

Optimizing Earthquake Early Warning Performance: ElarmS-3

by Angela I. Chung, Ivan Henson, and Richard M. Allen

ABSTRACT

The University of California Berkeley's (UCB) Earthquake Alert Systems (ElarmS) is a network-based earthquake early warning (EEW) algorithm that was one of the original algorithms developed for the U.S. west-coast-wide ShakeAlert EEW system. Here, we describe the latest update to the algorithm, ElarmS v.3.0 (ElarmS-3 or E3). A new teleseismic filter has been developed for E3 that analyzes the frequency content of incoming signals to better differentiate between teleseismic and local earthquakes. A series of trigger filters, including amplitude-based checks and a horizontal-to-vertical ratio check, have also been added to E3 to improve the quality of triggers that are used to create events. Because of its excellent performance, E3 is now the basis for EPIC, the only ShakeAlert point-source algorithm going forward. We can therefore also use the performance of E3 described here to assess the likely performance of ShakeAlert in the coming public rollout. We should expect false events with magnitudes between M 5 and 6 less than once per year. False events with $M \ge 6$ will be even less frequent, with none having been observed in testing. We do not expect to miss any $M \ge 6$ onshore earthquakes, though the system may miss some large offshore events and may miss one onshore earthquake between M 5 and 6 per year. Finally, in the metropolitan regions where the station density is on the order of 10 km, we expect users 20, 30, and 40 km from an earthquake epicenter to get 3, 6, and 9 s warning, respectively, before the S-wave shaking begins.

Electronic Supplement: Screenshot of the Earthquake Alert Systems (ElarmS) review tool, and example histograms and tables of algorithm performance created by the review tool.

INTRODUCTION

The principle behind earthquake early warning (EEW) is to detect earthquakes soon after they occur and then warn the nearby population that they are about to experience strong shaking. When an earthquake begins, seismic energy is radiated away from the hypocenter. The fastest of these seismic waves, P waves, travel at roughly 5–6 m/s and quickly reach any nearby seismic stations. The data can either be processed on site or at a

central processing center where the location, magnitude, and origin time of the event can be characterized. With the knowledge of the source, an alert can be disseminated to end users. Alerts may go directly to the public via warning sirens, TV and radio broadcasts, cell phone applications, or dedicated alerting devices that provide users with a warning to take precautionary action. Alerts may also be ingested by systems that have been designed to automatically perform certain actions upon receiving the alerts. Examples include slowing or stopping trains to prevent derailment, shutting off gas or water mains, and opening fire station doors to allow the engines to exit the building in the event of power loss following significant shaking.

There currently exist EEW systems in various stages of development and production in earthquake-prone countries around the world (Allen *et al.*, 2009) including Mexico (Espinosa-Aranda *et al.*, 2011; Cuéllar *et al.*, 2017; Suárez *et al.*, 2018), Switzerland (Cua *et al.*, 2009), Italy (Zollo *et al.*, 2014; Picozzi *et al.*, 2015), Romania (Böse *et al.*, 2007), Turkey (Wenzel *et al.*, 2014), Taiwan (Hsiao *et al.*, 2009), Japan (Nakamura, 1988; Hoshiba and Ozaki, 2014; Liu and Yamada, 2014; Kodera *et al.*, 2016), Israel (Nof and Allen, 2016), South Korea (Sheen *et al.*, 2014), and Chile (Lancieri *et al.*, 2011). Both the demand and the usefulness of these systems have been well established through numerous studies (Kamigaichi *et al.*, 2009; Gasparini *et al.*, 2011; Strauss and Allen, 2016).

In 2006, the California Integrated Seismic Network began developing the ShakeAlert EEW system in collaboration with participants from California Institute of Technology, University of California Berkeley, and the U.S. Geological Survey (Brown et al., 2011). University of Oregon and University of Washington joined the ShakeAlert project in 2011. Originally, three EEW algorithms contributed alerts to the ShakeAlert system: Earthquake Alert Systems (ElarmS), a network-based algorithm (Allen and Kanamori, 2003; Allen, 2007; Wurman et al., 2007); OnSite, a single-station algorithm (Böse et al., 2009, 2012); and Virtual Seismologist (VS), a Bayesian EEW algorithm (Cua and Heaton, 2007; Cua et al., 2009). As over time, it became apparent that alerts from VS were slower than those from ElarmS and OnSite, and its source estimates did not offer significant improvements over the other two algorithms (Table 1). VS was

Table 1 Performance Statistics for the 146 M ≥ 3 California Earthquakes Alerted on by All Three ShakeAlert Algorithms (Earthquake Alert Systems [ElarmS] v.2.0 [E2], OnSite [ON], and Virtual Seismologist [VS]) from 1 January 2014 through 1 April 2016					
Statistic	E2	ON	VS		
Median alert time (s)	6.7	8.2	13.2		
Standard alert time (s)	4.1	4.1	4.7		
Median location error (km)	2.3	10.2	2.3		
Standard location error (km)	16.7	12.5	6.8		
Median magnitude error	0.3	0.3	0.2		
Standard magnitude error	0.2	0.3	0.2		
Total number of $\mathbf{M} \ge 3$ events alerted during time period	410	181	277		
Missed ANSS M ≥ 3 events	213	442	346		
False $\mathbf{M} \ge 3$ alerts	33	21	58		
Statistics from the best performing algorithm in each category are bolded. ANSS, Advanced National Seismic System.					

retired from ShakeAlert in 2016 when the production prototype version of ShakeAlert officially came online (Kohler *et al.*, 2018).

Each of the two remaining ShakeAlert algorithms is unique and offers advantages and disadvantages. ElarmS is a networkbased EEW system that uses a minimum of four stations to alert, but only requires a small amount of data from those stations (0.2 s from four stations). OnSite uses a single-station approach and can detect earthquakes using just one station and a minimum of 3 s data from that station. However, to reduce false alerts, a minimum of two stations are currently required to send out an alert (Böse et al., 2014). The ShakeAlert decision module (DM) receives event notifications from each of the EEW algorithms, which include the estimated earthquake location and magnitude, and creates alerts using a weighted average of the reporting algorithms (Böse et al., 2014). On 1 February 2016, ShakeAlert was deployed on production-ready machines to prepare for its future public release. ShakeAlert originally only provided alerts for earthquakes within California, but in April 2017 ShakeAlert coverage became west-coast-wide with the addition of data from Oregon and Washington (Chestler, 2017). During this time, a number of test users have been receiving the alerts and developing applications. A summary of these can be found in Strauss and Allen (2016).

In 2010–2011, ElarmS was completely rewritten and rebuilt to make it a streamlined production code and to introduce new algorithms to improve its performance. These updates resulted in the release of ElarmS v.2 (ElarmS-2 or E2) in 2012 (Kuyuk *et al.*, 2014). Some improvements to the code included a new waveform processing (Waveform Processor [WP]) module, modifications to how the triggers are associated, a filter to prevent alerts due to teleseismic events, and updated magnitude scaling relationships.

