify a threshold alert message all on its own. Any trigger which is not associated with an existing event, or used to generate a new event, remains in the hopper. An isolated trigger may never be associated, and in that case may remain in the hopper until it expires and is deleted after 90 seconds.

This multi-stage association method ensures that (a) earthquakes generating multiple triggers will be found first, thus prioritizing large earthquakes, and (b) small spurious triggers do not bog down the system by generating unnecessary events.

P-Wave Parameter Thresholds
To support multi-stage association, E2 defines multiple thresholds for acceptable $\tau_P max$ and P_d values. Figure 4.3 shows the magnitude scaling relations for $\tau_P max$ and P_d in Northern and Southern California (see Brown, *et al.,* 2009, for more details about scaling relations and magnitude estimation). When associating a new trigger with an existing event, or creating a new event with three or more triggers, each trigger's $\tau_P max$ and P_d

values must be within the level 1 range, which covers the entire scaling relation and all observed magnitudes.

Level 1:	$-0.9 \leq \log_{10}(\tau_P max) \leq 1.0$	$-5.5 \leq \log_{10}(P_d) \leq 3.5$
Level 2:	$0.3 \leq \log_{10}(\tau_P max) \leq 1.0$	$0.5 \leq \log_{10}(P_d) \leq 3.5$
Level 3:	$0.45 \leq \log_{10}(\tau_P max) \leq 1.0$	$1.25 \leq \log_{10}(P_d) \leq 3.5$

When creating a new, two-station event, each trigger's $\tau_P max$ and P_d values must be within the level 2 range, corresponding to larger magnitude events. And for single-station events, E2 requires the $\tau_P max$ and P_d values to be within level 3 range, corresponding to extremely large magnitude events and prompting an immediate threshold alert message.

Station Density Adjustments

The new E2 associator also adjusts its association criteria to the local density of seismic stations. Rather than requiring a set number of stations (after the initial event creation), it requires 50 percent of near-source stations to trigger. In densely-instrumented areas, far more stations must trigger for the event to be created than in sparsely-instrumented areas. As new stations trigger, the system considers whether including those triggers would violate the 50% requirement before deciding whether to associate them with the event.

Split Event Prevention

One challenge for the original ElarmS associator was "split events", when a single earthquake would generate many triggers, some of which were associated with one system event, and some with a second system event. ElarmS-1 thus created two event numbers, two locations, two magnitudes, and two alert messages, when in reality there was only one earthquake.

The new E2 associator attempts to confront the problem of split events by defining a "blackout window" around existing events. When the associator has a set of triggers prepared for generating a new event, it checks all existing events. If any existing event epicenter is within 15 seconds and 90 kilometers of the proposed new event epicenter, the associator cancels the new event. Any triggers which were flagged for association with the new event are released back into the hopper of unassociated triggers. This simple prevention of redundant events prevents the vast majority of split events.

Figure 4.4 shows the San Leandro M3.6 event of August 24, 2011, processed by the E2 associator with and without the blackout window. Without the blackout window (Figure 4.4a), a coincident trigger north of the epicenter, seven seconds prior to the event, biases the event epicenter and origin time estimates. Additional triggers cannot be reconciled with the inaccurate epicenter, so they are formed into a second, erroneous event. Still later triggers do not match either event and are not associated at all. When the blackout window is enforced (Figure 4.4b), the second event is rejected and only the first event is created. As later triggers arrive they are associated with the first epicenter, gradually improving the accuracy of the estimated epicenter and origin time. By the time 9 stations have been associated with the event, the estimated epicenter is 1km from the ANSS catalog epicenter, and the estimated origin time is correct. Figure 4.4b shows the final event map 18 seconds after the event origin, when all stations have triggered and all but two have been associated with the event.

4.4.2 New Alert Filter

The new E2 event monitor also upgraded the Alert Filter Module at the end of processing, just prior to sending messages to the ShakeAlert Decision Module. In order to send an alert message, the ElarmS-1 event monitor required an event to have four or more triggers, an event magnitude of 2.0 or greater, and an estimated epicenter within the greater San Francisco Bay Area (regions mSA, SFBA, and nSA in Figure 4.5).

Number of Stations

The new event monitor requires four stations, rather than triggers. This seems like a minor technicality, but as the seismic network in California is expanded, more and more stations are installed with multiple instruments, such as co-located accelerometers and broadband seismometers. The old requirement of four triggers can now often be satisfied by just two stations, which are not sufficient for triangulating an accurate epicenter.

Multiple Magnitudes

The new event monitor also considers not just the averaged event magnitude, but the $\tau_P$max and P_d magnitudes separately, as well as their relationship to each other. The average event magnitude must still be greater than M2.0, but also the P_d magnitude must be greater than M1.5. Perhaps more significantly, the difference between the two magnitudes can not be more than 2.5 magnitude units. Agreement between both frequency and amplitude of the P-wave minimizes the effect of scatter in the magnitude scaling relations for each parameter and reduces the likelihood of false alarms.

Expansion of Alert Regions

The alert block also expands on the regional alert criteria defined by the ElarmS-1 event monitor, adding alert criteria for the remaining regions of the state (Table 4.2 and Figure 4.5). The new association and alert modules improve our confidence in ElarmS performance statewide, and we now send alerts for all regions.

Event History

Lastly, the new alert module stores a history of alert messages sent, so it can recognize whether changes to an event justify an updated alert message. The alert history has also proven useful for later event analysis. Thirty minutes after an event is first alerted, E2 sends an event history email to the development team, detailing all steps of event processing, such as when new triggers were associated and how the magnitude and location evolved with time.

4.4.3 Upgrades to Waveform Processing and CISN hardware

The new waveform processing module, which runs at each of the processing centers, has also been upgraded. WP2 is designed to read and process smaller packets of waveform data, and send the resulting parameters more promptly. Details about the new WP2 can be found in Kuyuk, *et al.,* 2012 (in prep).

Finally, in addition to the software upgrades to both the waveform processing and event monitor modules, the CISN network of seismic instruments has benefitted from the

recent American Recovery and Reinvestment Act (ARRA) funding, which paid for hardware upgrades. Newer instruments, dataloggers, and telemetry systems have greatly improved the speed of data transmission throughout the network.

4.5 Conclusion

The new E2 code upgrade improves both the speed and accuracy of the ElarmS system. It is robust, modern code, compatible with the production-grade ShakeAlert system. E2 is being further tested and implemented by the Berkeley Seismological Laboratory EEW group, and performance results will be published in Kuyuk, *et al.*, 2012 (in prep).

