

A Global Study of the Risk of Earthquakes to IXPs

Jorik Oostenbrink
Delft University of Technology
J.Oostenbrink@tudelft.nl

Fernando Kuipers
Delft University of Technology
F.A.Kuipers@tudelft.nl

Abstract—In this paper, we study the risk of earthquakes to global Internet infrastructure, namely Internet eXchange Point (IXP) facilities. Leveraging the CAIDA IXPs dataset and publicly available earthquake models and hazard computation tools, we find that more than 50% of the facilities have at least a 2% probability of experiencing potentially damaging levels of shaking, due to earthquakes, within a period of 50 years. Furthermore, we estimate that there is a 10% probability that at least 20 facilities will simultaneously experience potentially damaging levels of shaking within a period of 50 years. Fortunately, our analysis shows that IXPs that host many Autonomous Systems (ASes) tend to be located in less earthquake-prone areas, and that spreading out over multiple facilities significantly reduces the impact of earthquakes to IXPs. Following this observation, we propose a novel metric to help AS operators select peering facilities based on the probability of simultaneous facility failures. We show that applying our metric can significantly increase the resilience of individual ASes, as well as that of the Internet as a whole.

I. INTRODUCTION

The resilience of the Internet has been the focus of many studies. Yet, studies to the Internet's resilience to rare, impactful events, such as natural disasters, are rare themselves. Such events can inflict significant, concentrated damage to Internet infrastructure, disrupting local (and sometimes global) connectivity just when people need it most.

Many of the physical components and facilities making up the Internet may fail under intense levels of shaking [1]–[3]. In this work, we aim to take a global look at the risk of earthquakes to Internet eXchange Points (IXPs). An IXP is a physical infrastructure used by Autonomous Systems (ASes) to directly exchange traffic between their networks. Besides potentially reducing costs (by reducing the amount of traffic delivered via transit providers), IXPs have been shown to also increase Quality-of-Service (QoS) [4].

Given the presence of multiple ASes at each of their facilities, the destruction of IXP facilities could have severe consequences for the Internet as a whole. IXPs do take resiliency measures, such as distributing their services over multiple facilities, and/or rerouting traffic through other IXPs and ASes, in case of failures. But the loss of an IXP facility would certainly cause temporary issues and reduced QoS.

In this paper, using publicly available earthquake models and hazard computation tools, we estimate the hazard to individual IXP facilities, as well as the probability of simultaneous facility failures. Our main findings are:

- Many IXP facilities are at risk of potentially damaging levels of shaking: 32.4% (50.9%) of facilities have at

least a 10% (2%) probability of experiencing potentially damaging levels of shaking within 50 years.

- Facilities that host more ASes tend to be located in less earthquake-prone areas.
- In 50 years, we estimate that there is a 10% probability that at least 20 facilities will simultaneously experience potentially damaging levels of shaking (and a more than 6% probability for IXPs).
- Distributing IXPs over multiple facilities helps. We estimate that the median probability that an IXP with multiple facilities will simultaneously experience potentially damaging levels of shaking at *all* its facilities is well below 1%.

Furthermore, to help operators increase the resilience of their ASes to earthquakes, we propose a new metric for selecting IXP facilities, based on the probability of simultaneous facility failures. Applying our metric can increase the resilience of both individual ASes and the Internet as a whole.

II. RELATED WORK

There have been numerous studies on how to assess the risk of earthquakes and other natural disasters to *single* communication networks [5]–[9]. While crucial to our understanding of disaster risk, such studies focus on the resilience of single communication networks, and their methods and results do not necessarily scale well to the Internet as a whole.

There are few studies on the resilience of the Internet as a whole to disasters. Jyothi studied the risk of solar storms to the Internet by considering a number of risk factors (such as the geographical spread of ASes and datacenters) [10].

Anderson et al. analyzed the risk of wildfires to cellular infrastructure in the United States, by studying which cellular transceivers are under threat from wildfires [11].

