

Earthquake Early Warning ShakeAlert 2.0: Public Rollout

Monica D. Kohler^{*1}, Deborah E. Smith², Jennifer Andrews³, Angela I. Chung⁴, Renate Hartog⁵, Ivan Henson⁴, Douglas D. Given², Robert de Groot², and Stephen Guiwits²

Abstract

The ShakeAlert earthquake early warning system is designed to automatically identify and characterize the initiation and rupture evolution of large earthquakes, estimate the intensity of ground shaking that will result, and deliver alerts to people and systems that may experience shaking, prior to the occurrence of shaking at their location. It is configured to issue alerts to locations within the West Coast of the United States. In 2018, ShakeAlert 2.0 went live in a regional public test in the first phase of a general public rollout. The ShakeAlert system is now providing alerts to more than 60 institutional partners in the three states of the western United States where most of the nation's earthquake risk is concentrated: California, Oregon, and Washington. The ShakeAlert 2.0 product for public alerting is a message containing a polygon enclosing a region predicted to experience modified Mercalli intensity (MMI) threshold levels that depend on the delivery method. Wireless Emergency Alerts are delivered for M 5+ earthquakes with expected shaking of $\text{MMI} \geq \text{IV}$. For cell phone apps, the thresholds are M 4.5+ and $\text{MMI} \geq \text{III}$. A polygon format alert is the easiest description for selective rebroadcasting mechanisms (e.g., cell towers) and is a requirement for some mass notification systems such as the Federal Emergency Management Agency's Integrated Public Alert and Warning System. ShakeAlert 2.0 was tested using historic waveform data consisting of 60 M 3.5+ and 25 M 5.0+ earthquakes, in addition to other anomalous waveforms such as calibration signals. For the historic event test, the average M 5+ false alert and missed event rates for ShakeAlert 2.0 are 8% and 16%. The M 3.5+ false alert and missed event rates are 10% and 36.7%. Real-time performance metrics are also presented to assess how the system behaves in regions that are well-instrumented, sparsely instrumented, and for offshore earthquakes.

Cite this article as Kohler, M. D., D. E. Smith, J. Andrews, A. I. Chung, R. Hartog, I. Henson, D. D. Given, R. de Groot, and S. Guiwits (2020). Earthquake Early Warning ShakeAlert 2.0: Public Rollout, *Seismol. Res. Lett.* **XX**, 1–13, doi: [10.1785/0220190245](https://doi.org/10.1785/0220190245).

[Supplemental Material](#)

Introduction

The goals of the ShakeAlert earthquake early warning (EEW) system for the West Coast of the United States, like other similar systems around the world, are to automatically identify and characterize an earthquake rapidly after it begins, estimate the intensity of ground shaking that will result, and deliver alerts to people and systems that may experience shaking (Given *et al.*, 2014, 2018; Kohler *et al.*, 2018). Its objective is to issue alerts for a defined level of ground motion before that ground motion occurs at a user's location. The U.S. Geological Survey (USGS), together with partner organizations, has been developing and operating ShakeAlert for the highest-risk areas of the United States by leveraging the current earthquake monitoring capabilities of the Advanced National Seismic System (ANSS; USGS, 2017).

In 2018, ShakeAlert went live in a regional public test of the first phase of a general public rollout (hereafter, referred to as "ShakeAlert 2.0"). This event was preceded in 2016 by the

Production Prototype version (referred to as "ShakeAlert 1.0"; Kohler *et al.*, 2018), which went live in California only, and began providing notifications to a limited group of about 20 early adopter community participants through pilot implementations. In 2017, Production Prototype 1.2 went live to a small group of community participants on the entire West Coast (California, Oregon, and Washington). The ShakeAlert system is now providing alerts to institutional users such as transportation systems (e.g., Bay Area Rapid Transit, Los Angeles County Metropolitan Transportation Authority) and

1. Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, California, U.S.A.; 2. U.S. Geological Survey, Pasadena, California, U.S.A.; 3. Seismological Laboratory, California Institute of Technology, Pasadena, California, U.S.A.; 4. Seismological Laboratory, University of California, Berkeley, Berkeley, California, U.S.A.; 5. Department of Earth and Space Sciences, University of Washington, Seattle, Washington, U.S.A.

*Corresponding author: kohler@caltech.edu

© Seismological Society of America

commercial entities developing hardware for alert delivery in the three states of the western United States where most of the nation's earthquake risk is concentrated (Federal Emergency Management Agency [FEMA], 2008).

The overall goal of the ShakeAlert product is to issue warnings of potentially damaging earthquake shaking to the public (Given *et al.*, 2014, 2018). To that end, the 2018 ShakeAlert product for public alerting is a message containing a polygon enclosing a region predicted to experience modified Mercalli intensity (Wood and Neumann, 1931) (MMI) \geq IV for an earthquake of $M \geq 5.0$. Recall that whereas MMI is an empirical measure of ground-shaking intensity, magnitude is an empirical measure of the energy released by an earthquake. Estimated ground shaking in the form of MMI level is used as the basis of alert thresholds, in part because future algorithms may not explicitly estimate an earthquake's location or magnitude (e.g., ground motions only) and because this is a more direct way to estimate the potential for structural and non-structural damage. MMI is chosen because it combines the attributes of peak ground acceleration (PGA) and peak ground velocity (PGV), which both correlate with the strength of felt shaking and the expected damage (Wald *et al.*, 1999). An MMI threshold of IV is chosen to alert populations who might experience moderate to strong ground shaking (the "damage" threshold), to accommodate shaking estimate uncertainties in the system, and to maximize potential alert times. Alerts are only provided for earthquakes that have at least some minimal potential to cause structural or nonstructural damage, defined here as $M 5.0+$. To be effective, alerts must be provided quickly enough to arrive at most users' locations before the arrival of damaging ground motion; the threshold at which damage occurs depends on the specific application. By design, these alerting thresholds are monitored and may change based on system performance and user demand.

This article focuses on the software architecture of ShakeAlert 2.0, relative to its predecessor (ShakeAlert 1.0), which was described by Kohler *et al.* (2018), and on the performance of the system during testing with both real-time and historic event datasets. The testing process includes a more extensive quantitative measure of how well the system alerts for specific levels of ground motion, and assessments of these measures for the real-time and historic event datasets are presented. For a comprehensive description of ShakeAlert history and overall goals, particularly within the context of USGS infrastructure, see Given *et al.* (2018).

Phase 1 Rollout

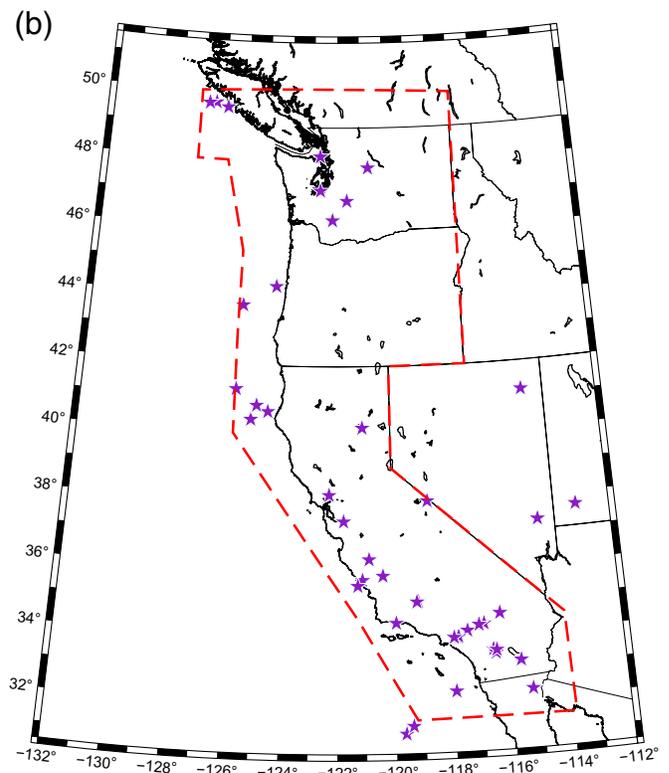
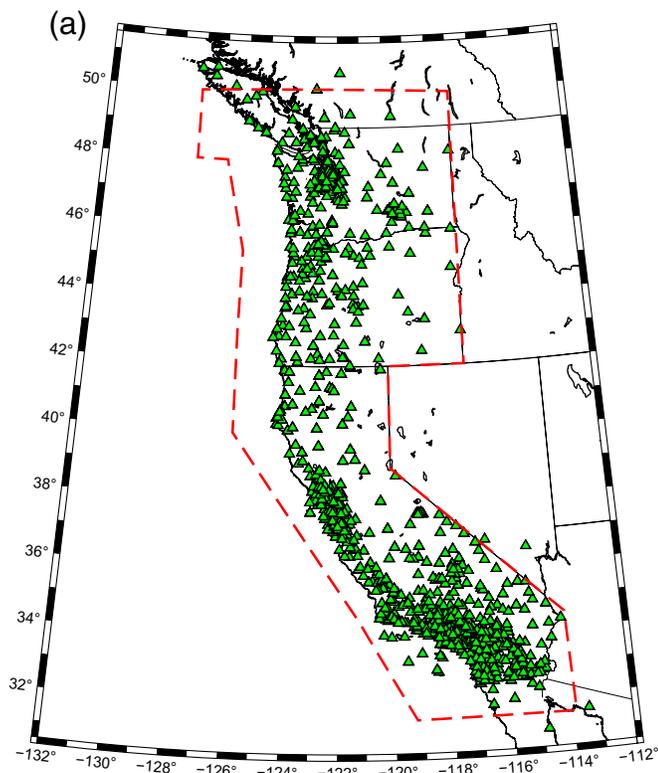
In fall 2018, ShakeAlert entered phase 1 of alerting in California, Oregon, and Washington. More than 60 institutional partners are now alerting personnel and taking automated actions, important steps in a strategy of phased rollout leading to full public operation. Concurrent with the fall 2018 release of ShakeAlert 2.0, it was announced that

