

Research Article

Processing and Communication Delays in EWS: On the Performance of the Earthcloud Prototype

Martin Klapez ^{1,2} Carlo Augusto Grazia ^{1,2} Maurizio Casoni ¹
Simone Zennaro,² and Matteo Cozzani²

¹Department of Engineering Enzo Ferrari, University of Modena and Reggio Emilia, via Pietro Vivarelli, 10, 41125 Modena, Italy

²Earthcloud Association, Italy

Correspondence should be addressed to Martin Klapez; martin.klapez@unimore.it

Received 19 November 2018; Accepted 5 March 2019; Published 26 March 2019

Academic Editor: Mauro Femminella

Copyright © 2019 Martin Klapez et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

A Seismic Alert System (SAS), also called Earthquake Warning System (EWS) or Earthquake Early Warning System (EEW or EEWS), represents one of the most important measures that can be taken to prevent and minimize earthquake damage. These systems are mainly used to detect P-waves and the faster seismic waves and to subsequently trigger an alarm about the incoming S-waves, the slower and most dangerous seismic waves. In some cases, distributed systems are also able to alert some locations before the impending P-waves strike them. This paper presents Earthcloud, a cloud-based SAS that aims to provide all the former capabilities while retaining financial accessibility. Earthcloud first results, generated from four months of data acquisition, are compared with those coming from other systems. In particular, the paper focuses on processing and communication delays, showing how the Earthcloud new detection strategy may minimize delays. Although a thorough test campaign with more sensor nodes is needed to assess performance reliably, especially for highly dense urban scenarios, initial results are promising, with total latencies for Earthcloud always kept under the 1-second mark, despite being at the expense of solid magnitude estimation.

1. Introduction

Due to the relatively fast propagation of earthquake waves, it is only possible to receive alerts from SASs in a timeframe that goes from a few seconds to tens of seconds before the strike. Nevertheless, they can potentially yield very significant benefits. If an earthquake occurs at night when people are usually sleeping, for instance, the primary risk for human health (primary in terms of the number of people affected) is represented by the possible collapse of buildings or parts of them. Potential victims usually wake up abruptly, most likely in a state of panic and with the risk of being overwhelmed by events. Risks for human health would be even higher if an earthquake occurs during the day when many people are working or traveling. Examples of these hazards, unfortunately, abound and can go from small injuries to large-scale nuclear disasters, for instance, derailment of trains on railroads and subways, vehicles on bridges and tunnels, dangerous machines and chemicals in work environments,

suspended loads and work at heights in building sites, fires, and more, plus the collapse of buildings, as mentioned earlier. In addition to all of these, some risks involve machinery that may or may not affect human health but that would inevitably cause economic damage. Receiving a seismic alert, even if only seconds before the arrival of destructive waves, can prevent human losses, injuries, and damage to machinery and infrastructures. For instance, it can trigger the enforcement of a previously approved emergency plan that people can consist in moving away from hazardous equipment, taking cover under desks or load-bearing structures, etc. Machinery, on the other hand, can be automatically slowed, shut down, and/or isolated, to prevent damage as well as to avoid the ignition of cascading threats, such as fires that can kindle in the aftermath of an earthquake.

Designing and operating an effective SAS, however, is often not trivial, because an earthquake has to be detected, processed, and notified significantly faster than the propagation speed of its seismic waves. Furthermore, environmental,

traffic, industrial, and other kinds of noise that may easily generate false positives are often present and have to be effectively classified in a very short time with a very low error rate. This paper presents the second iteration of Earthcloud, an Internet of Things (IoT) SAS designed to be low-cost, low-power, and cloud-based. From the detection of the first anomaly to an alarm, processing and communication delays of several systems is here analyzed and compared with the results obtained from the Earthcloud prototype.

To provide the reader with solid interdisciplinary bases, Section 2 pinpoints related work, while Section 3 summarizes the necessary geological concepts. The deployed prototype is described in Sections 4 and 5, where sensor units and the cloud system are detailed, respectively. Section 7 presents the outcomes of four months of processing, with an emphasis on processing and communication delays. Section 8 concludes the writing.

2. Related Work: An Overview on SASs

The goal of this Section is not to be exhaustive; it is instead to present an indicative group of works from academia and industry, to form a representative overview of the options and to point the reader to relevant sources.

Back in 1996, Leach et al. wrote one of the first works on SASs based on computer technology [1]. Authors designed a neural network and trained it with real data extracted from seismograms. Then, the neural network has been fed with other data from the same sources, and its outcomes have been studied. An actual system was not developed, but the algorithm was successful in the accurate estimation of hazard start, stop, and duration time. The algorithm was able to generate a warning signal 0.3 seconds after detecting the first ground motion.

The majority of SAS prototypes reported in the literature aim to be low-cost use MEMS (i.e., Microelectromechanical System) accelerometers as sensors. In [2], for instance, Hoque et al. built a prototype where sensors are composed of accelerometers attached to Arduino microcontrollers. The control system is a LabVIEW software on a laptop that communicates with sensors through the ZigBee standard. A ZigBee antenna can theoretically have a range up to hundreds of meters, but, generally, actual range is in the order of tens of meters. The latter applies to the system, as presented by the authors. With respect to Earthcloud, it is also moderately more complex to set up. From the paper, it is not clear whether the prototype has been deployed or not. Another work that leverages accelerometers and Arduino microcontrollers is that presented by Sherki et al. [3]. Instead of ZigBee, a GSM module is employed, and the system is tested against a vibration generator machine. The focus of the authors, however, is in this case on signal analysis.

Accelerometers embedded in smartphone have also been extensively used. For instance, Uga et al. [4] test these sensors specifically for SASs applications and study how to separate false positives from real ground motion, although tests are performed against artificial data obtained with a shakeboard, and smartphones are positioned on horizontal flat surfaces

for the entire duration of the experiments. Shakeboards are also employed to test prototypes designed to form an IEEE 802.11 (WiFi) Wireless Mesh Network [5, 6]. In these cases, however, authors focus specifically on the performance of wireless communications. They found that the sudden small-amplitude P-waves (see Section 3) shaking can have a huge impact on their performance, especially with no or scarce Line Of Sight between transceivers. Accelerometers in smartphones have also been coupled with cloud elaboration before. Heryana et al. [7] rely on 2G/3G/4G cellular technologies or IEEE 802.11 to pass data between devices and a cloud service. However, most of the paper is spent detailing the software development cycle of the proposed Android application, and very few information can be found about the effectiveness of the system.

Other works focus on unconventional systems. Shayo et al. [8] use the same low-cost accelerometer sensors seen in other solutions. However, sensors transmit data to a computer that, in the event of an emergency, transforms the alarm signal in SMS form and then send it through cellular networks. Another workstation has been set up to receive those SMS and measure the elapsed time. Authors found that, although the delays varied, on average these exceeded the maximum tolerable delay for a SAS and concluded that SMS is not a reliable platform for those means. An unorthodox sensor is presented by Heindl [9]. Author proposes to detect ground motions through the variations in read/write error rates of common electromechanical magnetic hard disks. The distributed system is designed in a P2P fashion. As in [7], the SAS would require substantial user collaboration to work correctly. Alternatively, although vendors would probably oppose such a choice for security and privacy reasons, operating systems might be modified to perform those SAS tasks in the background, without the users' consent. Anyway, the system may become weak in the long term, as magnetic hard drives are replaced by faster solid-state storage devices.