Eriksson et al. proposed RiskRoute, a routing framework that can configure routes based on both historical and forecasted outage threats [12].

Durairajan et al. and Mayer et al. used data from the Internet Atlas [13] to analyze the risk of, respectively, global warming [14] and earthquakes [15] to Internet infrastructure in the United States. Both of these works essentially analyze the risk to Internet infrastructure by determining the amount of infrastructure at risk.

The true danger of an earthquake to communication networks is not only the damage it can inflict to any individual point of presence, but also its ability to disrupt multiple points of presence at once. Any approach that only considers the risk

to individual network components in isolation only paints half the picture. For a more thorough analysis, we need to consider which components may be disrupted simultaneously, and with what probability. This requires a more complex approach that considers individual earthquake scenarios.

To the best of our knowledge, we are the first to analyze the risk of natural disasters to the Internet using a large number of realistic disaster scenarios generated based on actual disaster data, as well as the first to assess the risk of earthquakes to Internet infrastructure globally. We combine a set of 19 earthquake hazard models covering approximately 68.9% of global IXP facilities and generate a total of 902,134,602 earthquake scenarios to estimate the risk to individual facilities, as well as the risk of earthquakes to the Internet as a whole.

III. DATASETS

A. IXPs

We use the CAIDA IXPs Dataset [16]. This dataset has been constructed by combining information from PeeringDB, Hurricane Electric, and Packet Clearing House. The dataset gives the geographical locations of IXPs (from all three sources), the locations of facilities (i.e., datacenters) hosting these IXPs (only from PeeringDB), and the autonomous systems (ASes) peering at each IXP. A single facility can host multiple IXPs, and an IXP can be distributed over multiple facilities.

Our study will be on the level of individual facilities. Thus, as a first step, we create a singly facility for each IXP without assigned facilities. We place these facilities at the location of the IXP itself. IXPs without location information (country + city or lon+lat) are filtered out. We also filter out all facilities that do not host an IXP. The resulting dataset contains 1,887 facilities, hosting a total 1,162 IXPs.

Most facilities are already assigned precise geographical locations. For the 220 facilities missing coordinates, we assign the coordinates of their city, as given by Geonames [17]. Two facilities were assigned incorrect coordinates by PeeringDB, placing them in the middle of the ocean. In addition, there was a mismatch between the assigned city and country of some IXPs. We manually corrected the locations of these facilities.

B. OpenQuake Engine

We use the OpenQuake Engine [18] to estimate the earthquake hazard at each facility. The OpenQuake Engine is an open-source software tool for earthquake hazard and risk calculation. One of the key benefits of the OpenQuake Engine is the availability of hazard data for most of the world. This allows us to use largely the same process to determine earthquake risk, independent of the location of a facility.

To calculate earthquake hazards, the OpenQuake engine needs both a seismic source system and a ground motion system. In the remainder of this paper, we will refer to the combination of seismic source system and ground motion system as a *hazard model*.

TABLE I
HAZARD MODELS USED IN OUR CALCULATIONS.

Region	Version	Facilities
Euro-Mediterranean [19]	6.1	629
South America [20]	2016.0.0	158
Australia [21]	2018.032	81
Indian Subcontinent [22]	2.0.1	61
Southeast Asia [20]	2018.0.1	59
Canada [20]	2015.1.1	44
Indonesia [20]	2017.0.0	41
Western Africa [20]	2018.0.0	40
Middle East [23]	1.5.0-2016-10-31	35
Eastern Sub-Saharan Africa [20]	2018.0.0	33
The Caribbean and Central America [20]	2018.0.0	31
New Zealand ¹	04	27
South Africa [20]	2018.0.1	14
Central Asia [24]	1.1	10
The Philippines [20]	2018.1.1	10
The Arabian Peninsula [20]	2018.0.0	9
Taiwan [20]	2015.0.0	9
Northern Africa (2018) [20]	2018.0.0	7
Papua New Guinea [25]	NSHA_2019	3

1) *Hazard Models*: To attain global coverage, we need to combine results from multiple hazard models (see Table I). We only make use of publicly available models that are not under any NDA. With these models, we are able to estimate the hazard to 1,301 out of 1,887 facilities (68.9%).