ShakeAlert was "open for business," which included plans to accelerate the recruitment and development of technical partnerships. This officially marked the end of the production prototype (or pilot) phase in which, by design, the number of partners working on "proof-of-concept" pilot projects (e.g., using alerts to close a valve to protect a water supply or open a firehouse door) was kept small. In addition, the introduction of a ShakeAlert "license to operate" debuted, allowing partners in the advanced stages of their relationship with ShakeAlert to take their products to the market provided they continue to abide by restrictions on alerting thresholds and meet the latency requirement to deliver alerts to 95% of their end users within 5 s.

To participate in a ShakeAlert technical partnership, partners must meet some basic criteria. Alerts and automated actions must be fast enough to be effective, and the partners must develop a plan to address meeting and sustaining the latency requirement described earlier. Automated actions must be tolerant of system errors including false, missed, or late alerts and incorrect intensity estimates, and must make every effort not to have the potential to result in injury, damage, or loss. As of September 2019, >60 technical partners have been working with the ShakeAlert Joint Committee for Communication, Education, and Outreach (JCCEO) to develop, test, and implement appropriate personal protective actions and specialized responses that prioritize human safety. The end user-JCCEO partnership draws on existing industry best practices and social science research to optimize human response in taking a protective action such as drop, cover, and hold on or other mandated action as defined in a partner's standard operating procedure.

Three mobile phone apps (MyShake, QuakeAlert, and ShakeAlertLA) designed to deliver alerts derived from USGS-issued ShakeAlerts moved into the testing phase as part of the phase 1 rollout. All app partners are required to (1) have the institutional capacity to manage and sustain the app, (2) have appropriate server capacity, (3) demonstrate timely delivery of alerts, and (4) deliver postalert messages. All app development teams must provide appropriate education and training in partnership with the JCCEO. In addition, all developers must abide by USGS-mandated alerting thresholds. On 31 December 2018, the first large-scale test of an app began in Los Angeles County with the release of the ShakeAlertLA app by the City of Los Angeles. It had >400,000 downloads in the first month it was available.

Broader public alerting (e.g., via FEMA's Wireless Emergency Alert [WEA] system) at the $M 5.0$ and MMI IV levels will begin when existing mass alerting technologies are able to deliver ShakeAlerts at the speed or scale needed for effective EEW and with the concurrence of the state emergency management agency. ShakeAlert partners are working with both public and private mass alert system operators, including FEMA's Integrated Public Alert and Warning System, cellular carriers, mass notification companies, and others, to provide



this functionality. Mass notification companies (e.g., Regroup and Everbridge) have existing infrastructure to alert college campuses, large businesses, and cities about severe weather, active shooters, and so on using modalities such as reverse 911, SMS, email, and push notifications from apps. Both Regroup and Everbridge are ShakeAlert technical partners, and they are testing the viability of EEW alert delivery via several of the modalities that they currently offer. Finally, sufficient, broad public awareness, education, and training are being conducted in partnership with each state's emergency management agency. The USGS has invested resources in a partnership with Denver-based Nusura, Inc., to develop strategic communication resources (e.g., talking points, fact sheets, animations, engagement strategies, teaching resources) that are part of tool kits that will be used by ShakeAlert stakeholders in all three states. The Nusura tool kits emphasize appropriate protective actions, preparedness principles, and setting appropriate expectations from ShakeAlert. Parallel resources (but focused more on the technical side) are in development by the USGS Office of Communication and Publishing. Materials will be used for public meetings and other gatherings, and will also be available on the web (see e.g., [Data and Resources](#)). USGS maintains an active Twitter presence @USGS_ShakeAlert (8800+ followers as of November 2019). In addition, the USGS is collaborating with the Incorporated Research Institutions for Seismology to develop technical animations (e.g., the difference between magnitude and intensity), educational activities, and other resources that can be used by

Figure 1. (a) Advanced National Seismic System (ANSS) network stations (green triangles) in the western United States, making up the three tier 1 regional seismic networks used by ShakeAlert 2.0 as of May 2019. A tier 1 center is a seismic network covering a broad area with an established data processing center (U.S. Geological Survey [USGS], 2019), consisting here of the Pacific Northwest Seismic Network, the Northern California Seismic Network, and the Southern California Seismic Network. The alerting region is shown by the dashed red outline. Not all stations are required to be within the alerting region. In the Pacific Northwest, the alerting region extends to just west of the Cascadia subduction zone. (b) Earthquakes used in historic event testing. The color version of this figure is available only in the electronic edition.

classroom teachers, emergency managers, and in free-choice learning environments such as museums. There has been considerable effort to leverage and build on existing resources developed by partners such as the Earthquake Country Alliance (developers of The Great ShakeOut). In addition, on the USGS end, ShakeAlert is being marketed as a critical new offering of ANSS products that contribute to earthquake risk reduction.

Alert areas and thresholds

ShakeAlert issues alerts for earthquakes that fall within the U.S. West Coast reporting region, a polygon that includes the states of California, Oregon, and Washington and extends a short distance into Baja California and Canada (Fig. 1). The

ShakeAlert system publishes alerts for M 3.5+ earthquakes, and they are made available for all delivery partners. However, the USGS has mandated higher minimum thresholds that are tailored for specific types of alert delivery system. For example, WEA can only be delivered for M 5+ earthquakes and only to people (i.e., their wireless devices) who could potentially experience damaging shaking (MMI IV+). For cell phone apps, the mandated thresholds are slightly lower— M 4.5+ and weak shaking (MMI III) or greater. For earthquakes with magnitudes $\geq M$ 6.0, an estimated line (finite-fault) source will be included in the alert if available, providing an indication of both the strike and extent of the rupturing fault. Estimated ground motion is also provided.

Updates to regional networks (sensors and telemetry)

The ShakeAlert project plan calls for a network of 1675 high-quality, real-time seismic stations: 1115 in California and 560 in Oregon and Washington. At the beginning of 2019, >900 stations were contributing to the system (Fig. 1a), although some of them still require upgrades, primarily to accommodate faster and more reliable telemetry. Station construction has been focused on seismic high-risk areas; therefore, despite being incomplete, the ShakeAlert network is sufficient for rapid alerting in parts of Southern California; the San Francisco Bay area, including Silicon Valley; and the Seattle–Tacoma area. Alert delivery in areas where the network is not complete will experience additional latency plus latencies contributed by the delivery modality that was used. USGS is working with state emergency managers and alert delivery organizations to address this challenge, including communicating this limitation to end users.

Software and Central Processing Architecture

ShakeAlert 2.0 produces both point-source and line-source earthquake solutions, has added ground-motion estimation products, and has reduced the number of false and missed events. The ShakeAlert 2.0 system has also satisfied government cybersecurity requirements and includes improved operational procedures.

