THE COMMUNITY SEISMIC NETWORK FOR DENSE, CONTINUOUS MONITORING OF GROUND AND STRUCTURAL STRONG MOTION

Monica Kohler[1] and the Community Seismic Network Team

[1]Department of Mechanical and Civil Engineering
California Institute of Technology, Pasadena, CA

Abstract

The Community Seismic Network (CSN) is a cloud-based, strong-motion network of seismic stations deployed in the greater Los Angeles area. The sensors report three-component acceleration time series data and peak acceleration scalar data for use in assessments of earthquake shaking intensity in buildings and on the ground level, monitoring structural health of instrumented buildings, zonation maps of future shaking potential, and the ShakeAlert earthquake early warning system. The hardware and software behind CSN's client and server architecture are described, as well as network subarrays deployed at Los Angeles Unified School District campuses, the NASA-JPL campus, and in mid-rises and high-rises.

Introduction

This paper describes the architecture of the Community Seismic Network (CSN), a permanent strong-motion seismic network. CSN has been developed over the past 10 years by a team of scientists, listed in the Acknowledgments section of this paper, whose work is represented here. CSN hardware comprises commercially produced MEMS accelerometers that are coupled with processors, external storage, and power supply. CSN consists of over 700 accelerometers that are deployed in mid-rises and high-rises, school campuses, civic service buildings, and homes in California (Fig. 1). The CSN project has increased the number of 3D (ground level+all upper floors of buildings) seismic observations in greater Los Angeles by an order of magnitude, by taking advantage of advances in small-form-factor MEMS sensing technologies, on-site computing, and cloud infrastructure. The mission of CSN is to: 1) Provide high spatial resolution assessments of shaking intensity in buildings and on the ground following major earthquakes; 2) Monitor the health and safety of structures through detection and location of damage; 3) Create zonation maps of future shaking potential in populated areas; and 4) Provide data for the ShakeAlert earthquake early warning system (Given et al., 2014, 2018; Kohler et al., 2020).

The Community Seismic Network (CSN) currently comprises hundreds of stations located in southern California, most of which are in the greater Los Angeles area (Clayton et al., 2011, 2015, 2020; Kohler et al., 2013, 2014, 2018; Massari et al., 2017). The accelerometers are triaxial, and capable of recording accelerations up to twice the level of gravity. The primary product of the network is measurements of shaking of the ground as well as upper floors in buildings, in the seconds during and following a major earthquake. Each sensor uses a small, dedicated ARM processor computer running Linux, and analyzes time series data in real time at 250 sps, which then is downsampled to 50 sps for data storage purposes.

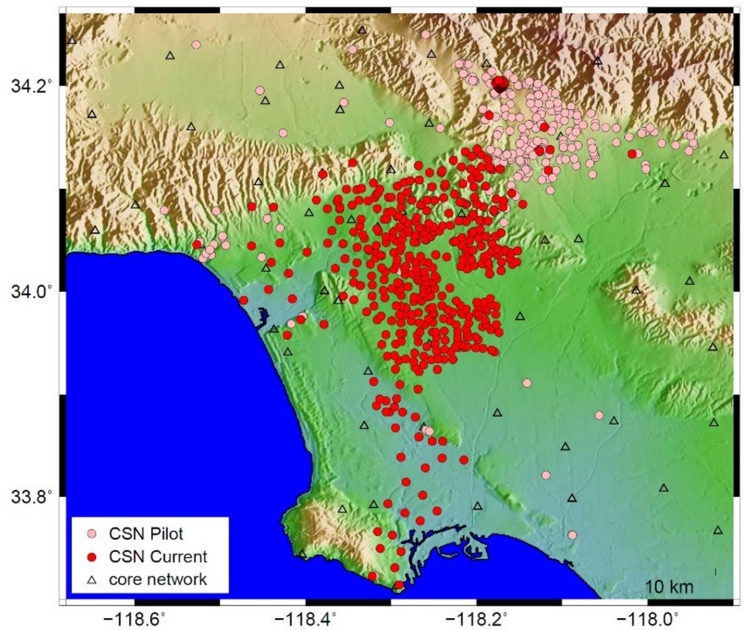

Figure 1. Seismic stations in southern California. CSN stations: circles. SCSN stations: triangles.

CSN client architecture

Hardware

CPU

CSN clients have primarily been deployed using low-power, "single-board-computer" (SBC) platforms since 2012. Such platforms typically have a physically small form factor and consume under 20 W of power, making them suitable for use with low-wattage battery backup units. They also have between one and four USB 1.0 ports, and at least one ethernet port with 10 Mbps or greater data transmission rates. Some platforms have a USB serial port suitable for on-site connectivity without disturbing an operating client. Other platforms have wireless radios, such as Bluetooth or WiFi, that provide on-site access without having to displace the active network connection.

