

Toward earthquake early warning in northern California

Gilead Wurman,¹ Richard M. Allen,¹ and Peter Lombard¹

Received 31 October 2006; revised 5 April 2007; accepted 30 April 2007; published 22 August 2007.

[1] Earthquake early warning systems are an approach to earthquake hazard mitigation which takes advantage of the rapid availability of earthquake information to quantify the hazard associated with an earthquake and issue a prediction of impending ground motion prior to its arrival in populated or otherwise sensitive areas. One such method, Earthquake Alarm Systems (ElarmS) has been under development in southern California and, more recently, in northern California. Event magnitude is estimated using the peak amplitude and the maximum predominant period of the initial P wave. ElarmS incorporates ground motion prediction equations and algorithms from ShakeMap for prediction of ground motions in advance of the S wave arrival. The first peak ground motion estimates are available 1 s after the first P wave trigger, and are updated each second thereafter for the duration of the event. The ElarmS methodology has been calibrated using 43 events ranging in size from $M_{\rm L}$ 3.0 to $M_{\rm w}$ 7.1 that occurred in northern California since 2001. We present the results of this calibration, as well as the first implementation of ElarmS in an automated, noninteractive setting and the results of 8 months of noninteractive operation in northern California. Between February and September 2006, ElarmS successfully processed 75 events between M_d 2.86 to M_w 5.0. We find that the ElarmS methodology processed these events reliably and accurately in the noninteractive setting. The median warning time afforded by this method is 49 s at the major population centers of the Bay Area. For these events the magnitude estimate is within an average of 0.5 units of the network-derived magnitude, and the ground motion prediction from ElarmS is within an average of 0.1 units of the observed modified Mercalli intensity.

Citation: Wurman, G., R. M. Allen, and P. Lombard (2007), Toward earthquake early warning in northern California, *J. Geophys. Res.*, *112*, B08311, doi:10.1029/2006JB004830.

1. Introduction

[2] Earthquake early warning (EEW) systems are combinations of instrumentation, methodology, and software designed to analyze rapidly an ongoing earthquake and issue real-time information about the hazard to persons and property before the onset of strong ground motions in populated areas. Japan, Mexico, and Turkey currently operate EEW systems, while Taiwan, Italy, Romania and Greece are testing EEW systems [Allen, 2006, and references therein]. Japan's EEW system, which has been providing warnings to a limited group of users, is anticipated to begin widespread public dissemination of warnings in the summer of 2007. The system operating in Mexico is a frontal detection system, which relies on the fact that the largest potential earthquake epicenters are 300 km from Mexico City in the Middle America Trench, such that an array of seismometers between the fault and the city can reliably detect and measure the intensity of an earthquake's S waves and issue a warning well before those waves arrive at the city [*Espinosa Aranda et al.*, 1995].

[3] In California, the proximity of major faults to population centers limits the utility of frontal detection systems for EEW. Under the conditions found in California a useful EEW system must be able to rapidly and reliably estimate the location, origin time and size of an earthquake based on the P wave alone. The system must then generate predictions of ground motion at multiple locations of interest and disseminate these predictions in the time between the P wave arrival and the S wave arrival. Such systems are being developed in Taiwan [Wu and Kanamori, 2005] and in Japan [Odaka et al., 2003] which rely on measurement of the amplitude of the P wave as a proxy for the magnitude of the earthquake. Such systems are effective for small- and moderate-size events but are susceptible to saturation in large events. Ground accelerations near the source of large earthquakes saturate at approximately $10-15 \text{ m/s}^2$, due in part to ground response becoming nonlinear under large stresses.

[4] The Earthquake Alarm Systems (ElarmS) methodology [*Allen and Kanamori*, 2003] has been tested using data from southern California, Taiwan, Japan and the Pacific Northwest of the United States [*Olson and Allen*, 2005; *Lockman and Allen*, 2007], and uses the maximum pre-

¹Department of Earth and Planetary Sciences, University of California, Berkeley, California, USA.

Copyright 2007 by the American Geophysical Union. 0148-0227/07/2006JB004830\$09.00



Figure 1. Map of California showing distribution of events used in the calibration process (white circles) and stations in the NCSN (gray) and BDSN (black) networks. Triangles signify high-gain, broadband velocity sensors. Inverted triangles signify low-gain strong motion accelerometers, and diamonds signify a station with collocated velocity sensor and accelerometer.

dominant period (τ_p^{max}) of the first 1 to 4 s of the P wave as an estimate of earthquake magnitude. The ElarmS methodology has been shown to be effective in these areas for M 3 and larger earthquakes [Allen and Kanamori, 2003; Lockman and Allen, 2005; Olson and Allen, 2005; Allen, 2006; Lockman and Allen, 2007; Allen, 2007]. In the process of testing ElarmS in northern California, we find that using both τ_p^{max} and the peak amplitude of the P wave improves the accuracy of the ElarmS magnitude estimate. We have been testing the effectiveness of the combined methodology since February 2006 and find that the system estimates the magnitude of earthquakes in northern California rapidly, accurately, and reliably.

[5] In addition to incorporating P wave peak amplitude in the magnitude determination, we have incorporated the attenuation relationships (hereafter referred to as ground motion prediction equations, GMPEs) and algorithms of ShakeMap [Wald et al., 2005] into the part of the methodology which predicts ground motions during an event. The GMPEs used by ShakeMap [Newmark and Hall, 1982; Boore et al., 1997; Wald et al., 1999a; Boatwright et al., 2003] replace the empirical attenuation relationships developed for southern California [Allen, 2004, 2007]. ShakeMap algorithms [Wald et al., 1999a] incorporate individual station corrections to observations as well as scaling of predicted ground motions based on local geology [Borcherdt, 1994; Wills et al., 2000] throughout northern California. In addition to making ElarmS ground motion predictions directly comparable to other products like ShakeMap itself,

we find the incorporation of these algorithms allows us to generate accurate and timely predictions of ground motion at seismic stations.

2. ElarmS Methodology

[6] Implementing earthquake early warning in northern California presents opportunities not seen in other places for improving the robustness of the ElarmS methodology across different networks. A functional EEW system in northern California must integrate data from both high-gain, broadband velocity instruments and from low-gain, strong motion accelerometer stations. The system must collect this data over the two networks currently operating in northern California: the Northern California Seismic Network (NCSN) operated by the US Geological Survey, and the Berkeley Digital Seismic Network (BDSN) operated by the University of California Berkeley Seismological Laboratory. Within each network, high-gain velocity instruments (channels HHE, HHN, and HHZ, which we address as HH henceforth) are more useful for measuring the small (M \leq 4.5) events on which we rely for calibration and routine validation of the method. However, these stations will clip quickly in the event of a nearby major earthquake. Lowgain, strong motion accelerometers (channels HNE, HNN, and HNZ; or HLE, HLN, and HLZ, which we address as HN and HL, respectively) will remain on-scale longer in the event of a nearby major earthquake, but are of limited use in measuring small events due to low signal-to-noise ratios. Between networks, differences in instrumentation may lead to different behavior within the same channel type (i.e., velocity or accelerometer). All of these differing behaviors must be accounted for by an EEW system which seeks to maximize the amount of usable data in a minimum amount of time.

[7] We use the Earthquake Alarm Systems (ElarmS) methodology developed by Allen and Kanamori [2003], with some modifications for the particular problems of northern California. The ElarmS methodology is built of two systems: a single waveform processing system extracts parameters of interest from a single channel of data, and sends these parameters to an event processing system. The latter integrates output from waveform processing of multiple channels into information about an event's size, time and location, and in fact whether there is an event at all. Given an event's size and location, ground motion predictions are issued for specific locations on a second-by-second basis during the event. Within both of these systems we encounter the need to modify the original ElarmS methodology to account for the specific challenges of northern California data. These will be discussed at length later.