E2 performed well with these improvements and only minor improvements to the code were made over the next few years. E2 alerts were usually the fastest of the ShakeAlert algorithms, and it created alerts for significantly more earthquakes than either OnSite or VS (Table 1). Between 1 January 2014 and 1 April 2016, E2 created alerts for 410 $M \ge 3$ alerts. In contrast, OnSite and VS alerted for only 181 and 277 $M \ge 3$ alerts, respectively. During this time period, E2 missed 213 M \ge 3 earthquakes (the majority of which were offshore or in areas without dense station coverage), OnSite missed 442 earthquakes, and VS missed 346 earthquakes. In this study, a matched event is defined as an alert that is correctly created for an earthquake that matches an Advanced National Seismic System (ANSS) catalog earthquake within 100 km and 30 s. A missed event is when the system fails to create an alert for an earthquake. A false alert is when the system creates an alert, but there is no matching earthquake in the ANSS catalog within 100 km and 30 s. Of the 146 $M \ge 3$ earthquakes for which all three algorithms created alerts during this time, the median alert time (time between the earthquake origin and when the alert was sent out) was 6.7 s for E2 ($\sigma = 4.1$ s), 8.2 s for OnSite ($\sigma = 4.1$ s), and 13.2 s for VS ($\sigma = 4.7$ s). The median location error of these alerts from both E2 and VS was just 2.3 km ($\sigma = 16.7$ km for E2; $\sigma = 6.8$ km for VS) and 10.2 km for VS ($\sigma = 12.5$ km). The median magnitude error was comparable among the algorithms, 0.3 magnitude units for E2 and OnSite ($\sigma = 0.2$ and 0.3 units for E2 and OnSite, respectively), and 0.2 magnitude units for VS ($\sigma = 0.2$ units).

From the above statistics, it is clear that E2 is a fast and accurate algorithm and that it was able to create alerts for the majority of earthquakes that occurred within California during that time. E2 continued to use the short-term average/long-term average (STA/LTA) method of triggering from the original version of ElarmS (Wurman *et al.*, 2007). This simple triggering method is very sensitive, which is advantageous as E2 is able to trigger on small and/or out-of-network earthquakes. E2 also created alerts for more than twice as many $M \ge 3$ earthquakes as OnSite from 1 January 2014 through 1 April 2016. Because of its sensitivity, however, E2 is also susceptible to creating false events due to teleseismic signals, anomalous signals, and other problematic events.

A single-large (M 7.8) deep (664 km depth) earthquake near Japan on 30 May 2015 caused E2 to generate 10

Table 2 List of Earthquakes Used in Teleseismic Dataset						
UTC Date (yyyy/mm/dd)	UTC Time (hh:mm:ss)	Magnitude <i>M</i> w	Depth (km)	Latitude (°)	Longitude (°)	
2002/08/19	11:08:22	7.7	649.9	-23.868	178.454	
2008/07/05	02:12:06	7.7	646.1	53.946	152.863	
2010/07/23	22:08:11	7.3	610.2	6.711	123.488	
2010/07/23	23:15:10	7.5	633.7	6.740	123.327	
2011/09/15	19:31:03	7.3	629.0	-21.593	-179.324	
2013/05/24	05:44:50	8.3	607.0	54.815	153.391	
2015/05/30	11:23:02	7.8	664.0	27.839	149.493	
Events were obtained by searching the Incorporated Research Institutions for Seismology Data Management Center earthquake catalog for earthquakes with $M > 7.2$ and depth > 600.0 km during the time period 23 July 2002 through 1 June 2015.						

 $\mathbf{M} \ge 3$ false alerts. Of these, seven of the events had estimated magnitudes $\mathbf{M} \ge 5$. The seismic waves first triggered the stations in the northwest corner of California, and then continued to trigger stations as the waves moved quickly to the southeast. As the seismic waves from the distant event progressed throughout the state, E2 grouped triggers from nearby stations and created false events. To differentiate between local and teleseismic earthquakes E2 has a teleseismic discriminator that calculates the average τ_p^{\max} and P_d for the event (Kuyuk *et al.*, 2014). Though this empirical filter worked well for the test suite used when developing E2, it was not as successful with events as large and as deep as the aforementioned Japanese earthquake.

The sensitivity of the STA/LTA picker also means that E2 has the unintended ability to trigger on spurious signals. E2 associated spikes, boxcar functions, and other small-amplitude noisy signals with either other similarly noisy signals or with small local events to create false alerts. The worst example of this was when E2 created two false alerts (one with magnitude **M** 8.2 and the other with magnitude **M** 6.3) during a network calibration procedure.

In this report, we describe the latest version of ElarmS, ElarmS v.3.0 (ElarmS-3 or E3), the primary objective of which is to reduce false alerts without compromising the overall excellent performance of E2. We will also demonstrate the performance of E3 by comparing replays of E2 and E3.

ELARMS IMPROVEMENTS

Teleseismic Filter

The first major updates that were made to E2 were chosen by focusing on the situations that caused E2 to create the most false alerts. As mentioned previously, a single-large teleseism can cause E2 to generate multiple false alerts within a very short time span. This made filtering out teleseismic signals of critical importance.

To create the E3 revised teleseismic filter, we follow the method of the probabilistic filter bank EEW algorithm by Meier *et al.* (2015). In contrast to that study, which aims to quickly estimate the magnitude and station-to-event distance, we use the filter bank method to distinguish local

and teleseismic earthquakes due to the decreased high-frequency content of teleseismic signals. As the waves travel long distances through the Earth, the high-frequency energy is attenuated significantly more than the lower frequency energy. We begin by obtaining data for all large earthquakes with M >7.2 and depth > 600.0 km during the time period 23 July 2002 through 1 June 2015 (Table 2). We convert all data in the teleseismic and local datasets to velocity (if needed) and then trim the data in windows from 30 s before the trigger through a time t after the trigger. We explored a range of tvalues including 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0, and 2.0 s. We then filter all waveforms using a second-order bandpass Butterworth causal filter with nine filter ranges n. The ranges depend on the sampling rate and are shown in Table 3. The distinction between teleseismic and local earthquake signals is easily seen by comparing amplitudes of the filtered traces at different frequencies (Fig. 1).

We next go on to calculate the narrowband peak ground velocity (PGV_{nb}), which is the maximum velocity for each narrow passband-filtered trace during the time window beginning at the time of the trigger and ending at time t and combine the suite of values resulting from the filter bank processed traces to create a fingerprint of each signal.