The most recent large CSN deployment in early 2020 used 100 Raspberry Pi Model 4B units, housed in passive cooling aluminum cases. The previous large CSN deployment in early 2019 used 200 Raspberry Pi Model 3B units, also housed in passive cooling plastic cases. Board failures have not been observed in the ~2.5 years after deployment, despite the reliance on small chip-size heat sinks used in the absence of aluminum cases to serve as heat sinks. A small number (~5) of Raspberry Pi Model 3B+ units have also been deployed; these were originally for laboratory use but were later migrated into the field to meet deployment goals. Each of the Raspberry Pi units allows for the use of SD cards for a moderate amount (about three months) of on-site storage.

CSN's first SBC platform beginning in early 2012 was the Global Scale Technologies SheevaPlug, which includes a single USB 2.0 port, a single RJ-45 1 Gbps ethernet port, 512 MB internal flash drive, and an external SD card slot. It also provides a mini-USB port for serial console access. Over 400 of these are still deployed and operating at CSN stations in the field. The SheevaPlug has a single USB mini-A port configured so that it can be connected to a laptop or desktop USB A port via USB serial interface. MacBook Pros and most Windows-compatible laptops need no special USB driver software to communicate with a SheevaPlug; however earlier MacOS systems may need a dedicated driver for the SheevaPlug.

Accelerometer

Through 2021, CSN has used only one type of accelerometer family to provide the base acceleration data - the phidgets.com Model 1056, based on the STMicroelectronics analog MEMS LIS344ALH triaxial inertial sensor with ± 2g range. Initial CSN stations used the original 1056 model that included a compass and gyroscope, but most CSN stations now use the model 1056-1 which only has the triaxial accelerometer.

Power, battery backup, and power boards

A combination of power options attempt to provide CSN stations with a few hours of backup power. In some cases, emergency power is available from the sensor host enterprise; in other cases, local enterprise-provided Uninterruptible Power Supply (UPS) units are available for use. When reliable backup power is not available, small battery backup units are added to the setup.

The canonical UPS deployed to date is a CyberPower CP350SLG unit which typically provides a couple of hours of standby power. This unit has a form factor that neatly fits into the external CSN box packaging that has been deployed since the outset, but the units lack an interface that would support status monitoring. One issue with these units is the limited lifetime (2-3 years) of the internal sealed lead acid battery. In addition, while a single battery replacement cycle works well, a second replacement cycle is less likely to succeed. At the third replacement point, the unit is scrapped.

Early SheevaPlugs were notorious for premature failure of their internal 120 VAC-to-5 VDC power supply board. Now, after 5-8 years of deployment, CSN's SheevaPlugs are experiencing increasing rates of power board failures. As of early 2021, these power boards are no longer available from the vendor, so a suggestion from the user community was adopted in which failed boards are replaced with a generic 2.5Amp power adapter wall wart, and the factory output plug is replaced with a harness obtained from a retired SheevaPlug power board.

CSN server architecture

Software

Operating system

The CSN client software is based on the Linux operating system running on the hardware platform with USB and ethernet interfaces. The USB interface is primarily used to connect to the accelerometer. The ethernet interface is the primary means of sending locally collected accelerometer data to the CSN server environment. The ethernet interface is also used to maintain an open reverse-ssh tunnel to the server environment. This permits remote access from the server to the client even if client-side enterprise firewalls are present.

Preparing the SBC's operating system platform involves a combination of tasks that include: ensuring that a base level of applications has been installed on the platform to support CSN tools; setting various configuration parameters specific files for each application of interest; and in some cases disabling conflicting applications. CSN practice is to retain the operating Linux system version originally installed on a hardware platform. This provides a degree of stability and predictability that is valuable over many years. As of 2021, several 8-year-old deployments are still in operation and are expected to continue in perpetuity. This also addresses early versions of hardware that do not readily support operating system version upgrades. However, the above practice also implies ongoing support for an ever-growing number of Linux versions, each typically customized for a particular hardware platform family and model.

Client application

The CSN client application is a Python script, currently written to Python 2.7. It is thread-based, in which the threads are used to handle several different tasks:

- Interfacing with the accelerometer over USB and receiving triaxial samples from the sensor.
- Processing the triaxial samples from the sensor (including decimation, mean removal, property assessment).
- Creating picks from incoming sensor samples.
- Monitoring the system clock, obtaining Network Time Protocol (NTP) time from a time server, computing a regression, and providing other threads with timestamp.
- Uploading 10-minute raw data files to Amazon S3, and uploading latest station configuration data file to Amazon S3.
- Implementing a web server interface for remote users to obtain data from the client including: a) uptime, b) version, c) latest accelerometer sample, d) 10-minute files for arbitrary periods of time, e) latest 2 minutes of data in the form of Google Charts for each of the three sensor axes.
- Ensuring that sufficient space is available on local storage by deleting older files when necessary.