About systems that rely on users collaboration through software deployed on common devices, the Earthquake Network Project by Finazzi [10] is one of the best examples, claiming 4 million downloads of its Android application, 750000 active users, and more than a 1000 early warnings sent. While the analysis of data coming from the network of smartphones is not trivial, statistical approaches are being studied to improve the system [11].

An example of SASs from industry is ShakeAlarm by Zaicenco et al. [12], developed by Weir-Jones Engineering Consultants, a system that detects P-waves (see Section 3) and claims to be able to determine in less than 0.5 seconds if following waves will be dangerous. ShakeAlarm is deployed in some regions of Canada and of the USA. There are also SASs developed by institutions and deployed on a large scale. In Japan, the Japan Meteorological Agency operates an EWS [13] mainly formed by about 300 single-function and multipurpose seismometers; the latter are equipped with satellite mobile phone communication capability for backup purposes and a power supply that can keep the whole system operational for about 72 hours in the event of power failure. The Japanese EWS fetches data also from seismometers managed by universities, by the National Research Institute




FIGURE 1: Basic pattern of an earthquake.

for Earth Science and Disaster Resilience, and by the Japan Agency for Marine-Earth Science and Technology. Taiwan also has a nationwide SAS [14], while Mexico has a partial EWS called SASMEX that comprises the SAS of Mexico City, in continuous operation since 1991, and the SAS of Oaxaca City that started its services in 2003. In the United States, the west coast is currently covered by an experimental system called ShakeAlert, presented by Burkett et al. and Given et al. in [15], [16], respectively.

Earthcloud is a low-cost, low-power, and cloud-based SAS at the prototypal stage. Klapez et al. previously presented a version [17] that had two processing layers. The first was represented by sensor nodes and the second by the cloud infrastructure. New data was continuously generated by the sensor devices and continuously sent to the cloud system. The same data was processed twice, once by the sensors, once by the cloud. Both were able to issue an alarm, on the basis of the respective processing results. As the processing power in the first layer is limited, and to keep processing times acceptable, the cloud infrastructure was responsible for issuing the most reliable warnings. In this paper, we present the second iteration of Earthcloud. As it is described in Section 5, it embeds a fundamental change in how warnings are triggered and on how data is handled.

3. Key Geological Concepts

Earthquakes are generated in an area called the nucleation zone. Most of the times, this area is located inside the seismogenic layer, that is the part of the Earth's crust with a depth ranging from 5 to 30 km. In the nucleation zone, multiple types of waves are generated, with different characteristics based on destructive potential and speed.

These waves mainly come in two arrivals. The time between the two is the opportunity to detect and confirm the earthquake and send an alarm signal. Each of these two arrivals consists of many individual elastic waves that have traveled from the epicenter to recombine at the recording site as a function of their respective velocities, focal distances, and propagation paths. As depicted in Figure 1, waves belong to two types: body waves and surface waves.

The first arrival is composed by the fastest of body waves, the P-waves (from Latin *prima unda*, i.e., primary waves). These are compressional or longitudinal and shake the ground in the direction of their propagation using compression or rarefaction, while their speed is between 4 and 8 km/sec. P-waves are usually not destructive.

Body waves and surface waves may compose the second arrival. They often produce both horizontal and vertical ground motion and their peak velocities, peak accelerations, and duration in time may cause significant damage to structures.

The body waves in the second arrival are called S-waves (from Latin *secunda unda*, i.e., secondary waves); they shake the ground in a direction that is perpendicular to the direction of propagation, while their speed is about 60% to that of the respective P-waves [2]. S-waves can be destructive in up to several kilometers from the epicenter due to the phenomenon of seismic amplification, linked to the local geology and morphology.

Surface waves are categorized into two types, called Love and Rayleigh, from the name of the scientists that modeled them. They are formed by constructive interference between P-waves and S-waves, and they are the most dangerous. Love waves travel slower than both P-waves and S-waves; Rayleigh waves are the slowest. However, while they are slower, surface waves decay much less with the distance than




FIGURE 2: An Earthcloud Sensor Subsystem.

body waves, as they mainly travel on one axis instead of three. The Earth's crust where they travel acts as a waveguide that provides little attenuation loss. In big earthquakes, they may circumnavigate Earth several times before being completely dissipated. In general, surface waves cause more movement than body waves. In addition, the interaction among them as they propagate can produce considerable amplification of ground motion near the surface, a phenomenon called the free surface effect that occurs when upgoing and downgoing reflected waves are in phase and of considerably greater wavelength than the thickness of the crust.

4. Earthcloud: Sensor Devices

An Earthcloud sensor system is composed of three elements: a Raspberry Pi 3 Model B V1.2, the Adafruit ADS1115 Analog-Digital Converter (ADC), and a set of three 4.5 Hz geophones. Figure 2 depicts and details elements and the wired connections among them. Size proportions are real. Raspberry Pis mount a Linux distribution as the operating system; specifically, they mount Raspbian Stretch 9.4 with Linux kernel version 4.14-34 (4.14.34-v7+).

As a whole, a single sensor system needs less than 1W. To estimate the power consumption of the Raspberry Pi boards, we based our calculations on the specifications [18] published by the Raspberry Pi Foundation, in particular, those regarding power consumption of the same specific board in idle and low-load conditions. From an average current of $\sim 300\text{mA}$, we removed 50mA for the HDMI port and 100mA for mouse and keyboard (we picked up the best case, the aggregated figure can go up to 1A), all components not used in our setup. The values in [18] include WiFi usage; we assumed a similar consumption for the Ethernet port, that is the one used in our setup. All this, when considering

that the board is powered with 5V, yields a worst-case power consumption of 0.75W. The GPIO pins may safely draw up to 50mA overall, while an individual GPIO pin can only safely draw 16mA. In our setup, only the cable pictured in yellow in Figure 2 draws a significant amount of current to power the ADC through the 3.3 VDC GPIO pin; the power needed for that is, therefore, equal to 0.0528W at most. The blue(ish) wires transmit data from the ADC to the Raspberry Pi board only, and, therefore, for the latter, those pins need a negligible amount of power. Geophones are instead passive devices that only need to be grounded. System-wide, this amounts to a power upper bound of $\sim 0.8\text{W}$; considering common power supply efficiency around 80%, we can safely claim a total power consumption value equal to or less than 1W.

The geophones are the Earthcloud sensor elements. They usually sit in between accelerometers and seismometers in function and price, the latter being the most accurate, delicate, and expensive system. A typical geophone has one working axis and has the function to convert ground movement that has the same direction of the working axis into voltage. Therefore, any voltage deviation from the baseline is data. A typical geophone consists of housing and, inside, of a mass suspended by means of mechanical springs. The spring-mass system has a resonance frequency, also called natural frequency, in the working axis of the geophone. Conceptually, when the sensor moves, if the velocity produces a frequency that is lower than the natural, both housing and mass move. If the generated frequency is higher, only the housing moves and the mass tends to hold its position.

In modern geophones, like those used by Earthcloud, the mass is formed by a cylindrical frame in which, externally, two coils are attached so as to surround the cylinder. The double coil allows reducing data distortion. The cylinder and coil assembly are usually supported by springs, directly




FIGURE 3: Inside an Earthcloud Geophone.