The ShakeAlert 2.0 algorithm and hardware architecture, shown schematically in Figures S1 and S2 (available in the supplemental material to this article), is an improvement over ShakeAlert 1.0 (Kohler *et al.*, 2018), which was in operation since 2016. In particular, each ShakeAlert 2.0 production server has two algorithms that can detect earthquake events. The first algorithm, Earthquake Point-Source Integrated Code (EPIC), triggers from seismic ground-motion data and generates a point-source estimate, which includes magnitude, location, and origin-time (OT) parameters. This algorithm is largely based on Earthquake Alarm Systems (ElarmS, one of the original ShakeAlert algorithms; Chung *et al.*, 2019). The

second algorithm, Finite-Fault Rupture Detector (FinDer), also triggers from seismic data and generates both a point-source and a line-source estimate; hence, this additionally yields fault length and strike information (Böse *et al.*, 2012b). The source parameters from these two independent algorithms are then combined in the Solution Aggregator (SA) algorithm. Next, the eqInfo2GM algorithm takes the information from the SA and generates ground-motion estimates for the earthquake in the form of contour messages and map messages. Last, the information is sent through the Decision Module (DM) algorithm, which runs on six alert servers maintained by the USGS. The DM runs a series of quality tests and checks if certain thresholds are exceeded to determine if the alert will be sent out and distributed to end users such as emergency responders, Bay Area Rapid Transit train system, the City of Los Angeles, and other users. A Heartbeat Monitor algorithm also running on the alert servers characterizes overall system health. Figure S1 summarizes the algorithms and message-sharing pathways among the four contributing regional networks, starting from the incoming seismograms (Fig. S1, bottom), to alerts sent to general public (Fig. S1, top).

EPIC algorithm

ShakeAlert originally consisted of three point-source algorithms: ElarmS, Onsite, and Virtual Seismologist (VS; Allen and Kanamori, 2003; Allen, 2007; Wurman *et al.*, 2007, Böse *et al.*, 2009, 2012a; Cua and Heaton, 2009; Cua *et al.*, 2009). Although VS was able to accurately locate and estimate the magnitude of earthquakes, it was consistently slower than the other two algorithms (Chung *et al.*, 2019) and was removed from ShakeAlert in 2016. In early 2017, it was decided that to streamline ShakeAlert processing, the remaining algorithms should be merged as the system's single point-source algorithm. Because ElarmS historically created more alerts for earthquakes and more accurate estimates of the earthquake location and magnitude than Onsite (Chung *et al.*, 2019), it was used as the basis for the new algorithm. The resulting algorithm, EPIC, came online in 2018 with the release of ShakeAlert 2.0.

EPIC was created from the most recent version of ElarmS: “ElarmS 3” (Chung *et al.*, 2019). EPIC is identical to ElarmS 3 in how it creates triggers (i.e., it uses the identical classic short-term average/long-term average picker), how it evaluates trigger quality, and how it determines the best location (minimizing travel-time difference via a grid search) and magnitude (using peak P -wave displacement up to 4 s after the trigger, averaging magnitude estimates from a minimum of four stations); see Chung *et al.* (2019) for details. One difference from ElarmS 3 is that the EPIC waveform processor applies a different discrete-time high-pass filter and then performs an additional baseline correction to the incoming seismograms, features taken from the Onsite code. ElarmS 3 used a first-order Butterworth filter with a filter coefficient that was the same for all sampling frequencies (Kanamori *et al.*, 1999),

which led to slightly different cutoff frequencies for channels with different sample rates (Chung *et al.*, 2019). The EPIC waveform processor uses a second-order Butterworth filter with a prescribed cutoff frequency of 0.075 Hz, that is, the filter coefficients depend on the sample rate. EPIC also uses a simple baseline correction by removing the average amplitude calculated using the previous 60 s of the waveform from each sample. If this length of data is not yet available, such as when the algorithm is first started or following a gap in data, the available data are used so as not to delay further processing.

EPIC is susceptible to magnitude saturation, performing well for $M < \sim 7$, because it uses ground acceleration and velocity as inputs. However, because of the sensitivity of the algorithm, it is able to consistently create faster and more accurate estimates of location and magnitude than other point-source algorithms (Cochran *et al.*, 2018; Chung *et al.*, 2019). Provided adequate station coverage, EPIC can create alerts for earthquakes as small as M 3.0. The combined use of EPIC and FinDer, described next, allows ShakeAlert to create rapid, accurate alerts for a large range of earthquake magnitudes, from M 3.0 and up.

FinDer algorithm

The FinDer algorithm (Böse *et al.*, 2012b, 2015, 2018) became a component of ShakeAlert in March 2018 after extensive development and testing. This algorithm eliminates magnitude saturation and extends ShakeAlert's alert magnitude range to M 9.0+. It also estimates finite-fault parameters leading to improved characterization of ground motion. Unlike EPIC and previous point-source algorithms tested and used by ShakeAlert, FinDer is not dependent on a picking algorithm. It instead uses a complementary methodology based on the spatial pattern of ground-motion amplitudes to produce the additional finite-fault parameter characterization.

The input data for FinDer are the maximum PGA values within 120 s time windows recorded by all available stations in the network; these values are gridded (5 km resolution, to match templates) and contoured using functions in the Generic Mapping Tools library (Wessel and Smith, 1998), and then converted to a binary map image that shows regions above and below a configured threshold. FinDer then uses template matching to compare this observation image with template images for different fault lengths and orientations (using functions in the openCV library; Bradski and Kaehler, 2013). FinDer characterizes the best-matching image in terms of fault centroid, length, and strike. Fault length is used to estimate magnitude using the relations of Wells and Coppersmith (1994) for crustal events. Because FinDer is primarily concerned with matching near-source, high-frequency ground motions, and scaling relations are necessarily averaged over earthquakes with a range of characteristics (e.g., stress drop), FinDer magnitudes can differ from catalog magnitudes (Böse *et al.*, 2018). In the future, it would be preferable for FinDer to report ground-motion estimates directly.

A second method of magnitude estimation is used when the template-based magnitude is $< M$ 5.5. In these cases, station amplitudes are assigned a phase type (P or S) depending on the expected arrival time, and the Cua and Heaton (2009) ground-motion prediction equation (GMPE) is used to find the predicted magnitude with the lowest misfit to match the amplitudes. Whereas the template-matching method assumes that observations are peak ground motions (i.e., S -wave or later) and tends to underestimate final earthquake magnitude in the earliest alerts, this amplitude regression method facilitates faster estimation of larger magnitudes.

The lowest acceleration amplitude threshold used by FinDer within ShakeAlert 2.0 is 2 cm/s/s, which permits detection of earthquakes with magnitude down to ~ 3.0 in areas of good network coverage. Event detection is controlled within FinDer by trigger parameters that include the number of stations needed to trigger an event and a minimum number of station connections (in which two triggered stations are considered "connected" if their separation is less than a specified maximum distance). In ShakeAlert 2.0, a maximum distance of 50 km permits good event detection performance in regions of moderate or good station coverage, such as the urban areas of Seattle, San Francisco, and Los Angeles, but can lead to a failure of detection in areas of low station coverage. The required number of triggered stations, currently set to four, controls a trade-off between the risk of false alerts and the delay to event detection.

In ShakeAlert 2.0, FinDer is configured to use a set of generic, symmetric ground-motion templates, computed using the GMPE of Cua and Heaton (2009) for a fixed depth of 10 km. This tailors FinDer for determining onshore hazard, for example that posed by major West Coast crustal faults such as the San Andreas and Hayward. However, for offshore seismic regions such as the Cascadia subduction zone or Mendocino triple junction, FinDer needs an extended template set that accounts for asymmetric monitoring and the larger depths of interface and slab events (Böse *et al.*, 2015). This has not yet been implemented and will be a priority for future development.

Because of the limited template set used by FinDer, which results in significant mislocation for earthquakes in out-of-network or sparsely instrumented regions, FinDer cannot generate an alert that is not also declared by EPIC. However, its results can still be used by the SA to determine the alert location, rupture parameters, and magnitude.

SA and DM algorithms

The SA algorithm receives alert messages from the EPIC and FinDer algorithms. The SA combines the point-source information from those two algorithms and sends out an alert message that is received by the DM algorithm. The DM algorithm makes the final decision whether or not to release the alert to end users, based primarily on the SA's determination of whether location and magnitude are within reporting thresholds.

The two algorithms that send alert messages to the SA (“senders”) assign uncertainty values to their point-source estimates. FinDer does this using fixed uncertainty values that were determined from the average algorithm performance for a catalog of real-time and historic events. EPIC’s uncertainty values are from empirical relationships between the EPIC event location errors and the number of stations that are contributing to the event, derived from a comparison of EPIC (ElarmS) event locations with the ANSS catalog for the time period May 2012 to December 2012. The SA algorithm screens the alert messages that it receives using bounds on the acceptable point-source parameters: location, depth, magnitude, origin time (OT), and their uncertainty values. These bounds are simply sanity checks, with the exception of location, which checks for an epicenter within the ShakeAlert reporting regions. The SA associates the event messages and identifies when the differences are within an acceptable limit. Namely, the messages from the two senders will be associated together when their locations are within 100 km and their OTs are within 30 s. The SA computes the combined point-source parameters as the average of the sender parameters weighted by the normalized sender parameter uncertainties. This is computed as $1/((\text{algorithm parameter uncertainty})^2)$, normalized so that the sum of weights for all algorithms equals 1.