We randomly select 50 waveforms from the teleseismic dataset and obtain the mean and standard deviation of the PGV_{nb} values from those waveforms for each of the 9 passbands and each of the 11 window lengths. We iteratively repeat this process 150 times and calculate the grand mean $(\bar{\mu})$ and the mean of the standard deviations of the iterations $(\bar{\sigma})$. For each passband filter range *n*, we create series of amplitude cutoffs using the following:

$$V_{\max}^{t,n} = \bar{\mu} + a \times \bar{\sigma} \tag{1}$$

$$V_{\min}^{t,n} = \bar{\mu} - a \times \bar{\sigma},\tag{2}$$

in which $V_{\text{max}}^{t,n}$ and $V_{\text{min}}^{t,n}$ (equations 1 and 2, respectively) are the maximum and minimum PGV_{nb} values for a teleseismic trace, for each time window length *t* and narrowband filter

	Table 3 Parameter <i>a</i> -Values Used for Each Time Window and Filter Bank Narrow Passband Filter Range Number <i>n</i>							
n	Narrow Passband Filter Range (HNZ/HHZ) (Hz)	Narrow Passband Filter Range (BHZ) (Hz)	a (HNZ)	a (HHZ)	a (BHZ)			
1	24–48	12–16	4	2.5	1.8			
2	12–24	8–12	4	2.2	1.8			
3	6–12	6–8	4	2.2	1.8			
4	3–6	3–6	4	2.0	1.8			
5	1.5–3	1.5–3	4	2.0	1.8			
6	0.75–1.5	0.75–1.5	4	2.0	1.8			
7	0.375–0.75	0.375–0.75	4	2.2	1.8			
8	0.1875–0.375	0.1875–0.375	4	2.2	1.8			
9	0.09375–0.1875	0.09375–0.1875	4	2.2	1.8			
<i>a</i> is ind	<i>a</i> is the coefficient in equations (1) and (2) and is multiplied by the standard deviation of the filtered waveforms. A larger <i>a</i> -value indicates a wider range of values that are classified as teleseismic and a smaller <i>a</i> -value indicates a narrower range.							

range n, respectively. a is obtained empirically and ranges from 1.8 through 3 (see Tables 3 and 4) depending on the channels (HNZ, HHZ, and BHZ) and severity of the filter needed to correctly identify the most signals. As we prefer to classify a teleseismic signal as being due to a local event and potentially send out a false alert rather than classify a local signal as teleseismic and possibly cancel an alert for a real earthquake, we create the filter in such a way that if a signal is not easily identifiable, we err on the side of caution and identify it as a signal from a local earthquake.

Though the best results were obtained when using longer time windows of data following the trigger, to maximize performance of the filter and minimize delays to the EEW system following repeated tests it was clear that using a minimum window of t = 0.2 was necessary. Final $V_{\text{max}}^{t,n}$ and $V_{\text{min}}^{t,n}$ values using t = 0.2 are shown in Table 4.

To determine whether an incoming signal is teleseismic, the following steps are taken (Fig. 2): The trigger is sent from the station and is received by the waveform processor. The data from 30 s before the trigger through t s after the trigger is filtered using the filter bank parameters described above. PGV is calculated for each of the nine narrow passband filtered traces (PGV_{nb}) . If all PGV_{nb} values for that signal lie within the $V_{max}^{t,n}$ and $V_{\min}^{t,n}$ values and also the ratio of $\mathrm{PGV_7}/\mathrm{PGV_3} > 0.90$ (Fig. 3), then the signal is determined to be teleseismic. The additional constraint requiring $PGV_7/PGV_3 > 0.90$ for a trigger to be classified as teleseismic was added to make the filter even more robust. As teleseismic signals usually lack high-frequency signals, this requirement verifies that the ratio of lowfrequency amplitudes to high-frequency amplitudes is relatively high. If any of those requirements is not satisfied, the signal is classified as nonteleseismic.

We repeat the calculations using the remaining t window lengths of data after the trigger. If other stations have not yet triggered or completed calculations, then more time can be used at the first station to increase robustness of the identification. Once a signal is classified as teleseismic, it is not allowed to revert back to nonteleseismic status. If at least two of the first three stations are flagged as teleseismic, then the earthquake is categorized as teleseismic, and no alert is sent out. Triggers continue to be aggregated into this event to prevent any extra triggers that may not have been correctly identified by the filter from creating another event. If at least two of the first three stations are not flagged as teleseismic, then the earthquake is classified as nonteleseismic, and the alert is sent out as usual.

Trigger Filters

The next update that was added to E3 was a trigger filter to address false alerts caused by a wide variety of problematic signals. Unlike the teleseismic filter mentioned in the Teleseismic Filter section, which still allows the flagged triggers to be associated into an event to sweep up any triggers that might not be flagged as teleseismic (noisy signals, etc.), the trigger filter described here discards any triggers that do not satisfy the requirements and does not allow that trigger to be associated with any other triggers.

Both E2 and E3 use the same STA/LTA picker to identify a trigger on the high-pass-filtered vertical component. Triggers are only identified using the velocity traces, using either the original velocity signal or signals converted from acceleration to velocity. After a trigger is detected, the algorithms perform a few basic amplitude checks to verify the trigger and prevent any exceptionally small- or large-amplitude triggers from being used. To remove problematic triggers, E3 employs additional filters that have been empirically obtained using a database of triggers. New parameters that have been added to E3 include revised minimum amplitude checks, a "range post-trigger" parameter, and a horizontal-to-vertical amplitude ratio check.

Amplitude Checks

Because of its sensitive STA/LTA picker, ElarmS is capable of associating triggers and identifying very small earthquakes. E2 and E3 both have a minimum cutoff of magnitude M > 2.0



▲ Figure 1. An example showing the difference in high-frequency amplitude between (a) a local M 4.0 earthquake and (b) a teleseismic earthquake. Green vertical lines show the trigger location, and red vertical lines are 2 s after the trigger. Note the lack of signal in the higher frequency band-pass filtered teleseismic signals.

for an alert to be sent out. Though allowing ElarmS to create alerts for such small events is useful for making sure that the system is working properly, it is not very practical to alert on such small events; the affected population is likely to be very small and the estimated shaking will be minimal. Having such a low threshold, however, has the potential to generate false events from triggers due to small blips or offsets in the signals. Some of the poorer quality stations send hundreds of bad triggers a day to the ShakeAlert servers. An example of a false alert caused by spurious triggers was when E2 associated triggers from a real M 2.7 earthquake near Humboldt County in California with a small-amplitude boxcar-shaped signal from a station further to the northeast and created a false alert for an M 6.1 earthquake just north of the California–Oregon boundary.

In E3, we introduce a simple amplitude check and require that the following amplitude restrictions are satisfied within the time window $0.1-0.2s:-0.9 < \log(\tau_p) < 1.0;-5.5 < \log(P_d) < 3.5;-5.5 < \log(P_v) < 3.0;-2.5 < \log(P_a)$, in which τ_p is the predominant period, and P_d , P_v , and P_a are the maximum displacement, velocity, and acceleration amplitudes, respectively.

As many of the nonseismic triggers are large spikes or boxcar-shaped signals, the initial offsets are often large enough to exceed the P_a threshold. Because of this, we also introduce a range post-trigger parameter R. We measure the amplitude of the original signal in units of either acceleration or velocity for strong-motion and broadband instruments, respectively, from 0.1 to 0.2 s after the trigger. By choosing this time window of data, we ensure that the triggering signal is not a single pulse or rapid offset, and that it has a duration of at least 0.2 s. If the maximum amplitude minus the minimum amplitude of the signal during that time period does not exceed 2.2×10^{-3} cm/s² for acceleration signals and 2.2×10^{-6} cm/s for velocity signals, then the trigger is discarded.