2.1. Calibration Data Set

[8] Prior to applying ElarmS in a real-time setting, we tested the method on 43 calibration events ranging in size from M_L 3.0 to M_w 7.1 which occurred in northern California since 2001. The calibration events are shown in Figure 1. We were restricted from using older events such as the 1989 Loma Prieta and 1992 Petrolia earthquakes in the calibration, because prior to 2001 the NCSN and BDSN networks did not have sufficient station coverage or the appropriate instrument types to measure these events. The

calibration events were used to establish the maximum predominant period versus magnitude and peak amplitude versus magnitude relations described below.

2.2. Triggering and Event Location

[9] The first step in the early warning process is to detect an event. This begins with the waveform processing system, which must detect the initial P wave of an event and issue a trigger at the onset of that P wave. We use a short-term/ long-term average method following Allen [1978]. The algorithm is applied to the vertical velocity trace with timescales of 0.5 s for the short-term average and 5 s for the long term, and a triggering threshold of 20. Triggering can be accomplished using any real-time algorithm, but cannot be done with any method which requires data after the trigger itself, as such data is by definition unavailable at the time of the trigger. Consequently, methods such as autoregressive pickers [Sleeman and van Eck, 1999] and pickers based on wavelet transforms [Zhang et al., 2003], while more precise than a simple short-term/long-term average method are not practical for this application. This also means generally that any filter applied to the data must be causal.

[10] When the first station triggers, the event processing system will provisionally locate the event beneath that station. When a second station triggers the provisional location moves to a point directly between the two stations, based on the timing of the arrivals. Once trigger times are produced at three or more locations, the event location and origin time is estimated using trilateration and a grid search algorithm to find the optimal solution. Although a depth can be estimated using more stations or more sophisticated algorithms, this is unnecessary for the geologic setting of northern California, where most events nucleate at less than 20 km depth [*Hill et al.*, 1990]. We currently fix the depth of the event to be 8 km.

[11] On the basis of the estimated event location and time, warning times can be calculated for any geographical locations of interest. These warning times are based on a moveout speed of 3.75 km/s, which is determined from observations of the onset times of significant ground motions in southern California. Again, although more sophisticated methods exist for the estimation of time until significant shaking, when one considers the computational requirements for greater sophistication against the need for rapid processing and notification, this simple moveout speed seems sufficient for the purpose of estimating the warning time.

2.3. Magnitude From Predominant Period

[12] The ElarmS methodology rests largely on the use of the maximum predominant period (τ_p^{max}) within the first 4 s of the P wave as an indicator of the size of the event [*Allen and Kanamori*, 2003; *Olson and Allen*, 2005]. The predominant period, τ_p of a single vertical channel (HHZ, HLZ or HNZ) is calculated in real time using the iterative relation

$$\tau_{p,i} = 2\pi \sqrt{\frac{X_i}{D_i}} \tag{1}$$

where $X_i = \alpha X_{i-1} + x_i^2$ and $D_i = \alpha D_{i-1} + (dx/dt)_i^2$. The constant α is a smoothing constant, and x_i is the ground velocity of the last sample. Because both velocity sensors

and accelerometers are used, the accelerometer traces must be integrated to velocity before τ_p can be calculated. In addition, a causal 3 Hz low-pass Butterworth filter is applied iteratively to the velocity data [*Allen and Kanamori*, 2003]. This calculation is done by the waveform processing system, and the maximum value of τ_p within the first 4 s of the P wave arrival is recorded and sent to the event processing system, which uses it to estimate magnitude according to a predetermined relationship.

[13] Our initial attempts to establish a relationship between magnitude and τ_p^{max} were frustrated by noise in the low-magnitude data (M < 4.5). This problem led us to adopt two significant additions to the ElarmS methodology. The first of these is a criterion for the disqualification of S wave data. Part of the low-magnitude scatter was due to many small events being located close to our stations in the San Francisco Bay Area. As a result, the S wave arrival occurs within 4 s of the P wave arrival, and since S waves generally have longer periods than the associated P waves, it is the S wave τ_p which gets recorded as τ_p^{max} . A simple criterion based on an S-minus-P moveout of 1 s per 8 km eliminates these false signals and cleans up the data somewhat, though we do apply a minimum S-minus-P time of 1 s, based on the assumption that the event is 8 km deep.

[14] In addition to this S wave criterion, we chose to incorporate a second criterion for the exclusion of data, based on the signal-to-noise ratio (SNR) of each waveform. As this was a particular problem for low-gain accelerometers (HL and HN channels), we chose to treat each channel type separately. The absolute noise level is calculated as a very long-term average from interevent data and is frozen when a trigger is detected. From the time of the trigger until the event is over, the signal level is calculated using a 0.05-s short-term average [Allen, 1978], and the ratio of these two is the SNR. In principle the higher we require the SNR to be, the better our results. However, we must consider the need for fast measurements as well as good ones, and the greater SNR we require, the fewer measurements of the first second of the P wave will be admitted. By weighing the reduction in scatter of small magnitude $\tau_p^{\rm max}$ against the number of excluded τ_p measurements in the first second of data, we arrive at the optimal minimum SNR: 100 for HH channels and 200 for HL and HN channels.

[15] The results of calibrating τ_p^{max} versus magnitude are shown in Figure 2. Note that low-gain accelerometer data still shows a significant scatter in spite of the two added criteria. We are investigating the root cause of this scatter, but presently the HN channels (Figure 2c) exhibit the largest scatter, and we have provisionally removed them from the τ_p^{max} determination until this can be resolved. The best fit relationship between τ_p^{max} and magnitude is

$$M = 5.22 + 6.66 \cdot \log_{10} \left(\tau_p^{\max} \right) \tag{2}$$

using only the HH and HL channels. This relationship is plotted in Figure 2. In determining magnitude from τ_p^{max} for any new event we only use data from HH and HL channels to be consistent with the calibration of this relationship.

2.4. Magnitude From P Wave Peak Amplitude

[16] While we have succeeded in reducing the scatter in measurement of τ_p^{max} at low magnitudes, the scatter is still



Figure 2. Plots of τ_p^{max} versus magnitude for all calibration events. Each point represents a single station measurement. Measurements are separated by channel code: (a) HH represents velocity sensors, while (b) HL and (c) HN represent accelerometers. The dotted line in all three plots is the same line of best fit using HH and HL data simultaneously.

sufficient to present us a problem in discriminating between small nonhazardous events and large hazardous events. Because of this scatter, there is a potential to misidentify a small event as a large one, leading to a false alarm. This is of critical importance in many early warning applications, as a high incidence of false alarms will drastically reduce the credibility and utility of the warnings. This is especially true in applications where the cost of false alarms is high, such as industrial process interruption. In order to further improve this discrimination, we have added a second, independent estimate for rapid magnitude determination. Using a method similar to that of Wu et al. [2006], we calculate the peak amplitude of the P wave, scaled by the logarithm of the epicentral distance. As with τ_p^{max} , we glean the peak amplitude from the first 4 s of the vertical record. Wu et al. used the peak displacement, P_d of the P wave, but we chose to analyze displacement, velocity and acceleration for each channel type independently. In theory, the displacement record has longer periods than the acceleration or velocity records, and will be less susceptible to random high-frequency excursions. For velocity instruments (HH channels) this holds true, and measuring the peak amplitude in displacement yields the lowest error. However, for accelerometer channels (HN and HL), the act of numerically integrating twice (from acceleration to velocity and then again to displacement) introduces errors to the point where using the velocity record rather than displacement yields a better magnitude estimate.