The SA algorithm is configured to only accept alerts from senders reporting alerts within an approved region for that sender. This is useful for sender algorithms that are sensitive to station density, such as FinDer, which is assigned a polygon region that approximately encloses the instrument network (i.e., excludes offshore regions). Different magnitude bounds can also be assigned to individual senders. There is the option to require the SA to wait for a confirmation message from a second sender (i.e., second algorithm) when the first sender magnitude is above a threshold value, but this is not currently being implemented.

The DM algorithm receives alert messages from the SA algorithm and from the eqInfo2GM module discussed in the [EqInfo2GM ground-motion estimation algorithm](#) section. The DM has its own set of bounds on the point-source parameters for issuing its alert messages. These bounds can differ from the SA bounds; currently, this is the case for magnitude. The SA is allowed to issue lower magnitude event messages, which are useful for monitoring the system performance, whereas the DM alerts have a higher magnitude threshold. The DM will also pass along the FinDer information that is included in the SA messages to the DM when the event magnitude meets or exceeds a threshold, currently set at **M** 6.0.

EqInfo2GM ground-motion estimation algorithm

The eqInfo2GM algorithm computes estimated ground motion for the earthquake parameters provided by the SA and is a new component for ShakeAlert 2.0. The core calculations involve GMPEs and ground-motion intensity conversion equations (GMICE) to provide estimated PGA, PGV, and shaking

intensity in MMI units. Messages incorporating ground-motion estimates are sent to the DM, which will in turn send out these alerts if they meet various criteria (see the [SA and DM algorithms](#) section).

GMPE and GMICE selections follow those currently used by the regional seismic networks to provide as much consistency as possible with other USGS products (e.g., ShakeMap): [Chiou and Youngs \(2008\)](#) and [Worden *et al.* \(2012\)](#) for the Pacific Northwest and [Boore and Atkinson \(2008\)](#) and [Wald *et al.* \(1999\)](#) for northern and southern California. These implementations closely follow those in ShakeMap v.3.5, as documented in [Worden and Wald \(2016\)](#), though because of the paucity of source characterization information available on the early warning timescale, many terms are unused or simplified (e.g., fault type, aftershock terms; see [Thakoor *et al.*, 2019](#), for all implementation details).

The eqInfo2GM algorithm provides ground-motion information in two formats, contour and grid, which differ in their computation time, message size, resolution, and accuracy. The contour message is the faster, more simplified product, and the system is currently configured to compute a contour for MMI levels of II and above. For a given earthquake magnitude, the region-appropriate GMPE and GMICE are used to compute the distance from the line source (if available) or epicenter at which the specified MMI level is expected, assuming a site condition (V_{S30} value, or average shear-wave velocity in the upper 30 m of soil) of 500 m/s. At regular angular intervals around the source, this distance is converted to coordinates describing a closed, eight-vertex polygon. The maximum size of this message type is on the order of tens of kilobytes, and computation time is on the order of 0.001–0.01 s.

The grid or map message is computed using a fixed grid that covers the ShakeAlert reporting region and has a grid point spacing of 0.2° (~ 20 km). For each grid point, the V_{S30} value (derived from regional knowledge of geologic units, topography, or both) is combined with the distance to source (for line source if available, otherwise epicenter) and earthquake magnitude, to compute the estimated ground motion. The maximum size of this message type is on the order of hundreds of kilobytes, and the computation time is 0.01–1.5 s. This computation time increases with increasing earthquake magnitude because it depends on the grid extent and the number of segments describing the line source; the upper time limit of 1.5 s corresponds to an **M** 7.8 scenario line-source earthquake using the maximum number of grid points (~ 6000).

ActiveMQ message broker and server structure

Communication involving data or message passing between ShakeAlert servers (Fig. S1) is provided by the Apache ActiveMQ open-source software. ActiveMQ acts as a message broker and uses ShakeAlert-defined “topics” to which a server can subscribe to either provide or obtain a message (data) on the topic. For example, the ground-motion information from

the eqInfo2GM algorithm is forwarded to users via a separate ActiveMQ topic from the point-source and line-source alert messages.

In addition to the pair of prealert production servers at each regional seismic network (i.e., Pacific Northwest Seismic Network based in Seattle, Northern California Seismic Network based in Menlo Park and Berkeley, and Southern California Seismic Network based in Pasadena), there is an overlying alert layer that consists of two virtual servers at three of the four sites: USGS–Pasadena, USGS–Menlo Park, and USGS–Seattle (Figs. S1 and S2). As with production, there are two identical servers for failover capability. Although the algorithms and the SA run on the prealert production server layer, the DM runs at the alert layer. The advantages of the alert layer are that it simplifies that part of the system that is responsible for sending alerts directly to the public. Because the alert layer only runs the DM, the ActiveMQ broker, and a heartbeat monitor, the servers are lightweight when it comes to processing and memory requirements. This allows for easy virtualization, which in turn allows for horizontal scaling to accommodate a growing number of user subscriptions. Another benefit of separating the alert layer from the production servers is increased security. By moving the alert servers away from the production servers, external connections to the production servers can be shut down, eliminating all but a few subnets from seeing them on the Internet. To further distance the production servers from the alert servers, the alert servers are configured to become subscribers to ActiveMQ SA brokers on the production hosts. From the production standpoint, the alert server is just an external user subscribing to its data feed; all knowledge of the alert servers and how they relate to production is removed. External alert users only have access to the alert layer. In addition, because the alert servers run only a few required processes, they can be hardened against constant public exposure.

Within the coming year, a proxy layer will be developed and deployed that will replace the ActiveMQ infrastructure at the alert layer. Currently, each ActiveMQ broker can reliably handle ~1000 concurrent connections. This will not be sufficient going forward as more users subscribe to the ShakeAlert data feeds. The new proxy layer will be able to handle the increasingly large number of connections ($\leq 10^6$) and will give users an interface with which to manage their own accounts. Currently, the accounts are created and managed by the ShakeAlert staff, an inefficient way to manage an increasing number of subscribers. The new proxy layer will most likely be established using the Neural Autonomic Transport System open-source, cloud-native messaging system, or some other resilient, high-performance, messaging system.

Security

ShakeAlert uses two-factor authentication for administrative access to the ShakeAlert servers. Each ShakeAlert server

handles access control through Pluggable Authentication Modules for Linux and does not rely on a central authentication server. An independent firm conducts a periodic penetration test that includes vulnerability scanning, brute force login attempts, port scans, and social engineering tests. In addition, the ShakeAlert staff performs routine vulnerability scans and has an aggressive patch management plan. Production servers have limited access and are set up with strict firewall configurations shielding them from the external world. No ports are opened to the public. Connections are only allowed from subnets that contain other ShakeAlert production servers.

System Testing and Performance

Performance of ShakeAlert 2.0—False alerts and missed events

In preparation for public rollout, testing of the ShakeAlert 2.0 algorithms was conducted in ShakeAlert's Testing and Certification Platform (TCP). TCP was developed as a major component of ShakeAlert to test new candidate code and new operating system environments before they are implemented in production (Cochran *et al.*, 2018). Figure S2 shows schematically how the testing process fits into the ShakeAlert architecture. To assess the quality of new candidate code, it is tested against a baseline, which nearly always consists of the performance given by the system running in production at that point in time. This is a quantitative methodology for assessing potential code improvements. The discussion next describes the performance of the ShakeAlert 2.0 test (specifically v.2.0.1) and includes results for historic events, real-time performance, and performance of the system that was in place for significant 2018 and 2019 events. See [Data and Resources](#) for details on data source.

Following the Cochran *et al.* (2018) methodology, the first component of testing involves comparing the ShakeAlert point-source solution with the ANSS comprehensive catalog (ComCat) solution for earthquake source parameters (hypo-centers, magnitudes, and OTs). During testing, an earthquake solution match is declared (as opposed to a false alert) if the magnitude difference between ShakeAlert and ComCat is within 2.0 magnitude units, the location estimate is within 100 km, and the OT estimate is within 15 s. In addition, timeliness is quantified by looking at the alert time associated with the largest distance traveled by an *S* wave to a contour with MMI = IV to assess whether the alert was early enough to be useful. This metric depends on the distance between the MMI = IV contour and the epicenter, and constant *S*-wave velocity. It is thus complementary to additional assessments of ground motion discussed in the [Ground-motion assessment](#) section.