The third E3 trigger check is designed to prevent S-wave triggers from coming into the system. Triggers occurring on S-wave arrivals can be associated with other S-wave triggers or with a combination of P-wave and S-wave triggers to create poorly constrained and/or false events following an earthquake. For this parameter, we take the velocity amplitude range from the time of the trigger through 0.05 s after the trigger of either of the two horizontal components, whichever is greater, and compare with the amplitude range of the vertical component as follows:

$$\frac{H}{V} = \frac{\max(|\max(N) - \min(N)|, |\max(E) - \min(E)|)}{|\max(Z) - \min(Z)|}.$$
(3)

If the ratio of these values exceeds 0.95, then the trigger is discarded.

Table 4 $V_{max}^{t,n}$ and $V_{min}^{t,n}$ Values Using $t = 0.2$ s after the Trigger for Each Filter Bank Narrow Passband Filter Range Number n and Channel HNZ, HHZ, and BHZ						
n	$V_{\min}^{t,n}$ (HNZ)	$V_{\max}^{t,n}$ (HNZ)	$V_{\min}^{t,n}$ (HHZ)	$V_{\max}^{t,n}$ (HHZ)	$V_{\min}^{t,n}$ (BHZ)	$V_{\max}^{t,n}$ (BHZ)
1	-2.962	-6.641	-3.693	-6.644	-3.889	-5.906
2	-2.899	-6.225	-3.741	-6.270	-3.892	-5.788
3	-3.012	-5.598	-3.776	-5.956	-3.790	-5.602
4	-2.475	-5.223	-3.586	-5.491	-3.648	-5.438
5	-2.265	-4.518	-3.075	-5.140	-3.175	-5.034
6	-2.029	-4.355	-3.115	-4.782	-3.139	-4.753
7	-2.001	-4.874	-3.327	-5.185	-3.640	-5.198
8	-2.238	-5.272	-3.751	-5.225	-3.801	-5.126
9	-2.201	-5.163	-4.005	-5.289	-4.024	-5.198

ELARMS FLOW

ElarmS Waveform Processor

In this section, we describe in detail the workflow of E3 beginning first with the WP, and then moving on to the Event Associator (EA). The WP (Fig. 4) is the module that reads in and processes incoming waveforms. Data are fed into the WP either from an Earthworm ring (see Data and Resources) or a tracebuf2-message file. These are mostly 1 s duration packets. Data can also be read directly from tankplayer files for faster-than-real-time replays, which we used extensively when testing E3, and the results of which are described in the Replay Results section. These data packets are then sorted into queues by network, station, and location. Individual packets are processed with a separate thread per network, station, location identifier—up to approximately 200 threads, with more than one network, station, and location packet per thread, if necessary. These data packets are then processed resulting in a RawPacket object containing the channel description (network.station.location.channel), start time, sample rate, number of samples in the packet, an integer data array (in counts), the UTC time that the packet was read from the ring, and the time that the packet spent in the queue (up to 40 s of data are saved in the channel buffer).

The packets are next fed into the "To Ground Motion" module, which further processes the data to obtain the integer data in counts, the predominant period (τ_p) , and signal and noise levels. The "To Ground Motion" module also converts the incoming waveform to displacement, velocity, acceleration, and a high-pass velocity used by the picker. The STAs/LTAs are also calculated in this module.

Once the SampleStruct object is created, two steps take place. In the first, a peak ground-motion (PGM) message is made including the channel description (network, station,



Figure 2. Filter bank teleseismic filter logic flowchart.



▲ Figure 3. (a) Narrowband peak ground velocity (PGV_{nb}) values for HHZ recordings of local earthquakes of various magnitudes (rainbow colors), and the sharply contrasting PGV_{nb} values for teleseismic events (purple). Yellow bands in the middle of and bounding above and below the purple teleseismic lines are the median and ± 2.5 standard deviation values, respectively, of the teleseismic values. To classify as a teleseismic event, the PGV_{nb} values for an event must lie within this yellow band. (b) Plotting the ratio of two key passband values: 6– 12 and 0.375–0.75 Hz. Note the distinct separation between local and teleseismic earthquakes. We also require that the ratio of the log(PGV_(0.375-0.75 Hz)) / log(PGV_(6-12 Hz)) > 0.90 to classify as a teleseism. Example local (circle) and teleseismic (triangle) events are plotted on both (a,b).

location, channel, latitude, and longitude), the time of the first sample in the packet, the number of samples in the packet, the displacement/velocity/acceleration maxima for this packet, and the latency of this packet. The Sender module then packages the PGM message from multiple channels within a 0.5 s window and transmits them as a single message. The ActiveMQ topic eew.alg.elarms.trigger.data is updated by the WP process with post trigger measurements of displacement, velocity, acceleration, τ_p , signal level, and noise level for the Z, N, and E components that are used by the EA to determine the acceptability of the trigger and its amplitude for event magnitude calculations. The latency information in the PGM message is also stored in a Postgres database (see Data and Resources), where it is available to a web page that monitors station latencies in near-real time.

The second process that begins is the triggering algorithm. E3 checks whether or not this network.channel.location is currently in a triggered state. If not, the STA/LTA picker runs. When the picker detects a trigger, the WP checks to see if the trigger is a new trigger. If it is not, the packet is discarded. If it is new, the WP continues its processing. At this point, the triggering processing is either started (new triggers from network.channel.locations that are in a newly triggered state), or continued (existing triggers from a previously triggered state). The WP then calculates the maximum amplitude for multiple data window lengths and frequency bands (used for the filter bank teleseismic filter), and then the Sender module sends a message for each packet, with updates through 4 s after the initial trigger. The ActiveMQ topic eew.alg.elarms.trigger.data is updated with the trigger information, and a final message is sent out by the WP. This message, the Production System message, includes data from 1 s before the trigger to 4 s after the trigger with displacement, velocity, acceleration, τ_p , signal level, noise level, and the STA/LTA for Z, N, and E components at the full sample rate.

ElarmS Event Associator

The EA (Fig. 5) consists of two main processes: the Message Receiver and the Module Manager. The Message Receiver first receives trigger messages from the WP and, if the message is not a duplicate trigger, saves the trigger in the trigger buffer. The Message Receiver ingests the PGM message and updates the station active time (used to create a list of active stations, which is in turn used to calculate the percentage of active nontriggering stations), and the 10-min average latency is logged. If a Telestifle Warning Message has been created by the ANSS Quake Monitoring System, it saves the calculated arrival times of incoming teleseismic waves to block triggers from those stations during that time (P. Hellweg, personal comm., 2018).

Elarms Waveform Processor



▲ Figure 4. ElarmS v.3.0 (E3) Waveform Processor flowchart.

Elarms Event Associator



▲ Figure 5. E3 Event Associator flowchart.



▲ Figure 6. Recent event dataset replay results for (a) Earthquake Alert Systems (ElarmS) v.2.0 (E2) and (b) E3 showing alerts for ANSS or ElarmS $M \ge 3.0$ events. On the maps, the following are plotted: yellow circles, ANSS locations; green circles, ElarmS location estimates; orange circles, missed events; and red circles, false events. The ShakeAlert reporting region is shown by the white boxy outline encompassing Washington, Oregon, and California, and extending into the Pacific Ocean.