[17] We also investigated the merit of using between 1 and 5 s of P wave data for determining P_d or P_v (peak velocity, for HN and HL channels). Using less than 4 s yielded greater errors, and between using 4 and 5 s there was little difference in performance (4 s performed slightly better for HH, slightly worse for HL and HN channels). We chose to use 4 s for the sake of internal consistency with our τ_p^{max} measurements, which also use 4 s of P wave data.

[18] The results of calibrating P_d and P_v (which we henceforth abbreviate $P_{d/v}$) versus magnitude are shown in Figure 3. The amplitudes are plotted as a function of magnitude, after being scaled to an epicentral distance of 10 km using the best fit relations in equations (3), (4) and (5) below. These plots do not show nearly the scatter at low magnitudes that the τ_p^{max} versus magnitude plot does in Figure 2. However, the $P_{d/v}$ of the largest (M_w 7.1) event is significantly lower than predicted. This is due to the fact that this event incorporates data from more distant stations than is normally allowed, as will be discussed in detail later. Because of this effect, the $P_{d/v}$ measurements for the M_w 7.1 event were not used in the best fit lines plotted in Figure 3.

[19] Although the variability of the HN data seen in τ_p^{max} is visible to a lesser degree in P_v , (Figure 3c) the data are not unusable. However, we chose to fit HL and HN data separately to minimize the error of measurements on the HL channels. The best fit relationships between magnitude and $P_{d/v}$ are

$$M = 1.04 \cdot \log_{10} (P_d) + 1.27 \cdot \log_{10} (R) + 5.16 \text{ (HH channels)}$$
(3)

$$M = 1.37 \cdot \log_{10} (P_v) + 1.57 \cdot \log_{10} (R) + 4.25 \text{ (HL channels)}$$
(4)

$$M = 1.63 \cdot \log_{10} (P_{\nu}) + 1.65 \cdot \log_{10} (R) + 4.40$$
 (HN channels) (5)

where *R* is the distance in kilometers from the station to the epicenter. When the source of scatter in the HN data is found and controlled for, it may be beneficial to unify the relationships for HL and HN channels. Since the HH channels use P_d rather than P_v the relationship for HH channels must remain separate from the other two.

2.5. Data Integration and Magnitude Determination

[20] The waveform processing system sends any new information available to the event processing system every



Figure 3. Plots of $P_{d/v}$ versus magnitude for all calibration events. Each point represents a single station measurement, corrected to a distance of 10 km using the empirical best fit equation. Measurements are separated by channel code: (a) HH, (b) HL, and (c) HN. The dotted line in each plot is the line of best fit using data from the corresponding channel only.

tenth of a second for 4 s after a trigger. This includes the maximum value of τ_p^{max} or $P_{d/\nu}$ only if that maximum has changed since the last tenth of a second. The τ_p^{max} data is accompanied by the value of the SNR at the time of the measurement. This low data volume has the advantage of being easily transmissible over existing station telemetry, so that the waveform processing system can potentially be implemented at each station independent of the rest of the network. The advantage of this approach, in turn, is that the waveform processing happens much sooner and much more reliably, as there is no delay for telemetry of data over the network, and no risk of data dropout leading to errors in processing. Instead, the large volume of data being produced by the sensors is reduced on site to a few parameters of interest which can be cheaply transmitted over the network.

[21] The event processing system gathers the transmitted data from the waveform processing systems at each station within 100 km of the estimated epicenter. This is the distance within which frequency-dependent attenuation (Q) has a minimal effect on the predominant period measurement [Allen and Kanamori, 2003]. For the two largest calibration events, the Mw 6.5 San Simeon earthquake and the Mw 7.1 earthquake in the Gorda Plate, this cutoff distance is increased to 150 km and 200 km, respectively, due to the lack of stations within 100 km of these events. We justify this in particular for the Gorda Plate event by asserting that the intervening crust between the event and the stations is mostly oceanic, and has higher Q than continental crust [Vera et al., 1990]. The system integrates the data from the stations once per second to determine a magnitude estimate for the event as it progresses. Each time a new maximum τ_p^{max} or $P_{d/v}$ value is reported, the event processing system checks it for validity by examining whether the S wave may have arrived at that station, as described earlier in this section. It also checks that the SNR at the time of a τ_p^{max} measurement exceeds the minimum required value. If any of these checks fail, the event processing system ignores that measurement and proceeds as if it were never reported.

[22] The event processing system makes one more check, in which it looks for an indication that a given channel has clipped. This indication is actually given by the waveform processing system in the form of a negative SNR beginning when the channel's output first exceeds a particular threshold, and extending for a fixed duration after the last sample which exceeds that threshold. This duration represents the time required for the channel to recover from the clipping event and become usable. The clipping threshold and recovery time vary from channel to channel, and are encoded in the waveform processing system at each station, so the event processing system does not know anything about the value of the data, only whether it has clipped. If the event processing system receives a clipping indication, it immediately stops updating information from that station for the duration of the event. The τ_p^{max} value at the time of clipping is recorded as the final τ_p^{max} for that station, and the $P_{d/v}$ value for the station is stricken. The reason for treating the two estimates differently is that we often find that the time at which τ_p^{max} is taken is not the same time as the peak amplitude of the P wave, so the τ_p^{max} value before the clipping occurred is still potentially valid. In contrast, the fact that the sensor has clipped means a priori that the previous $P_{d/v}$ value has been exceeded, and is therefore invalid. In this respect τ_p^{max} is more robust, as it can tolerate clipping of a channel and still represent a valid estimate.

[23] If the data passes all the checks, the quality of the data can be reasonably assured, and the event processing system uses the updated information to produce a magnitude estimate for the event. It takes the log10 average of τ_p^{\max} from each available channel, and calculates a magnitude from the average value. The results of magnitude estimation for the calibration events using τ_p^{\max} alone are shown as gray triangles in Figure 4. Note the significant scatter in the magnitude estimate below $M \approx 4.5$, consistent with the calibration of τ_p^{\max} versus magnitude from Figure 2.

[24] The event processing system also takes the average value of $P_{d/v}$ from each station and calculates a magnitude



Figure 4. Plot of ElarmS magnitude estimate versus network-derived magnitude (usually M_L or M_w) for the calibration events. Gray triangles are magnitude estimates using τ_p^{max} only, gray squares are estimates using $P_{d/v}$ only, and black circles are a linear average of the two magnitudes for each event.

from that average value. The results for the calibration events are shown as gray squares in Figure 4. Note the comparatively low magnitude assigned to the largest event in the calibration data set, consistent with Figure 3. As discussed earlier, this is the result of incorporating data from more distant stations for this event. However, it highlights a potential limitation of the $P_{d/v}$ estimate. The $P_{d/v}$ estimate is susceptible to saturation near the fault for very large events. This is because at fault-normal distances less than the length of the rupture the distance to the farthest point of the rupture is significantly greater than the distance to the nearest point. As a result, the effective distance between the station and the rupture (i.e., the average distance between the station and all points on the rupture) is greater than the actual faultnormal distance, leading to lower P wave amplitude than predicted for a given epicentral distance.