The public datasets for New Zealand and Central Asia only contained seismic source input models. For these regions, we used the ground motion system specified by GEM² instead.

IV. FACILITIES AT RISK

A. Methods

One of the more common intensity measures in use today is Peak Ground Acceleration (PGA). As the name implies, PGA measures the peak acceleration of the ground during an earthquake. It is seen as a good indicator of earthquake hazard for short buildings (of up to 7 floors) [27]. We have chosen to focus on PGA in this study as it is one of the more intuitive intensity measures, and because we assume most IXP facilities are located in short buildings.

We are interested in (1) the level of shaking we can expect in a given investigation period and (2) how often we can expect potentially damaging levels of shaking at each facility. Both of these objectives can be achieved through a classical Probabilistic Seismic Hazard Analysis (PSHA). Simply stated, a classical PSHA considers all specified earthquake ruptures together with ground motion prediction equations, to compute a hazard curve for each location [28], [29]. A hazard curve gives the probability of exceeding given levels of shaking at a location (or *site*) within a specified investigation time. These curves can be reduced to a hazard map, which shows the level of shaking with a given probability of exceedance (e.g., the PGA with a 2% probability of exceedance) for each site.

¹The Earthquake Rates – National Seismic Hazard Model is owned by GNS Science and is based on the model explained in [26]. The model is held under licence from GNS Science.

²<https://hazard.openquake.org/gem/models>

OpenQuake incorporates epistemic uncertainties within a logic tree. Each path through this logic tree (called a realization in the OpenQuake Engine) constitutes a different combination of ground motion prediction equations and source model. This means that instead of computing a single hazard curve for each site, the engine needs to compute a hazard curve for each realization. Thus, when we discuss a probability of exceedance within this paper, we are actually referring to a *mean* probability of exceedance over all hazard curves.

1) *Damaging Levels of Shaking*: PGA is an objective measure of ground-motion due to an earthquake; it is not a direct measure of the damage to buildings and infrastructure. In contrast, a macroseismic intensity scale, such as the Modified Mercalli Intensity scale (MMI), measures the observable effects of an earthquake. In some papers and hazard maps (e.g., [8], [15], [30]), a macroseismic intensity of 6 (in MMI or the Mercalli-Cancani-Sieberg (MCS) scale) is used as a sort of lower-bound for potentially damaging levels of shaking³.

Unfortunately, it is not straightforward to convert PGA to a macroseismic intensity. For one, there are inherent regional differences in the relationship between ground motion and macroseismic intensity. Caprio et al. quantified some of these regional differences, and constructed global ground motion to intensity conversion equations (for a combined MMI/MCS intensity scale) [31]. While one would preferably use regional conversion equations, the global scope of our study makes the global equations a practical, albeit imperfect, alternative.

A global macroseismic intensity of 6 roughly corresponds to a PGA of 0.086g, which we will use as a threshold for potentially damaging levels of shaking. For comparison, using conversion equations for California [32] would result in a threshold of 0.11g (or 0.084g if we round up from an intensity of 5.5), and Mayer et al. assumed infrastructure is potentially damaged if the PGA exceeds 0.092g [15].

2) *Calculation Setup*: We run a classical PSHA with an investigation time of 50 years on each hazard model⁴. The configuration of these calculations is described in the appendix. We compute the probability of exceeding a PGA of 0.086g in 50 years, as well as the PGA with a probability of 10% and 2% of being exceeded in 50 years. The 10% and 2% probabilities of exceedance in 50 years are two common choices for seismic hazard maps.

B. Results

Fig. 1 shows the PGA with a 2% probability of exceedance of each unique location of the facilities covered by one of the hazard models, as well as those in the conterminous United States. In this section, we discuss the facilities covered by the hazard models. For a more complete analysis, we will briefly discuss the hazard of US facilities in Section IV-D.