The second main process is the Module Manager. The Module Manager reads in the trigger buffer, identifies whether the trigger is new or an update, and then checks whether or not the trigger criteria are passed. These trigger criteria include the following: (1) check if τ_p , displacement, velocity, and acceleration are within acceptable ranges; (2) check the zero-crossing rate (this experimental trigger criterion is currently not used by E3); (3) check the horizontal-to-vertical amplitude ratio; (4) check whether the maximum post-trigger change in acceleration or velocity exceeds threshold; and (5) to contribute to an event magnitude, the τ_p , P_d , and P_v signal-to-noise ratio (SNR) must be within acceptable ranges, and the station must be within 200 km of the calculated epicenter.

If the trigger passes these checks, the EA then checks if the trigger associates with an existing event. If it does not, the EA creates a new event. If it does associate with an existing event, the trigger is added to that event. The magnitude is then either computed (new event) or updated (existing event), and the event is then checked against the alert criteria. To send out an alert, E3 requires that (1) there are triggers from four stations; (2) 40% of nearby stations are triggered (using the list of active stations, mentioned above); (3) the magnitude is

between (2.0 < M < 10.0) and that there is at least one station reporting an acceptable magnitude; and (4) the event is not on the edge of the grid, and the 4-station root mean square (rms) < 1.0.

If the event does not pass the alert criteria check, E3 does not send out an alert. If the event does pass the alert criteria check, E3 creates the alert, sends a log message, updates the Postgres database, and finally, updates the webpages.

REPLAY RESULTS

To verify that E3 was able to reduce the number of false alerts caused by E2 without sacrificing the performance of the system, we ran replay tests using both E2 and E3 algorithms. For these replays, we used two distinct datasets. The first dataset used was the ShakeAlert Testing and Certification (T&C) group's dataset of historic earthquakes and anomalous events (Cochran *et al.*, 2017). This dataset was put together by the T&C group to provide a test suite of events that the ShakeAlert team could use to assess the performance of algorithms and any changes made to those algorithms. The T&C test suite is composed of 40 moderate-to-large (M 3.8-7.2)



▲ Figure 7. Recent event dataset replay results for E2 and E3 showing histograms of (a) the location errors and (b) absolute magnitude errors of alerts created by E2 (blue) and E3 (green) replays for Advanced National Seismic System (ANSS) or ElarmS $M \ge 3.0$ events.

local earthquakes, 30 teleseismic and regional earthquakes, and 40 calibration signals.

As the majority of the events in the T&C test suite are potentially troublesome signals that have been collected over the last few years of development and testing of ShakeAlert (teleseisms, regional earthquakes, and anomalous signals) and each of the teleseismic and regional earthquakes has the potential to create numerous false alerts by ElarmS, the dataset is particularly useful for evaluating an algorithm's performance during problematic events. To test the performance of E3 in normal operating conditions, we also create a second database of 726 $M_{\rm w} \ge 3$ earthquakes throughout the west coast from 1 February 2016 through 31 October 2017, and also any events that occurred outside of the ShakeAlert reporting region boundary but which caused real-time E2 to create a false alert. This "Recent Events" dataset allows us to compare the performance of E2 and E3 for more typical events than those found in the T&C test suite. Using these two datasets, we are able to obtain as complete a picture as possible of how ElarmS performs in a variety of circumstances.

As can be seen in Figures 6–8 and Table 5, replay results of the two algorithms show that E3 is a dramatic improvement over E2. When replaying the recent events dataset, E3 creates just 6 $M \ge 3$ false alerts, significantly fewer than the 53 $M \ge 3$ false alerts created by E2. Of the six false events that E3 created from this dataset, two are poorly located (distance error between 100 and 200 km) real earthquakes, leaving just four false alerts that are due to poor triggers (anomalous signals, *S*wave triggers, etc.) and/or mislocated events. Because of the new amplitude filters in E3, it is inherently less sensitive than E2. As expected, E2 creates alerts for more earthquakes than E3 (61 earthquakes missed by E2; 87 missed by E3); however, most of these events are smaller (**M** 3.0–4.0) and in regions with sparser station coverage (Fig. 6).

In Figure 7, we compare the earthquake location and magnitude estimates from E2 and E3 for the recent events dataset. As the objective of E3 was to reduce the number of false alerts without negatively impacting the performance of ElarmS during local earthquakes and improving location and magnitudes was not a focus for this study, as expected there is very little difference in the accuracy of the location and magnitude estimates from the two versions of ElarmS. The majority of the events with poor locations (distance errors of 50 km or more) are offshore and out-of-network earthquakes, which are notoriously difficult for ElarmS to locate due to lack of azimuthal station coverage for these events.

When replaying the T&C test suite, E2 created 54 false events with $\mathbf{M} \ge 3$ (Fig. 8). Most of these false events were due to teleseismic earthquakes, though a few were either split events from larger earthquakes within the reporting region or poorly located events. In contrast, E3 created just three $\mathbf{M} \ge 3$ false alerts from the same dataset. Of these three events, two



Figure 8. (a) E2 replay performance for the Testing and Certification group's test suite showing alerts for ANSS or ElarmS $\mathbf{M} \ge 3.0$ events. (b) E3 replay performance for the same dataset. (E) Images from screenshots of the ElarmS review tool (available in the electronic supplement to this article).

were very poorly located real earthquakes that had occurred outside of the network. The third false event was due to the M 7.3 Kuril Islands teleseism that occurred on 17 November 2002.

In addition to these large replay datasets, we also ran a replay of a single day (10 October 2017) of continuous data with data from all available EEW stations. During this 24-hr period, the real-time E2 system had created 42,383 triggers. As the station list that was available for the replay was not exactly the same as the original station list, the number of triggers created by E3 was slightly lower (36,564). 34,845 triggers were found in both the original E2 log and the E3 replay log. The E3 replay filtered out 11,413 more of those triggers created by both versions of the algorithm than E2.

REAL-TIME RESULTS

E3 began running on the ShakeAlert production systems on 25 January 2018. Even though it is only possible to look at performance statistics for a relatively short time window of just a few months, E3 appears to be running well thus far (Fig. 9). Between 1 February and 20 June 2018, E3 created 31 matched

Table 5 E2 and ElarmS v.3.0 (E3) Replay Results of the "Recent Events" Dataset					
Algorithm	Matched	Missed	False (within 100 km)	False (within 200 km)	
E2	402	61	53	46	
E3	375	87	6	4	
This dataset is composed of 726 earthquakes with ANSS $M \ge 3.0$, and 87 known problematic events. The numbers of matched, missed, and false events where either the ElarmS or ANSS magnitude is M 3.0 or greater are shown here. E2 creates alerts for 27					

more earthquakes than E3; however, most of these events are smaller (\mathbf{M} 3.0–4.0) and in regions with sparser station coverage. New amplitude filters in E3 prevent alerts for these events. False events are reduced by 47 events with E3. Of the six false events that E3 created from the dataset, two are poorly located (distance error between 100 and 200 km).