[25] The value of τ_p^{max} does not appear to be susceptible to this effect [*Olson and Allen*, 2005], but is much more susceptible to noise pollution at lower magnitudes than peak amplitude measurements. Thus the two estimates of τ_p^{max} and $P_{d/v}$ are particularly complementary, with the strengths of one compensating for the weaknesses of the other, and using some combination of the two magnitude estimates from τ_p^{max} and $P_{d/v}$ produces a more robust single estimate. Currently, the two estimates are combined in a linear average, the results of which are shown as black circles in Figure 4. Note the superior performance at both ends of the magnitude scale as a result of this combined approach. The large events are not underestimated, and the scatter in the small events has been greatly reduced. [26] A more sophisticated scheme may be conceivable for the combination of the τ_p^{\max} magnitude with the $P_{d/\nu}$ magnitude. In particular, since we are interested in the low-magnitude performance of $P_{d/\nu}$ and the high-magnitude performance of τ_p^{\max} , it makes sense to consider a progressive weighting scheme in which the latter is more heavily weighted at low magnitudes and the former is more heavily weighted at high magnitudes. We investigated a scheme by which the weighting changes linearly with the magnitude of the event, but found that the data does not bear out the use of such a scheme. At this time the simple linear average appears to be as good as any weighted average, so we use the linear average.

2.6. Ground Motion Prediction

[27] The final step in an early warning system is to predict the severity of imminent ground motions from an ongoing earthquake and to issue warnings based on those predictions. We do not address the question of when and how to issue warnings. For a discussion of this aspect of EEW, see V. F. Grasso and R. M. Allen (Uncertainty in real-time earthquake hazard predictions, submitted to Bulletin of the Seismological Society of America, 2007, hereinafter referred to as Grasso and Allen, submitted manuscript, 2007). For this study, we observe that the 1σ error in magnitude estimate reduces to a reasonable level (0.5 magnitude units) when 4 s of data are available from 4 channels. We define this criterion of 4 s of data in four channels as the "alarm time" for the purposes of performance evaluation in the next section. However, we arrive at this definition somewhat arbitrarily, and different users would require different levels of uncertainty or timeliness, depending on their tolerance for false or missed alarms (Grasso and Allen, submitted manuscript, 2007).

[28] ElarmS is capable of generating ground motion predictions through the incorporation of algorithms from ShakeMap [Wald et al., 2005]. These algorithms, which have been developed for and tested extensively in California, incorporate empirically derived GMPEs [Newmark and Hall, 1982; Boore et al., 1997; Wald et al., 1999a; Boatwright et al., 2003], as well as geological amplification correction and corrections for site conditions at seismic stations [Borcherdt, 1994; Wills et al., 2000]. The ground motion predictions are initially calculated using only the estimated magnitude and location of the event, as no observations of peak ground motion are yet available. We use the GMPE for the given magnitude to compute the predicted ground motion on a regular grid of points with a spacing of 0.1° around the source. The prediction at each point is then corrected for local geological effects. The result is a coarsely spaced grid of points with predictions of peak ground motion based solely on the magnitude and location of the event. This grid can be interpolated to create predictions at finer resolution, and to generate predictions for discrete locations of interest, such as urban centers or seismic stations.

[29] As the event progresses and the S wavefield expands outward from the source, peak ground motion observations become available at each station. The observations are first corrected for the site condition at the station, and then the GMPE curve, based on magnitude and location, is linearly scaled up or down to best fit the corrected observations. The



Figure 5. Map of California showing distribution of events (white circles) processed noninteractively by ElarmS from February through September 2006. Stations in the NCSN (gray) and BDSN (black) networks are plotted as triangles for high-gain, broadband velocity sensors. Inverted triangles signify low-gain strong motion accelerometers, and diamonds signify a station with collocated velocity sensor and accelerometer.

resulting equation is used to generate ground motion predictions on a regular grid as before, with the addition of grid points representing the individual station observations available at the time. This irregular grid is interpolated to produce a finer, regular grid of peak ground motion incorporating magnitude, location and station observations. This grid is predictive at all points ahead of the S wavefront. The process is similar to that used to produce ShakeMaps after an earthquake, but here it is done once per second. Initially, there is very little information to incorporate and the ground motion predictions are correspondingly rough, but with each second that passes the information becomes more complete and the ground motion predictions are refined in real time. ElarmS produces predictions of PGA and PGV at all points, which are combined to produce a prediction of modified Mercalli intensity (MMI) using the relationship of Wald et al. [1999b].

2.7. Simulating ElarmS

[30] It is not practical to implement the ElarmS methodology online as outlined in the first part of this section without first testing it offline to ensure its functionality. This is because a full implementation requires the investment of time and money to emplace the waveform processing system at each station in the network. Therefore we test the performance of the methodology offline using a program that simulates the causality of information after the event has completed. While this simulation may differ somewhat from the final implementation of ElarmS, the behavior of the methodology will not change appreciably from the results of the simulation. Henceforth, when referring to "ElarmS" we refer to the simulation unless otherwise stated.

3. ElarmS Performance

[31] Since February 2006, we have been operating ElarmS automatically following every event of M 3.0 or larger in northern California. This processing is initiated 10 min after notification of a new event, in order to allow the requisite data to be collected at the network data center for retrieval. The processing is performed automatically with no human input or oversight. We have been using the results of this automatic processing to make improvements to the ElarmS methodology, and consequently it is necessary on occasion to reprocess these events after the fact, when a significant change is made in the methodology. This reprocessing is prompted by a human operator, but without any added input from the operator. The process is identical to the automatic processing and uses the same data which was gathered 10 min after each event. We call this "noninteractive" processing, and we use it to indicate how a real-time implementation of ElarmS might perform.

3.1. Performance of Noninteractive Processing

[32] Between February and September 2006, a total of 85 instances of noninteractive processing were initiated. Of these, one is a duplicate event, a result of the email notification system posting an update to an existing event. One instance was a false event. This was not the result of a false detection by ElarmS, but of an erroneous email notification.

[33] The geographic distribution of the remaining 83 events is shown in Figure 5. Of these, one event was offshore Mendocino, with no stations within 100 km of the source. This is the cutoff distance for usable stations in the ElarmS methodology, so the event produced no output. Seven events suffered errors as a result of maintenance of the operating system. Five of these occurred consecutively, due to the extraordinary misfortune of an update to the operating system coinciding with a temporal (not spatial) cluster of small events (all $M_L \leq 3.6$). The system maintenance prevented the acquisition of data 10 min after the earthquake. Acquiring data at a later time invalidates the noninteractive procedure, so these events are not considered in this analysis.

[34] The remaining 75 events range in magnitude from $M_d 2.86$ to $M_w 5.0$. The results of magnitude estimation for these 75 events are presented in Figure 6. Figure 6 shows the magnitude errors (with respect to network-based magnitudes, usually M_w or M_L) produced by ElarmS at three different times for each event. The initial magnitude error (Figure 6a) refers to the magnitude estimation based on only the first second of P wave data at the first station or stations to detect the event. This is the earliest possible magnitude determination, which can be used to give the maximum warning time. The initial magnitude has a significant scatter ($\sigma = 0.72$ magnitude units) due to its reliance often on a single station's data.



Figure 6. Histograms showing errors in magnitude estimate for all noninteractive events (a) at 1 s after detection (75 events), (b) at alarm time when 4 s of data are available in four channels (66 events), and (c) 60 s after the event began for those events which achieved the alarm condition (66 events).

[35] Figure 6b shows the errors at alarm time, which we define as in the previous section to be the time at which at least 4 s of P wave data are available from at least four different channels. The magnitude error at this time is considerably less than in the first second ($\sigma = 0.54$ magnitude units). There are fewer events represented in this plot (66 events versus 75 in Figure 6a), because not all of the events are ever detected in enough channels to meet the alarm criteria. This is primarily due to the weak signal from small (M \approx 3) events, and in some cases results from a lack of enough stations within 100 km of the epicenter.