³Note that the building itself does not need to be damaged to disrupt an IXP facility. A facility could also be disrupted if equipment inside the building is damaged or falls down, or if infrastructure in the surrounding area is damaged.

⁴The hazard models of the Caribbean and Central America and the Philippines are fixed at an investigation time of 1 year. We convert their results to a 50-year investigation time by assuming Poissonian occurrences.

TABLE II
THE NUMBER OF FACILITIES WITH GIVEN PROBABILITIES OF EXCEEDING POTENTIALLY DAMAGING LEVELS OF SHAKING (PGA OF 0.086) WITHIN A PERIOD OF 50 YEARS.

Probability of Exceedance	facilities
≤ 0.01	496
0.01 - 0.02	143
0.02 - 0.1	240
0.1 - 0.2	116
0.2 - 0.5	113
0.5 - 0.8	130
0.8 - 1	63

The hazard models cover a total of 1,301 facilities spread out over 1,135 unique locations. Together, these facilities host 849 unique IXPs. Fig. 2 shows the PGA versus the number of facilities with at least a 10% (respectively 2%) probability of exceeding this PGA in our investigation time of 50 years. A significant number of facilities are at risk of potentially damaging levels of shaking. While the median PGA with a 10% probability of exceedance is only 0.0333g, the median PGA with a 2% probability of exceedance is 0.0928g - just above our threshold of 0.086g.

Table II shows the number of facilities with given probabilities of exceeding potentially damaging levels of shaking within a period of 50 years. 422 (32.4%) facilities have at least a 10% probability of experiencing potentially damaging levels of shaking within a period of 50 years and 662 (50.9%) facilities at least a 2% probability.

Of course, not every facility is equally important. To measure the importance of each facility, we count the number of ASes at each IXP. Although the dataset does not contain all ASes that peer at every IXP, we expect this number to be proportional to the real number of ASes at an IXP. We set the weight of each facility to the sum of the number of ASes of each of the IXPs it hosts. Most IXPs host few ASes: the median number of ASes at an IXP is 11, and there are only 138 IXPs (out of 1162) with at least 100 ASes.

Fig. 3 shows the weight and probability of exceeding potentially damaging levels of shaking of each facility covered by one of the hazard models. Overall, facilities with a larger weight have a lower probability of experiencing damaging levels of shaking: the median probability of exceeding potentially damaging levels of shaking within a period of 50 years is respectively 0.0355, 0.0134, and 0.00576 for facilities with a weight below 100, at least 100, and at least 1,000. However, there are a number of high-weight facilities in higher-risk areas: there are 215 (of 502) facilities with a weight of at least 100 that have at least a 2% probability of exceeding potentially damaging levels of shaking, and 26 (of 84) facilities with a weight of at least 1,000 that have at least a 2% probability of exceeding potentially damaging levels of shaking.

C. Country-Level Analysis

In this section, we analyze the risk of earthquakes to IXP facilities on a country-level by mapping each facility to the region denoted by its ISO 3166 two-letter country code [33].