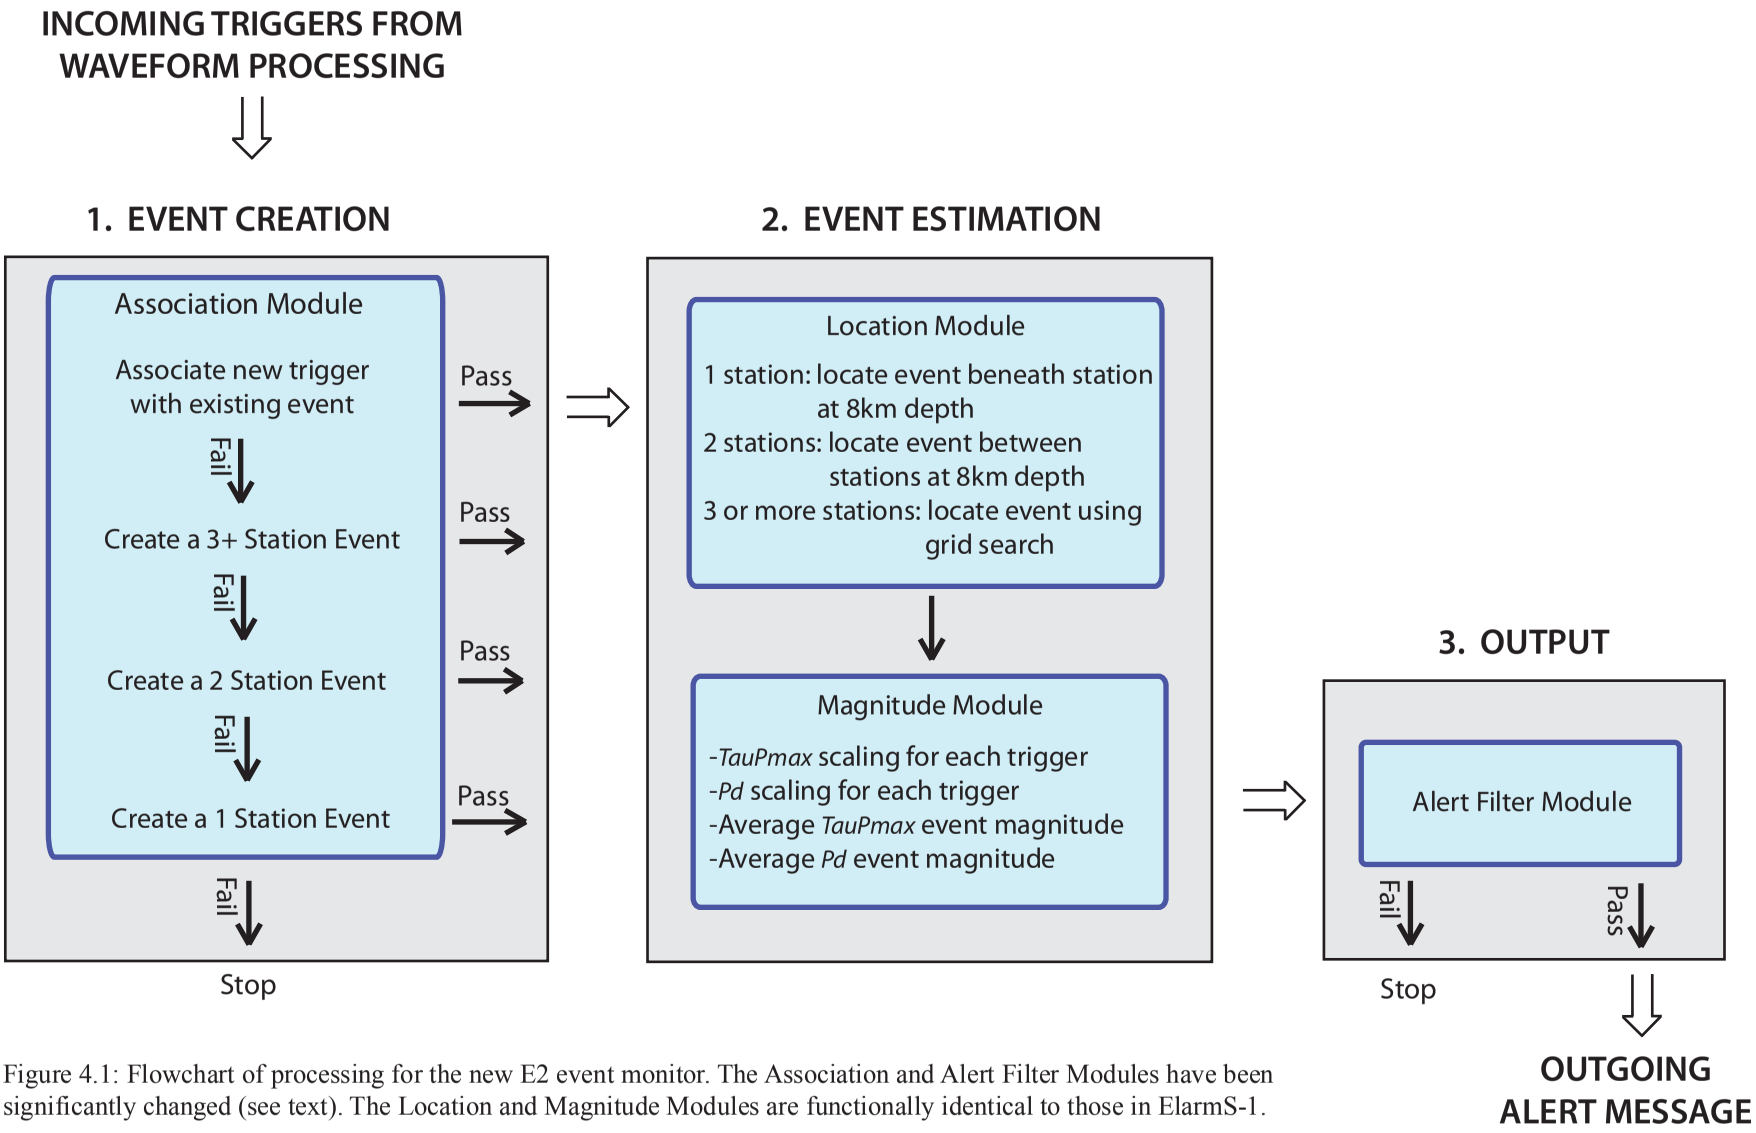

Figure 4.1: Flowchart of processing for the new E2 event monitor. The Association and Alert Filter Modules have been significantly changed (see text). The Location and Magnitude Modules are functionally identical to those in ElarmS-1.

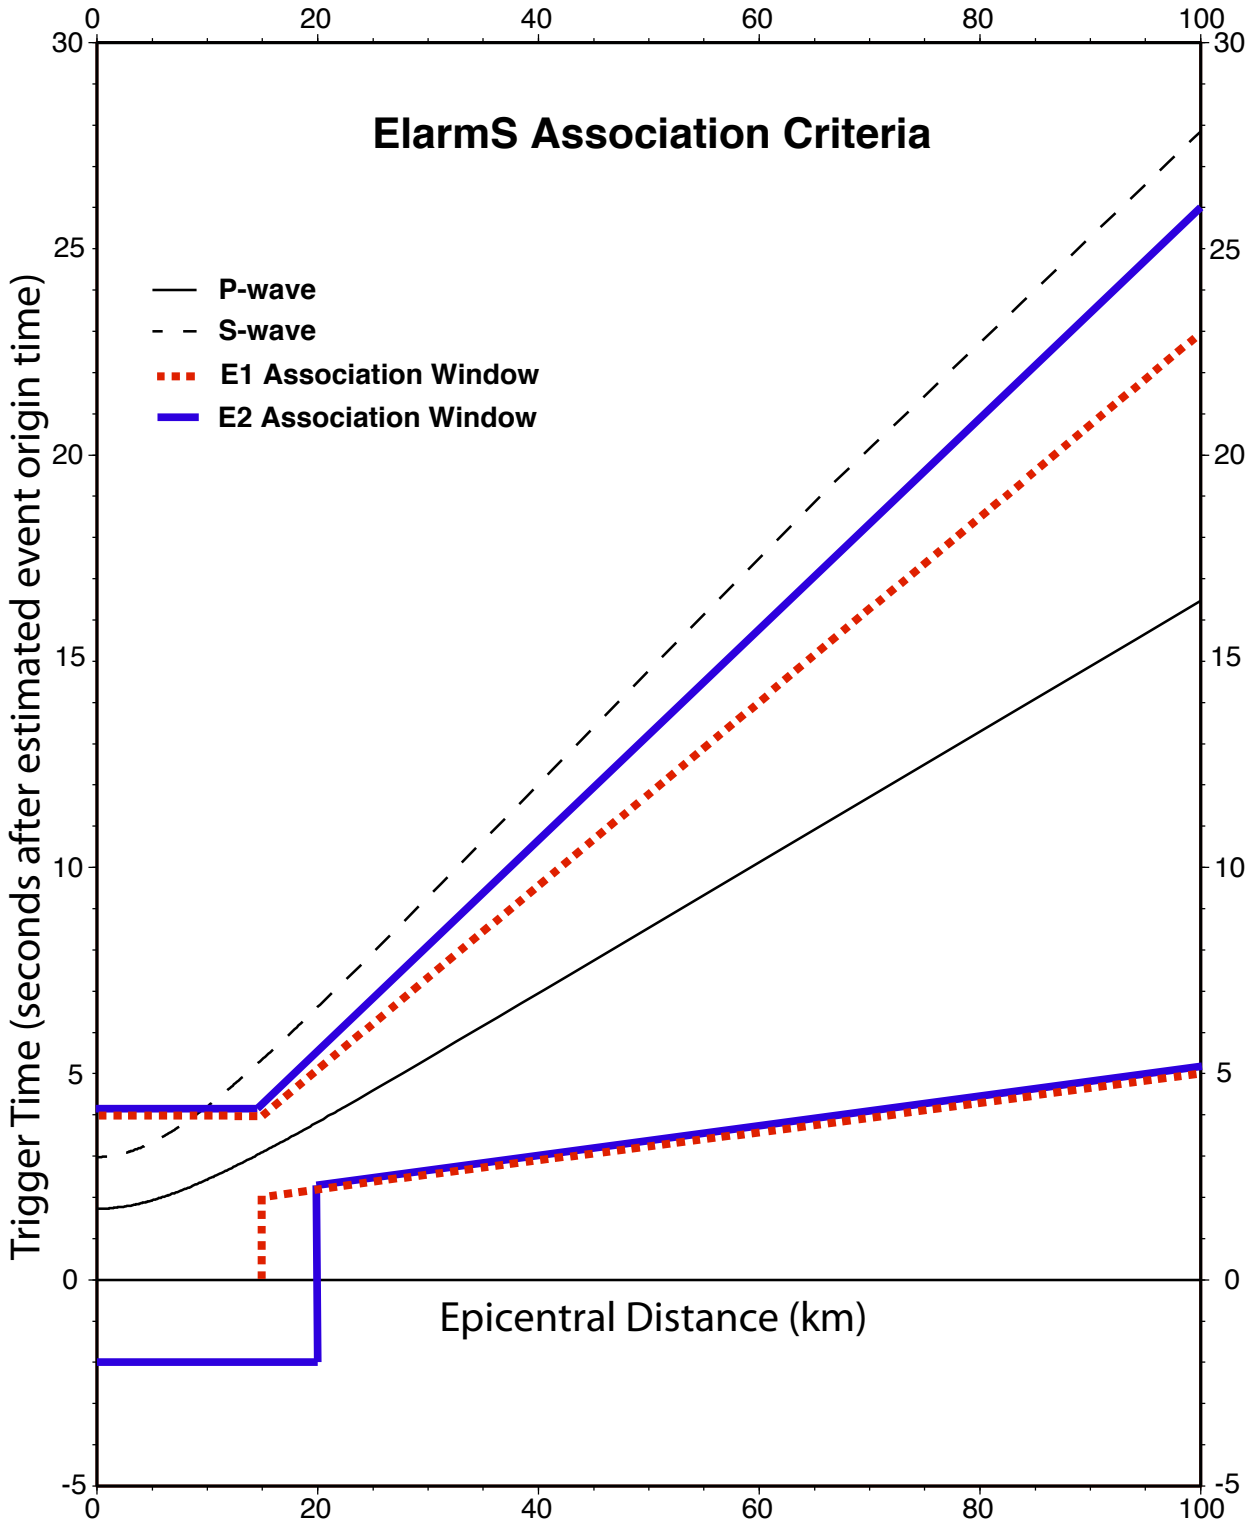

Figure 4.2: Space/Time Association boxes for ElarmS-1 (red dashed line) and E2 (blue solid line).