[36] Figure 6c shows the error in the final magnitude determination for events that met the alarm criteria, using all available data from stations within 100 km of the source. The scatter has decreased slightly ($\sigma \approx 0.48$ magnitude units) due to the incorporation of more station information. In all three of these plots, the magnitude estimate is biased slightly downward (mean of -0.57 magnitude units in the initial estimate, -0.13 at alarm time and -0.02 in the final magnitude estimate). This is due to events beyond the physical edge of the network, which can be mislocated by

tens of kilometers due to poor azimuthal coverage. This does not affect τ_p^{\max} -based magnitude estimates, but $P_{d/v}$ -based magnitude estimates are strongly affected by epicentral distance errors. Location errors can also cause the system to set the S wave arrival time earlier than the true arrival time because it considers some events to be closer to the station than it is. This causes the system to discard valid data when measuring both τ_p^{\max} and $P_{d/v}$ because it considers the signal contaminated by the S wave. This biases the estimates downward because it prevents both τ_p^{\max} and $P_{d/v}$ values from being revised at later times, and these revisions are always upward.

[37] Figure 7 summarizes the error in ground motion prediction at seismic stations in the NCSN and BDSN networks for all events. The logarithm of the observed ground motions is subtracted from the logarithm of the predicted ground motions for PGA (Figures 7a and 7b) and PGV (Figures 7c and 7d). An error of 1 signifies overprediction by a factor of 10. Figures 7a and 7c show the errors in the first second of data, and Figures 7b and 7d show the errors at alarm time. At alarm time the 1σ error in PGA and PGV is approximately a factor of 4 (0.6 log units). For the MMI errors (Figures 7e and 7f) no logarithm is necessary as MMI incorporates logarithms of PGA and PGV [Wald et al., 1999b]. The number of predictions in the center bin (less than ± 0.17 MMI unit error) is off scale in Figures 7e and 7f. There are 435 observations in the center bin in the first second (Figure 7e) and 398 at alarm time (Figure 7f). The scatter in ground motion prediction is significantly reduced by waiting for the alarm condition to be met ($\sigma = 0.42$ MMI units in the first second, versus $\sigma =$ 0.08 at alarm time). Figure 7e has a positive bias (0.12 MMI units). This is because the ShakeMap MMI scale ranges from 1 to 10, and for many of the smaller events MMI cannot be significantly underpredicted simply because the actual MMI is only 1 or 2. There is also a slight negative bias (-0.01 MMI units) in Figure 7f reflecting the same effects as seen in Figure 6. The errors reported in Figures 6 and 7 are valid only for events in the validation data set. We cannot evaluate the error for larger earthquakes, as none occurred in the time of the study.

[38] The times between event origin and detection and the achievement of the alarm condition for all the events are summarized in Figure 8. Initial detection occurs an average of 8.0 ± 4.8 s after event origin (Figure 8a). The alarm condition is reached an average of 14.9 ± 4.6 s after origin (Figure 8b). These results are presented in Figure 9 in terms of time until largest ground shaking at three major Bay Area metropolitan centers: San Francisco, San Jose and Oakland. The median warning time in these cities at initial detection is 56 s in San Francisco (Figure 9a) or Oakland (Figure 9e), and 48 s in San Jose (Figure 9c). If we wait for the alarm condition to be reached, the median warning times reduce to 39.5 s in San Francisco (Figure 9b), 40 s in San Jose (Figure 9d), and 39 s in Oakland (Figure 9f). This analysis does not show the warning time for any future earthquakes, but the distribution of event locations in Figure 5 does coarsely reflect the potential locations of future large earthquakes. For a more detailed analysis of warning time for potential damaging earthquake scenarios, see Allen [2006].

[39] The two event parameters which have not been discussed are the epicenter estimates and the origin time



Figure 7. Histograms showing ground motion prediction errors at all stations for all noninteractive events. Only stations which had not already observed peak motions at the initial and alarm times are included in the data. (a) Errors in log(PGA) at 1 s after detection, and (b) at alarm time when 4 s of data are available in four channels; (c) errors in log(PGV) at 1 s after detection and (d) at alarm time; (e) errors in MMI at 1 s after detection and (f) at alarm time. In Figures 7e and 7f, the 0-centered bin is off scale. The value of this bin is 435 in Figure 7e and 398 in Figure 7f.

estimates. When only one station has triggered, ElarmS assumes the event origin time and epicenter correspond to the time and location of that first trigger. When two stations have triggered, the epicenter is located at a point between the first two triggered stations, based on timing. Consequently location and origin time errors are significant when fewer than three stations are used. However, when three or more stations are used we find the error is small, and by the time the alarm condition is met both of these estimates have insignificant error. This characterization does not hold as well for events beyond the edge of the network. At alarm time, the mean absolute error in epicenter location is $13.7 \pm$ 23.4 km and the mean absolute error in origin time is 2.3 \pm 3.1 s, including events beyond the edge of the network. For both of these measures, the mean error is within a standard deviation of zero.

3.2. Two Bay Area Scenario Events

[40] Among the 75 events processed noninteractively by ElarmS, two moderate events represent likely hazardous earthquake scenarios for the Bay Area (Figure 10), and thus provide some insight into what can be expected of ElarmS. For these two events we use M_L as a reference, even though

 M_w values exist for both. This is because M_L is sensitive to the same frequencies ($\sim 1-2$ Hz) as ElarmS, and because M_L is more directly related to the severity of the event in terms of damage to persons and property.

[41] The first event is a M_L 4.7 event near Gilroy, California, on 15 June 2006. This event is located near the southern Calaveras fault, in a geographic location where a Calaveras or Southern Hayward fault rupture might nucleate (Figure 10). Figure 11a shows the magnitude estimate for this event as a function of time in relation to the arrival time of significant shaking at San Francisco, Oakland and San Jose (vertical lines). The time at which the alarm condition was reached is also plotted, and the dashed horizontal line represents M = 4.7, the actual local magnitude of the event.

[42] Figure 11b shows the error in predicted PGA versus time, over all stations which have not yet reported peak ground motion at the given time. The solid line is the mean error at each time, and the dashed lines are $\pm 1\sigma$ error margins. Figure 11c shows the error in predicted PGV versus time for the event in a similar fashion. Both of these factor into the predicted MMI for the event [*Wald et al.*, 1999b].



Figure 8. Histograms of time (a) between event origin and initial detection and(b) between event origin and the alarm time at which 4 s of data are available in four channels for all noninteractive events.

[43] Figure 11d shows the error in predicted MMI versus time in a manner similar to Figures 11b and 11c. After 13 s the only stations which have not reported peak ground motion are those far from the fault, which experience an intensity of 1. The predicted MMI at all these stations is 1 (the ShakeMap algorithm does not produce MMI less than 1), and this is reflected in Figure 11d by the fact that the mean and $\pm 1\sigma$ lines converge to zero after 13 s.

[44] Figure 12 shows the predicted peak ground shaking from ElarmS (which we call the "AlertMap") for 7 s following event detection for comparison with the Shake-Map for this event (Figure 12h). The time since event origin and magnitude estimate are given above each AlertMap. Stations are plotted as small white symbols in the same fashion as the maps in Figures 1 and 5: triangles for velocity sensors, inverted triangles for accelerometers, and diamonds for a collocated installation. Stations which have triggered are plotted as larger gray symbols, and stations which are experiencing peak ground motions are plotted in black. When the peak observations become available, the stations are colored according to their observed peak MMI, following the scale at the bottom of the plot. The circular contours radiating from the epicenter represent the estimated time until the onset of the largest ground motions at all locations, based on the current epicenter location and a moveout speed of 3.75 km/s. The color field in Figures 12b–12g represents estimated peak MMI following the scale at the bottom of the plot, which is the same as for the ShakeMap in Figure 12h. At all points outside the 0-s warning time contour, this peak MMI estimate is predictive.