FIGURE 4: Geophone Housing.

connected to the geophone housing. Inside the cylinder frame, there is a permanent magnet also directly attached to the housing that does not move. Figure 3 illustrates a typical geophone of this kind. Figure 4 shows the geophones in an Earthcloud sensor system, in which they are enclosed in a rigid structure firmly mounted upon a surface like a wall.

Therefore, if the movement frequency is lower than the natural, there is conceptually no internal movement at all if considering the relationship between coils and magnet; i.e., the variation in the position of the coils with respect to that of the magnet is zero. If the frequency exceeds the natural, coils tend to maintain their position while the magnet moves with the housing; therefore, there is a shift in location between coils and magnet. As the coils are in the field of the magnet, this shift produces a variation in voltage that is proportional to the movement velocity. If the geophone is fixed to the ground, its response is (if above natural frequency) proportional to ground velocity. For comparison, MEMS accelerometers respond to ground acceleration.

Geophones also have what is called a spurious frequency, usually due to the small movements that the mechanical springs can have in directions perpendicular to the working axis of the sensor. These springs, in fact, are designed to move linearly in the working axis but also to have the possibility of small movements in the plane perpendicular to the working axis. Motion in this plane, either transverse or rotational, is essential to allow freedom of mobility in the working axis for the cylinder and coil assembly. As there is a resonance frequency in the working axis, called natural

frequency, there is also a resonance frequency in the axis perpendicular to the former, that is called spurious frequency (actually, there are spurious frequencies, the lowest is taken as the spurious). The springs are, of course, very stiff in the direction perpendicular to their working axis, hence the high-frequency nature of the spurious resonance [19]. As shown in Figure 5, the natural frequency represents the lower bound of the usable bandwidth, while the spurious frequency the upper bound. Therefore, the two frequencies together determine the useful frequency range of the sensor. For all frequencies in this bandwidth, the output sensitivity (i.e., the smallest variation in voltage that can be measured by the sensor) remains approximately constant.

As it can be seen from Figure 5, the actual response of a geophone is not an on-off Boolean function. The mass still moves below the resonance frequency, but output resolution (i.e., the smallest variation in space that can be measured by the sensor) drops fast (output resolution is calculated by dividing sensitivity over noise, S/N). Inverse filtering can compensate by flattening the response below the natural frequency of the geophone, but it is only useful if there is an adequate signal-to-noise ratio. On the other hand, for frequencies higher than the spurious one, springs introduce additional resonances that generate noise and therefore lower resolution. Inconsistencies in materials and manufacturing processes have a substantial effect on the determination of the spurious frequency, hence its common specification as $SF \geq x$, which indicates that the actual frequency band in which spurious resonance will occur may be x or a $y > x$.

Geophones usually maintain signal digitization as a separate process, as the analog output is generated in the sensor first and then is sent to an external digitizer. For comparison, MEMS accelerometers derive their feedback from within the digitizing process instead. In general, accelerometers have lower noise at high frequencies, while geophones have lower noise at low frequencies. The signal itself actually degrades slower with MEMS accelerometers, usually around a 6-dB as the frequency halves. For comparison, the same figure for standard geophones is about 12-dB [20]. This should make those accelerometers good candidates for low-frequency recording; if the signal is strong, this is indeed the case. The problem with MEMS accelerometers is the intrinsically

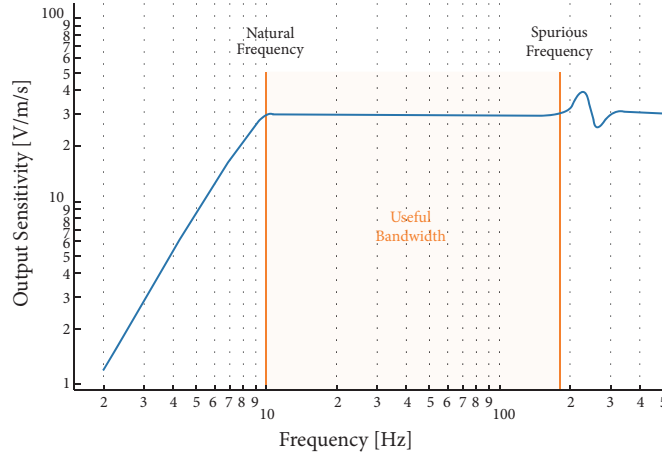


FIGURE 5: A transfer function for a typical 10 Hz geophone, with $\sim 60\%$ damping and 180 Hz spurious frequency.

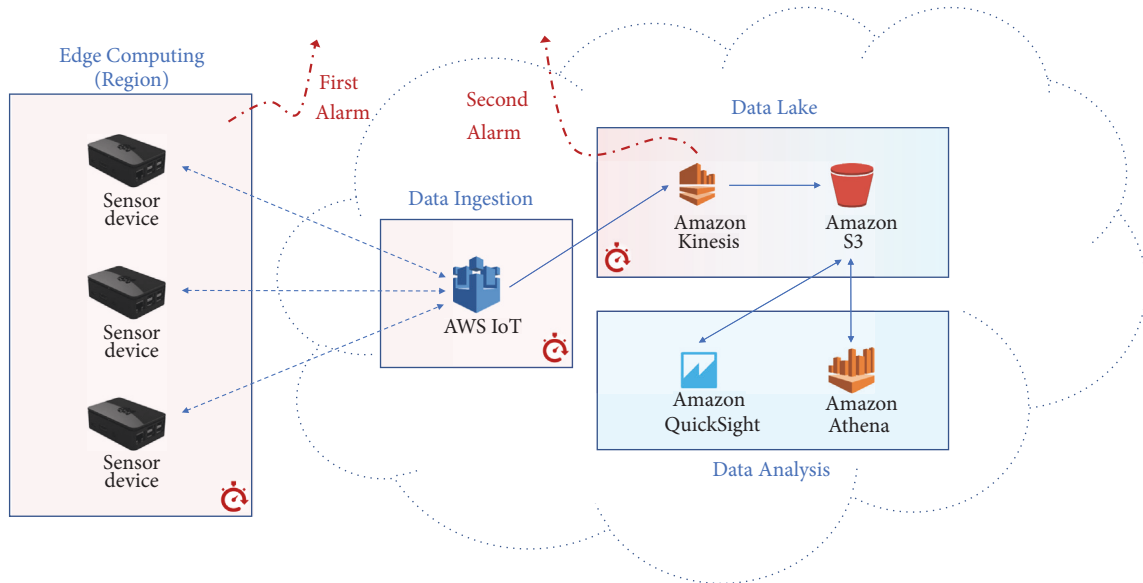


FIGURE 6: Earthcloud V2 Architecture.

high self-noise which makes them unsuitable for accurate low-frequency signals detection. Table 1 lists the technical specifications of the geophones currently used by Earthcloud. As shown in Figures 2 and 4, an Earthcloud sensor includes three geophones for three different axes that permit to gather data in a 3D domain.

5. Earthcloud: Cloud Infrastructure

Figure 6 shows the overall architecture of Earthcloud together with the composition of cloud services. At the moment, the latter are based on the Amazon AWS platform, in which Earthcloud services mainly operate from Ireland, EU.