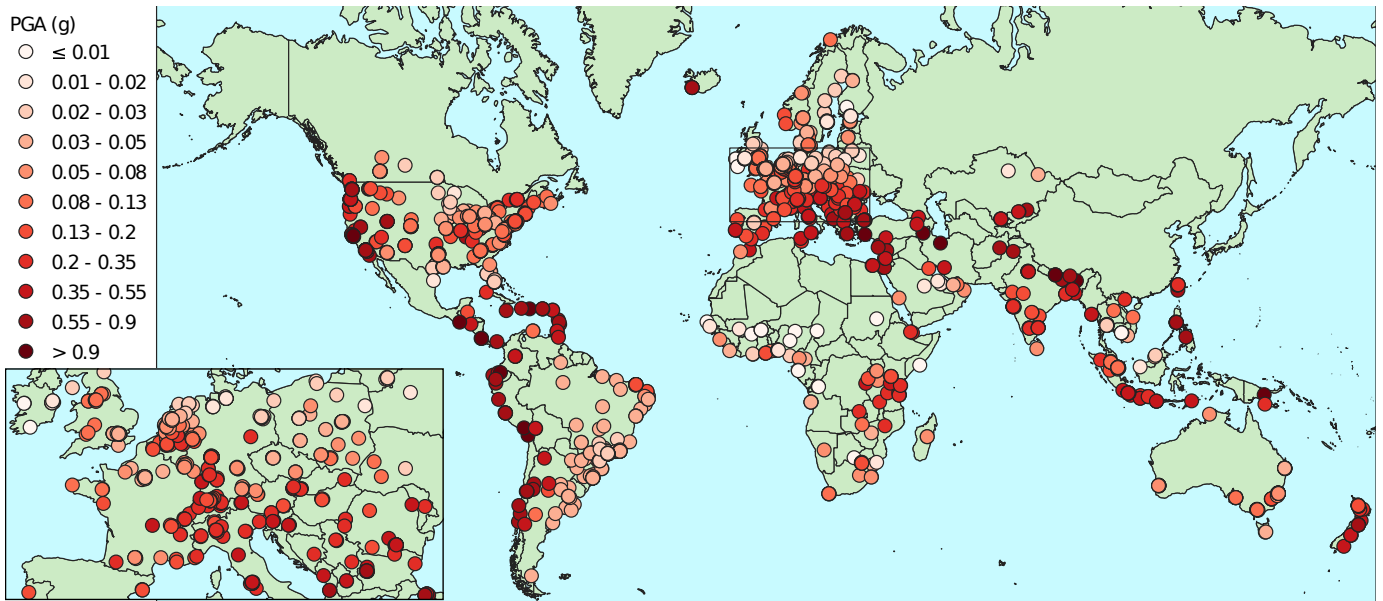


Fig. 1. Locations of facilities covered by a hazard model from Table I, and the local PGA with a 2% probability of being exceeded in 50 years. Results for the conterminous US were added by extracting PGA values from the 2018 USGS long-term seismic hazard map [30].

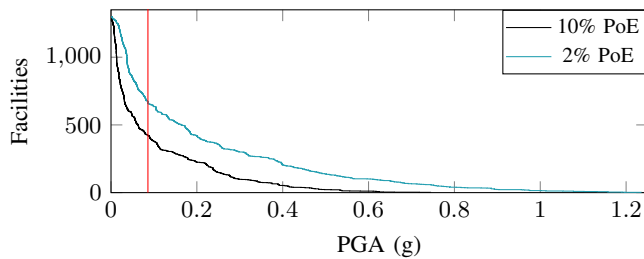


Fig. 2. The number of facilities with at least a 10% (2%) probability of exceeding a given PGA in 50 years. The red line indicates our threshold of potentially damaging levels of shaking.

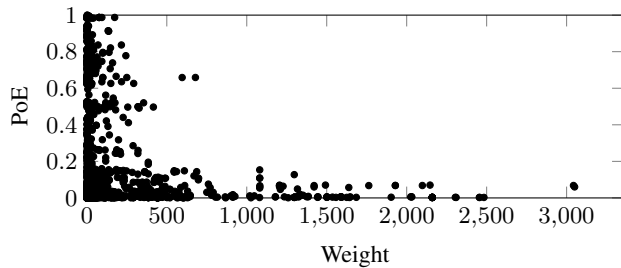


Fig. 3. Weight of each facility versus the probability of exceeding potentially damaging levels of shaking within a period of 50 years.

Fig. 4 shows the median probability of exceeding potentially damaging levels of shaking for each country. This essentially shows the earthquake hazard that an average facility in each country faces. These values are affected by both the frequency and intensity of earthquakes in each country, as well as the exact placement of facilities within the country.