The main client program contains the credentials for accessing both the Amazon S3 service that will store the sensor data, and the location and credentials for the ActiveMQ broker which will receive the picks from the client. (ActiveMQ is open-source messaging software that is employed by the distributed algorithms and applications that require messaging, and the broker is an application that validates, translates and routes a message from a sender to a receiver). The CSN client file is in a human-readable format that contains CSN station metadata including the sensor's latitude and longitude, building identification, floor number, and client name. This is routinely edited to contain the required details.

Timing and accuracy

Accurate timestamps are required in the client to ensure that the data points are all synchronized with accurate external clocks and with those from other clients. This is especially critical for deployments of several clients within the same building, for example all sensors from sequential floors in a building for which inter-story drift or propagating wave property calculations are desired. To achieve accurate timestamps, the client applies a linear regression to the reported times obtained from one of several possible NTP servers on the network. Typically, the servers are at Caltech, UCLA, and USC, but sometimes the local router is used for failover capability. In some cases, a CSN SBC in the field serves as the NTP server (e.g. a separate CSN NTP server in a building).

The client polls the NTP server every minute, and the offset between the system clock and the NTP time is added to a sliding 10-minute window of offsets. A linear regression is then performed on the latest 10 minutes of data, which allows an estimate to be made for the true time at any data point (i.e. time series sample) over the coming minute before the NTP server is polled again. An example of how the NTP offsets and predicted offsets look over a 1-hour period is shown in Fig. 2.

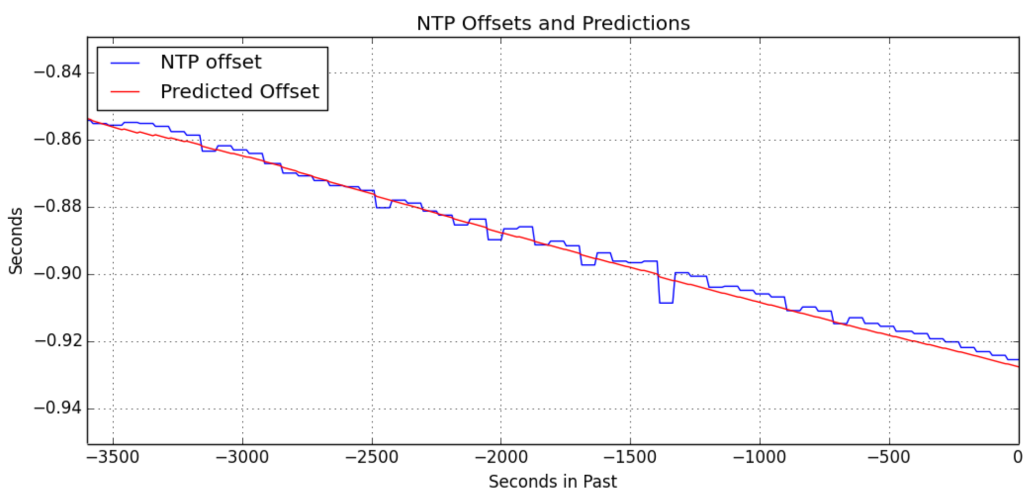

Figure 2. Example of NTP offsets (blue curve) and predicted time offsets (red line) over a 1-hour period of waveform data.

Client server architecture

CSN provides both real-time and near-real-time access to client-generated data. Limited real-time data flows instantly from CSN clients to a CSN cloud server data broker, to which two subscribers currently connect – one for earthquake early warning test applications and the other for a ShakeMap (Wald et al., 2008; Worden et al., 2020) service test instance. Near-real-time data flows in short bursts (currently, 10-minute-long time series) from CSN clients through a CSN cloud server to a local server-based archive at Caltech.

The method by which these parallel data flows are carried out is as follows. The CSN client running on the Linux-based processor at each station has two different but concurrent modes for uploading data. In the first mode (referred to as "pick mode"), the station software instantly detects anomalously large waveform events in the data stream, such as high-amplitude accelerations, and immediately sends the computed peak amplitude scalar values and their associated time stamps to the Amazon cloud (AWS). In this mode, the CSN python application assesses the incoming stream of samples from the phidget sensor device library, and it forwards the selected samples to a server running on the CSN cloud infrastructure. The pick mode data are those that are currently designated for ShakeMap and the ShakeAlert earthquake early warning system.