In the historic event test of ShakeAlert 2.0, four instances of the algorithms including EPIC, FinDer, SA, DM, and eqInfo2GM were run in parallel with each other. The historic data suite consists of 60 earthquakes, both mainshocks and aftershocks, that occurred in California and the Pacific

Northwest, mostly inside the ShakeAlert reporting region. It also includes 10 events, most of which were regional earthquakes, that had occurred outside the ShakeAlert reporting region; 25 sensor calibration or recentering events; 10 extremely noisy channel events; and 20 deep or large-magnitude teleseismic events. See Cochran *et al.* (2018), supplementary section, for a complete list of historic events used. Figure 1b shows the earthquakes that comprise the West Coast earthquake subset of the historic event test suite. For each earthquake, the algorithms produce multiple alerts that evolve with time and are updated as new station data arrive at the production servers during the earthquake. The testing process examines both the first (earliest) alert, as well as the later “best” alert (best overall performance in terms of accuracy of estimating source parameters). The false alert and missed event analysis is unchanged whether we use the first alert or best alert. Averages are computed for four parallel test instances because there are slight performance differences among the instances due to nondeterministic behavior related to multithreading; see Cochran *et al.* (2018), for more details on this. Thus, the results are sometimes noninteger because they are averages from the multiple instances.

For the 60 M 3.5+ earthquakes in the test suite, ShakeAlert 2.0 produced an average of 4.5 and a maximum of 6.0 false alerts. Therefore, for the M 3.5+ earthquakes, the DM false alert rate was 7.5%–10%. For the 25 M 5+ earthquakes, ShakeAlert 2.0 produced an average of 1.25 and a maximum of 2.0 false alerts, yielding a 5%–8% false alert rate. The DM missed event alert rate was based on the parallel test instance that had the largest number of missed events. This instance missed 22 M 3.5+ events (of 60), resulting in a DM missed event rate of 36.7% (18 of the missed events were aftershocks or secondary events in an earthquake sequence, and some were offshore or near the borders of the alerting region). This same instance also missed four M 5+ events (of 25), resulting in a DM M 5+ missed event rate of 16%. Thus, in summary, for the historic event ShakeAlert 2.0 test, the average DM M 5+ false alert and missed event rates are 8% and 16%. The DM M 3.5+ false alert and missed event rates are 10% and 36.7%.

Upon looking more in depth at the false alerts and missed events, two of the DM false alerts arose from EPIC incorrectly locating an earthquake. This occurred for the 10 March 2014 M 6.8 offshore Eureka, California, earthquake, and the 6 May 2005 M 5.2 earthquake near Anza, California. For the offshore Eureka event, three of the four parallel EPIC instances resulted in poor source parameters, generating two small-magnitude false alerts. This was not unexpected; locating offshore earthquakes is notoriously difficult because of poor azimuthal seismic ray-path coverage. For the Anza event, EPIC produced an incorrect OT and location because of a combination of poor azimuthal coverage and sparse station spacing near the epicenter.

Three of the four missed M 5+ events occurred at the boundaries of the ShakeAlert reporting region. This includes two

offshore events: the 10 March 2014 M 6.8 offshore Eureka, California, earthquake, and the 24 April 2014 offshore Vancouver Island, British Columbia, event. For the 4 April 2010 M 7.2 El Mayor–Cucapah event, an M 5.3 aftershock occurred a little more than 2 min after the mainshock but was missed by ShakeAlert 2.0. Moreover, the El Mayor–Cucapah SA/DM alerts were delayed 5 s compared with EPIC’s alert time because the earthquake initiated outside the reporting region. This illustrates in general how only using the epicenter in the triggering criteria can delay alerts for large earthquakes that may initiate outside the reporting region. The fourth missed M 5+ event was the Anza, California, earthquake discussed earlier. Finally, an M 3.6 DM false alert was produced from a calibration test signal, and an M 4.8 false alert was caused by a drift in OT in source parameters produced by FinDer because it incorporated noise and aftershock energy into updates of the mainshock alert.

The historic event test suite includes many older events recorded by generally lower density seismic networks. It also includes more recent events that have occurred along the entire West Coast in areas of varying station density, which is currently still the case in many regions. Sparser station coverage can contribute toward the missed event rate.

Historic teleseismic earthquakes also produced false alerts in the ShakeAlert 2.0 tests. Three DM false alerts were caused by the 17 November 2002 M 7.3 Kuril Islands earthquake, which generated an M 5.3 false alert in EPIC; the 29 May 2015 M 6.8 Alaska Peninsula earthquake, which generated an M 3.8 false alert in EPIC; and the 30 May 2015 M 7.8 deep earthquake south of Japan, which generated an M 5.3 false alert in EPIC. For ShakeAlert 2.0, these false alerts might be prevented in real time using the algorithm “teletistle,” which works by accessing real-time teleseismic earthquake information from the USGS National Earthquake Information Center (NEIC) to suppress teleseismic triggers (Chung *et al.*, 2019). Because it is a real-time process, it cannot be run in the historic event test, but it is run in ShakeAlert 2.0 production.

In preparation for ShakeAlert 2.0 deployment, real-time tests were run for the six-week time period 3 August 2018 to 17 September 2018. The tests were run on real-time test servers located in Pasadena, Menlo Park, Berkeley, and Seattle (Fig. S2). There was a total of eight ANSS catalog M 3.5+ events. The results show a distinct decrease in the number of false alerts for ShakeAlert 2.0 compared with ShakeAlert 1.0. The number of false alerts decreased from 22 to 2, primarily because of the removal of Onsite. However, there were also two missed events that arose from the slightly higher DM configuration minimum threshold of M 3.5 (the minimum was 3.0 for ShakeAlert 1.0). There are, however, only seven matched events with M 3.5+ in common between ShakeAlert 2.0 and 1.0 and very few earthquakes in total for this time period (eight with M 3.5+), so these comparisons are based on small numbers of events.

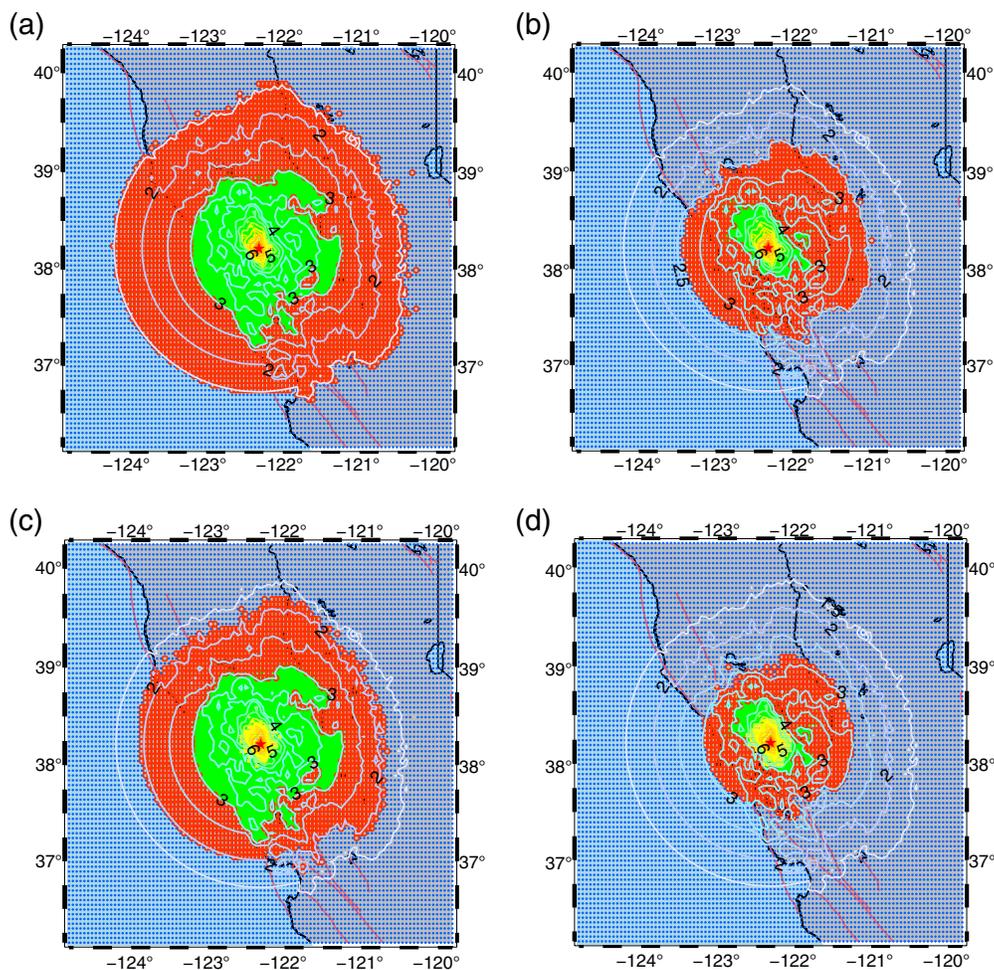


Figure 2. Ground-motion assessment test results for Earthquake Point-Source Integrated Code (EPIC) and Finite-Fault Rupture Detector (FinDer) compared with observed ShakeMap for the 24 August 2014 M 6.0 South Napa, California, earthquake. (a) EPIC output for modified Mercalli intensity (MMI) = III, (b) EPIC output for MMI = IV, (c) FinDer output for MMI = III, and (d) FinDer output for MMI = IV. Maps show the distribution of false-negative alerts (yellow region centered on the epicenter), true-positive alerts (green region), false-positive alerts (red region), and true-negative sites (region indicated by small blue dots). Contours show ShakeMap MMI from observed data. Epicenter is shown by red star. Tests do not include latencies due to data transmission or alert delivery. The color version of this figure is available only in the electronic edition.