IoT is a paradigm based on the connection between devices in the physical world, often sensors and data-gathering devices, and the Internet. Each Earthcloud sensor

TABLE 1: Technical specifications of the geophones used.

| | |
|----------------------------------|---|
| Name | EG-4.5-II |
| Natural frequency | $4.5 \pm 10\%$ Hz |
| Damping | $0.6 \pm 5\%$ |
| Coil resistance | $375 \pm 5\%$ Ω |
| Open circuit sensitivity (V/m/s) | $28.8 \pm 5\%$ |
| Harmonic distortion | $\leq 0.2\%$ |
| Typical spurious Frequency | ≥ 150 Hz |
| Moving mass | 11.3 g |
| Max. coil excursion p-p | 4 mm |
| Element diameter | 25.4 mm |
| Element height | 36 mm |
| Element weight | 86 g |
| Operating Temperature | -40°C to $+100^\circ\text{C}$ |

represents the edge of the cloud system, and it continuously scans the data incoming from its geophones. The only processing mechanism is a threshold-based filter; i.e., if data coming from a geophone exceeds a certain threshold (e.g., $v \geq 0.5mV$), a prewarning is issued to the cloud, together with the data that generated it. The principle on which Earthcloud warnings are based is purely probabilistic. In order to minimize to the maximum processing delays, sensors do not attempt to determine if signals from geophones are due to earthquake waves or noise. The system relies instead on the number of prewarnings received. As the number of sensors issuing a warning in the same region and in the same timeframe increases, the probability of contemporaneous false positives decreases.

Data generated by the Earthcloud sensors are encapsulated in MQTT messages continuously fed to AWS IoT, which preprocess and route them to Amazon Kinesis. When the data coming from its geophones exceed a set threshold, the sensor device sends a warning to AWS IoT, which in turn broadcasts it to the other sensor devices located in the same region. If a sensor device in a region issues a warning and receives a certain number of warnings with which it recognizes other n sensor devices of the same region as the source, it emits an alarm to registered entities. Subsequent alarms from other sensor devices located in the same region can be then safely ignored for a certain timeframe. Sensor data is small in size but can be continuously generated, from, potentially, thousands of sources. For Earthcloud, Kinesis has a dual function: to process data and issue alarms in real-time, and to act as a bridge between AWS IoT, the data entry point, and Amazon S3, the data endpoint. Clearly, the *real-time* part (depicted in Figure 6 with a light red background) is critical for Earthcloud. After receiving data for several warnings from the same region, Kinesis also issues an alarm to registered entities, serving as the confirmation that alarms are relevant and not false positives. Due to the significantly higher processing power of Kinesis with respect to sensor devices, in fact, the former can process more data and extrapolate more precise information, all in a short timeframe. The Kinesis alarm is supposed to follow those arriving at registered entities from sensor devices. However, there might be corner cases in which the latter fail to identify an earthquake; in these cases, the Kinesis alarm would be the first that is received by registered entities. Through an assessment of an ever-increasing data set, we expect to determine the probability with which Kinesis would issue a false alarm while being the only entity raising the issue. If we will find it to be sufficiently low, the alarms coming from the cloud system alone could be considered final even without the explicit support of sensor devices.

If needed, data can be converted into different formats before leaving the Kinesis bridge. Amazon S3 is a data storage service, well integrated with the other services from the same vendor. Data in S3 can be then further analyzed, manually and automatically. Amazon Athena allows querying S3 data through regular SQL (S3 data are stored as objects, not as SQL rows), while Amazon QuickSight, or other solutions, eases or performs data analysis. All the passages from Amazon

S3 forward are regarded as batch processing (depicted in Figure 6 with a light blue background); more sophisticated analyses can be performed, but it can take anywhere from few seconds to few minutes to obtain the results. Although it is in the works, we do not have a system based on machine learning implemented at this stage. However, as the data set grows, all relevant data collectable from S3 becomes more valuable to train a suitable neural network adequately.

5.1. Network Performance Considerations. Cloud services are provided with already built-in redundancy features, at least as long as they are deployed on premises of established vendors. However, as strong earthquakes are likely to disrupt infrastructures [21–24], redundancy should be introduced to sensor devices too. As we used flexible, low-cost, and low-power devices, it should be feasible to power them with small rechargeable batteries as well as to enable on them multiple connection channels [25, 26], all while retaining financial convenience. Within shared wireless networks, it is important to maintain the latency low. To do so, the signal-to-noise ratio of the wireless network chain (i.e., sensor devices, access point, repeaters, routers, etc.) must be kept high, while the saturation of channels belonging to devices of the same chain must be avoided. To have high-throughput capabilities is not crucial per se for sensor devices, as the data rate they produce is very low. However, if a wireless channel starts to be congested, jitter and latency increase, even to the order of seconds. It is therefore important to either prevent [27, 28] or manage [29–33] congestion.

6. Earthcloud: The Prototype Setup

The first prototype of Earthcloud consists of the cloud system (in beta) and three devices. These have been deployed in Modena, an Italian city in the Emilia-Romagna region, a mid-north area of Italy, EU, classified as a moderate-to-high seismic zone according to the 1999 Global Seismic Hazard Map [34]. In 2012, a seismic swarm struck Modena's area with intense earthquakes [35]. The first strong earthquake was registered on May 20, 2012, at 02:03:52 UTC (04:03:52 local time, i.e., at night). It had a 5.9 Richter magnitude (M_L , for local magnitude) and a hypocenter located 6.3 km underground. Other subsequent earthquakes followed, two with 5.1 M_L . Seven people died, 50 were injured, and 5000 lost their houses, while many historical buildings collapsed, together with several farms and factories. Consequential to the disaster, soil liquefaction also caused collapses of recently built structures. On May 29, 2012, a second strong earthquake hit the same region with a 5.8 ML and a 10.2 km-deep hypocenter, at 07:00:03 UTC (09:00:03 local time, i.e., in the morning). Other subsequent earthquakes followed, among which one with 5.3 M_L and another with 5.1 M_L . Consistently with what is exposed in Section 1, damages were more significant. 20 people died, 350 were injured, and 10000 more lost their houses. Among the aftershocks that followed, on June 3, 2012, another one with 5.1 M_L hit the area. There were




FIGURE 7: Comparison of first warning delays among different SASs (both prototypes and in production).

building damage and collapse, but no casualties. The two strongest earthquakes were also felt in nearby European countries, in particular in southeast France, Switzerland, south Germany, Austria, Slovenia, and Croatia. Damages were classified on the Mercalli scale (EMS-98) that classifies earthquake damages instead of its released energy, with a value of 8/12.

Two sensors have been initially installed on June 19, 2018, at the Department of Engineering Enzo Ferrari of the University of Modena and Reggio Emilia, inside a temperature-controlled server room. Devices are connected to the Internet through Ethernet cables, directly attached to the department switches for minimum latency, and are powered through electrical outlets. The geophone housings have been firmly mounted directly on walls concrete, with an industrial-grade wood-concrete glue. One has been installed on a load-bearing wall, while the other on a wall next to the room door, which is opened and closed very infrequently. The intent was to have two sensors in the same place but with one more sensitive to false positives. A week later, a third sensor has been deployed in a storage room belonging to a dismissed industrial building, located in a different area of the same city. The geophone housing was mounted on a load-bearing wall. This device was also connected to the Internet through an Ethernet cable, but with the difference of the latter being attached to a commercial off-the-shelf router manufactured for domestic purposes. Power was again provided by an electrical wall outlet. For the first prototype, unrefined ad hoc containers for geophones have been built by creating several rectangular prisms of approximately the same size from solid wood. Three holes to house the geophones have been drilled in each container, two for horizontal axes and one for the vertical axis, as in Figure 4. The holes allow for very limited movement of the geophones inside the housing, around 1 mm. As data analysis gets refined, more sophisticated housings will be created, in order to further minimize geophone allowances inside their containers or cancel them altogether.