The CSN station uploads all waveform data in the second near-real-time mode ("continuous mode"). The station's processor accumulates 10 minutes of three-component, minimally processed, unfiltered, acceleration data into a file. Once the file is complete, it is compressed and uploaded to the cloud, while data are being accumulated into the next file. The 10-minute data files are retained by the station's system until the local storage use reaches a maximum threshold value, at which time the files are aged out; typically this amounts to a few weeks of sensor data with the flash storage cards. This storage system would need to be accessed for accumulated data in the case of power outages or communications problems.

CSN's software client computes peak accelerations which are reported as picks (the "pick mode" introduced above) if values are > 0.5% g. The picks are computed from the time series' deviations from a long-term mean calculated over 10-second sliding windows. Each orthogonal axis is treated independently and the minimum repick interval is 1 second on the same axis; thus the maximum pick rate is 3 picks/second. Timestamps associated with the picks are calculated using a continuous regression on the NTP offsets to the computer's system clock, as discussed earlier. Only ground-level station picks are sent to the cloud for the ShakeMap and ShakeAlert applications.

In the current implementation of CSN's pick distribution method, the CSN client (running at the station) directly generates an ActiveMQ message for each pick locally on the station's processor, in the required format for ShakeMap or the earthquake early warning applications FinDer (Böse et al., 2012, 2015, 2018) and PLUM (Cochran et al., 2019). The CSN server sends that message to an ActiveMQ broker running on an AWS virtual machine. The CSN client includes an NTP-based corrected pick timestamp in the ActiveMQ message. ShakeAlert operates its own ActiveMQ brokers, whose topics are subscribed to by the various algorithms, including FinDer and PLUM. A channel between the CSN broker and the ShakeAlert broker used by development versions of FinDer and PLUM allows it to receive all CSN client picks. Both FinDer and PLUM use all reported CSN client picks associated with an earthquake, since they always exceed 0.5% g. At the server side where PLUM is running, MMI values are computed continuously for the incoming accelerometer measurements. MMI is computed from the incoming PGA pick values on all three components, and sent when an MMI threshold for PLUM is exceeded. The maximum rate of MMI messages being sent by each client is one message per second, for the duration of the shaking. Similar to FinDer, these MMI values are relayed to the ShakeAlert ActiveMQ broker to which the development version of PLUM is subscribed.

As of August 2021, the data stream provided by the CSN python client application to the CSN ActiveMQ broker sufficiently matches the expectations of the ShakeAlert infrastructure to support a direct subscription. Future enhancements to the stream from the CSN client to the CSN ActiveMQ broker may require an additional level of processing and assessment within the cloud, prior to making the data available to the ShakeAlert production system.

The data flow architecture is illustrated in Fig. 3. The key motivation for this setup is to prevent large latencies that can arise in part from the current ShakeAlert requirement that each station sends full waveforms to the central processing site.

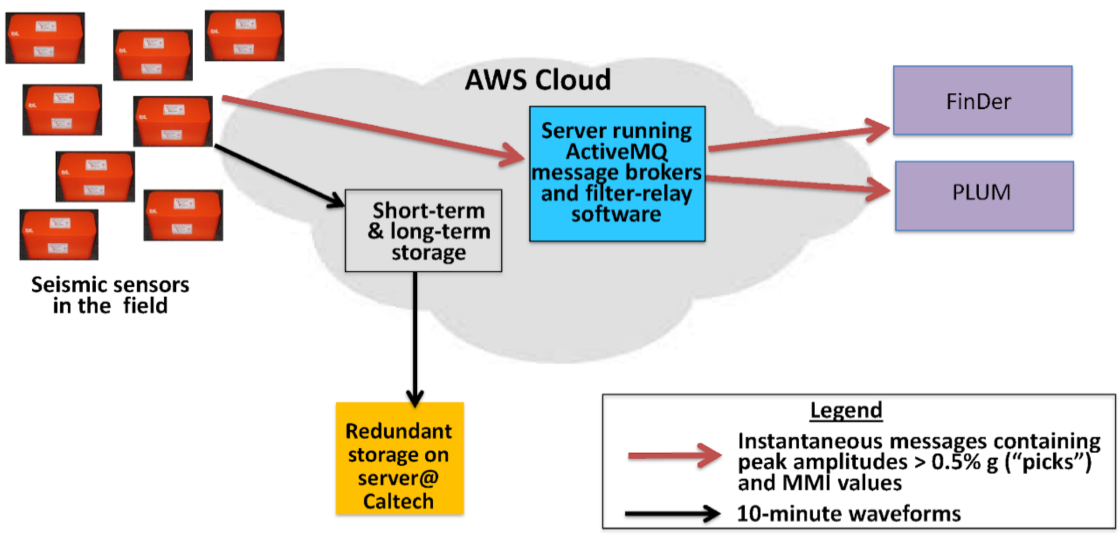

Figure 3. Data flow architecture schematic. Sensors in the field (left side of diagram) send two types of data to the message broker and storage servers in the Amazon cloud (AWS cloud). Data are then passed on to development instances of FinDer and PLUM (right side of diagram).