Ground-motion assessment

The second major component of testing includes assessments of how well the ground motion is estimated by the different ShakeAlert algorithms. This component examines the combination of algorithm source parameters, GMPE uncertainty, and alert thresholds on the performance of EEW alerting. For this test, we followed the method of Cochran *et al.* (2018) to convert EPIC and FinDer source parameter estimates from the test with historic earthquakes (excluding all data latencies) into ground-motion estimates using the ShakeMap 3.5 software (Worden and Wald, 2016). An example is shown in Figure 2 for the 24 August 2014 M 6.0 South Napa, California, earthquake, which shows a gridpoint-by-gridpoint comparison between

predicted and observed ShakeMaps that have been calculated using the same GMPEs, with grid points classified by their alerting success. “Predicted” ShakeMaps are computed using the ShakeAlert algorithm’s source parameters and do not include ground-motion observations. “Observed” ShakeMaps are computed using catalog source parameters and do include all available ground-motion observations. Performance is assessed at grid points with 0.05° spacing and is a combination of the first alert and all updates. This test examines the accuracy of the predicted shaking intensity, accounting for whether the alert arrived with enough time to be useful by comparing the alert time with the estimated S -wave arrival time for each location. Figure 2 illustrates the performance of ground shaking alerting at different MMI threshold levels individually for EPIC and FinDer. Figure 2a and 2b shows alerting performance for EPIC for thresholds MMI = III and IV. Figure 2c,d shows the analogous results for FinDer, which are similar because the point-source estimates produced by FinDer for this dataset are similar to EPIC. In Figure 2, “true-positive” alert regions are those

where the alerts were both timely (before the estimated S -wave arrival time) and accurate for the specified MMI level. “False-negative” alert regions are where the alerts were either late because they occurred after the estimated S -wave arrival time or were missed because the predicted MMI level was too low (underpredicted). The “false-positive” regions are where the alerts overpredicted the MMI level. Finally, “true-negative” regions are where the system correctly generated no alert (neither the estimated shaking nor observed shaking exceeded the specified MMI level). In this test, EPIC provided the first alert at 5 s after OT and FinDer at 5.5 s after OT.

Secondary tests include the validation of the eqInfo2GM GMPE/GMICE implementations (comparisons of

ShakeMaps that do not consider observations against the eqInfo2GM-produced ground-motion map) and validation of the full ShakeAlert ground-motion map accuracy (a comparison of ShakeMap with observations, e.g., from USGS-NEIC, against the ShakeAlert-produced ground-motion map). For a more detailed analysis of these additional validation tests, see the supplemental material and Figures S3 and S4. Ideally, these approaches would eventually be combined so that thresholding uses ShakeAlert-produced maps and not ShakeMaps.

Performance during 2018 and 2019 earthquakes

We selected six representative earthquakes that occurred in 2018–2019 to give an indication of how ShakeAlert performs, and can be expected to perform in the future under certain conditions. Some of these earthquakes resulted in alerts produced only by ShakeAlert 1.0 because ShakeAlert 2.0 was not in production until late 2018. However, the behavior of the system is illustrative of the various ways in which alerts can be expected to be produced, depending on where the earthquakes occur with respect to station density. The alerts, described next in reverse chronological order, indicate how the system performs when the network station density is high, when it is sparse, and when the earthquake occurs offshore.

The most recent assessment of system performance for a reasonably well-instrumented region of the Mojave Desert in southern California was demonstrated by the 6 July 2019 M 7.1 Ridgecrest, California, earthquake. ShakeAlert 2.0 reported a first alert 8.0 s after OT. The radius of the zone shaken by the S wave before receiving the first alert was 26 km. The ANSS catalog source depth was 8 km, but note that for all ShakeAlert solutions, the depth is fixed to 8 km by EPIC and 10 km by FinDer. The first reported magnitude was 5.5, reaching a magnitude of 6.3, reported by both EPIC and FinDer. The evolution of magnitude estimates produced by each algorithm is shown in Figure 3a. Other observations included a maximum observed instrumental intensity MMI level of IX in the epicentral region and MMI VII over a 40 km wide region near the epicenter. More details on this earthquake’s performance will be presented in future studies because the data analysis of this event is still ongoing.

Less than a month prior, the production system was tested during the 23 June 2019 M 5.6 earthquake that occurred near Petrolia, California, on the northern California coast, where station density is lower and the source–station azimuthal gap is large. ShakeAlert 2.0 produced its first alert at 7.0 s after OT, created by EPIC, but the initial magnitude was overestimated at M 6.3. FinDer created an alert at 8.7 s after OT with a more accurate magnitude of M 5.5 but with larger location error. The evolution of magnitude estimates is shown in Figure 3b. The ANSS catalog depth was 9.4 km.

An example of system behavior for well-instrumented regions is illustrated by the 29 August 2018 M 4.4 earthquake

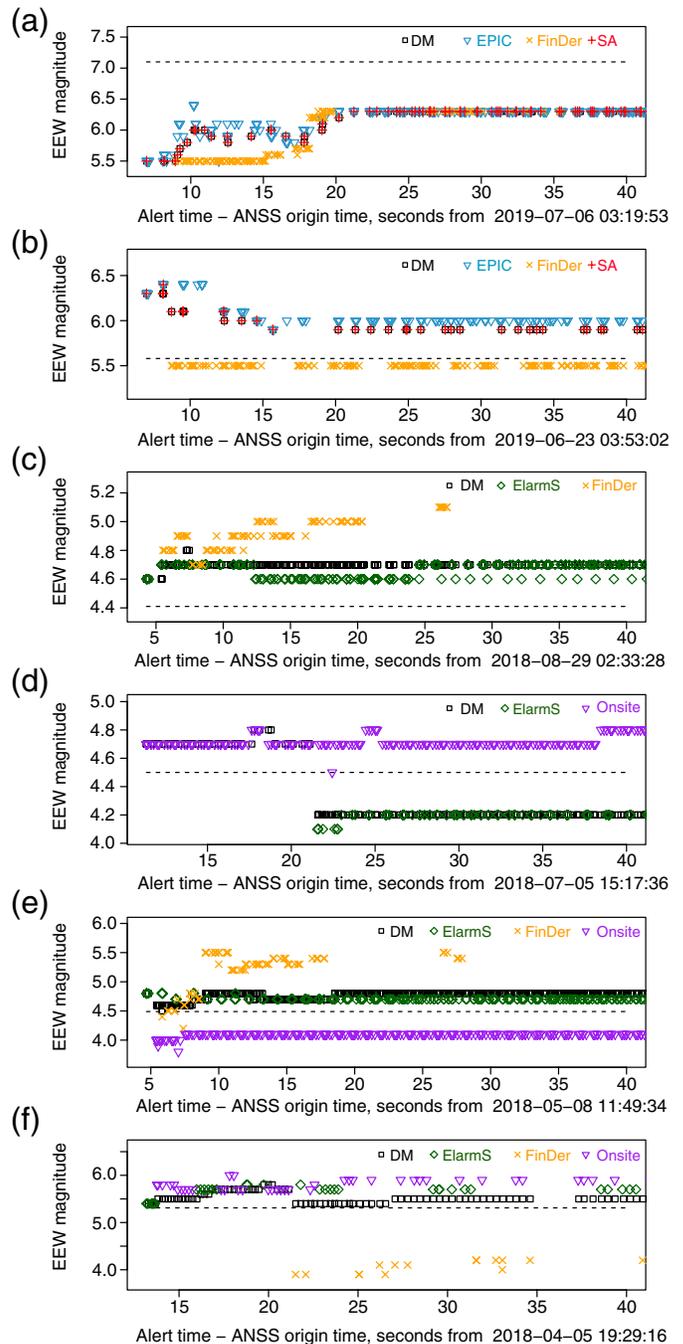


Figure 3. Evolving magnitude estimates made by EPIC, FinDer, Earthquake Alarm Systems (ElarmS), Onsite, Solution Aggregator (SA), and Decision Module (DM) for the following earthquakes discussed in the [Performance during 2018 and 2019 earthquakes](#) section: (a) 6 July 2019 M 7.1 Ridgecrest, California; (b) 23 June 2019 M 5.6 near Petrolia, California; (c) 29 August 2018 M 4.4, near La Verne, California; (d) 5 July 2018 M 4.5 earthquake 17 km west-northwest of Sandy Valley, Nevada; (e) 8 May 2018 M 4.5 earthquake 11 km north of Cabazon, California; (f) 5 April 2018 M 5.3 earthquake 27 km southwest of Santa Cruz Island, California. Algorithms shown are those that were in testing or production during the time of the earthquake. EEW, earthquake early warning. The color version of this figure is available only in the electronic edition.

that occurred 4 km north of La Verne, California. This part of southern California is relatively densely instrumented. ShakeAlert 1.0 reported a first alert 4.30 s after the OT. The radius of the zone shaken by the *S* wave before receiving the first alert was 14.0 km, and the first reported magnitude was 4.6. The time for the *P* wave to reach the surface was 1.0 s (ANSS catalog source depth was 5.5 km). On the test system where ShakeAlert 2.0 was being tested, EPIC reported an **M** 4.6 at 4.5 s after OT, and FinDer reported an **M** 4.3 at 4.7 s after OT. The evolution of magnitude estimates is shown in Figure 3c. Other observations included a maximum observed instrumental intensity MMI level of VI (14.3%g).