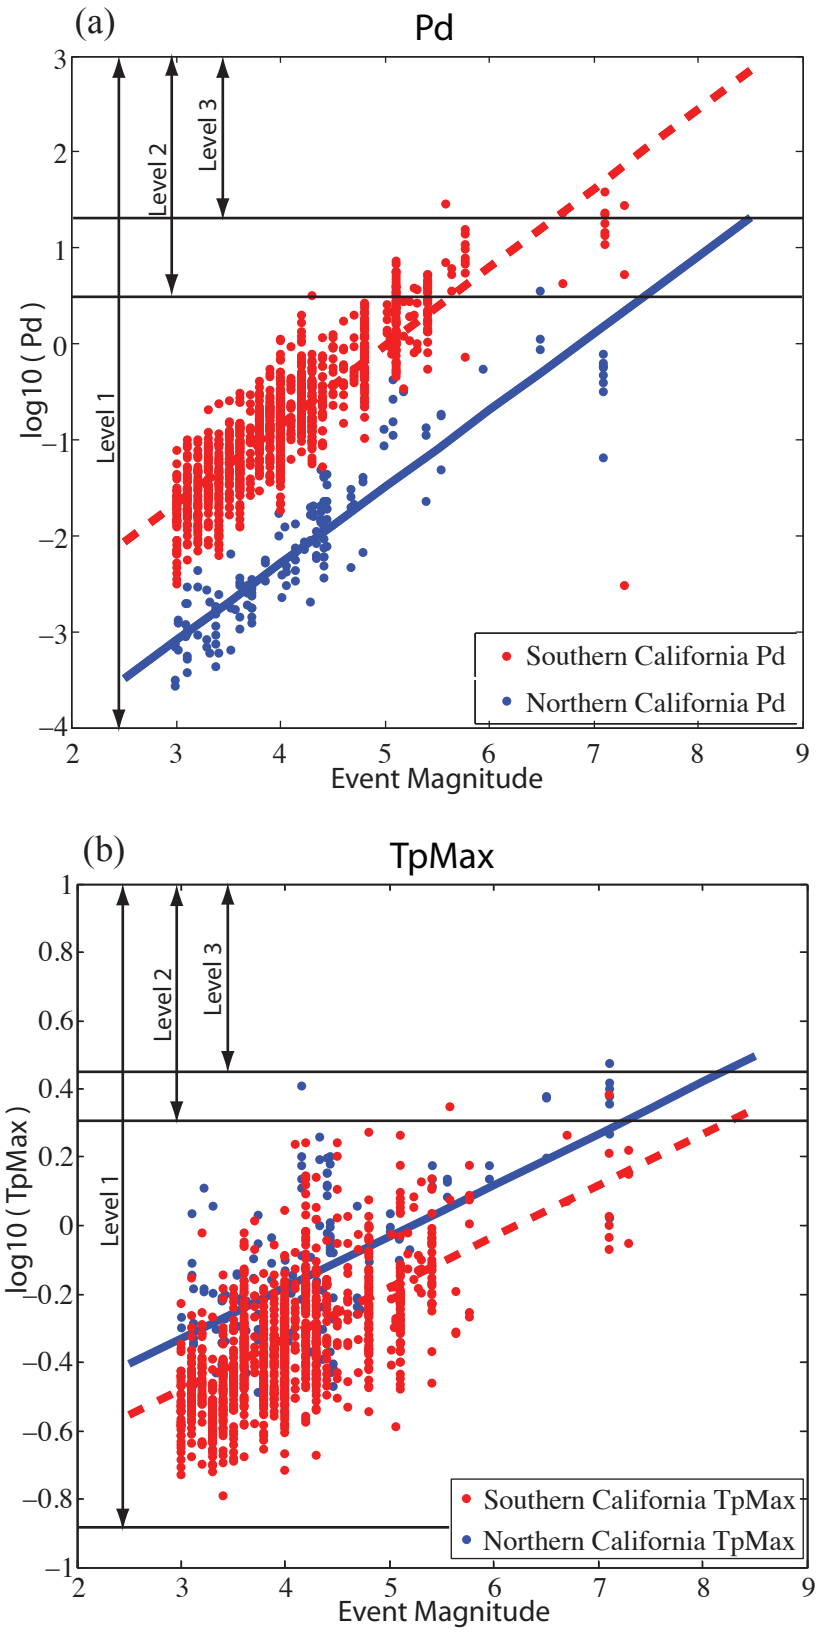

Figure 4.3: Threshold values for (a) Pd and (b) TpMax P-wave parameters. Values must be within level 1 for a trigger to be associated with an existing event or used to generate a new, 3+ station event. Values must be within level 2 to generate a 2 station event, and within level 1 to generate a single station event. More information about the Pd and TpMax magnitude scaling relations can be found in Brown, et al., 2009.