The ShakeMap infrastructure requires a different format and content for peak amplitude data than is provided by the CSN cloud ActiveMQ broker for ShakeAlert. Therefore, a service running on the local server platform subscribes to the CSN ActiveMQ broker, provides very low latency reformatting of the incoming peak amplitude (pick) data stream, and then sends those data onward to the ShakeMap infrastructure.

A number of waveform data retrieval applications operate on the local server to provide continuous waveform time series files to researchers. The main application is based on the Seismogram Transfer Program (STP) client (STP, 2007) and can deliver data for all stations rapidly for recent months, and with a small latency for data older than that. Customized CSN STP-based clients serve a subset of sensors such as all stations at Los Angeles Unified School District campuses, NASA-JPL, and several instrumented mid-rise and high-rise buildings. The applications rely on three file types for their operation: 1) the station file which contains the metadata information for each station in CSN, 2) the waveform metadata files which contain metadata information about each waveform segment, and 3) the waveform segment files themselves which contain the digital samples in SAC format. The data used by the applications is refreshed by restarting them; this is done on a regular basis several times a day.

Los Angeles Unified School District (LAUSD) stations

The majority of CSN stations are deployed at campuses of the Los Angeles Unified School District (LAUSD). LAUSD is a public, general-community stakeholder and partner, with approximately 1000 campuses that span the City of Los Angeles (2000 km^2). The school campus buildings are typically 1-3 story wood frame or reinforced concrete structures built after 1950. These structures include types that are known to be prone to severe damage due to seismic hazards, especially for older construction and soft first-story construction. Approximately 400 of the LAUSD campuses are instrumented with CSN sensors (majority of red circles shown in Fig. 1). Average CSN station spacing at LAUSD campuses is about 0.5 km.

An experimental ShakeMap-ShakeCast-like setup has been developed for the LAUSD campuses at which CSN has deployed a sensor (Kohler et al., 2018). The setup is generally based on features of ShakeCast (Wald et al., 2008), including the use of the Hazus Earthquake Model Loss Estimation Methodology (Hazus, 2020) to classify structures and supply fragility curves. However, key differences are that it uses CSN data recorded at the actual structure as shaking intensity input into the fragility functions for the structure, and CSN-developed web-based tools for its user interface. The LAUSD campuses used in this installation consist of only low-rise structures across a lateral dimension spanning about 20 km. All sensors are located in communication or utility closets; none are in classrooms. The ShakeCast application is installed in the central LAUSD office in downtown Los Angles and communication is modeled on a centralized decision engine setup in which information could be subsequently sent via formal channels to local principals and campuses.

As mentioned above, CSN's current software client computes broadband peak accelerations which are reported to ShakeMap if values are > 0.5% g, obtained from the time series' deviations from the long-term mean, on any axis. Many CSN stations on school campuses are in locations with frequent human activity that influences noise levels. For example, many LAUSD stations exhibit noticeably higher noise levels during school hours. In future work, station-specific noise models taking into account time of day and day of the week could be trained, allowing for more reliable picking and signal-to-noise estimation at stations with predictable human-generated noise.

Mid-rise and high-rise instrumentation

Several mid-rise and high-rise buildings are currently instrumented by CSN with at least one triaxial sensor deployed on most floors. All are located in the downtown or greater Los Angeles region. The buildings include a 52-story dual system (concentrically braced steel frames at core with outrigger moment frames (with 63 sensors); 15-story steel moment frame and concrete shear wall (with 34 sensors); 9-story reinforced concrete (with 10 sensors); and two 9-story steel moment-frame with trusses and girders (one with 31 sensors and the other with 15 sensors).

Several of the instrumented buildings have two or three triaxal sensors deployed on most floors, for the purpose of measuring torsion, rotations or rocking. Although the sensors are not usually located at the edges of the floors, their locations relative to the center of mass on the building floor can be determined from structural engineering drawings obtained for most of the instrumented buildings. The majority of sensors are located in electrical or IT closets. One of the CSN-instrumented buildings – the 52-story high-rise – also has CSMIP instrumentation that has recorded significant earthquakes since the 1992 M7.3 Landers earthquake, and could be used for data comparison of the July 2019 M7.1 and M6.4 Ridgecrest earthquakes, the September 2020 M4.5 South El Monte earthquake, and the April 2021 M4.0 Lennox earthquake.