An example of system behavior for an epicentral region that is onshore but not densely instrumented is the 5 July 2018 **M** 4.5 earthquake that occurred 17 km west-northwest of Sandy Valley, Nevada. This earthquake occurred on the California–Nevada border in a region that is sparsely instrumented. ShakeAlert 1.0 reported a first alert 11.4 s after OT. The radius of the zone shaken by the *S* wave before receiving the first alert was 39.4 km because of the sparsity of the seismic network in the source region. The first reported magnitude was 4.7; the evolution of magnitude estimates is shown in Figure 3d. The time for the *P* wave to reach the surface was 1.4 s. (ANSS catalog source depth was 8.6 km.) The point-source algorithms produced source parameter estimates, but FinDer did not produce a solution because its checks for “connected” (i.e., nearby) trigger stations failed.

Similarly, on 8 May 2018, an **M** 4.5 earthquake occurred 11 km north of Cabazon, California, in a region of southern California where the seismic network becomes sparser. ShakeAlert 1.0 reported a first alert 4.1 s after the OT. The first reported magnitude was 4.8, and the time for *P* waves to reach the surface was 2.4 s (ANSS catalog source depth, 12.9 km). The evolution of magnitude estimates is shown in Figure 3e. The radius of the zone shaken by *S* waves before receiving the first alert was 1.9 km. This small zone benefited from the fact that the 12.9 km depth delayed the receipt of the *S* waves at the closest seismic stations.

Finally, insight into system behavior for offshore events is illustrated by the 5 April 2018 **M** 5.3 earthquake that occurred 27 km southwest of Santa Cruz Island in the offshore region of southern California. ShakeAlert 1.0 reported a first alert 13.1 s after OT. This relatively long time is due to the lack of enough reporting stations (including those on islands) between the hypocenter and the mainland and the ElarmS requirement of four stations to trigger. The radius of the zone shaken by *S* waves before receiving the first alert was 40.5 km. The first reported magnitude was 5.4, and the time for the *P* wave to reach the surface was 1.7 s (ANSS catalog source depth, 9.9 km). The evolution of magnitude estimates is shown in Figure 3f. There was a simultaneous false **M** 3.8 event mislocated into central California due to a split (i.e., a second simultaneous associated event) produced by EPIC.

ShakeAlert User Engagement

In July 2016, the ShakeAlert JCCEO was formed to provide well-informed feedback regarding the ShakeAlert system’s human interface. Since its founding, JCCEO has coordinated the development of the resources that have been elemental to successful and sustained user uptake of the system to maximize life-safety and property protection impacts. Some of the primary roles of JCCEO are to recruit ShakeAlert technical users, and to provide a consistent and accessible forum for the discussion and coordination of implementation plans. These plans are customized for various uses (automated, organizational, and public) across the ShakeAlert system and for the general public. The intention of the technical engagement program is to work with partners to develop real-world ShakeAlert applications that trigger automated actions and/or trigger human-based protective actions (e.g., drop, cover, and hold on). In addition, JCCEO is recruiting a class of partners that are collectively called “Enablers” to develop hardware, software, or processes that other users can use “off the shelf” rather than having to develop unique solutions for each application.

In 2017, JCCEO and ShakeAlert stakeholders chose five high-priority and high-impact focus sectors for technical engagement. These focus sectors represented ShakeAlert technical partners in the corporate, governmental, and educational sectors that would use ShakeAlerts to facilitate public safety, emergency response, training, and education but without actually issuing alerts to the general public. These focus sectors were emergency management, transportation, utilities, educational institutions, and health care, with the goal of bringing numerous ShakeAlert Pilot Projects and Enablers to Market by 2018. As of September 2019, ShakeAlert had >60 technical partners in >12 different sectors (e.g., transportation, mass notification, and health care), and several of them have conducted large-scale tests of slowing trains (Bay Area Rapid Transit, Los Angeles Metro), automatically controlling valves, gates, and power generators in water and power systems (RH2 Engineering in the Pacific Northwest; Eugene Oregon Water and Electric Board), alerting residents of a 224-unit high-rise apartment (in Marina Del Rey, California), and release of emergency doors (Regatta Seaside in Marina Del Rey, California). In the Pacific Northwest, numerous water and sewer districts in Washington and Oregon are developing ShakeAlert Pilot Projects. The first completed project developed by northeast Sammamish Water and Sewer district with support from RH2 Engineering was completed in March 2018. In California, AT&T collaborated with the City of Los Angeles to develop an app (ShakeAlertLA) that was rolled out at the end of 2018 and is currently undergoing testing with the residents of Los Angeles County. JCCEO is providing recommendations for the content and messaging (e.g., the short message on a cell phone screen prompting an immediate protective action) and is providing ongoing support as the app is updated

and improved. The Bay Area Rapid Transit system is working on messaging in transit stations.

Several technical partners have developed apps that are delivering alerts to test populations within the general public. These include SkyAlert, UC Berkeley (MyShake), and Early Warning Labs (QuakeAlert). All of the app developers are conducting systematic and coordinated testing to reduce latencies. In addition to addressing technical aspects of ShakeAlert delivery, partners are working with the JCCEO to provide appropriate education and training to those receiving an alert. The JCCEO has partnered with all of the app developers to develop, review, and evaluate education and training content on each of the apps, which includes the actual alert message(s), recommended protective actions, and pre-education that users get when they first download the app and start to use it. JCCEO is not in the position to mandate that any particular content is used on apps but draws its authority from the literature and from experienced practitioners.

The ShakeAlert community is genuinely interested in promoting practices that will save lives and protect property. The major goal of the JCCEO-led ShakeAlert technical engagement program is to promote the development of a sustainable EEW industry. Current technologies for mass notification via cell phones, TV, and radio have not been demonstrated to be fast enough (i.e., delivery of alerts to 95% of users in 5 s or less) for any region in California or cannot handle the volume of messages needed for effective early warning. The USGS is working with alert delivery organizations to address this issue. Speeding up mass delivery is being addressed by FEMA, cellular telephone companies, and others with encouraging developments on this front. Despite the need to improve latency for alert delivery, the USGS and the State of California proceeded with ShakeAlert Rollout Phase 3: Public Alerting in California, when alerts can be distributed statewide via the WEA system and apps. On 17 October 2019, California Governor Gavin Newsom announced that statewide alerting would commence while also making it clear that EEW was still in a testing phase. The JCCEO in partnership with the USGS Office of Communications and Publishing will continue to message what end users can expect from WEA and apps. At the time of the announcement of public alerting in California, the emergency management agencies in Oregon and Washington announced their intention to implement statewide alerting by ShakeOut day 2020 (15 October 2020).

Conclusions and Future Directions

ShakeAlert will continue to evolve as components of the system are improved, expanded on, and modified based on information provided by new earthquakes that produce alerts. There are now 900 stations in parts of Southern California, the San Francisco Bay area (including Silicon Valley), and Seattle–Tacoma that are contributing to the system. New construction has focused on seismically high-risk areas, but

existing stations and telemetry will continually be upgraded as well on an as-needed basis. Future additions may include new geodetic data to provide tighter constraints on large earthquake rupture processes through the contribution of large-amplitude displacement data. Improvements in telemetry will not only add redundancy in data pathways but will also increase the speed with which data are streamed from the field to the central processing sites.