▲ Figure 9. Figures showing real-time performance of E3 from 1 February (when E3 was added to the ShakeAlert production system) through 20 June 2018 for ANSS $M \ge 3.5$ earthquakes. (a) Map showing E3 matched, missed, and false alerts during this time. (b) Map showing alert times of the correctly matched E3 alerts. In general, alerts are faster in regions where the station density is higher. (c) Histogram showing ANSS-E3 location error for the matched events during this time. The larger location differences are, in general, due to out-of-network (e.g., offshore) events. (d) ANSS-E3 magnitude difference for matched events during this time. Blue line shows 1:1 ratio.

events for earthquakes with ANSS magnitudes $M \ge 3.5$. The initial E3 estimates for these events had a median distance error of 3.89 km (standard deviation: 17.0 km), and a median

magnitude error of 0.30 magnitude units (standard deviation: 0.19 magnitude units). Median alert times for these events were 7.7 s (standard deviation: 6.5 s). Similar to E2 alerts, the E3

Table 6 Summary of E3 Replay and Real-Time Performance						
Replay and Magnitude Range Used	Matched	Missed	False	Time Span	Region	
T&C test suite $\mathbf{M} \ge 4$	34	20	2	1999/10-2015/05	West Coast	
T&C test suite $\mathbf{M} \ge 5$	20	7	2	1999/10-2015/05	West Coast	
T&C test suite $\mathbf{M} \ge 6$	9	2	0	1999/10-2015/05	West Coast	
Recent event replay $\mathbf{M} \ge 4$	20	6	1	2016/02-2016/12	California only	
Recent event replay $\mathbf{M} \ge 5$	3	1	0	2016/02-2016/12	California only	
Recent event replay $\mathbf{M} \ge 6$	0	0	0	2016/02-2016/12	California only	
Real time M ≥ 4	12	2	1	2018/02-2018/06	West Coast	
Real time M ≥ 5	1	0	0	2018/02-2018/06	West Coast	
Real time M ≥ 6	0	0	0	2018/02-2018/06	West Coast	
Columns show number of matched (alerts created for events that match ANSS events within 30 s and 100 km) missed and false						

Columns show number of matched (alerts created for events that match ANSS events within 30 s and 100 km), missed, and false alert created by E3. Performance statistics are shown for the Testing and Certification (T&C) test suite (significant west coast and teleseismic earthquakes between October 1999 and May 2015), the recent events test suite (all $\mathbf{M} \ge 3$ earthquakes in California between 1 February and 31 December 2016), and real-time performance for E3 throughout the west coast from 1 February through 20 June 2018).

alert times are later and the location estimates are not as good for events that occur outside of the network. As it is more difficult for E3 to accurately estimate the location of offshore earthquakes due to the lack of azimuthal coverage for these events, to understand how the system performs under more ideal station coverage conditions, it can be helpful to look at the statistics for events that occur within the network only. The median distance error of initial E3 alerts for the 20 innetwork events that occurred during this time was 2.78 km (standard deviation: 9.8 km), though the magnitude error was almost identical to the magnitude error using both in-network and out-of-network events (median magnitude error: 0.30 magnitude units; standard deviation: 0.15 magnitude units). Median alert times for in-network events only were 6.0 s (standard deviation: 2.0 s).

During this same time window, E3 failed to create alerts for five $M \ge 3.5$ earthquakes. Two of these missed events occurred offshore northern California, one was an M 3.7 in the Salton Sea area, one was an M 3.5 in the East Bay area, and one was an M 3.6 earthquake near Mammoth Lakes.

Because it began running on the production system, E3 created just one $M \ge 4$ false alert and two false alerts with magnitudes between M 3.5 and 4. The M 4.0 false alert was due to the combination of a 562-km-deep M 6.8 earthquake in Bolivia with unusually large high-frequency amplitudes and a previously unknown bug in the code. The bug allowed E3 to create an alert even though only 4% of nearby stations were reporting. The bug has since been fixed. The other two smaller false alerts were due to noisy stations.

EXPECTED PERFORMANCE OF E3

In this section, we analyze both the replay and the real-time performance of E3 to better understand the expected

performance of the algorithm. E3 has only been running continuously on the production machines since early 2018, during which time there have been relatively few problematic events encountered by the ShakeAlert system. We combine statistics from the T&C test suite, the recent events test suite, and realtime performance. Each of these datasets spans a different time period, from more than 15 yrs for the T&C test suite to one year for the recent events replay, and just four months for the real-time performance. In addition, the recent events is for events within California, whereas both of the other two datasets show performance for the entire west coast. Nevertheless, with these performance statistics, we can extrapolate how we expect E3 will perform.

As shown in Table 6, E3 creates no $\mathbf{M} \ge 6$ false alerts for any of the three datasets. E3 also does not create any $\mathbf{M} \ge 5$ false alerts in either the recent events (1 yr of data) or the realtime datasets (4 months continuous). E3 does create two $\mathbf{M} \ge$ 5 false alerts from events in the T&C test suite; one of which was a teleseismic earthquake, and the other was an offshore event for which E3 produced a very poor location estimate. That represents $25 < \mathbf{M} < 6$ events over 15 yrs, but the T&C dataset does not include all possible false events of course. Combining this information, we should therefore expect E3 to create false events with magnitudes between \mathbf{M} 5 and 6 less than once per year. False events with $\mathbf{M} \ge 6$ will be even less frequent, with none having been observed in testing.

E3 failed to create alerts for just two $\mathbf{M} \ge 6$ earthquakes in the T&C test suite replays, both of which were far offshore. As again this test suite spans approximately 15 yrs and included all large events, which results in a missed rate of one missed offshore earthquake every 7 yrs. E3 missed one $\mathbf{M} \ge 5$ earthquake in the recent events test suite replays (E3 created a poor location estimate). As this dataset spans roughly 1 yr, this results in a missed rate of ~1 earthquake per year. Summarizing these results, we do not expect that E3 will miss any $M \ge 6$ onshore earthquakes, though it may miss some large offshore events. We expect that E3 may miss one onshore earthquake between M 5 and 6 per year.

To obtain the estimated alert rate for E3, we can look at the combined performance of the E2 and E3 real-time systems over the past 2 yrs. As the goal of E3 was to reduce the number of false alerts, the matched alerts performance of E3 should not differ drastically from that of E2. During the past 2 yrs, there were five $M \ge 5$ earthquakes throughout the west coast. The production versions of E2 (from June 2016 through January 2018) and E3 (from February through June 2018) created alerts for all five earthquakes, three of which were offshore. During the same time period, there were 54 $M \ge 4$ earthquakes (2–3 per month). Approximately 80% of these $M \ge 4$ events were detected by E2 or E3.

Median alert times for these $\mathbf{M} \ge 4$ events in regions where station spacing is close to the target 10 km station spacing (Given *et al.*, 2014; San Francisco and Los Angeles) are just 3.7 s. As earthquakes in California typically occur at depths of around 8 km, then at the time of the alert the *S* wave will have traveled less than 9 km from the epicenter. People within this region will receive little or no warning. If the initial alert is issued across the metropolitan regions, then people 20, 30, and 40 km from the event would get 3, 6, and 9 s warning, respectively, before the *S*-wave shaking begins.