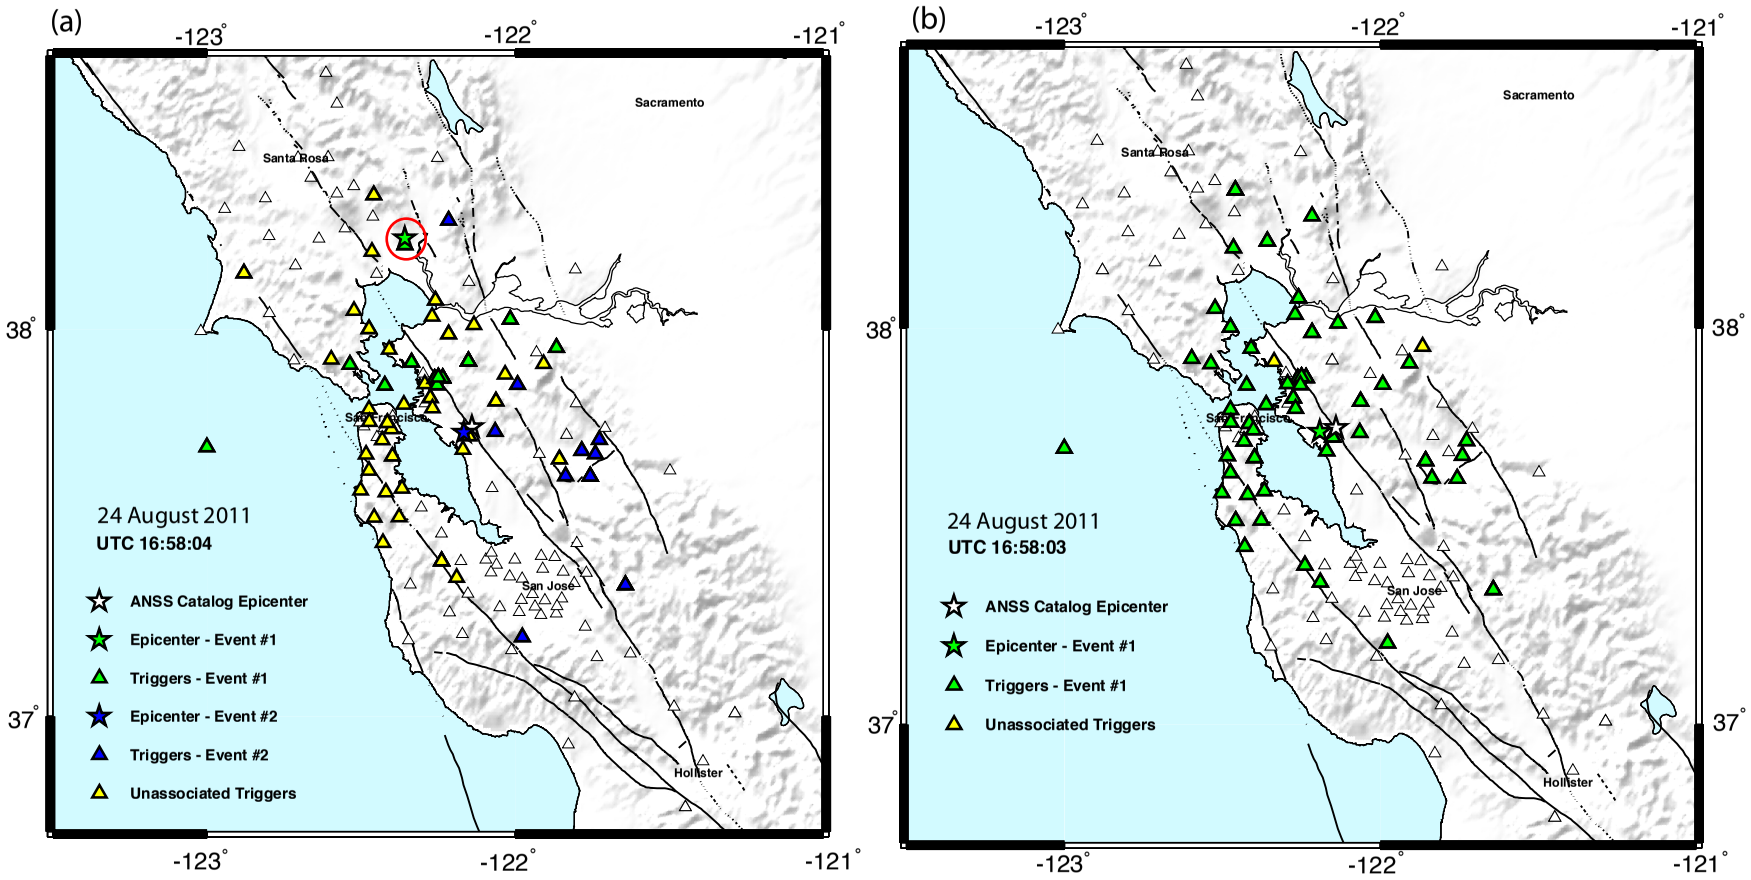

Figure 4.4: San Leandro Mw 3.6 event, 24 August 2011. (a) E2 processing of event without blackout box. A coincidental trigger to the north (red circle) seven seconds prior to the event biases the epicenter location (green star). The system then generates a second event (blue), trying to account for all the triggers to the south. Many additional triggers (yellow) cannot be associated with either event, because their trigger times do not match either estimated event origin. (b) E2 processing after application of the blackout box. The second event is rejected. As more triggers are added to the first event, the accuracy of the event estimation improves, canceling out the effects of the one bad trigger. The final estimated epicenter (green star) is 5km from the ANSS catalog epicenter (white star), and almost every triggered station is associated with the event.

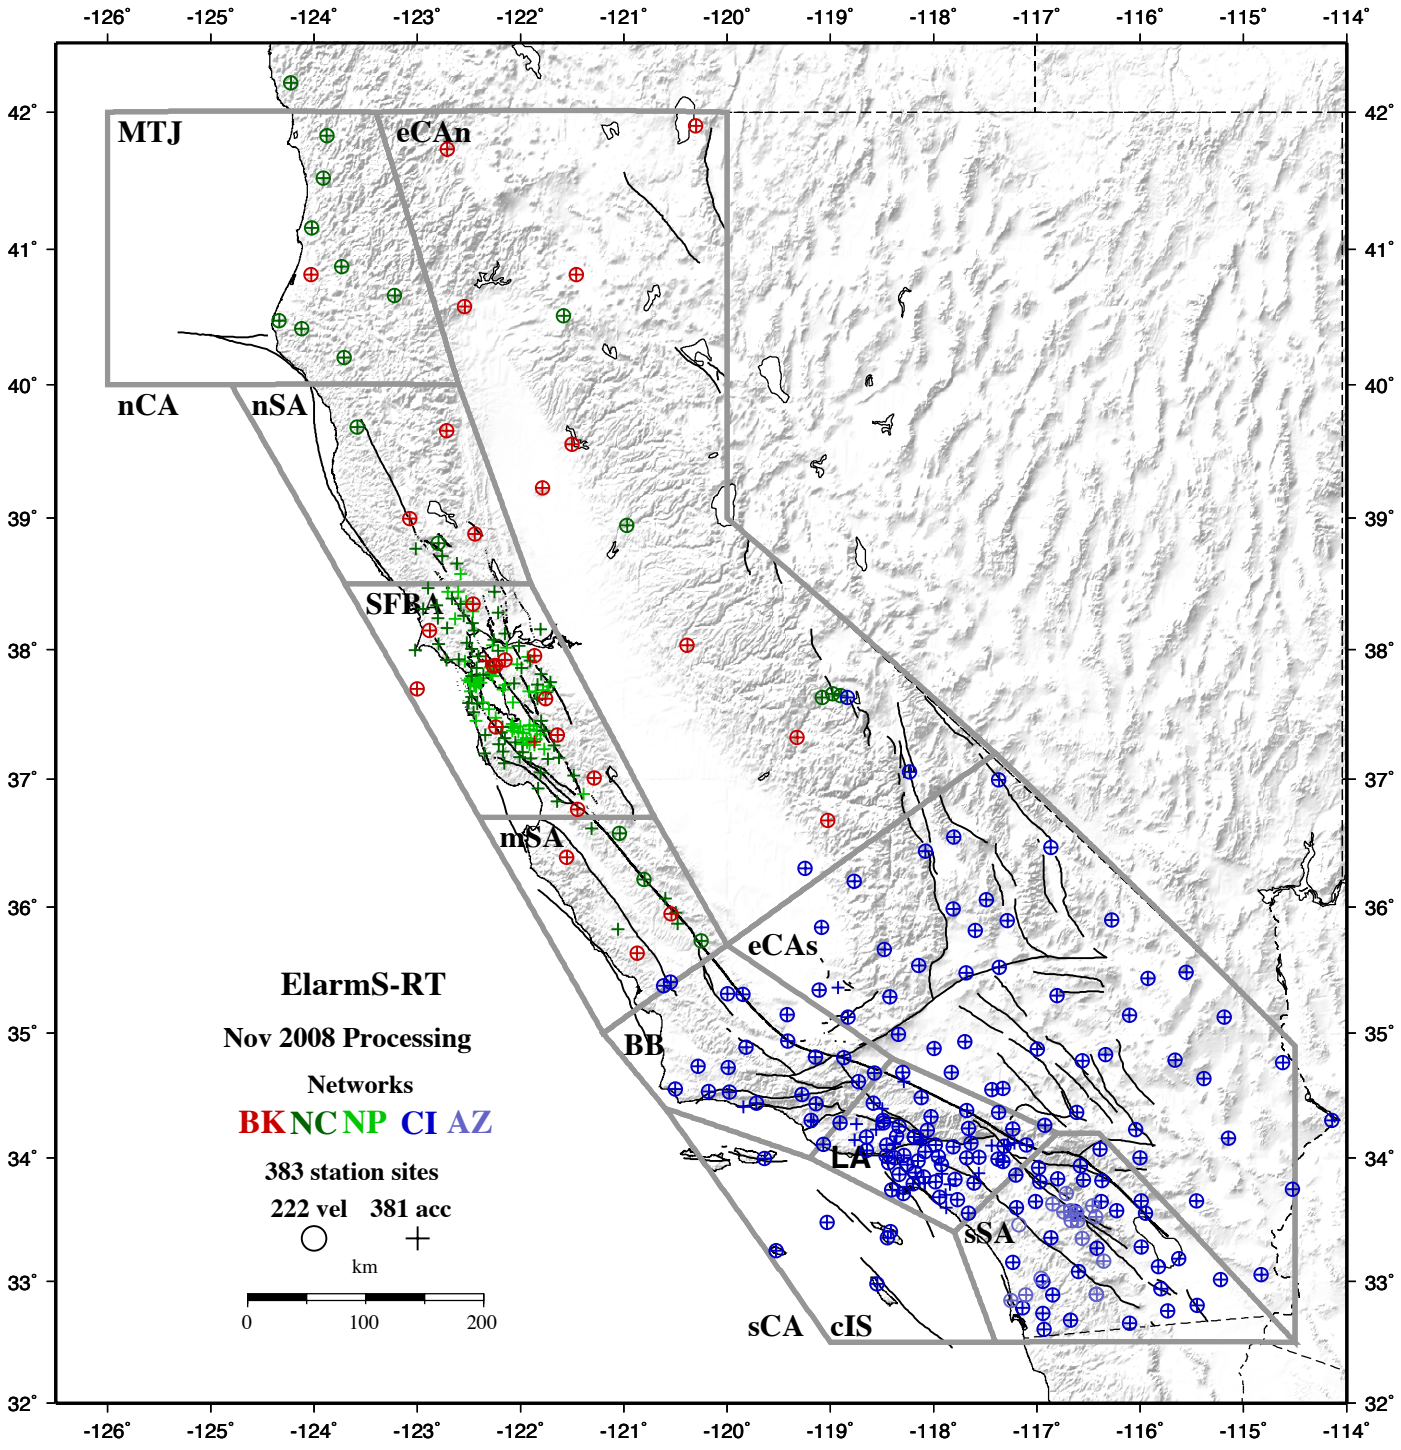

Figure 4.5: Map of alert regions in California. ElarmS-1 sent realtime alerts for events in the nSA, SFBA, and mSA regions. E2 is designed to send alerts for all regions.

Table 4.1: Comparison of Association criteria for ElarmS-1 and E2.