The quantitative testing procedure of ShakeAlert also continues to evolve and has new ground-motion metrics that involve comparisons with ShakeMaps. The test events database will eventually include a comprehensive suite of M 3.5+ West Coast earthquake data for the past few years, as well as large-magnitude Japanese earthquake datasets to test algorithms on larger events in a dense network. Finally, synthetic ground-motion data covering San Andreas fault, Hayward fault, and Cascadia scenario earthquakes is a long-term goal for inclusion in the test event suite.

The new ShakeAlert 2.0 algorithms produce both point-source and line-source earthquake solutions, as well as new ground-motion metrics products. ShakeAlert 2.0 testing shows that both the false alert rate and missed alert rates have dropped relative to the prototype version. The algorithms continue to be subject to modifications that increase the accuracy of alert estimates, decrease the number of false alerts, and improve the accuracy of ground-motion estimates.

Data and Resources

All historic event test data used here can be found at <http://scedc.caltech.edu/research-tools/eewtesting.html>. Some plots were made using the Generic Mapping Tools v.4.2.1 (www.soest.hawaii.edu/gmt/; Wessel and Smith, 1998). Open-source Neural Autonomic Transport System (NATS) messaging system software can be downloaded from <https://nats.io>. The ShakeMap scenario data are available from <https://earthquake.usgs.gov/scenarios/catalog>. The Hayward scenario data can be found at https://earthquake.usgs.gov/scenarios/eventpage/nclegacyhaywardroddgerscreekcrhnhsm7p3_se#shakemap. The San Andreas scenario data are from https://earthquake.usgs.gov/scenarios/eventpage/sclegacyspsanandreasbbnmsmnsbssbbgcom7p9_se#shakemap. The Cascadia scenario data are from https://earthquake.usgs.gov/scenarios/eventpage/bssc2014cascadia_sub0_m9p34_se#shakemap. The other relevant data are from ShakeAlert.org and MyShakeAlert.org. All websites were last accessed in March 2020. Supplemental material contains four figures and a detailed description of ground-motion assessment tests.

Acknowledgments

Funding for ShakeAlert has been provided primarily by the U.S. Geological Survey (USGS), California Governor's Office of Emergency Services (Cal OES), the State of Oregon, and the Gordon and Betty Moore Foundation. The authors are grateful for the useful and comprehensive reviews of this article provided by Sarah Minson and Debi Kilb. Credit goes to Elijah Marchese for creating Figure S4. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

References

- Allen, R. M. (2007). The ElarmS earthquake early warning methodology and application across California, in *Earthquake Early Warning Systems*, Vol. 2003, 21–43.
- Allen, R. M., and H. Kanamori (2003). The potential for earthquake early warning in Southern California, *Science* **300**, 786–789.
- Boore, D. M., and G. M. Atkinson (2008). Ground-motion prediction equations for the average horizontal component of PGA, PGV, and 5%-damped PSA at spectral periods between 0.01 s and 10.0 s, *Earthq. Spectra* **24**, no. 1, 99–138, doi: [10.1193/1.2830434](https://doi.org/10.1193/1.2830434).
- Böse, M., C. Felizardo, and T. H. Heaton (2015). Finite-fault rupture detector (FinDer): Going real-time in Californian ShakeAlert Warning System, *Seismol. Res. Lett.* **86**, no. 6, 1692–1704.
- 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.
- Böse, M., T. Heaton, and E. Hauksson (2012a). 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.
- Böse, M., T. Heaton, and E. Hauksson (2012b). Real-time finite fault rupture detector (FinDer) for large earthquakes, *Geophys. J. Int.* **191**, 803–812.
- Böse, M., D. E. Smith, C. Felizardo, M.-A. Meier, T. H. Heaton, and J. F. Clinton (2018). FinDer v.2: Improved real-time ground-motion predictions for M2-M9 with seismic finite-source characterization, *Geophys. J. Int.* **212**, 725–742.
- Bradski, G., and A. Kaehler (2013). *Learning OpenCV: Computer Vision in C++ with the OpenCV Library*, O'Reilly Media, Inc., 575 pp.
- Chiou, B. S.-J., and R. R. Youngs (2008). An NGA model for the average horizontal component of peak ground motion and response spectra, *Earthq. Spectra* **24**, no. 1, 173–215, doi: [10.1193/1.2894832](https://doi.org/10.1193/1.2894832).
- Chung, A. I., I. Henson, and R. M. Allen (2019). Optimizing earthquake early warning performance: ElarmS-3, *Seismol. Res. Lett.* **90**, no. 2A, 727–743.
- Cochran, E. S., M. D. Kohler, D. Given, S. Guiwits, M.-A. Meier, M. Ahmad, I. Henson, J. Andrews, and R. Hartog (2018). Earthquake early warning ShakeAlert system: Testing and certification platform, *Seismol. Res. Lett.* **89**, no. 1, 108–117, doi: [10.1785/0220170138](https://doi.org/10.1785/0220170138).
- Cua, G. B., and T. H. Heaton (2009). Characterizing average properties of southern California ground motion amplitudes and envelopes, *Earthquake Engineering Research Laboratory Report No. 2009-05*, Pasadena, California, available at <http://resolver.caltech.edu/CaltechEERL:EERL-2009-05> (last accessed January 2019).
- 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.
- Federal Emergency Management Agency (FEMA) (2008). *HAZUS-MH Estimated Annualized Earthquake Losses for the United States*, Vol. 366, FEMA, Washington, D.C.
- Given, D. D., R. M. Allen, A. S. Baltay, P. Bodin, E. S. Cochran, K. Creager, R. M. de Groot, L. S. Gee, E. Hauksson, T. H. Heaton, et al. (2018). Revised technical implementation plan for the ShakeAlert system—An earthquake early warning system for the West Coast of the United States, *U.S. Geol. Surv. Open-File Rept. 2018-1155*, 42 pp., doi: [10.3133/ofr20181155](https://doi.org/10.3133/ofr20181155).
- Given, D. D., E. Cochran, T. Heaton, E. Hauksson, R. Allen, M. Hellweg, J. Vidale, and P. Bodin (2014). Technical implementation plan for the ShakeAlert production prototype 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](https://doi.org/10.3133/ofr20141097).
- Kanamori, H., P. Maechling, and E. Hauksson (1999). Continuous monitoring of ground-motion parameters, *Bull. Seismol. Soc. Am.* **89**, 311–316.
- 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](https://doi.org/10.1785/0220170140).
- Thakoor, K., J. Andrews, E. Hauksson, and T. H. Heaton (2019). From earthquake source parameters to ground motion warnings near you: The ShakeAlert earthquake information to ground motion (eqInfo2GM) method, *Seismol. Res. Lett.* **90**, no. 3, 1243–1257, doi: [10.1785/0220180245](https://doi.org/10.1785/0220180245).
- U.S. Geological Survey (USGS) (2019). ANSS Participation Policy, Version 2.0, available at https://prd-wret.s3.us-west-2.amazonaws.com/assets/palladium/production/atoms/files/ANSS_Participation_Policy_2019_1.pdf (last accessed March 2020).
- U.S. Geological Survey (USGS) (2017). Advanced National Seismic System—Current status, development opportunities, and priorities for 2017–2027 (ver. 1.1, July 2017), *U.S. Geol. Surv. Circular 1429*, 32 pp., doi: [10.3133/cir1429](https://doi.org/10.3133/cir1429).
- Wald, D. J., V. Quitoriano, T. H. Heaton, and H. Kanamori (1999). Relationships between peak ground acceleration, peak ground velocity, and modified Mercalli intensity in California, *Earthq. Spectra* **15**, no. 3, 557–564, doi: [10.1193/1.1586058](https://doi.org/10.1193/1.1586058).
- Wells, D. L., and K. J. Coppersmith (1994). New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement, *Bull. Seismol. Soc. Am.* **84**, 974–1002.
- Wessel, P., and W. H. F. Smith (1998). New, improved version of the Generic Mapping Tools released, *Eos Trans. AGU* **79**, 579.
- Wood, H. O., and F. Neumann (1931). Modified Mercalli intensity scale of 1931, *Bull. Seismol. Soc. Am.* **21**, 277–283.
- Worden, C. B., and D. J. Wald (2016). *ShakeMap Manual Online: Technical Manual, User's Guide, and Software Guide*, U.S. Geological Survey, doi: [10.5066/F7D21VPQ](https://doi.org/10.5066/F7D21VPQ).
- Worden, C. B., M. C. Gerstenberger, D. A. Rhoades, and D. J. Wald (2012). Probabilistic relationships between ground-motion parameters and modified Mercalli intensity in California, *Bull. Seismol. Soc. Am.* **102**, no. 1, 204–221, doi: [10.1785/0120110156](https://doi.org/10.1785/0120110156).
- Wurman, G., R. M. Allen, and P. Lombard (2007). Toward earthquake early warning in northern California, *J. Geophys. Res.* **112**, no. 8, 1–19.

Manuscript received 4 September 2019