[45] The first AlertMap in Figure 12a, 3 s after event origin, represents the initial detection time for this earthquake. There is no magnitude estimate yet, since ElarmS requires a full second of P wave data before making the initial estimate. The larger gray station is the first station to trigger for this event, and the red star represents the epicenter, currently located at the station as described above. The initial magnitude estimate of M 5.0 is available 1 s later, 4 s after the origin (Figure 12b). With the magnitude estimate ElarmS begins predicting ground motions based only on the GMPE. The epicenter has also been relocated at this time due to a second station triggering. As described above, the epicenter is now located directly between the two triggered stations based on the trigger times. At 5 s after origin (Figure 12c), a third station triggers and from this point forward the location is fit to the trigger times using a grid search. At this time, data is being collected from 4 channels (two channels at the station southeast of the epicenter, one at each of the other two stations), so 4 s later (i.e., 9 s after origin) the alarm condition will be reached. At 6 s (Figure 12d) two more stations have triggered, but by now the epicenter location is good and does not move noticeably.

[46] At 7 s after origin (Figure 12e), the magnitude estimate has dropped to 4.7, and the first peak ground motion observation is available. At this time, the GMPE curve is biased to pass through that single observation, which is why the predicted MMI field for the event changes so drastically between Figures 12d and 12e. The station southeast of the epicenter is plotted in black to indicate that it is currently experiencing peak ground motion, consistent with its location within the 0-s warning time contour. At 8 s (Figure 12f) two more stations have triggered and another has entered the peak ground motion window. At 9 s (Figure 12g) the alarm condition is reached since the fourth channel triggered at 5 s, and the magnitude estimate is 4.3, only 0.4 magnitude units below the actual M_L . At this time, 3 additional stations have triggered, bringing the total to 13 triggered channels at 10 different stations. Another station has entered the peak ground motion window, and the station southeast of the epicenter has reported its peak ground motion observation. When there are multiple observations, ElarmS biases the GMPE curve to best fit the available observations. In this case, this results in a slight increase in the predicted ground motion over Figure 12f.

[47] The reason for the low intensities far from the event in the final AlertMap (Figure 12g) versus the ShakeMap (Figure 12h), is that the ShakeMap incorporates peak ground motion observations from stations on the Peninsula and in the East Bay. At 9 s after the origin, the S wavefront has not arrived at many of these stations, so that information is not used to bias the GMPE curve in the AlertMap. However, in Figure 11d it is apparent that the ElarmS ground motion predictions at 9 s are accurate to within a standard deviation of 0.3 MMI units of the actual observed



Figure 9. Histograms of warning time until onset of largest ground motions at San Francisco, San Jose, and Oakland for all noninteractive events. Onset times are estimated using a moveout speed of 3.75 km/s. Warning times are calculated (a, c, and e) from 1 s after initial detection and (b, d, and f) from alarm time, defined as having 4 s of data in four channels.

intensities. The vertical lines in Figure 11 represent the arrival of the largest ground motions at the three major urban centers in the Bay Area. San Jose experienced peak ground shaking only 12 s after event origin, meaning San Jose would have had about 3 s warning time in this event, not considering telemetry and dissemination delays. However, Oakland and San Francisco would have had 20 and 22 s of warning, respectively, for this event. These warning times depend primarily on the disposition of stations around the epicenter, so they would be comparable for a magnitude 7 event. In this case, the distance between the epicenter and the major cities is comparable to that of the M_w 6.9 Loma Prieta earthquake. Thus, in the case of a Loma Prieta repeat, we would expect comparable warning times in the major cities.

[48] The second scenario event is a M_L 4.7 event near Santa Rosa on 2 August (local time) 2006. This event is located near the Rodgers Creek fault, near possible epicenter locations for a southward rupturing Rodgers Creek/Hayward fault event (Figure 10). Figures 13 and 14 show the history of this event in the same manner as for the Gilroy event.

[49] Initial detection of this event occurs 3 s after event origin, as shown in the AlertMap in Figure 14a. The epicenter at this time is collocated with the only triggered station. At 4 s (Figure 14b), two other stations have triggered, so the epicenter is located using a grid search method. In addition, one of the stations has both an accelerometer channel and a velocity channel, bringing the triggered channel count to 4. Thus the alarm condition will be reached in 4 s, at 8 s after origin. The initial magnitude estimate at this time is 5.8, over one magnitude higher than the actual size of the event. The high-magnitude estimate in turn causes the ground motion predictions to be high, as these are produced using the GMPE alone, with no station observations to bias the curve. This is reflected both in the AlertMap, which shows significantly higher MMI than the ShakeMap (Figure 14h), and in Figure 14b, which shows that the first MMI predictions exceed actual observations by as much as 2 MMI units.

[50] The comparatively large magnitude error highlights the utility of waiting for more data to become available rather than issuing the alarm immediately, based on information from a single station only. In this case, the large error is due to the first triggered station being a strong motion accelerometer. Most of the stations to the north of the Bay Area are strong motion accelerometers, which are susceptible to noise pollution below $M \approx 5$. For large earthquakes this is not a problem, but in smaller events high-gain broadband velocity sensors yield superior data.



Figure 10. Map of the San Francisco Bay Area showing the location and focal mechanisms of two hazardous scenario earthquakes processed noninteractively by ElarmS. The black symbols are stations in the NCSN and BDSN networks. Stations are plotted as triangles for high-gain, broadband velocity sensors, inverted triangles for low-gain strong motion accelerometers, and diamonds for stations with collocated velocity sensor and accelerometer. Focal mechanisms are from the BDSN regional moment tensor catalog [*Pasyanos et al.*, 1996].



Figure 11. Plot of ElarmS output for the first 40 s of the Gilroy event. Vertical lines represent (as labeled) the time of alarm condition, or the onset of largest ground motions at San Francisco, San Jose or Oakland based on a moveout of 3.75 km/s. (a) Magnitude estimate. The dotted horizontal line represents the network-based M_L of 4.7. (b) Error in the logarithm of predicted peak ground acceleration. At each time interval, the plot incorporates all stations which have not yet observed peak ground motions. Dotted lines represent the 1σ error envelope. (c) Error in the logarithm of predicted peak ground velocity. (d) Error in the predicted modified Mercalli intensity.

[51] At 5 s (Figure 14c) after origin the magnitude drops to 4.3 due to the incorporation of 1 s of data from the three channels which triggered in the previous second. The magnitude estimate is now 0.4 units lower than the actual M_L for this event, but the error is nearly a third of that in the previous second, again suggesting that it is better to wait one second for multiple stations to provide data rather than relying on a single-station estimate. At 6 s after origin (Figure 14d) the first peak ground motion observations become available, and the GMPE curve is biased to minimize the errors at these stations. Two more stations to the southwest have triggered at this time, leading to a small revision in the epicenter location. At 7 s after the origin (Figure 14e) the magnitude estimate drops slightly to 4.2. Two more stations have triggered, but the location does not change noticeably after 6 s from origin. The station directly south of the epicenter now reports an additional peak ground motion observation, further informing the ground motion predictions at this time.