CONCLUSIONS

The ElarmS EEW algorithm is the fastest and most accurate of the ShakeAlert algorithms, and it has detected more $M \ge 3$ earthquakes than either of the other original ShakeAlert algorithms (OnSite and VS). Though E2 performed well for the majority of events, it would still create an unacceptable number of false events from teleseismic earthquakes, spikes in the data, and other spurious signals. E3 addresses the false-alert problem by requiring that incoming triggers pass a series of quality checks before they can be used to create an event. Though E3 requires the system to wait for a fraction of a second of additional data (~0.2 s) from each trigger, the resulting alerts are significantly more robust than those from E2. E3 significantly reduces the number of false alerts created by E2 while continuing to provide fast and accurate alerts.

Though we tried to make ElarmS as robust as possible, it is always possible that the system will encounter a new and unexpected challenge. Whether the cause is a unique earthquake, an abnormal signal, or the coincidence of several highly unusual events, it is not possible to guarantee that the performance of ElarmS will be flawless. Nevertheless, we believe that the benefits of the ElarmS and ShakeAlert EEW alerts greatly outweigh the cost of any rare false alerts. A reconnaissance team traveling to Mexico following the **M** 8.1 Chiapas earthquake on 7 September 2017 found that people there were more tolerant of false alerts than they were of missed alerts (Allen *et al.*, 2017). False alerts to a degree were even considered valuable because they gave people an opportunity to practice taking precautionary action and helped build awareness of earthquake risk.

E3 is now running on the fully functional ShakeAlert production system. Because of its excellent performance as part of the ShakeAlert system, E3 was chosen as the basis for EPIC, the single point-source algorithm that ShakeAlert will use in the future. Although improvements will continue to be made to EPIC in the coming years, the only significant change to E3 in the current version of EPIC is the waveform filter, the implementation of which has been modified to make E3 more compatible with modules that may be added in the future.

DATA AND RESOURCES

Waveform data, metadata, or data products for this study were accessed through the Northern California Earthquake Data Center (NCEDC, doi: 10.7932/NCEDC) and through the Southern California Earthquake Center (SCEDC). The facilities of Incorporated Research Institutions for Seismology (IRIS) Data Services, and specifically the IRIS Data Management Center, were also used for access to waveforms, related metadata, and/or derived products used in this study. Earthquake Alert Systems (ElarmS) uses both the PostgreSQL (http://postgresql.org) and Earthworm (http://www.earthwormcentral.org) software. All websites were last accessed in November 2018. **≦**

ACKNOWLEDGMENTS

This work was funded by the Gordon and Betty Moore Foundation through Grant GALA Number 3024 to U.C. Berkeley, and the U.S. Geological Survey (USGS) Cooperative Agreement G17AC00346. The Southern California Earthquake Center (SCEDC) and Southern California Seismic Network (SCSN) are funded through USGS Grant Number G10AP00091, and the SCEDC, which is funded by National Science Foundation (NSF) Cooperative Agreement EAR-0529922 and USGS Cooperative Agreement 07HQAG0008. Incorporated Research Institutions for Seismology (IRIS) Data Services are funded through the Seismological Facilities for the Advancement of Geoscience and EarthScope (SAGE) Proposal of the NSF under Cooperative Agreement EAR-1261681.

REFERENCES

- Allen, R. M. (2007). The ElarmS earthquake early warning methodology and application across California, in *Earthquake Early Warning Systems*, P. Gasparini, G. Manfredi, and J. Zschau (Editors), Springer-Verlag, Heidelberg, Germany, 21–43, doi: 10.1007/978-3-540-72241-0_3.
- Allen, R. M., and H. Kanamori (2003). The potential for earthquake early warning in southern California, *Science* 300, 786–789, doi: 10.1126/science.1080912.
- Allen, R. M., E. S. Cochran, T. Huggins, S. Miles, and D. Otegui (2017). Quake warnings, seismic culture, *Science* 358, no. 6367, 1111, doi: 10.1126/science.aar4640.

- Allen, R. M., P. Gasparini, O. Kamigaichi, and M. Böse (2009). The status of earthquake early warning around the world: An introductory overview, *Seismol. Res. Lett.* 80, no. 5, 682–693, doi: 10.1785/ gssrl.80.5.682.
- Böse, M., R. M. Allen, H. Brown, G. Gua, M. Fischer, E. Hauksson, T. Heaten, M. Hellweg, M. Liukis, D. Neuhauser, *et al.* (2014). CISN ShakeAlert: An earthquake early warning demonstration system for California, in *Early Warning for Geological Disasters*, F. Wenzel and J. Zschau (Editors), Springer-Verlag, Heidelberg, Germany, 49–69, doi: 10.1007/978-3-642-12233-0_3.
- Böse, M., E. Hauksson, K. Solanki, H. Kanamori, Y. M. Wu, and T. H. Heaton (2009). A new trigger criterion for improved real-time performance of onsite earthquake early warning in Southern California, *Bull. Seismol. Soc. Am.* 99, no. 2A, 897–905, doi: 10.1785/0120080034.
- Böse, M., T. Heaton, and E. Hauksson (2012). Rapid estimation of earthquake source and ground-motion parameters for earthquake early warning using data from a single three-component broadband or strong-motion sensor, *Bull. Seismol. Soc. Am.* 102, no. 2, 738– 750, doi: 10.1785/0120110152.
- Böse, M., C. Ionescu, and F. Wenzel (2007). Earthquake early warning for Bucharest, Romania: Novel and revised scaling relations, *Geophys. Res. Lett.* 34, no. 7, 1–6, doi: 10.1029/2007GL029396.
- Brown, H., 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 Dynam. Earthq. Eng.* **31**, no. 2, 188–200, doi: 10.1016/j.soildyn.2010.03.008.
- Chestler, S. (2017). PNSN announces launch of integrated West Coast Earthquake Early Warning System Pacific Northwest Seismic Network, retrieved on 12 January 2018 from https://pnsn.org/ blog/2017/04/14/pnsn-announces-launch-of-integrated-west-coastearthquake-early-warning-system.
- Cochran, E. S., M. D. Kohler, D. D. Given, S. Guiwits, J. Andrews, M. Meier, M. Ahmad, I. Henson, R. Hartog, and D. Smith (2017). Earthquake early warning ShakeAlert system: Testing and certification platform, *Seismol. Res. Lett.* 89, no. 1, 108–117, doi: 10.1785/0220170138.
- Cua, G. B., and T. Heaton (2007). The Virtual Seismologist (VS) method: A Bayesian approach to earthquake early warning, in *Earthquake Early Warning Systems*, P. Gasparini, G. Manfredi, and J. Zschau (Editors), Springer-Verlag, Heidelberg, Germany, 97–132.
- Cua, G. B., 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, no. 5, 740–747, doi: 10.1785/gssrl.80.5.740.
- Cuéllar, A., G. Suárez, and J. M. Espinosa-Aranda (2017). Performance evaluation of the earthquake detection and classification algorithm $2(t_S - t_P)$ of the seismic alert system of Mexico (SASMEX), *Bull. Seismol. Soc. Am.* **107**, no. 3, 1451–1463, doi: 10.1785/ 0120150330.
- Espinosa-Aranda, J. M., A. Cuéllar, F. H. Rodriguez, B. Frontana, G. Ibarrola, R. Islas, and A. Garcia (2011). The seismic alert system of Mexico (SASMEX): Progress and its current applications, *Soil Dynam. Earthq. Eng.* **31**, no. 2, 154–162, doi: 10.1016/j.soil-dyn.2010.09.011.
- Gasparini, P., G. Manfredi, and J. Zschau (2011). Earthquake early warning as a tool for improving society's resilience and crisis response, *Soil Dynam. Earthq. Eng.* **31**, no. 2, 267–270, doi: 10.1016/j.soildyn.2010.09.004.
- Given, D. D., E. S. Cochran, T. H. Heaton, E. Hauksson, R. M. Allen, M. Hellweg, J. Vidale, and P. Bodin (2014). Technical implementation plan for the ShakeAlert production system—An earthquake early warning system for the West Coast of the United States, U.S. Geol. Surv. Open-File Rept. 2014-1097, 25 pp., doi: 10.3133/ ofr20141097.