The median probability does not give the full picture, and

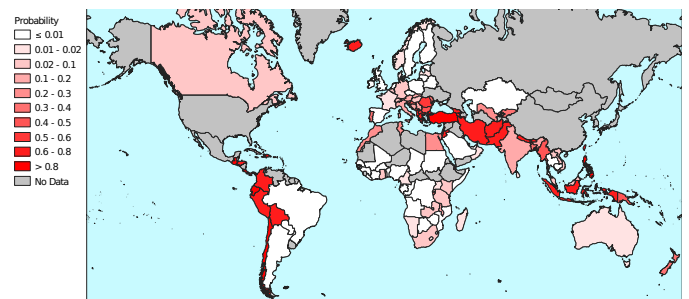


Fig. 4. The median probability of exceeding potentially damaging levels of shaking within a period of 50 years at each facility of every country. Countries with either (1) no facilities or (2) facilities that were not assigned to a hazard model are excluded (No Data).

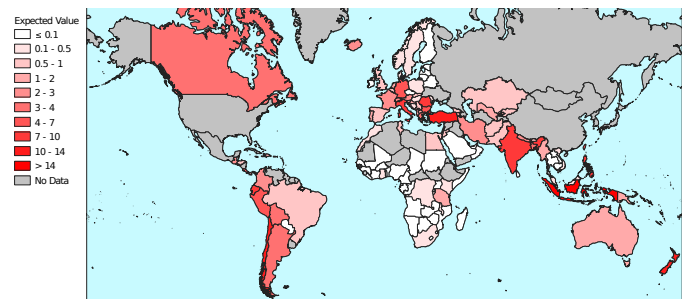


Fig. 5. The expected number of facilities that will experience potentially damaging levels of shaking within a period of 50 years in each country. Countries with either (1) no facilities or (2) facilities that were not assigned to a hazard model are excluded (No Data).

even hides the influence of any outliers within a country. Risk is a combination of probability and impact. Thus, what we are more interested in is the number of facilities that could be disrupted by earthquakes in each country.

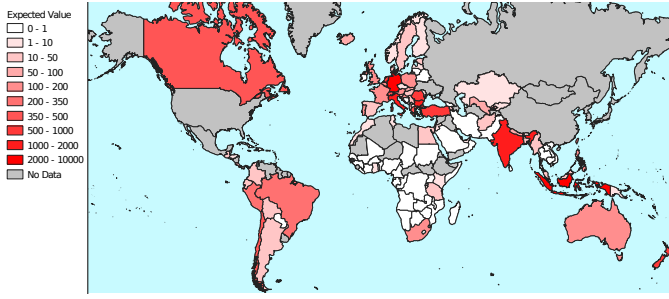


Fig. 6. The expected total weight of facilities that will experience potentially damaging levels of shaking within a period of 50 years in each country. Countries with either (1) no facilities or (2) facilities that were not assigned to a hazard model are excluded (No Data).

Fig. 5 shows the expected number of facilities in each country that will experience potentially damaging levels of shaking at least once within a period of 50 years. We can see that the risk in countries with a low median probability of exceeding potentially damaging levels of shaking can still be relatively high, simply due to the number of facilities. Similarly, some countries with a high median probability of exceeding potentially damaging levels of shaking are at lower risk than expected, because they do not host many facilities.

Indonesia is both prone to large earthquakes, and hosts a reasonably high number of IXP facilities (38). As such, it is the country with the highest expected number of facilities that will experience potentially damaging levels of shaking (20.8). Out of all countries covered by our hazard models, Germany hosts most IXP facilities (101). While it is not the most earthquake-prone country we have studied, it still ranks as the country with the 14th highest expected number of facilities that will experience potentially damaging levels of shaking (4.16).

As we discussed in the previous section, not every facility is equally important. Fig. 6 shows the sum of the product of weight and exceedance probability of each facility for each country. That is, the total expected weight of the facilities in each country that will experience potentially damaging levels of shaking at least once within a period of 50 years. Following this metric, Indonesia ranks as the second-most country at risk (with an expected weight of 2,021). Due to its concentration of high-weight facilities, Germany has the highest total expected weight of facilities that will experience potentially damaging levels of shaking (2,843⁵). Clearly, different weight functions may lead to a different ranking.