7. Processing and Communication Delays

This section compares the latency performance of Earthcloud with the delays reported by different SASs in the literature. When possible, it is reported data from real earthquakes. Where this kind of data is not available and if simulation data is instead present, the latter is reported here.

Figure 7 depicts the comparison, where the vertical axis renders the delay that exists between the first wave arrival and the issue of the first warning, while the horizontal axis reports the SASs. In the label identifying a SAS, the round brackets in the second line inform about the kind of sensors employed by the system. The square brackets instead contain whether the delay values include processing delays (P) and communication delays (C). The order of these delays in the square brackets matters. The processing delay, in particular, can be found before or after the communication delay, depending on whether the processing performed on sensors is more significant or that on a remote system. In the case of Earthcloud, the symbol P is italic, to represent an almost negligible contribution. In the case of Hualien V3, the acronyms FD and OS stand for front-detection and on-site, respectively. A front-detection system aims to detect earthquakes in one place and give early warnings for more distant locations. On-site systems aim to recognize P-waves and to issue an alarm, in the same place where they are located, before S-waves arrive. The boxes of the box plot are composed as follows. The full horizontal line inside the perimeter of the boxes is the median, while the dotted line is the arithmetic mean. The lower and upper bounds of the boxes mark the first and the third quartiles, respectively. The whiskers include instead the values between $Q1 - (1.5 * IQR)$ and $Q3 + (1.5 * IQR)$, where IQR is the interquartile range, calculated with $IQR = Q3 - Q1$. We considered the other values, i.e., the dots outside the whiskers bounds, as outliers.

7.1. Earthquake Network. The Earthquake Network is a research project [10, 11] that leverages the MEMS accelerometers embedded in common smartphones to create a

crowdsourced front-detection EEWS. The principles on which the Earthquake Network operates are similar to those of Earthcloud. For the former, sensor nodes are represented by smartphones where a specific application is voluntarily installed by the user. If the smartphone is monitoring the output of its accelerometers, it sends an “alive” message to a central server every 30 minutes. In this way, it is possible to estimate how many sensors are active in each moment and where are they located. Similarly to Earthcloud, the processing in sensor nodes is kept to a minimum. In the Earthquake Network case, every vibration detected by a sensor appears to be logged, filtered, and then eventually sent to the central server together with the position of the smartphone. The sensor nodes, therefore, use the resources needed to operate their embedded accelerometers, log their data, filter it, and send it to the central server. The delay between the beginning of the ground shaking and the vibration detection from a smartphone is reportedly equal to 1.5s [11]. The detection algorithm is completely deployed server-side. In general, the authors report several figures about the delay experienced by various subnetworks during real earthquakes, including also false alarms. In particular, it is reported in detail the system response in the case of the M7.3 earthquake of May 12, 2015, in Nepal. The detection happened after approximately 6.5s from the moment the seismic waves reached the first smartphone. Firstly, the aforementioned 1.5s are assumed. As communication delay, 0.5s are assumed of having been needed for each one-way communication between smartphone and server, yielding a RTT (Round Trip Time) of 1s. The server was able to confirm the earthquake after 4s from when it received the first notification. In general, the authors report total delay values in the range from 2 to 17 seconds, with a median of 7s and a mean value of 8.15s. The paper [11] also contains delays resulting from simulations, not included here as per the guidelines defined in the incipit of Section 7.

7.2. SMS. Authors in [8] evaluate Short Message Service (SMS) messages as a platform to transmit seismic alerts. They created a prototype able to deliver SMSs through three mobile operators in Tanzania, finding that the communication delays fluctuate strongly, almost always exceeding the threshold of 5s took as reference. The delays resulting from the study have been deliberately isolated as communication delays only. Data is generated by a MEMS accelerometer connected to a sender computer and is then transmitted as SMSs to a receiver computer, every 20 minutes, for 21 days. It is worth to note that the same SMS was always simultaneously transmitted via the three operator networks. The study specifies that values greater than 15s have been discarded, as considered too high to be useful. The data reported here is composed by the values registered in one of the days of the experiment; specifically, the day with average values overall among the days of the trial. In this 24 hours, authors report total delay values in the range from 8 to 14 seconds, with a median of 10s and a mean value of 10.34s.

7.3. WSN. In [6], Wireless Sensor Networks (WSNs) are proposed as low-cost tools to realize front-detection EEWSs.

The authors refer to WSNs as “computer networks whose nodes communicate wirelessly using a license-free spectrum in a self-organized manner”. This work focuses on the optimization of the network routing protocol, as it is determined that shakings representative of P-waves of M6 or greater at a distance of 40km or less can easily result in severe multipath and shadow fading effects that considerably affect performance of the WSN wireless communications right when the EEWS is needed the most. The paper mainly presents outcomes of simulations. Delays range from 0.1 to 3 seconds, growing together with the number of WSN nodes. As the latter increases, the reliability of the warnings decreases. The median results equal 0.3s while the mean 0.9s. In the study, values “as low as 4s” are cited for real implementations. The simulated values are very low. This is due to the fact that, like Earthcloud at its current stage, sensor nodes do not perform substantial processing. Unlike the Earthcloud prototype that can estimate the epicenter location, both the magnitude and the epicenter location estimations are not performed by WSN. There are some additional remarks that need to be made. Firstly, the reported figures are only correlated to the time interval between a generic P-wave arrival and a generic positive alarm decision, but information about which node does what is missing. Negative delays are also reported in the paper, clearly representing a warning arriving to a node still extraneous to the P-wave detection. This might strongly skew this system results in the framework of the comparison presented here. To mitigate this problem, only the positive values have been considered, although it had not been possible to correlate them correctly. On the other hand, the values cited to be “as low as 4s” for real implementations have not been included here. It has also to be considered that the focus of the article is the routing protocol and not the SAS. Lastly, the performance parameters of the simulated network are unknown. As a consequence, it is not possible to evaluate, nor mitigate, eventual offsets with systems that would be physically deployed. Because of all the above-mentioned issues, the presented values should be validated by a real-world implementation before considering them reliable.

7.4. MyShake. MyShake [36] is another example of front-detection SAS based on the MEMS accelerometers embedded in common smartphones. The study identifies earthquakes that happen within a 10 km radius and that have a magnitude that is at least 5, as those that can be detected by the average smartphone. When the MyShake application is installed and active, and ad hoc algorithm continuously monitors the accelerometer and communicates shaking data to a central server if a certain condition is triggered. The central server uses a detection algorithm to confirm that an earthquake is underway and, if the output is positive, it calculates location and magnitude and issues an alarm. The paper [36] does not outline the details of the triggering algorithm on the smartphones, but it appears that the system performs an on-phone detection that is subsequently validated by a central server. The MyShake proof-of-concept has been validated by simulations, resulting in a combined delay of 5 seconds

after the origin time. Unfortunately, from the paper it is not possible to identify the delay components, although the overall figure is in line with similar systems [15, 16]. In Figure 7, data of 4 and 6 seconds for MyShake have been added for graphic purposes only.

7.5. Hualien. Hualien is a highly seismic area in Taiwan. Following a strong 6.8 M_L earthquake in 1986, the Central Weather Bureau of the country developed and tested several EEWSs [37], in a time span of 25 years. With reference to Figure 7, we present here the main outcomes of the various SAs presented in [14].