Association Requirements

1. Associate triggers with existing events	ElarmS-1	E2
Tp/Pd range requirement	none	level 1
Association Space/Time Box	Red dashed line in Fig 1	Blue solid line in Fig 1
Required percentage of triggers	none	50%
2. Create new events with 3+ triggers	N/A	
Tp/Pd range requirement		level 1
Association Space/Time Box		65 km, 10 sec
Blackout Box around existing events		90 km, 15 sec
3. Create new events with 2 triggers	N/A	
Tp/Pd range requirement		level 2
Association Space/Time Box		30 km, 5 sec
Blackout Box around existing events		90 km, 15 sec
4. Create new events with 1 trigger		
Tp/Pd range requirement	none	level 3
Blackout Box around existing events	none	90 km, 15 sec
Forced event creation?	yes	no

Table 4.2: Alert Criteria for ElarmS-1 and E2.

Alert Requirements

Regional Criteria	ElarmS-1	E2
SFBA	4 triggers w/in 35km	4 stations w/in 35km
mSA	5 triggers w/in 100km	5 stations w/in 100km
nSA	5 triggers w/in 100km	5 stations w/in 100km
LA	-	4 stations w/in 35km
sSA	-	5 stations w/in 50km
All others: MTJ, eCAn, BB, eCAs, cIS	-	10 stations w/in 100km
Magnitude Criteria	**ElarmS-1**	**E2**
Minimum Event Mag	2.0	2.0
Minimum Pd Mag	-	1.5
Maximum Mag Difference: Pd Mag - Tpmax Mag	-	2.5

Chapter 5

Faster Earthquake Alerts with Assistance from an Artificial Neural Network

This chapter has been submitted to the *Bulletin of the Seismological Society of America* as "Faster Earthquake Alerts with Assistance from an Artificial Neural Network", by Brown and Allen, October 2012.

5.1 Introduction

Earthquake Early Warning (EEW) systems scan realtime seismic waveforms for P-wave arrivals, to rapidly identify earthquakes in progress, estimate the location and intensity of the event, and warn surrounding communities of the danger a few seconds before severe shaking begins (Allen and Kanamori, 2003). Several EEW systems are in place throughout the world; some warn the general public via cellphone, some are limited to government and industrial applications, and some are still in development (Allen, *et al.*, 2009a). Earthquake Alarm Systems, or ElarmS, is one of the algorithms currently contributing to the ShakeAlert demonstration EEW system in California (Böse, *et al*., 2012). ShakeAlert now sends alerts to test users throughout the state, and is in the process of continually being evaluated and improved. ElarmS is a network-based EEW system, requiring several adjacent seismometers to confirm a P-wave arrival before it will process the event and send an alert message. Requiring multiple seismometers greatly improves the accuracy of the alerts, but potentially decreases the speed of the alert as the P-wave has to arrive a more sensors.

ElarmS receives realtime seismic waves from 572 seismic instruments (digital seismometers and accelerometers), located at 397 unique station sites, across California (Allen, *et al.*, 2009b). It associates P-wave arrivals, referred to as triggers, at neighboring instruments together to detect earthquakes in progress. Trigger times and locations are used to triangulate an event epicenter. The P-wave amplitudes and frequency content are then used to estimate event magnitude (Brown, *et al.*, 2009; Brown, *et al.*, 2011). Once an event is detected, and its location and magnitude estimated, it must pass the ElarmS alert criteria before an alert message will be sent to users. The alert criteria specifies minimum magnitude, regional boundaries and, most importantly for this study, the number of seismic stations which must be reporting P-wave triggers. (There may be more than one seismic instrument at any given station.) The current methodology and performance of ElarmS is detailed by Kuyuk, *et al.* (in prep).

When ElarmS requires four or more seismic stations to trigger (register a P-wave arrival) before sending an alert message, it has a statewide false alert rate of only 4%. If we relax

the alert requirements and allow ElarmS to send alerts with only 2 or 3 stations, we observe a prohibitively high false alert rate of 21% or 15%, respectively. These false alerts are mostly due to spurious noise triggers at individual stations. The majority of the false alerts are small magnitudes (M~2.5), but still we want to minimize them. Therefore the alert criteria in the current realtime version of ElarmS requires 4 stations before sending an alert to users. Waiting for four stations to trigger exacts a cost of one to two seconds. When the warning is only a few seconds to begin with (for all types of early warning systems), waiting an extra second or two for additional stations to trigger is a high cost. The purpose of this study is to determine how we can improve ElarmS performance with fewer triggered stations, and thus send faster, earlier alerts.

5.2 ANN Approach

In 2012 we began testing an Artificial Neural Network (ANN) at the end of ElarmS offline processing, to catch false alerts before they are released to users. The ANN reads input data (an earthquake alert message) and the desired output (true or false) for a large dataset of sample events, and optimizes a mapping function between inputs and outputs (Bishop, 1995). That function can then be utilized on future events to filter out alerts with a high probability of being false.

5.2.1 Training Method

We utilized the MATLAB Neural Network Toolbox (Demuth, *et al*., 2009) to generate each ANN, with ten neurons in the hidden layer and one neuron in the output layer. During training, the ANN is provided a random subset of 70% of the inputs and their outputs, and iteratively finds the best function to describe the relationship between the sets. Each neuron has an associated weight and bias (Bishop, 1995), and the training stage utilizes feed-forward back propagation to adjust the neural weights and biases until the output error is minimized (Derras, *et al.*, 2012). A separate 15% of the dataset is used to repeatedly test the evolving function and recognize when it has stopped improving. Finally, the last 15% of the dataset is used as an independent test of the function's performance.

5.2.2 Dataset

Our dataset was all ElarmS alert messages between December 6th, 2011, and May 14th, 2012. Within this time period the Advanced National Seismic System (ANSS) catalog lists 1061 real earthquakes of magnitude 2 or greater in California, up to a maximum observed magnitude of 5.6. An ElarmS alert message was designated true if its estimated location and origin time were within 100km and 30 seconds of an ANSS catalog event.

We configured ElarmS in three ways, to generate events and send alert messages using 2 or more stations, 3 or more stations, and 4 or more stations. Depending on the configuration, ElarmS produced between 337 and 512 alert messages in this time period. Each of these configurations produced a separate dataset, used to train a unique ANN function.

5.2.3 Variety of Input Configurations

Each of these datasets contained a great deal of information about the events. We varied the parameters offered to the ANN, to see which event details were most useful. A human

seismologist, glancing at an incoming alert, would depend most on location, magnitude, number of triggered stations, and distribution of stations to predict whether the event is true or false. What factors would the ANN depend on? The input options offered to the ANN included location (latitude, longitude), residual of the location estimate which is the average absolute difference between the observed trigger times and the predicted trigger times (referred to as location misfit), the number of triggered stations (nS), the number of triggered instruments (nT, of which there might be several at any given station), the estimated magnitude, two numerical descriptions of trigger distribution, and different combinations of these. The trigger distribution was described by counting the number of triggers as a function of azimuth (in 60 degree bins) and distance from the epicenter (in 20km radius bins).

5.3 Results

5.3.1 Statewide Results, 2-Station Configuration

We began by determining which input variables were of most value to the ANN to discriminate between true and false events. We configured ElarmS to generate alerts after two or more stations have triggered, since this configuration generates the most alert messages. The dataset contained 512 alert messages between December 6, 2011, and May 14, 2012, with a statewide false alert rate of 21%. We selected an assortment of input variables (latitude, longitude, magnitude, number of triggered stations, number of triggered instruments, and location misfit) and trained an ANN for each input variable or combination of variables. Table 5.1 shows the resulting false alert rate when each ANN function was tested.

The best results (rows H and I) are when we provide the ANN with all the information contained in a typical alert message: latitude, longitude, location misfit, number of triggered instruments, and number of triggered stations. The inclusion or exclusion of magnitude does not affect the results. Adding trigger distribution bins (rows J and K) shows no evidence of improvement, likely due to the heterogeneous distribution of stations in California, including the absence of stations offshore.

Dividing the input variables provides insight into which variables are most useful for predicting false alerts. Training an ANN with only latitude and longitude (row A) and no other event characteristics results in a false alert rate of 19%. This is a slight improvement over the original rate of 21%. Training an ANN with nT and nS only (row C), or with location misfit only (row B), results in a increased (worse) false alert rate of 23%. And training an ANN with magnitude alone (row D) gives the worst result of all, a 27% false alert rate.

This means that when we train an ANN with all available data (row H), it is depending most heavily on event location, with a lesser dependence on number of triggered stations/instruments and location misfit. This is not surprising, given that a human seismologist often considers event location first, as well. Events in areas with historically high seismicity are likely to be true. Events in historically quiet regions are more likely to be false. The ANN mapping function grants a similar *a priori* historical knowledge to the ElarmS algorithm.

For all subsequent ANNs, we use the five input parameters from row H in Table 5.1: latitude, longitude, location misfit, number of triggered instruments, and number of triggered stations. We choose to not include the magnitude as this does not improve performance and its inclusion may result in an increased likelihood that the ANN will declare an event false if the magnitude is unusual (i.e. large) which is precisely when the alert is most important.

5.3.2 Statewide Results, 2-, 3-, and 4-Station Configurations

Having determined the best combination of input variables, we moved to trying different ElarmS configurations. ElarmS can be configured to send alerts with 2 or more stations, 3 or more stations, or 4 or more stations. Table 5.2 shows the false alert rate before and after application of the optimized ANN mapping function.