D. Conterminous United States

Out of the 1,887 facilities in the dataset, 390 are located in the conterminous United States. The US has the 4th highest total weight of all countries. Although we lack a hazard model for the United States, we would be remiss if we completely ignore it. In this section, we give a brief analysis of the seismic hazard to IXPs in the United States based on the 2018 USGS long-term seismic hazard maps [30].

⁵Note that this is only 5.72% of Germany's total weight of 49,710.

The hazard maps give hazard data for a grid of points spread out over the conterminous United States. To determine the hazard for each facility, we map it to its closest grid point. We first extract the PGA with a 2% probability of exceedance for site class B/C⁶. It seems the average hazard at US facilities is only slightly higher than that of the rest of the world; the median PGA with a 2% probability of exceedance is 0.103g, compared to 0.0928g in the rest of the world.

The USGS includes a map of the chance of “slight (or greater) damaging earthquake shaking in 100 years” (i.e., the probability of MMI of 6 or higher), which can be easily converted to 50-year probabilities. Note that, while very similar, these probabilities were computed in a different manner than our probability of experiencing potentially damaging levels of shaking, and thus are not perfectly comparable.

As in the rest of the world, a large number of US facilities are at risk of earthquakes; the median probability of experiencing slight (or greater) damaging earthquake shaking in 50 year is 0.0312. Out of the 390 facilities, 89 (22.8%) have at least a 10% probability of experiencing damaging earthquake shaking, and 279 (71.5%) at least a 2% probability.

The expected number of US facilities that will experience damaging earthquake shaking in 50 years is 68.9. While this is indeed more than any other country we analyzed, the United States also contains by far the most facilities of all countries. The expected total weight of US facilities that will experience damaging earthquake shaking in 50 years is 11,384 (much more than any other country!). In the US, more than in the rest of the world, a large number of facilities with relatively high number of ASes are at high risk of damaging earthquakes.

V. COMBINED FAILURES

Whereas in the previous section we considered facilities individually, in this section we study the risk of simultaneous facility outages. In other words, we study the potential disruption of multiple IXP facilities due to a single earthquake. To this end, we first run an event-based PSHA in OpenQuake. In contrast to a classical PSHA, an event-based PSHA randomly generates sets of earthquake events, called stochastic event sets, as well as ground motions at each site during each of these events. A single stochastic event set is a realisation of potential earthquakes during the full duration of the investigation time. By generating multiple event sets, and processing the resulting ground motion fields, we can estimate which facilities could potentially be disrupted simultaneously.

A. Disruption

In this section, we say a facility is *disrupted* by an earthquake if it experiences shaking with a PGA of at least 0.086g. In addition, we say an IXP is disrupted if at least one of its facilities is disrupted, and is *fully* disrupted if all of its facilities are disrupted. Since our threshold of 0.086g is a lower bound on potentially damaging levels of shaking, this gives us a pessimistic view of the potential impact of an earthquake.

⁶Roughly equivalent to the site class used for our own analysis.

We run an event-based PSHA with almost exactly the same settings as we did for the classical PSHA. To reduce computation time and memory usage, we sample logic trees with more than 200 realizations 200 times. For each sampled realization, we generate 200 seismic event sets⁷. In total, we generate 902,134,602 events, out of which 8,615,935 disrupt one or more facilities.

Analogously to the probability of exceedance, our goal will be to compute the mean complementary cumulative distribution function (CCDF) of the worst-case impact of an earthquake within a period of 50 years. Since the Open-Quake Engine assumes earthquake occurrences are Poissonian, estimating these probabilities for a single hazard model is straightforward. Unfortunately, combining results from multiple hazard models is more complex. Fortunately, under some conditions, we can combine mean probabilities.

Lemma 1: Let n be the number of hazard models, and let X_1, \dots, X_n be random variables measuring the number of events of interest in each hazard model. Furthermore, let R_i be the realizations of hazard model i , and w_r the weight of realization $r \in R_i$.