7.5.1. Hualien V0. The first deployed system, not depicted in Figure 7, was a front-detection prototype based on 10 non-MEMS force-balance accelerometer stations that continuously transmitted data to a central mainframe. During the 2-years test period, the system could provide quake warnings in approximately 10 seconds or more. As the system was also designed to provide earthquake localization and to determine quake magnitudes, its scarce reliability in these respects was enough to shut down the system.

7.5.2. Hualien V1. The subsequent iteration of the front-detection system was based on similar principles but included 110 non-MEMS force-balance accelerometer stations organized in virtual subnetworks. This is the system still currently operated by the Taiwan Central Weather Bureau. Every 10 seconds of recordings is processed by each virtual subnetwork in order to determine if an earthquake is occurring and, in the case, its magnitude and hypocenter. V1 is significantly more accurate and precise in earthquake rate detection, magnitude estimation, and hypocenter determination, at the expense of a higher average warning time. This is on average around 19s, although large variations exist on the reported earthquakes, ranging from 13 to 27 seconds. Times like these may be useful for locations 70 km or more away from the epicenter. While the estimation of the hypocenter can be computed in less than 10s after the arrival of the P-waves, magnitude estimation requires higher times, as it needs data from the S-waves [38]. Tests have been also made by incorporating the signals from distant sensor stations; however, it has been found that the addition in warning time is significant, for an almost negligible improvement in hypocenter and magnitude estimation [39].

7.5.3. Hualien V2. To provide warnings to locations within 70 km from the epicenter of an earthquake, an on-site prototype based on a seismographic network has been developed. Seismic signals are continuously transmitted to a central station with IP-based networks. The system tries to identify the peak magnitude of the initial P-wave displacement instead of the magnitude of the earthquake, in order to shorten processing times. Total warning time is not clear. It is claimed that this on-site alert system may offer warnings in less than 10s; among 54 earthquakes detected by the system, however, on average almost 19s have been needed to issue a warning.

7.5.4. Hualien V3. To further reduce the warning times without densifying the network with a lot of high-cost non-MEMS force-balance accelerometers, a hybrid system named *Palert*, based on MEMS accelerometers, was developed by a consortium of industry and academia. *Palert* is both an on-site and a front-detection system. Sensor devices employ a local full-fledged algorithm to detect P-waves and calculate their peak magnitudes. Sensors also send, each second, all the acceleration signals to a central server. If a P-wave is detected on-site, sensor devices start an alert with a warning sound. If the central server recognizes that a certain number of *Palert* stations are triggered, it considers it an earthquake and it starts to compute hypocenter, magnitude, and issues an alarm. For the front-detection slice, a mean value of 15.5s, compared to ~ 19 s, and a median of approximately 14s, instead of ~ 19 s, indicate the V3 ability to issue warnings faster, at the expense of a slightly higher uncertainty regarding magnitude and hypocenter estimation. It was concluded that *Palert* is able to function as an EEWS, for regions located at 60 km or more with respect to the quake epicenter. Regarding the on-site subsystem, it is shown how the system can issue local warnings much faster. Apart from a value of 0.01s, that is probably an error, the majority of stations issued warning times between 1s and 7s.

7.6. Earthcloud. On July 1, 2018, a low-power earthquake of magnitude 3.6 on the Richter scale struck the region at 07:32:16 UTC (09:32:16 local time), with a hypocenter 14 km deep and epicenter 59.11 km (36.73 mi) and 59.69 km (37.09 mi) from the first two devices and the third one, respectively, in a straight line. The farther distance where the earthquake was perceived, instead, has been reported to be approximately 58 km in a straight line from the epicenter. The Earthcloud system determined that an earthquake was underway in times between 0.7s and 0.8s. Each Earthcloud sensor, instead, detected a probable earthquake and communicated a warning in times ranging from 0.1 to 0.2 seconds. All these delays do not include magnitude estimation that is not currently performed by the prototype. The variation in time between sensor devices is due to the different network paths to the cloud system. The processing time is conservatively assumed to be ~ 50 ms for each device (each device has the same processing power, and it is configured exactly the same way as the others). Transmissions between each of the two collocated devices and the AWS servers have a RTT of 106ms, therefore a one-way delay of 53ms. The same figures for the third device are instead 343ms and 172ms, respectively. As the region is the same and the number of devices is low, we considered a positive alarm only the case in which all 3 devices issued a warning in the same timeframe. The two collocated devices needed 0.275s to be aware that all 3 sensor nodes issued a warning, while the third device needed 0.393s instead. Kinesis knew of the three warnings in 172ms, and we assume that it may need up to 600ms to issue an alarm to farther sensor devices. This is confidently a worst-case scenario, as even the delays of geosynchronous satellite communications are usually lower

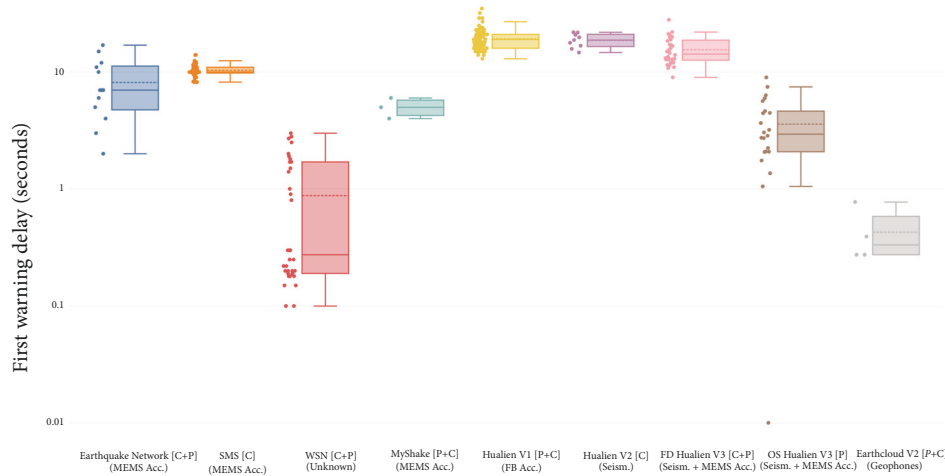


FIGURE 8: Comparison of first warning delays among different SASs, logarithmic scale.

[21]. Figure 8 depicts the same values of Figure 7 but on a logarithmic scale, and it is shown here to clearly compare the results of Earthcloud V2 with those of the other fastest systems.