When ElarmS is configured to send alerts with two or more stations, applying an ANN mapping function improves the false alert rate from 21% to 17%, as we saw previously. When ElarmS sends alerts with three or more stations, the ANN has a greater effect, improving the false alert rate from 15% to 8%. When ElarmS sends alerts with four or more stations, as the current realtime version does, the ANN can offer essentially no improvement. The false alert rate is 4%, with or without the ANN. When we examine the results from this trial (row M) more closely, we find that the ANN mapping function identifies every single ElarmS alert as true. There are so few false alerts in the 4-station dataset that the ANN can find no pattern at all, and the resulting function is simply a one-to-one relationship.

5.3.3 Regional Specifics

Noting that the ANNs depend most heavily on location, we examine which regions of California are most effected. Figures 5.1 and 5.2 show maps of alerted events, before and after application of an ANN mapping function.

The blue stars in Figures 5.1b and 5.2b represent blocked false alerts. Without the ANN mapping function, ElarmS would send alert messages for these nonexistent events. When we applied the mapping function, the alerts were correctly identified as false, and were not sent. The red stars represent newly missed events. These are real events, which ElarmS correctly sent alerts for without the ANN. However, when we apply the mapping function, the events are incorrectly identified as false, meaning alert messages would not be sent. In Figures 5.1b and 5.2b note that almost every red and blue star is located in Southern California. In other words, most of the false alerts are in the south and the mapping function is therefore blocking alerts primarily in the south and making almost no changes in the north.

We returned to our ANN outputs and divided them at the 36^{th} parallel, to see the false alert rates in each half of the state. Table 5.3 shows the false alert rates before and after the application of ANN mapping functions, for Northern California (latitude >= 36 degrees), Southern California, and total statewide.

The difference between performance in Southern and Northern California is stark. In Northern California, the ANN mapping function does not noticeably affect output. False alert rates are the same before and after the application of the ANN filter. Even when requiring only two stations to send an alert, ElarmS demonstrates a false alert rate of merely 5% in Northern California.

In Southern California, however, the false alert rates are six to eight times higher. When ElarmS is configured to send alerts with two stations, it has a Southern California false alert rate of 29%. The ANN decreases the false alert rate to 24%. When ElarmS is configured for three stations, it has a baseline false alert rate of 23% in Southern California, dropping to 12% when we apply the ANN mapping function. When we require four or more stations, as the current released version of ElarmS does, the Southern California false alert rate is 4%, equivalent to that of Northern California and unaffected by the ANN.

5.3.4 Effect of Earlier Alerts on Timeliness and Magnitude Accuracy

Table 5.4 shows the time gained by sending earlier alerts, and the change in accuracy of the magnitude estimate. We calculated the average time gained by subtracting the arrival time of the second triggering station from that of the fourth triggered station for each of the events in our dataset. On average, we find that ElarmS receives the second trigger 2.2 seconds before the fourth trigger (Table 5.4). The third trigger arrives 1.8 seconds before the fourth trigger on average.

An ElarmS four-station alert has an average magnitude error of -0.1 ± 0.2 magnitude units. Alerts sent with two or three stations have an average magnitude error of 0.2 ± 0.4 magnitude units. So by sending two- or three-station alerts we would be able to warn users about two seconds earlier than when we delay alerts until four stations have triggered, at a cost of about 0.1 magnitude units of accuracy.

5.3.5 Effect of Large Magnitudes

When we looked at the different possible input parameters (Table 5.1) that we could offer to the ANN, we determined that the row H combination was most effective. Each subsequent ANN was given event latitude, longitude, number of triggered stations, number of triggered instruments, and location misfit. The ANN mapping function and its results are thus completely independent of magnitude. There is then no difference between the small to moderate magnitude events used for the development of the ANN and large (magnitude > 6) events for which warning is most important. The number of triggered instruments and stations would be larger for a large magnitude event, but we are only considering using the mapping function for the initial 2- and 3-station events. When additional stations have triggered, ElarmS would bypass the mapping function anyway.

5.4 Conclusion

We find that applying a ANN-derived mapping function at the end of the ElarmS processing stream can improve the accuracy of ElarmS alert messages for Southern California events when only two or three stations have triggered. The number of false alerts generated by ElarmS in Northern California, and statewide after four or more

stations have triggered, is already very low and cannot at this time be improved by applying an ANN mapping function.

We recommend the following application of the ANN to ElarmS realtime processing:
1. When only two or three stations are providing trigger information, apply the ANN-derived mapping function to decrease the false alert rate. This study also provides the basis for likelihood estimates giving the probability that an alert is a true earthquake. Users can then have the option of receiving 2- or 3-station alerts earlier, with lower likelihoods, or not.
2. The ANN can be applied to all events in California for simplicity of implementation. However, it will only have an effect on the number false two- or three-station event alerts for events south of latitude 36, where the false alert rate is higher. North of this latitude the number of false alerts is already low and is not improved by the ANN.
3. Once an event has four or more stations reporting, bypass the ANN mapping function completely and issue all alerts generated by ElarmS as is currently done. Applying the ANN at the 4-station triggered stage provides essentially no improvement. Not using the ANN at this stage also removes the concern that the ANN will block an event because it is unusual compared to the training dataset.