[52] At 8 s after origin (Figure 14f) the alarm condition is reached. Two more stations have triggered in this second, bringing the total to 10 channels at 9 stations. The magnitude estimate is 4.2, half a magnitude lower than the actual M_L of 4.7. However, the peak ground motion predictions are biased up by the available station observations of peak ground motion, and match both the ShakeMap (Figure 14h) and the final observed peak ground motions at the stations outside the 0-s warning contour. Figure 14b shows that the MMI predictions are accurate to within a standard deviation of 0.3 MMI units at alarm time. For consistency with Figure 12 the AlertMap at 9 s, 1 s after the alarm time, is shown in Figure 14g. There is no noticeable change from Figure 14f, other than one additional station having triggered.

[53] When comparing the ElarmS AlertMap for this event (Figure 14f) with the ShakeMap (Figure 14h) the performance appears better at alarm time than for the Gilroy event, even in light of the low magnitude estimate. The two maps are almost identical in terms of peak intensities, though the AlertMap does underpredict the intensity near the epicenter as a result of the low-magnitude estimate. At alarm time, both San Francisco and Oakland have 11 s until the arrival of the largest ground shaking. The magnitude estimate is low, and will only rise to 4.6 at 13 s after the origin (Figure 11b) leaving 6 s of warning for San Francisco and Oakland, but this is not an issue when considering ground motion predictions, which are accurate at this time. San Jose experiences its largest ground motions 37 s after origin, so even with the additional 5 s delay for the magnitude estimate to rise, it still has 24 s of warning in this instance.

4. Improving ElarmS Performance

[54] The performance of ElarmS in the noninteractive processing arena is promising. At alarm time the 1σ magnitude error for these events is 0.5 magnitude units, which is consistent with the results using the calibration events, and also consistent with previous work with the ElarmS methodology in southern California [*Allen*, 2007]. This indicates that we now have a good understanding of how ElarmS behaves in a real setting, at least for earth-quakes less than $M \approx 5$, and suggests that we can expect the same behavior in the future in northern California and in other locations as well.

[55] Although the results are largely favorable, there are a few things to consider for full online implementation. For one thing, the warning times obtained in the noninteractive



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	-	IV	V	VI	VII	VIII	IX	X+

Figure 12. (a-g) ElarmS AlertMap output for 3-9 s after the origin of the Gilroy event. Time since event origin and the magnitude estimate are shown above each AlertMap. The epicenter is plotted as a red star. Stations which have triggered are shown in gray, stations which are experiencing peak ground motions are black, and those that have reported peak ground motion are color coded according to the scale at bottom. Stations are plotted as triangles for high-gain broadband velocity sensors, inverted triangles for low-gain strong motion accelerometers, and diamonds for stations with collocated velocity sensor and accelerometer. The circular contours represent time until onset of strong ground motion based on the location and origin time of the event and a moveout of 3.75 km/s. The color field is the ElarmS prediction of modified Mercalli intensity, according to the scale at bottom. (h) ShakeMap for the Gilroy event. The color field represents actual instrumentally observed modified Mercalli intensity for the event, processed after the event occurred.



Figure 12. (continued)

processing are maximum warning times, as they do not take into account telemetry, processing or dissemination delays. The processing time delay should not exceed 1 s if we are to run ElarmS with continuous updates every second. In terms of station telemetry, the actual transmission delays are negligible (R. Uhrhammer, personal communication, 2006), but currently, data are telemetered in up to 10-s-long packets for stations in the BDSN and NCSN networks (D, Neuhauser, personal communication, 2006). This packetization can be reduced to 1 s or less, but the overhead associated with transmitting each packet gets proportionally larger as the packet itself gets smaller. Finally, these warning times do not account for delays in disseminating the data and taking action at the user end. We cannot quantify these delays as no dissemination system exists at this time. However, the results reported here show that a dissemination delay of less than 1 s is ideal. When such a system is designed, the minimization of dissemination delays must be a primary design goal.

[56] One way to significantly improve warning times is to improve the disposition of seismometers around northern California. At present the broadband velocity and strong motion accelerometer stations in the NCSN and BDSN networks are distributed somewhat unevenly, and not



Figure 13. Plot of ElarmS output for the first 40 s of the Santa Rosa event. Vertical lines represent (as labeled) the time of alarm condition or the onset of largest ground motions at San Francisco, San Jose, or Oakland based on a moveout of 3.75 km/s. (a) Magnitude estimate. The dotted horizontal line represents the network-based M_L of 4.7. (b) Error in the logarithm of predicted peak ground acceleration. At each time interval, the plot incorporates all stations which have not yet observed peak ground motions. Dotted lines represent the 1σ error envelope. (c) Error in the logarithm of predicted peak ground velocity. (d) Error in the predicted modified Mercalli intensity.

always in optimal locations for observing an earthquake near its epicenter. In particular, most of the stations are located in and around the Bay Area, with comparatively few stations along the northern coast of California. The mitigating circumstance here is that most of the population of northern California lives in the Bay Area, and the preponderance of faults capable of major earthquakes is in and around the Bay Area, where most of the instruments are located. However, the comparative lack of instruments along the northern portions of the San Andreas Fault mean first that any people living north of the Bay Area would receive limited benefits from ElarmS or a similar EEW system. The second consequence of this station distribution is that an event nucleating on the northern portions of the San Andreas Fault and propagating southward would take a long time to achieve the alarm condition due to the dearth of stations in the vicinity of the epicenter. This means drastically reduced warning times for the Bay Area and likely a larger error margin due to the comparatively few measurements that would be incorporated in the event estimates. Given that a southward propagating rupture on the northern San Andreas Fault may cause as much as \$90 to \$120 billion in losses in the Bay Area [Kircher et al., 2006] it is worthwhile to instrument the northern reaches of the San Andreas more thoroughly.

[57] Even within the Bay Area there is some room for improvement. There are few broadband or strong motion stations around the southern segment of the Hayward Fault, which is considered the most hazardous fault in the region *[Working Group on California Earthquake Probabilities*, 2003]. Given this fault's proximity to the major urban centers in the Bay Area, as little as 1 or 2 s additional warning could make a significant difference. A few wellplaced seismometers along the Hayward Fault would go a long way toward attaining those extra 1 or 2 s.

5. Conclusions

[58] The ElarmS methodology incorporates two independent measurements to estimate the magnitude of an event: the peak amplitude of the P wave and its maximum predominant period, both within 4 s of the onset of the P wave. These measurements are controlled both for sufficient signal-to-noise ratio and to prevent pollution by the S wave arrival. The magnitude estimate produced by ElarmS is a linear average of the estimates determined from peak amplitude and maximum predominant period. The epicenter of an event is located using a grid search algorithm based on the trigger times at three or more stations. The epicenter and magnitude estimates are used to generate predicted ground motions at all points around the epicenter, using algorithms similar to ShakeMap. The predictions are updated each second based on updated epicenter and magnitude information. As observations of peak ground motion become available at stations near the epicenter, the prediction is corrected to conform to these observations, further refining the prediction.

[59] The ElarmS methodology has been applied in an offline simulation to every event greater than M 3 in northern California since February 2006. The methodology has been applied automatically and without human interaction 10 min after each event, to simulate how ElarmS performs without human assistance, as it would in a real-time application. Eight months of noninteractive operation of the ElarmS simulator have shown that the ElarmS methodology can reliably deliver accurate earthquake information within a few to a few tens of seconds of event origin.



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	1	-	IV	V	VI	VII	VIII	IX	X+

Figure 14. (a-g) ElarmS AlertMap output for 3-9 s after the origin of the Santa Rosa event. Time since event origin and the magnitude estimate are shown above each AlertMap. The epicenter is plotted as a red star. Stations which have triggered are shown in gray, stations which are experiencing peak ground motions are black, and those that have reported peak ground motion are color coded according to the scale at bottom. Stations are plotted as triangles for high-gain broadband velocity sensors, inverted triangles for low-gain strong motion accelerometers, and diamonds for stations with collocated velocity sensor and accelerometer. The circular contours represent time until onset of strong ground motion based on the location and origin time of the event and a moveout of 3.75 km/s. The color field is the ElarmS prediction of modified Mercalli intensity, according to the scale at bottom. (h) ShakeMap for the Santa Rosa event. The color field represents actual instrumentally observed modified Mercalli intensity for the event, processed after the event occurred.