NASA-JPL stations

A total of 220 CSN triaxial accelerometers are deployed on the ~ 1 km × 1 km NASA-JPL campus. This subset of sensors could be considered an "array within an array" due to their smaller but approximately equidistant station spacing. Sensors are installed on both ground-level and upper-level floors of several buildings. The ground-level stations have an average spacing of about 100 m.

The NASA-JPL sensor deployment can be viewed and tested as a prototype mini-city strong-motion deployment, as there are one or more CSN accelerometers installed in about 90 buildings (mostly single or two-story structures) on the campus. The building types comprise wood frame, steel sheds, modular trailers, steel-moment frame, and reinforced concrete. Of the total 220 stations deployed at JPL, only the 100 ground-level stations are contributing maximum shaking peak acceleration pick data for the experimental ShakeMap and earthquake early warning algorithms. As with the LAUSD subarray, a ShakeMap-ShakeCast-like setup has been configured for JPL (Massari et al., 2017). Each of the buildings uses fragility curves supplied by the Hazus Earthquake Model Loss Estimation Methodology (Hazus, 2020). The ShakeCast configurations for the JPL sites are set up so that they use the CSN ShakeMap as input for localized and customized building performance assessment. Several buildings have either two or three sensors located on the ground level floors because the structures are long or they contain a significant element joint halfway down the longitudinal axis of the building.

2019 Ridgecrest earthquake

The July 2019 Ridgecrest earthquake sequence that occurred in southeastern California, was recorded on hundreds of CSN sensors in the Los Angeles basin (Kohler et al., 2020; Filippitzis et al., 2021). In particular, CSN captured variations in ground-level motion and upper floor deformation within mid- and high-rise buildings and showed unexpected patterns of large spatial variations in shaking amplification, as was envisioned as a primary purpose of CSN. Work with CSN recordings of the M7.1 mainshock revealed amplified shaking in the instrumented 52-story high-rise in downtown Los Angeles lasting over two minutes. In addition, ground-level accelerations showed increases in the amplification of long-period motions (> 1 s) from the northern Los Angeles sedimentary basin (Fig. 4). In Fig 4, several locations show nearly co-located CSN and SCSN or CSMIP instrumentation (circles and diamond symbols in close proximity), indicating the consistent response between stations of the different networks for this earthquake. High-rises experienced unusually strong long-period shaking in the east-west direction as a result of excitation by a complex train of scattered shear waves inside the basin, including surface waves propagating in the basin. The density of the CSN observations demonstrated that the behavior of structures (e.g. buildings) with long natural periods does not follow long-standing expectations for how sedimentary basins affect amplification (Filippitzis et al., 2021).