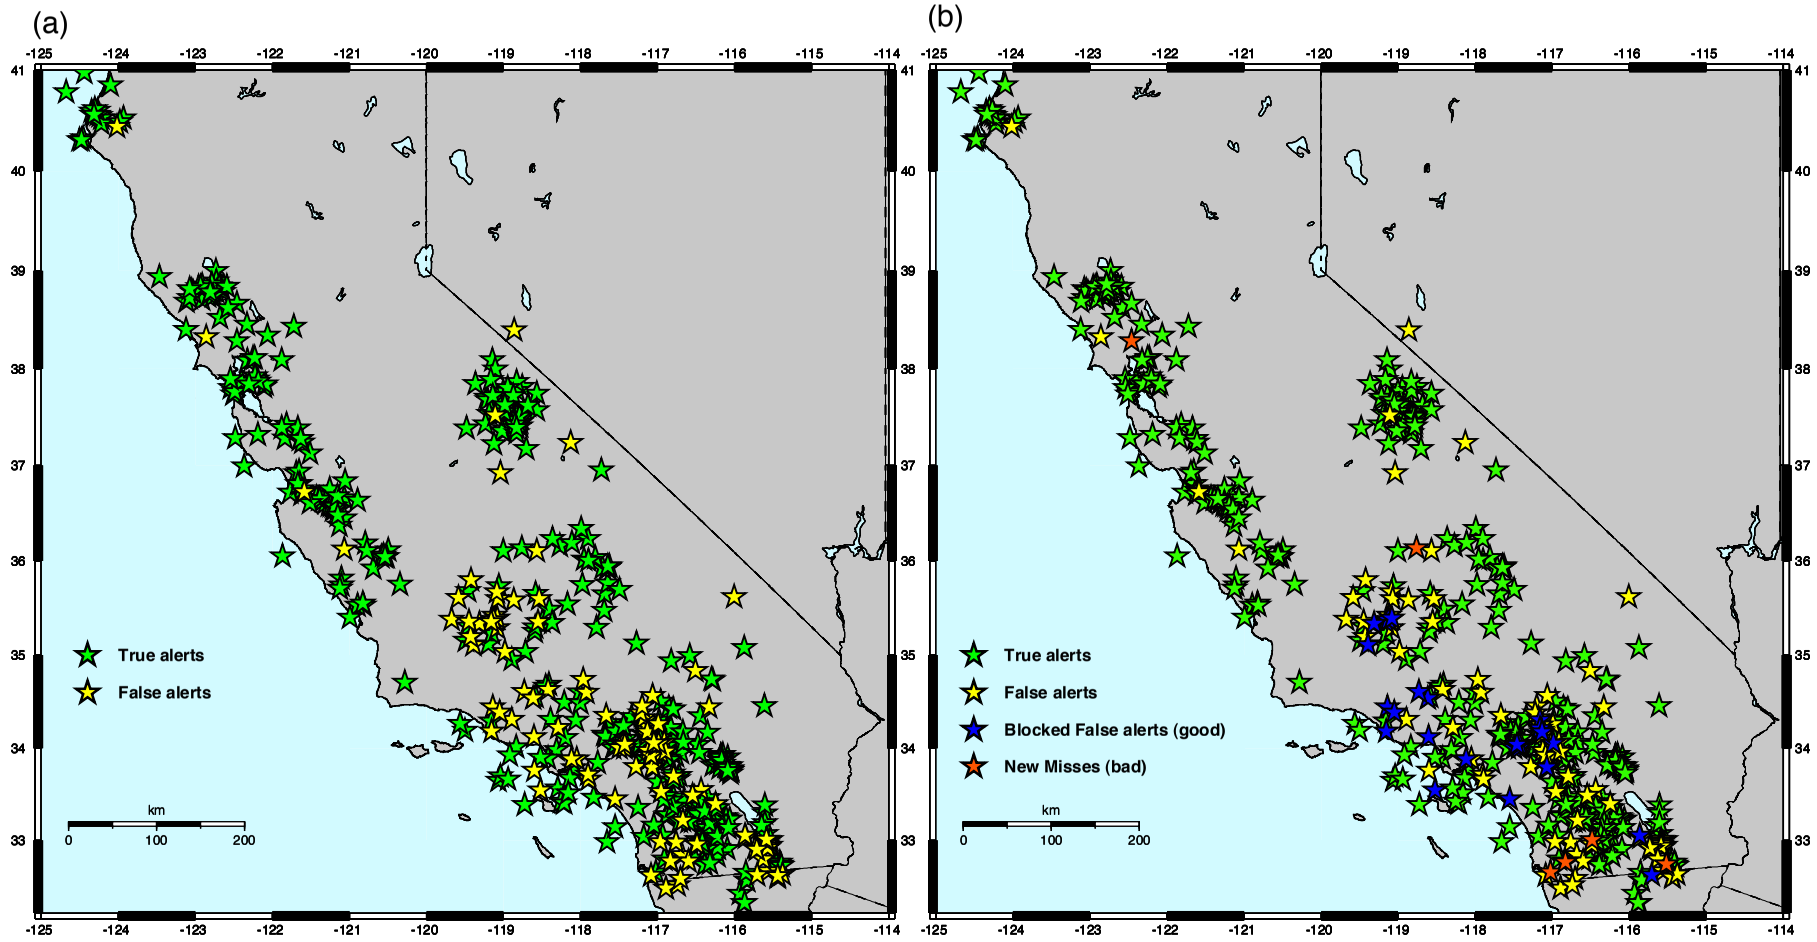

Figure 5.1: Maps of alerted events when ElarmS is configured to send alerts with 2 or more stations. For time period Dec. 6, 2011, to May 14, 2012. (a) Output without Mapping Function. (b) Output after application of Mapping Function. Green stars are true (correct) alerts. Yellow stars are false alerts. Blue stars are blocked false alerts. Red stars are newly missed alerts (incorrectly blocked true alerts).

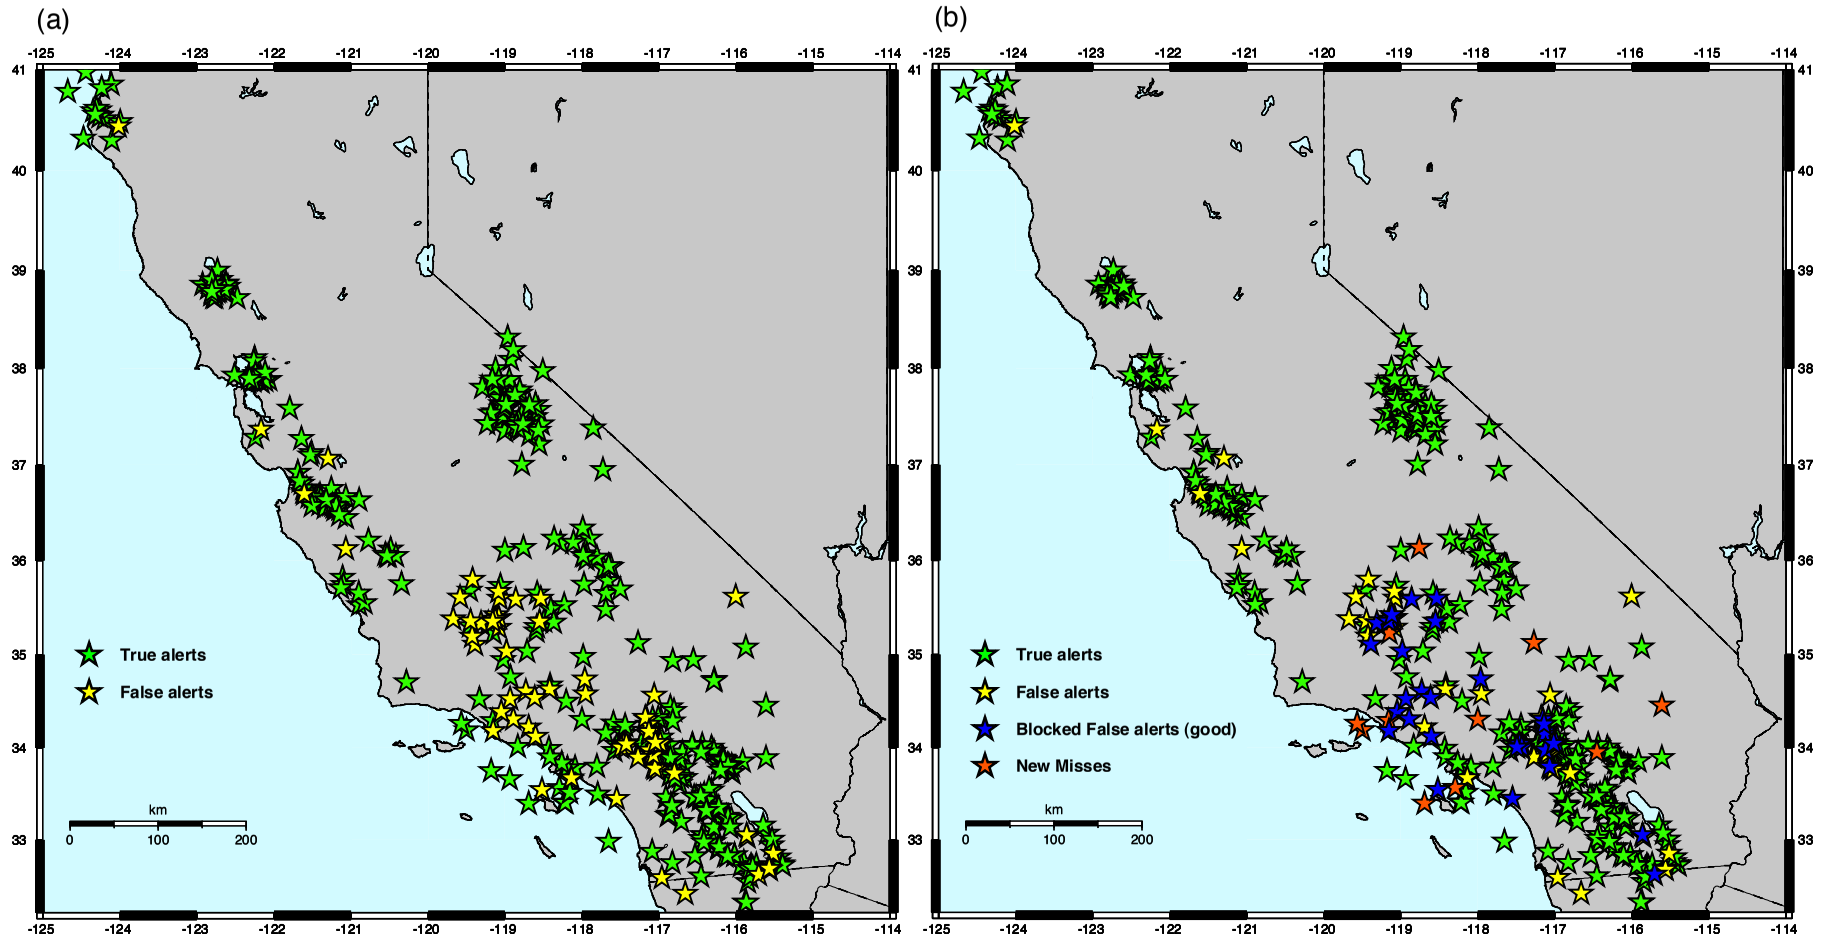

Figure 5.2: Maps of alerted events when ElarmS is configured to send alerts with 3 or more stations. For time period Dec. 6, 2011, to May 14, 2012. (a) Output without Mapping Function. (b) Output after application of Mapping Function. Green stars are true (correct) alerts. Yellow stars are false alerts. Blue stars are blocked false alerts. Red stars are newly missed alerts (incorrectly blocked true alerts).

Table 5.1: Comparison of input variable options. The first row shows the initial false alert rate, before application of any ANN. Each ANN (A through K) is trained on the same dataset of 512 alert messages, produced when ElarmS is configured to send alerts with 2 or more stations. The combination of input variables is unique to each ANN. The false alert rate shown is that during the final testing stage of the ANN function. Abbreviations are as follows: Lat = latitude; Lon = longitude; Loc. Misfit = the residual of the location estimate; nT = number of triggered instruments; nS = number of triggered stations.

	Input Variables	% False
	None (no ANN)	21%
A	Lat, Lon	19%
B	Loc. Misfit	23%
C	nT, nS	23%
D	Magnitude	27%
E	Lat, Lon, Loc. Misfit	21%
F	Loc. Misfit, nT, nS	21%
G	Lat, Lon, nT, nS	21%
H	Lat, Lon, Loc. Misfit, nT, nS	17%
I	Lat, Lon, Loc. Misfit, nT, nS, Magnitude	17%
J	Row H + Azimuth Bins	18%
K	Row H + Radius Bins	18%

Table 5.2: Comparison of ANN performance for different ElarmS configurations. A unique ANN (H, L, M) is trained for each ElarmS configuration (requiring at least two stations, three stations, or four stations to send an alert). The input variables are the same for each ANN: latitude, longitude, location misfit, number of triggered instruments, and number of triggered stations. Magnitude is not provided to the ANN. Row H is the same configuration as row H in Table 5.1.