Data harvested by Earthcloud on July 1, 2018, has been batch processed and cross-referenced with data of the same day from the Italian National Institute of Geophysics and Volcanology (INGV), which in turn sources data from the Italian National Seismic Network and other local, regional and national networks belonging to other national or international institutions. From cloud data gathered through Athena, it was confirmed that the two colocated devices detected the arrival of P-waves on all three directional axes, with peak values of $|0.5| mV$, at 07:32:31 UTC (09:32:31 local time) of the device clocks. The third device did not produce useful data, as it did not manage to differentiate earthquake waves from background noise. Its warning, however, contributed to the system earthquake determination. Several reasons may account for this noise in data; most likely, either there were some other vibration sources, the device was misconfigured, its assembly on the wall poorly performed, or it had some component (e.g., the analog-to-digital converter) failure. Considering a straight line distance of $59.110 km$ at earth level and a hypocenter with a $14.000 km$ depth, an approximation of the actual source point of the earthquake can be calculated as the hypotenuse of the right triangle formed through the previous figures, that is equal to $60.745 km$. Counting a time difference of $15.000 s$ and assuming that device clocks were synchronized with INGV time, the resulting P-wave average velocity between the source and the detection point would be equal to $4.050 km/h$. This figure is slightly higher than the lower end of the often-reported speed interval for P-waves of $4-8 km/h$. Actually, as reported in [40], P-waves can travel from anywhere between 300 and $6500 m/s$, depending on the terrain composition. Considering the presence of significant marshy deposits in the region (Modena, in the past, was a very swampy terrain), which can hamper wave propagation, it is plausible to expect wave propagation velocity attenuations when comparing it

to average values. Furthermore, no other relevant data could be extracted from Athena in a slightly larger timeframe. Therefore, the detection was considered positive.

At 07:32:38 UTC (09:32:38 local time), the *secunda unda* was identified. Earthcloud detected S-waves until 07:32:40 UTC (09:32:40 local time), through data with peak values of $|1.125| mV$ from the colocated devices. Detection started at second 38 and ended at second 40, resulting in a traveling time of $22.000 s$ and $24.000 s$, respectively, that give average velocities of $2.761 km/h$ and $2.531 km/h$, each, calculated following the same logic used for P-waves. The mean value between the two results is $2.646 km/h$. The latter is 65.333% of the P-wave average velocity, an amount that is consistent to the typical approximate figure of 60% that relate S-wave speed with the preceding P-wave speed, when considering the same direction and the same traversed materials. As a consequence, the detection was deemed to be positive.

7.6.1. Challenges. The number of sensor devices deployed in the Earthcloud network is still too low to definitely draw conclusions about the performance of the system. The same applies to the number of earthquakes experienced by the prototype. In particular, crucial information to extract is the relationship between the timing of definitive alerts and the accuracy of earthquake detection.

The time needed to perform the very basic processing functions in sensor devices is fixed. Being very scalable, a very similar property can be expected from the cloud. However, processing times might grow if algorithms for magnitude estimation are employed.

Due to the presence of Content Delivery Networks, middle boxes, etc., in the paths between sensor devices and cloud infrastructure, communication delays will most probably be defined by a certain range. In this regard, the most aleatory factor is instead the connection between sensor devices and the Internet. As described above, we experienced RTTs ranging from 106 to 343 milliseconds. Both include cloud processing time related to the routing of packets.

Being the infrastructure and the entry point the same, the 237ms difference is with certainty attributable to the network path between devices and AWS servers. The third sensor installed, in fact, instead of being directly connected to a university-grade network switch, is connected to a regular home router. As all sensors have been up to now connected to a router/switch through Ethernet cables, higher delays are expected if connecting devices through WiFi.

8. Conclusions

This article presented the second iteration of Earthcloud, a SAS designed to be low-cost, low-power, and cloud-based. Earthcloud processing and communication delays have been analyzed and compared with those of several other systems. Although more data is needed to draw solid conclusions, the Earthcloud prototype provided encouraging results. Soon after the system went operative, a small yet noticeable earthquake hit with a 14 km-deep hypocenter and an epicenter located approximately 60 km in a straight line from where sensor devices were deployed. Earthcloud emitted a number of warnings, validated through the collected data that indicate a significant correlation between the elaborated material and the seismic data published by the relevant national authority. All warnings issued by the prototype have been signaled in less than 1 second. This delay is in stark contrast with traditional SASs, as well as with later low-cost alternatives based on MEMS accelerometers and local processing; in both cases, these systems usually give alerts in approximately 5 seconds when on-site, or in 10 to 20 seconds when front-detection. Part of the reason is that they usually cover a very specific area while employing some form of magnitude and epicenter location estimation. While the latter would be, in perspective, given also by the Earthcloud system as it is, the former would need the employment of additional processing algorithms that would presumably increase the prototype warning times.

Data Availability

The data used to support the findings of this study are available from the corresponding author upon request.

Conflicts of Interest

The authors declare that there are no conflicts of interest regarding the publication of this paper.

Acknowledgments

The authors would like to express their gratitude to both the Department of Engineering Enzo Ferrari and the Ovestlab Center (<http://ovestlab.it>) together with its head, Silvia Tagliacuzzi, for allowing us to use their respective premises and resources for our devices without charge.

References

- [1] R. R. Leach Jr. and F. U. Dowla, "Earthquake early warning system using real-time signal processing," in *Proceedings of the IEEE Signal Processing Society Workshop*, pp. 463–472, September 1996.
- [2] R. Hoque, S. Hassan, M. A. Sadaf, A. Galib, and T. F. Karim, "Earthquake monitoring and warning system," in *Proceedings of the 3rd International Conference on Advances in Electrical Engineering, ICAEE 2015*, pp. 109–112, Bangladesh, December 2015.
- [3] Y. Sherki, N. Gaikwad, J. Chandle, and A. Kulkarni, "Design of real time sensor system for detection and processing of seismic waves for earthquake early warning system," in *Proceedings of the IEEE International Conference on Power and Advanced Control Engineering, ICPACE 2015*, pp. 285–289, India, August 2015.
- [4] T. Uga, T. Nagaosa, and D. Kawashima, "An emergency earthquake warning system using mobile terminals with a built-in accelerometer," in *Proceedings of the 2012 12th International Conference on ITS Telecommunications, ITST 2012*, pp. 837–842, Taiwan, November 2012.
- [5] J. Nachtigall, A. Zubow, R. Sombrutzki, and M. Picozzi, "The challenges of using wireless mesh networks for earthquake early warning systems," in *Proceedings of the 2009 2nd International Conference on Advances in Mesh Networks, MESH 2009*, pp. 155–162, Greece, June 2009.
- [6] J. Nachtigall and J.-P. Redlich, "Wireless alarming and routing protocol for Earthquake Early Warning Systems," in *Proceedings of the 4th IFIP International Conference on New Technologies, Mobility and Security, NTMS 2011*, pp. 1–6, France, February 2011.
- [7] A. Heryana, E. Nugraheni, B. Kusumo, A. F. Rojje, and B. Setiadi, "Applying agile methods in designing an earthquake and landslide early warning system application for Android," in *Proceedings of the 5th International Conference on Computer, Control, Informatics and its Applications, IC3INA 2017*, pp. 80–84, Indonesia, October 2017.
- [8] E. Shayo, A. Mwambela, and S. Kawambwa, "Evaluating quality of service for short message service as a platform to transmit seismic data in earthquake early warning system," in *Proceedings of the 2018 IST-Africa Week Conference, IST-Africa 2018*, Botswana, May 2018.
- [9] E. Heindl, "Peer-to-peer (P2P) earthquake warning system based on collaborative sensing," in *Proceedings of the 2009 3rd IEEE International Conference on Digital Ecosystems and Technologies, DEST '09*, pp. 174–176, Turkey, June 2009.
- [10] F. Finazzi, "The earthquake network project: toward a crowdsourced smartphone-based earthquake early warning system," *Bulletin of the Seismological Society of America*, vol. 106, no. 3, p. 1088, 2016.
- [11] F. Finazzi and A. Fassò, "A statistical approach to crowdsourced smartphone-based earthquake early warning systems," *Stochastic Environmental Research and Risk Assessment*, vol. 31, no. 7, pp. 1649–1658, 2017.
- [12] A. Zaicenco, S. Huffman, and I. Weir-Jones, "Seismic P-wave polarization in the context of on-site early warning system," in *Proceedings of International Conferences on Recent Advances in Geotechnical Earthquake Engineering and Soil Dynamics*, May 2010.