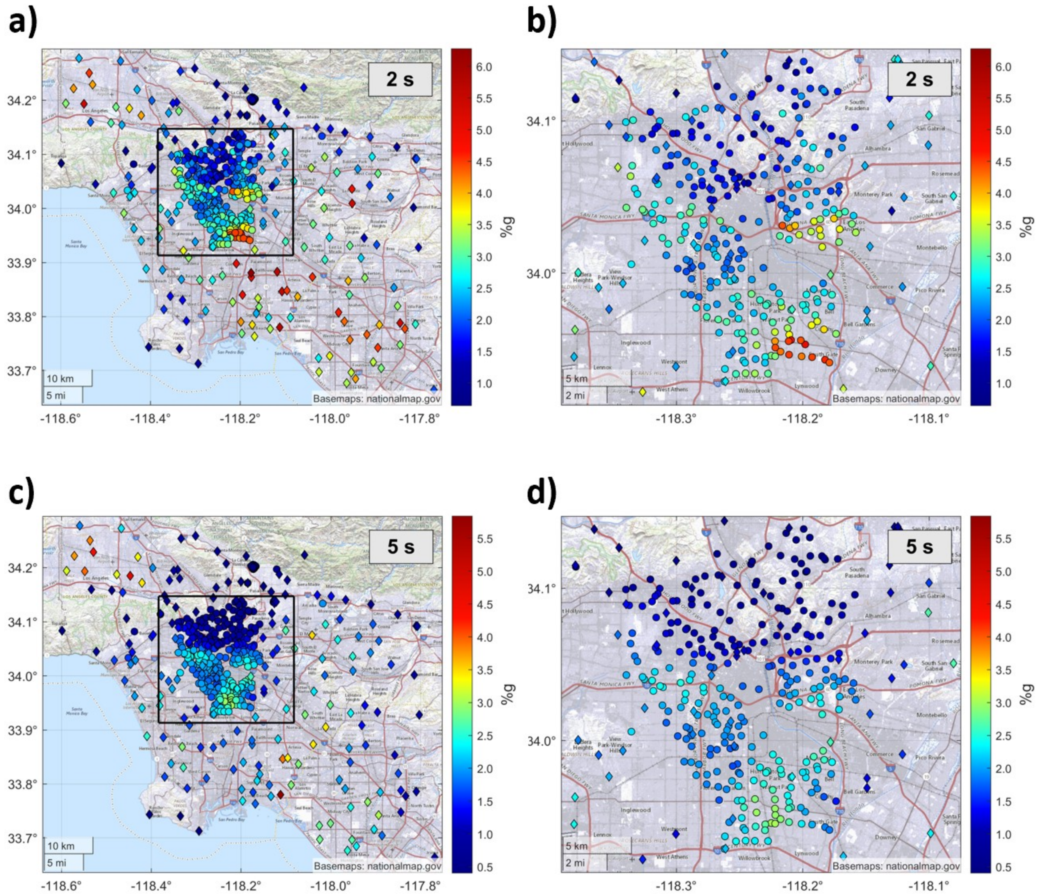

Figure 4. Pseudo-spectral acceleration amplitudes using 5% damping for the 2019 M7.1 Ridgecrest earthquake in southern California. Periods shown are: (a,b) 2 s, and (c,d) 5 s. Left column: Greater urban Los Angeles region. Right column: Blow-up of the region inside the marked squares on the left, showing detail. CSN stations=circles. SCSN+CSMIP stations=diamonds.

The ground-level accelerations recorded by CSN from the Ridgecrest earthquakes also exhibited coherent, gradational variations in the spectral amplitudes of high-frequency motions across the NASA-JPL campus that suggest correlations with geomorphological features (ridges, canyons, and foothills). The variations in spectral amplitudes are most pronounced for frequencies between 1 and 3 Hz. For the M7.1 mainshock, the overall maximum amplification occurred in the highest elevation, on top of the bedrock mesa bounding the campus to the north. The M6.4 and M5.4 foreshocks show a similar amplification pattern. The amplification pattern changes with the frequency as energy components of various wavelengths interact with surface and subsurface features of different characteristic lengths. For this higher-frequency range, the incoming waves likely interacted with numerous small-size features at the site to generate a complex, rapidly-varying amplification pattern.

The integration of site-specific and structure-specific instrumentation provided by the Community Seismic Network enables the deployment of large numbers of seismic sensors for dense spatial sampling. The developed framework offers a path forward for city-scale and regional-scale seismic network operations, and could serve as a scalable and reconfigurable tool for monitoring structures such as tall buildings, bridges and dams, as well as lifeline infrastructure.

Acknowledgements

The Community Seismic Network is the product of years of work by many people who have made up the CSN team over the years. The author is a part of this team, and in representing this team wishes to acknowledge their past and ongoing contributions to the CSN project. The current CSN team members are Robert Clayton (Seismological Lab; California Institute of Technology), Thomas Heaton (Department of Mechanical and Civil Engineering, and Seismological Lab, California Institute of Technology), Richard Guy (Department of Civil and Environmental Engineering, University of California, Los Angeles; and Seismological Lab, California Institute of Technology), Julian Bunn (Center for Data-Driven Discovery, Division of Physics, Math and Astronomy, California Institute of Technology), Filippos Filippitzis (Department of Mechanical Engineering, University of Western Macedonia, Kozani, Greece; and Department of Mechanical and Civil Engineering, California Institute of Technology), Yousef Bozorgnia (Department of Civil and Environmental Engineering, University of California, Los Angeles) and Ertugrul Taciroglu (Department of Civil and Environmental Engineering, University of California, Los Angeles). The CSN team is grateful to Caltech, UCLA, the Conrad N. Hilton Foundation, and Computers & Structures Inc., the Gordon and Betty Moore Foundation, and the National Science Foundation for providing support for the Community Seismic Network over the past decade.

References

Böse, M., T. Heaton, and E. Hauksson (2012). Real-time finite fault rupture detector (FinDer) for large earthquakes, *Geophys. J. Int., 191*, 803-812.

Böse, M., C. Felizardo, and T. H. Heaton (2015). Finite-fault rupture detector (FinDer): Going real-time in Californian ShakeAlert Warning System, *Seis. Res. Lett.,* 86(6), 1692-1704.

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.

Clayton, R., T. Heaton, M. Chandy, A. Krause, M. Kohler, J. Bunn, R. Guy, M. Olson, M. Faulkner, M. H. Cheng, L. Strand, R. Chandy, D. Obenshain, A. Liu, and M. Aivazis (2011). Community Seismic Network, *Annals of Geophysics*, 54, 6; doi: 10.4401/ag-5269.