Figure 14. (continued)

75 events were successfully processed noninteractively in that time frame. We define an alarm time of 4 s of data at four stations, and at this alarm time the 1σ magnitude errors for the noninteractive processing are half a magnitude unit. This is as expected given past performance of the methodology on calibration events and in southern California. At alarm time the ground motion predictions have a 1σ error within approximately a factor of 4 in PGA and PGV, or within 0.1 MMI of the actual observed MMI at seismic stations in the BDSN and NCSN networks. These errors are valid for the 75 events in the noninteractive data set, and it is difficult to estimate the errors for events larger than those

in that data set. The alarm time occurs an average of 15 s after event origin. Of the 75 events processed since February 2006, 66 achieve the alarm condition. For these 66 events, ElarmS provides a median warning time of 49 s in major Bay Area metropolitan centers.

[60] Two events since February 2006 represent hazardous scenario earthquakes for the Bay Area. In both cases, the magnitude estimate at alarm time is within 0.5 of the network-determined M_L of 4.7, and ground motion predictions are within 0.3 MMI units of the actual peak ground motions observed for the events. Warning times in the Bay Area achieved by the system for the two scenario earth-

quakes range from about 3 to 30 s, depending on the location of the epicenter relative to the city in question.

[61] A functional earthquake early warning system in northern California has the potential to save both lives and money in the event of a major earthquake. On the basis of the results of simulating the operation of ElarmS, we find the methodology in a condition in which we can move forward to a real-time, online implementation of ElarmS in northern California in the near future.

[62] Acknowledgments. We wish to thank Doug Neuhauser and Bob Uhrhammer for discussions relating to station equipment and networks and Dave Wald for discussions about ShakeMap. This paper is Berkeley Seismological Laboratory contribution 07-05. This work was funded by USGS/NEHRP grants 05HQGR0074 and 06HQAG0147.

References

- Allen, R. M. (2004), Rapid magnitude determination for earthquake early warning, in *The Many Facets of Seismic Risk*, edited by G. Manfredi et al., pp. 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*, edited by P. Gasparini et al., pp. 21–44, Springer Ital., Milan, Italy.
- Allen, R. M., and H. Kanamori (2003), The potential for earthquake early warning in southern California, *Science*, *300*, 786–789.
- Allen, R. V. (1978), Automatic earthquake recognition and timing from single traces, *Bull. Seismol. Soc. Am.*, 68(5), 1521–1532.
 Boatwright, J., H. Bundock, J. Luetgert, L. Seekins, L. Gee, and P. Lombard
- 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.
- Borcherdt, R. D. (1994), Estimates of site-dependent response spectra for design (methodology and justification), *Earthquake Spectra*, 10(4), 617– 654.
- Espinosa Aranda, J. M., A. Jimenez, G. Ibarrola, F. Alcantar, A. Aguilar, M. Inostroza, and S. Maldonado (1995), Mexico City seismic alert system, *Seism. Res. Lett.*, 66(6), 42–53.
- Hill, D. P., J. P. Eaton, and L. M. Jones (1990), Seismicity, 1980–86, in *The San Andreas Fault System, California*, edited by R. E. Wallace, U.S. Geol. Surv. Prof. Pap., 1515, 115–152.
- Kircher, C. A., H. A. Seligson, J. Bouabid, and G. C. Morrow (2006), When the big one strikes again: Estimated losses due to a repeat of the 1906 San Francisco earthquake, *Earthquake Spectra*, 22(S2), S297– S339, doi:10.1193/1.2187067.

- Lockman, A. B., and R. M. Allen (2005), Single-station earthquake characterization for early warning, *Bull. Seismol. Soc. Am.*, 95(6), 2029– 2039.
- Lockman, A. B., and R. M. Allen (2007), Magnitude-period scaling relations for Japan and the Pacific Northwest: Implications for earthquake early warning, *Bull. Seismol. Soc. Am.*, 97(1B), 140–150.
- Newmark, N. M., and W. J. Hall (1982), Earthquake spectra and design, *Geotechnique*, 25(2), 139–160.
- Odaka, T., K. Ashiya, S. Tsukada, S. Sato, K. Ohtake, and D. Nozaka (2003), A new method of quickly estimating epicentral distance and magnitude from a single seismic record, *Bull. Seismol. Soc. Am.*, 93(1), 526–532.
- Olson, E. L., and R. M. Allen (2005), The deterministic nature of earthquake rupture, *Nature*, 438(7065), 212–215, doi:10.1038/nature04214.
- Pasyanos, M. E., D. S. Dreger, and B. Romanowicz (1996), Toward realtime estimation of regional moment tensors, *Bull. Seismol. Soc. Am.*, 86(5), 1255–1269.
- Sleeman, R., and T. van Eck (1999), Robust automatic P-phase picking: an on-line implementation in the analysis of broadband seismogram recordings, *Phys. Earth Planet. Inter.*, 113(1–4), 265–275.
- Vera, E. E., J. C. Mutter, P. Buhl, J. A. Orcutt, A. J. Harding, M. E. Kappus, R. S. Detrick, and T. M. Brocher (1990), The structure of 0- to 0.2-m.y.old oceanic crust at 9°N on the East Pacific Rise from expanded spread profiles, J. Geophys. Res., 95, 15,529–15,556.
- 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., V. Quitoriano, T. H. Heaton, and H. Kanamori (1999b), Relationship between peak ground acceleration, peak ground velocity, and modified Mercalli intensity for earthquakes in California, *Earthquake Spectra*, 15(3), 557–564.
- Wald, D. J., B. C. Worden, V. Quitoriano, and K. L. Pankow (2005), ShakeMap[®] Manual: Technical manual, users guide, and software guide, *Tech. Methods* 12–A1, U.S. Geol. Surv., Reston, Va.
- Wills, C. J., M. D. Petersen, W. A. Bryant, M. S. Reichle, G. J. Saucedo, S. S. Tan, G. C. Taylor, and J. A. Treiman (2000), A site-conditions map for California based on geology and shear wave velocity, *Bull. Seismol. Soc. Am.*, 90(6B), S187–S208.
- Working Group on California Earthquake Probabilities (2003), Earthquake Probabilities in the San Francisco Bay Region: 2002–2031, U.S. Geol. Surv. Open File Rep., 03–214.
- Wu, Y.-M., and H. Kanamori (2005), Experiment on an onsite early warning method for the Taiwan early warning system, *Bull. Seismol. Soc. Am.*, 95(1), 347–353.
- Wu, Y.-M., H.-Y. Yen, L. Zhao, B.-S. Huang, and W.-T. Liang (2006), Magnitude determination using initial P waves: A single-station approach, *Geophys. Res. Lett.*, 33, L05306, doi:10.1029/2005GL025395.
- Zhang, H. J., C. Thurber, and C. Rowe (2003), Automatic P-wave arrival detection and picking with multiscale wavelet analysis for singlecomponent recordings, *Bull. Seismol. Soc. Am.*, 93(5), 1904–1912.

R. M. Allen, P. Lombard, and G. Wurman, Department of Earth and Planetary Sciences, University of California Berkeley, 215 McCone Hall, Berkeley, CA 94720, USA. (gwurman@seismo.berkeley.edu)