- [13] “Earthquake and Tsunamis - disaster prevention and mitigation efforts,” Japan Meteorological Agency, March 2018, http://www.jma.go.jp/jma/kishou/books/jishintsunami/en/jishintsunami_en.pdf.
- [14] Y.-M. Wu, N.-C. Hsiao, T.-L. Chin, D.-Y. Chen, Y.-T. Chan, and K.-S. Wang, “Earthquake early warning system in taiwan,” in *Encyclopedia of Earthquake Engineering*, pp. 1–12, Springer, Berlin, Germany, 2013.
- [15] E. R. Burkett, D. D. Given, and L. M. Jones, “ShakeAlert - an earthquake early warning system for the United States west coast,” U.S. Geological Survey Open-File Report 2014–1097, USGS Earthquake Science Center, February 2017.
- [16] D. Given, E. Cochran, T. Heaton et al., “Technical implementation plan for the ShakeAlert production system - an earthquake early warning system for the west coast of the united states,” U.S. Geological Survey Fact Sheet 2014–3083, USGS Earthquake Science Center, May 2014.
- [17] M. Klapez, C. A. Grazia, S. Zennaro, M. Cozzani, and M. Casoni, “First experiences with earthcloud, a low-cost, cloud-based iot seismic alert system,” in *Proceedings of the 14th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pp. 262–269, Limassol, Cyprus, October 2018.
- [18] “Power: what are the power requirements?” Web page, Raspberry PI Foundation, 2019, <https://www.raspberrypi.org/documentation/faqs/#pi-power>.
- [19] K. Faber and P. W. Maxwell, “Geophone spurious frequency: what is it and how does it affect seismic data?” in *Proceedings of the SEG Technical Program Expanded Abstracts 1996*, pp. 79–80, 2005.
- [20] M. Peter and M. Lansley, “What receivers will we use for low frequencies?” in *Proceedings of the SEG San Antonio Annual Meeting*, pp. 72–76, 2011.
- [21] M. Casoni, C. A. Grazia, M. Klapez, N. Patriciello, A. Amditis, and E. Sdongos, “Integration of satellite and LTE for disaster recovery,” *IEEE Communications Magazine*, vol. 53, no. 3, pp. 47–53, 2015.
- [22] C. A. Grazia, M. Klapez, N. Patriciello et al., “Integration between terrestrial and satellite networks: The PPDR-TC vision,” in *Proceedings of the 2014 10th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications, WiMob 2014*, pp. 77–84, Cyprus, October 2014.
- [23] M. Klapez, C. A. Grazia, and M. Casoni, “Towards massively multipath transmissions for public safety communications,” in *Proceedings of the 2016 IEEE 12th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pp. 1–7, New York, NY, USA, October 2016.
- [24] M. Casoni, C. A. Grazia, and M. Klapez, “A software-defined 5G cellular network with links virtually pooled for public safety operators,” *Transactions on Emerging Telecommunications Technologies*, vol. 28, no. 3, p. e3092, 2017.
- [25] M. Casoni, C. A. Grazia, and M. Klapez, “SDN-based resource pooling to provide transparent multi-path communications,” *IEEE Communications Magazine*, vol. 55, no. 12, pp. 172–178, 2017.
- [26] M. Klapez, C. A. Grazia, and M. Casoni, “A hybrid algorithm to combine redundancy and concurrency in virtual network resource pooling,” *Internet Technology Letters*, vol. 1, no. 2, p. e18, 2018.
- [27] C. A. Grazia, M. Klapez, N. Patriciello, and M. Casoni, “Pink: Proactive injection into ack, a queue manager to impose fair resource allocation among tcp flows,” in *Proceedings of the 11th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications, WiMob 2015*, pp. 132–137, UAE, October 2015.
- [28] M. Casoni, C. A. Grazia, M. Klapez, and N. Patriciello, “How to avoid TCP congestion without dropping packets: an effective AQM called PINK,” *Computer Communications*, vol. 103, supplement C, pp. 49–60, 2017.
- [29] A. A. Gebremariam, M. Usman, R. Bassoli, and F. Granelli, “SoftPSN: software-defined resource slicing for low-latency reliable public safety networks,” *Wireless Communications and Mobile Computing*, vol. 2018, Article ID 7253283, 7 pages, 2018.
- [30] M. Casoni, C. Augusto Grazia, and P. Valente, “Achieving a high throughput and a low latency through a modular packet scheduler,” *International Journal of Communication Networks and Distributed Systems*, vol. 20, no. 1, pp. 82–109, 2018.
- [31] M. Casoni, C. A. Grazia, M. Klapez, and N. Patriciello, “A congestion control middleware layer with dynamic bandwidth management for satellite communications,” *International Journal of Satellite Communications and Networking*, vol. 34, no. 6, pp. 739–758, 2016.
- [32] M. Casoni, C. A. Grazia, M. Klapez, and N. Patriciello, “QRM: a queue rate management for fairness and TCP flooding protection in mission-critical networks,” *Computer Networks*, vol. 93, pp. 54–65, 2015.
- [33] N. A. Surobhi and A. Jamalipour, “A MANET-based semantic traffic management framework for ubiquitous public safety networks,” *Wireless Communications and Mobile Computing*, vol. 14, no. 12, pp. 1127–1142, 2014.
- [34] D. Giardini, G. Grünthal, K. M. Shedlock, and P. Zhang, “The gshap global seismic hazard map,” *Annals of Geophysics*, vol. 42, no. 6, 1999.
- [35] L. Arcoraci, M. Berardi, F. Bernardini et al., “Rapporto macrosismico sui terremoti del 20 (ML 5.9) e del 29 MAGGIO 2012 (ML 5.8 E 5.3) nella pianura padano-emiliana,” Official Report, Istituto Nazionale di Geofisica e Vulcanologia, July 2012, http://terremoti.ingv.it/images/pdf/QUEST_Emiliana2012_RapportoFinale.pdf.
- [36] Q. Kong, R. M. Allen, L. Schreier, and Y. Kwon, “MyShake: a smartphone seismic network for earthquake early warning and beyond,” *Science Advances*, vol. 2, no. 2, Article ID e1501055, 2016.
- [37] Y.-M. Wu, N.-C. Hsiao, T.-L. Chin, D.-Y. Chen, Y.-T. Chan, and K.-S. Wang, “Earthquake early warning system in Taiwan,” in *Encyclopedia of Earthquake Engineering*, pp. 1–12, Springer, Berlin, Germany, 2014.
- [38] Y.-M. Wu and H. Kanamori, “Experiment on an onsite early warning method for the Taiwan early warning system,” *Bulletin of the Seismological Society of America*, vol. 95, no. 1, pp. 347–353, 2005.
- [39] Y.-M. Wu and T.-L. Teng, “A virtual subnetwork approach to earthquake early warning,” *Bulletin of the Seismological Society of America*, vol. 92, no. 5, pp. 2008–2018, 2002.
- [40] O. Coussy and B. Zinszner, *Acoustics of Porous Media*, Editions Technip, 1987.



Hindawi

Submit your manuscripts at
www.hindawi.com