Clayton, R. W., T. Heaton, M. Kohler, M. Chandy, R. Guy, and J. Bunn. (2015). Community Seismic Network: a dense array to sense earthquake strong motions, *Seis. Res. Lett*., 86, 1354-1363, doi: 10.1785/0220150094.

Clayton, R., M. Kohler, R. Guy, J. Bunn, T. Heaton, and M. Chandy (2020). CSN/LAUSD network: A dense accelerometer network in Los Angeles schools, *Seis. Res. Lett.,* 91(2A), 622-630, doi:10.1785/0220190200.

Cochran, E. S., J. Bunn, S. E. Minson, A. S. Baltay, D. L. Kilb, Y. Kodera, M. Hoshiba (2019). Event detection performance of the PLUM earthquake early warning algorithm in southern California, *Bull. Seis. Soc. Am.,* 109 (4), 1524–1541, doi:10.1785/0120180326.

Filippitzis, F., M. D. Kohler, T. H. Heaton, R. W. Graves, R. W. Clayton, R. G. Guy, J. J. Bunn, and K. M. Chandy, Ground motions in urban Los Angeles from the 2019 Ridgecrest earthquake sequence, *Earthquake Spectra*, version available online, doi:10.1177/87552930211003916, 2021.

Given, 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. Geological Survey Open-File Report 2014–1097*, 25 pp.

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, M. Hellweg, J. R. Murray, V. I. Thomas, D. Toomey, and T. S. Yelin (2018). Revised technical implementation plan for the ShakeAlert system—An earthquake early warning system for the West Coast of the United States, *U.S. Geological Survey Open-File Report 2018–1155*, 42 pp., doi:10.3133/ofr20181155.

Hazus (2020). *Hazus Earthquake Model Technical Manual.* Federal Emergency Management Agency, Washington, D.C. https://www.fema.gov/sites/default/files/2020-10/fema_hazus_earthquake_technical_manual_4-2.pdf.

Kohler, M. D., T. H. Heaton, and M. H. Cheng (2013). The Community Seismic Network and Quake-Catcher Network: enabling structural health monitoring through instrumentation by community participants, *Proceedings of the SPIE Smart Structures/Non-destructive Evaluation Conference*, San Diego, CA, March 10-14, 2013.

Kohler, M. D., T. H. Heaton, M. H. Cheng, and P. Singh (2014). Structural health monitoring through dense instrumentation by community participants: the Community Seismic Network and Quake-Catcher Network, *10th U.S. National Conference on Earthquake Engineering* (*10NCEE*), Anchorage, Alaska, July 21-25, 2014.

Kohler, M. D., R. Guy, J. Bunn, A. Massari, R. Clayton, T. Heaton, K. M. Chandy, H. Ebrahimian, and C. Dorn (2018). Community Seismic Network and localized earthquake situational awareness, *11th U.S. National Conference on Earthquake Engineering* (*11NCEE*), Los Angeles, CA, June 25-29, 2018.

Kohler, M. D., D. Smith, J. Andrews, A. Chung, R. Hartog, I. Henson, D. Given, R. de Groot, and S. Guiwits (2020). Earthquake Early Warning ShakeAlert 2.0: public rollout, *Seis. Res. Lett*., 91(3), 1763–1775, doi:10.1785/0220190245, 2020.

Kohler, M. D., F. Filippitzis, T. H. Heaton, R. W. Clayton, R. G. Guy, J. J. Bunn, and K. M. Chandy (2020). 2019 Ridgecrest earthquake reveals areas of Los Angeles that amplify shaking of high-rises, *Seis. Res. Lett*., 91(6), 3370–3380, doi:10.1785/0220200170.

Massari A., M. Kohler, R. Clayton, R. Guy, T. Heaton, J. Bunn, K. M. Chandy, D. Demetri (2017). Dense building instrumentation application for city-wide structural health monitoring and resilience, *16th World Conference on Earthquake Engineering* (*16WCEE*), Santiago, Chile, January 9-13, 2017.

Seismogram Transfer Program (STP) Reference Manual (2007). Prepared by the Southern California Earthquake Data Center, Seismological Laboratory, California Institute of Technology, Version: 1.01, https://scedc.caltech.edu/data/stp/STP_Manual_v1.01.pdf.

Wald, D., K.-W. Lin, K. Porter, and L. Turner (2008). ShakeCast: Automating and improving the use of ShakeMap for post-earthquake decision-making and response, *Earthquake Spectra*, 24, 2, 533-553.

Worden, C. B., E. M. Thompson, M. Hearne, and D. J. Wald (2020). ShakeMap Manual Online: Technical Manual, User's Guide, and Software Guide, U. S. Geological Survey, http://usgs.github.io/shakemap, doi: 10.5066/F7D21VPQ.