- Hoshiba, M., and T. Ozaki (2014). Earthquake early warning and tsunami warning of the Japan Meteorological Agency, and their performance in the 2011 off the Pacific Coast of Tohoku earthquake (M_w 9.0), in *Early Warning for Geological Disasters: Scientific Methods and Current Practice*, F. Wenzel and J. Zschau (Editors), Springer-Verlag, Heidelberg, Germany, 1–28, doi: 10.1007/978-3-642-12233-0_1.
- Hsiao, N. C., Y. M. Wu, T. C. Shin, L. Zhao, and T. L. Teng (2009). Development of earthquake early warning system in Taiwan, *Geophys. Res. Lett.* 36, no. 5, 3–7, doi: 10.1029/2008GL036596.
- Kamigaichi, O., M. Saito, K. Doi, T. Matsumori, S. Tsukada, K. Takeda, T. Shimoyama, K. Nakamura, M. Kiyomoto, and Y. Watanabe (2009). Earthquake early warning in Japan: Warning the general public and future prospects, *Seismol. Res. Lett.* **80**, no. 5, 717– 726, doi: 10.1785/gssrl.80.5.717.
- Kodera, Y., J. Saitou, N. Hayashimoto, S. Adachi, M. Morimoto, Y. Nishimae, and M. Hoshiba (2016). Earthquake early warning for the 2016 Kumamoto earthquake: Performance evaluation of the current system and the next-generation methods of the Japan Meteorological Agency, *Earth Planets Space* 68, no. 1, 202, doi: 10.1186/s40623-016-0567-1.
- Kohler, M. D., E. S. Cochran, D. Given, S. Guiwits, D. Neuhauser, I. Henson, R. Hartog, P. Bodin, V. Kress, S. Thompson, *et al.* (2018). Earthquake early warning ShakeAlert system: West Coast wide production prototype, *Seismol. Res. Lett.* 89, no. 1, 99–107, doi: 10.1785/0220170140.
- Kuyuk, H. S., R. M. Allen, H. Brown, M. Hellweg, I. Henson, and D. Neuhauser (2014). Designing a network-based earthquake early warning algorithm for California: ElarmS-2, *Bull. Seismol. Soc. Am.* 104, no. 1, 162–173, doi: 10.1785/0120130146.
- Lancieri, M., A. Fuenzalida, S. Ruiz, and R. Madariaga (2011). Magnitude scaling of early-warning parameters for the M_w 7.8 Tocopilla, Chile, earthquake and its aftershocks, *Bull. Seismol. Soc. Am.* 101, no. 2, 447–463, doi: 10.1785/0120100045.
- Liu, A., and M. Yamada (2014). Bayesian approach for identification of multiple events in an early warning system, *Bull. Seismol. Soc. Am.* 104, no. 3, 1111–1121, doi: 10.1785/0120130208.
- Meier, M. A., T. Heaton, and J. Clinton (2015). The Gutenberg algorithm: Evolutionary Bayesian magnitude estimates for earthquake early warning with a filter bank, *Bull. Seismol. Soc. Am.* 105, no. 5, 2774–2786, doi: 10.1785/0120150098.
- Nakamura, Y. (1988). On the urgent earthquake detection and alarm system (UrEDAS), *Proc. of the 9th World Conf. on Earthquake Engineering*, 2–9 August, 673–678, ISBN: 4-89580-010-5.
- Nof, R. N., and R. M. Allen (2016). Implementing the ElarmS earthquake early warning algorithm on the Israeli seismic network, *Bull. Seismol. Soc. Am.* 106, no. 5, 2332–2344, doi: 10.1785/ 0120160010.
- Picozzi, M., A. Zollo, P. Brondi, S. Colombelli, L. Elia, and C. Martino (2015). Exploring the feasibility of a nationwide earthquake early warning system in Italy, *J. Geophys. Res.* **120**, no. 4, 2446–2465, doi: 10.1002/2014JB011669.
- Sheen, D.-H., I.-S. Lim, J.-H. Park, and H.-C. Chi (2014). Magnitude scaling relationships using P waves for earthquake early warning in South Korea, *Geosci. J.* 18, no. 1, 7–12, doi: 10.1007/s12303-013-0066-3.
- Strauss, J. A., and R. M. Allen (2016). Benefits and costs of earthquake early warning, *Seismol. Res. Lett.* 87, no. 3, 765–772, doi: 10.1785/ 0220150149.
- Suárez, G., J. M. Espinosa-Aranda, A. Cuéllar, G. Ibarrola, A. García, M. Zavala, S. Maldonado, and R. Islas (2018). A dedicated seismic early warning network: The Mexican Seismic Alert System (SASMEX), *Seismol. Res. Lett.* 89, no. 2A, 382–391, doi: 10.1785/0220170184.
- Wenzel, F., M. Erdik, N. Köhler, J. Zschau, C. Milkereit, M. Picozzi, J. Fischer, J. P. Redlich, F. Kühnlenz, B. Lichtblau, *et al.* (2014). EDIM: Earthquake disaster information system for the Marmara

region, Turkey, in *Early Warning for Geological Disasters: Scientific Methods and Current Practice*, F. Wenzel and J. Zschau (Editors), Springer-Verlag, Heidelberg, Germany, 103–116, doi: 10.1007/978-3-642-12233-0_6.

- Wurman, G., R. M. Allen, and P. Lombard (2007). Toward earthquake early warning in northern California, J. Geophys. Res. 112, no. 8, 1– 19, doi: 10.1029/2006JB004830.
- Zollo, A., S. Colombelli, L. Elia, A. Emolo, G. Festa, G. Iannaccone, C. Martino, and P. Gasparini (2014). An integrated regional and onsite earthquake early warning system for southern Italy: Concepts, methodologies and performances, in *Early Warning for Geological Disasters: Scientific Methods and Current Practice*, F. Wenzel and J. Zschau (Editors), Springer-Verlag, Heidelberg, Germany, 117–137, doi: 10.1007/978-3-642-12233-0_7.

Angela I. Chung Ivan Henson Richard M. Allen U.C. Berkeley Seismological Laboratory University of California, Berkeley 307 McCone Hall #4760 Berkeley, California 94720-4760 U.S.A. aichung@berkeley.edu

Published Online 16 January 2019