		Before ANN		After ANN	
	ElarmS Configuration	# alerts	% false	# alerts	% false
H	2+ stations	512	21%	503	17%
L	3+ stations	470	15%	434	8%
M	4+ stations	340	4%	340	4%

Table 5.3: Comparison of Northern and Southern California false alert rates, before and after application of an ANN, for different ElarmS configurations (generating alerts with two or more stations, three or more stations, and four or more stations).

	% False Alerts, before ANN			% False Alerts, after ANN		
	Alert Requirement			Alert Requirement		
	2+ stations	3+ stations	4+ stations	2+ stations	3+ stations	4+ stations
North (lat >= 36)	5%	3%	3%	5%	3%	3%
South (lat < 36)	29%	23%	5%	24%	12%	5%
All California	21%	15%	4%	17%	8%	4%

Table 5.4: Average time saved (in seconds) compared to four-station alert time, and average magnitude error, for alerts sent with varying numbers of triggered seismic stations.

	2 stations	3 stations	4 stations	5 stations	6 stations
Average Time Saved	2.2 s	1.8 s	0	-1.8 s	-2.7 s
Average Magnitude Error	0.2 +/- 0.4	0.2 +/- 0.4	-0.1 +/- 0.2	-0.2 +/- 0.2	-0.2 +/- 0.1

Chapter 6

Conclusions

In 2007 ElarmS was a research prototype, consisting of solid foundations but still running offline in outdated code. During the subsequent five years we transformed ElarmS into a production grade system.

In 2007 and 2008 we tested it with large earthquakes in Japan to forecast its future performance for large events in California, and found that the magnitude scaling, location and ground motion prediction algorithms remain valid for large events. We thoroughly analyzed the uncertainties in all aspects of event estimation and developed a statistical error model to describe changing uncertainties during processing.

The Berkeley Seismological Laboratory EEW group then adapted ElarmS to run continuously, analyzing hundreds of seismic waveforms second by second from instruments throughout the state. A single installation of ElarmS at UC Berkeley coordinated all the incoming data and identified earthquakes in progress, sending immediate email notification to the development group. We analyzed latencies in the system and continually tested and improved system performance.

In 2010 and 2011 we overhauled the code. We transferred the high performance magnitude and location algorithms directly into C++, and significantly renovated the association and alert modules. The resulting E2 code is fast, efficient and reliable, running continuously in realtime and sending alert messages to CISN ShakeAlert.

Finally, in 2012 we further improved performance and speed by applying an Artificial Neural Network (ANN) at the end of ElarmS processing, to detect and block false alerts prior to transmission to the ShakeAlert Decision Module. The ANN study allowed us to investigate and improve uncertainties with varying numbers of seismic stations.

Now, at the end of 2012, what was a research prototype has become a robust, streamlined product. ElarmS processes simultaneous seismic waveforms from hundreds of seismic instruments throughout California, identifies earthquakes in progress, estimates the event magnitude, location and damage potential, and sends immediate alerts to the CISN ShakeAlert system. ShakeAlert, in turn, sends alert messages to government, industrial, and academic test users throughout the state. As ShakeAlert is further tested and trusted, the warnings will reach more Californians, providing precious seconds of warning before disaster strikes, and potentially saving countless lives in future earthquakes.

References

Allen, R. M. (2004). Rapid magnitude determination for earthquake early warning, in *The Many Facets of Seismic Risk*, G. Manfredi et al. (Editors), 15– 24, Univ. degli Studi di Napoli ‘‘Federico II’’, Naples, Italy.

Allen, R. M. (2006). Probabilistic warning times for earthquake ground shaking in the San Francisco Bay Area, *Seismol. Res. Lett*. **77**(3), 371–376.

Allen, R. M. (2007). The ElarmS earthquake early warning methodology and application across California, in *Earthquake Early Warning*, P. Gasparini, et al. (Editors)., 21– 44, Springer Ital., Milan, Italy.

Allen, R.M., P. Gasparini, O. Kamigaichi, M. Böse (2009a). The Status of Earthquake Early Warning around the World: An Introductory Overview, *Seismo. Res. Lett.* **80**(5) 682-693.

Allen, R. M., and H. Kanamori (2003). The potential for earthquake early warning in southern California, *Science* **300**, 786– 789.

Allen, R.M., H. Brown, M. Hellweg, O. Khainovski, P. Lombard, and D. Neuhauser (2009b). Real-time earthquake detection and hazard assessment by ElarmS across California, *Geophys. Res. Lett.* **36**, L00B08.

Bishop, C. M. (1995). Neural Networks for Pattern Recognition, Oxford University Press.

Boatwright, J., H. Bundock, J. Luetgert, L. Seekins, L. Gee, and P. Lombard (2003). The Dependence of PGA and PGV on Distance and Magnitude Inferred from Northern California ShakeMap Data, *Bull. Seismol. Soc. Am*. **93**(5), 2043-2055.

Boore, D. M., W. B. Joyner, and T. E. Fumal (1997). Equations for estimating horizontal response spectra and peak accelerations from western North American earthquakes: A summary of recent work, *Seismol. Res. Lett.* **68**(1), 128– 153.

Böse, M., E. Hauksson, K. Solanki, H. Kanamori, and T. H. Heaton (2009). Real-time testing of the on-site warning algorithm in southern California and its performance during the July 29 2008 M_w5.4 Chino Hills earthquake, *Geophys. Res. Lett.* **36**.

Böse, M., R. Allen, H. Brown, G. Cua, M. Fischer, E. Hauksson, T. Heaton, M. Hellweg, M. Liukis, D. Neuhauser, P. Maechling, and the CISN EEW Group (2012). CISN ShakeAlert – An earthquake early warning demonstration system for California, in *‘Early Warning for Geological Disasters - Scientific concepts and current practice’*, F. Wenzel and J. Zschau (Editors), Springer, in press.

Brown, H.M., R. M. Allen, and V. F. Grasso (2009). Testing ElarmS in Japan, *Seismol. Res. Lett.* **80**(5), 727-739.

Brown, H.M., R.M. Allen, M. Hellweg, O. Khainovski, D. Neuhauser, and A. Souf (2011). Development of the ElarmS methodology for earthquake early warning: Realtime application in California and offline testing in Japan, *Soil Dyn. and Earthquake Engin.* **31** 188-200.

Brown, H.M, and R.M. Allen (2012). Faster Earthquake Alerts with Assistance from an Artificial Neural Network, *Bull. Seis. Soc. Amer.*, submitted October 2012.

Cua, G., M. Fischer, T. Heaton, and S. Wiemer (2009). Real-time performance of the Virtual Seismologist earthquake early warning algorithm in southern California, *Seismol. Res. Lett.* **80**(5), 740-747.

Demuth, H., M. Beale, and M. Hagan (2009). Neural Network Toolbox TM 6: User's Guide. MATLAB: The MathWorks, Inc., Natick, Massachusetts.

Derras, B., P. Bard, F. Cotton, and A. Bekkouche (2012). Adapting the Neural Network Approach to PGA Prediction: An Example Based on the KiK-net Data, *Bull. Seis. Soc. Amer.* **102**(4) 1446-1461.

Hill, D. P., J. P. Eaton, and L. M. Jones (1990). Seismicity, 1980– 86, in *The San Andreas Fault System, California*, R. E. Wallace (Editor), U.S. Geol. Surv. Prof. Pap., 1515, 115–152.

Iervolino, I., M. Giorgio, C. Galasso, and G. Manfredi (2009). Uncertainty in early warning predictions of engineering ground motion parameters: What really matters?, *Geophys. Res. Lett.* **36**, L00B06.

Kuyuk, H.S., H. Brown, R.M. Allen, I. Henson, M. Hellweg, D Neuhauser (in prep). ElarmS-2: Designing a network-based earthquake early warning system for California.

Lockman, A. B., and R. M. Allen (2005). Single-station earthquake characterization for early warning, *Bull. Seismol. Soc. Am.* **95**(6), 2029 – 2039.

Olson, E. L., and R. M. Allen (2005). The deterministic nature of earthquake rupture, *Nature* **438**(7065), 212– 215.

Tsang, L., R.M. Allen and G. Wurman (2007). Magnitude scaling relations from P-waves in southern California, *Geophys. Res. Lett*. **34** L19304.

Wald, D. J., V. Quitoriano, T. H. Heaton, H. Kanamori, C. W. Scrivner, and B. C. Worden (1999a). TriNet ShakeMaps: Rapid generation of instrumental ground motion and intensity maps for earthquakes in southern California, *Earthquake Spectra* **15**(3), 537–556.

Wald, D. J., B. C. Worden, V. Quitoriano, and K. L. Pankow (2005). *ShakeMap1 Manual: Technical manual, users guide, and software guide*, Tech. Methods 12–A1, U.S. Geol. Surv., Reston, Va.

Wurman, G., R. M. Allen, and P. Lombard (2007). Toward earthquake early warning in northern California, *J. Geophys. Res.* **112**, B08311.

Youngs, R.R., S.-J. Chiou, W.J. Silva, and J.R. Humphrey (1997). Strong ground-motion relationships for subduction zones, *Seismol. Res. Letters* **68**, No.1, 58-73

Zollo, A., M. Lancieri, S. Nielsen (2006). Earthquake magnitude estimation from peak amplitudes of very early seismic signals on strong motion records, *Geophys. Res. Lett.* **33**, L23312