We define the mean probability

$$\begin{aligned} \bar{P}\left(\sum_{i=1}^n X_i \geq 1\right) &= \\ \sum_{r_1 \in R_1} w_{r_1} \cdots \sum_{r_n \in R_n} w_{r_n} P\left(\sum_{i=1}^n X_i \geq 1 | r_1, \dots, r_n\right) \end{aligned} \quad (1)$$

If X_1 to X_n are mutually independent, then

$$\bar{P}\left(\sum_{i=1}^n X_i \geq 1\right) = 1 - \prod_{i=1}^n \bar{P}(X_i = 0) \quad (2)$$

where

$$\bar{P}(X_i = 0) = \sum_{r \in R_i} w_r P(X_i = 0 | r) \quad (3)$$

Proof: Since

$$P\left(\sum_{i=1}^n X_i \geq 1 | r_1, \dots, r_n\right) = 1 - P\left(\sum_{i=1}^n X_i = 0 | r_1, \dots, r_n\right) \quad (4)$$

and the realization weights of each hazard model sum to 1, we can reformulate Equation 1 as

$$\begin{aligned} \bar{P}\left(\sum_{i=1}^n X_i \geq 1\right) &= \\ 1 - \sum_{r_1 \in R_1} w_{r_1} \cdots \sum_{r_n \in R_n} w_{r_n} P\left(\sum_{i=1}^n X_i = 0 | r_1, \dots, r_n\right) \end{aligned} \quad (5)$$

⁷For hazard models with an investigation time of 1 year, we generate 10,000 seismic events sets per realization instead.

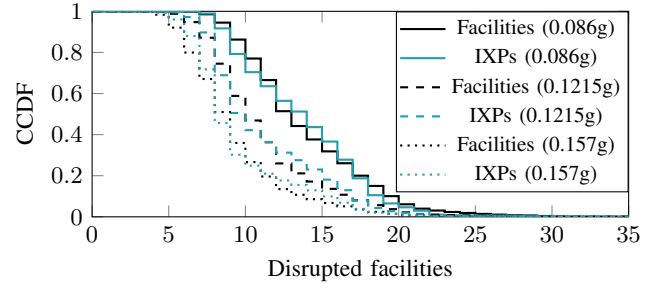


Fig. 7. The complementary CDF of the maximum number of facilities (IXPs) that are simultaneously disrupted by a single earthquake within a period of 50 years.

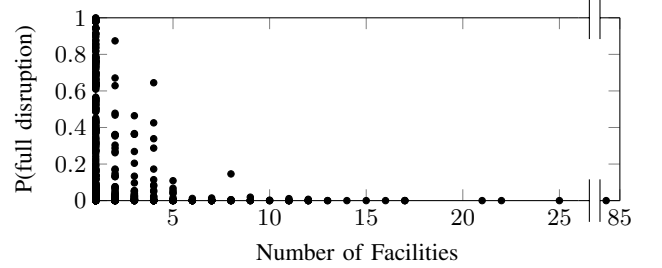


Fig. 8. The total number of facilities of each IXP, and their probabilities of full disruption within a period of 50 years.

Now, since we have mutual independence and hazard model i only depends on realization r_i :

$$\begin{aligned} 1 - \sum_{r_1 \in R_1} w_{r_1} \cdots \sum_{r_n \in R_n} w_{r_n} P\left(\sum_{i=1}^n X_i = 0 | r_1, \dots, r_n\right) &= \\ 1 - \sum_{r_1 \in R_1} w_{r_1} \cdots \sum_{r_n \in R_n} w_{r_n} \prod_{i=1}^n P(X_i = 0 | r_i) &= \\ 1 - \prod_{i=1}^n \bar{P}(X_i = 0) \end{aligned} \quad (6)$$

Equation 2 allows us to estimate any overall mean CCDF, by separately computing the mean estimated probability of zero events of interest for each hazard model.

Our approach ignores the potential overlap between different hazard models. Consider the ESHM13 and EMME14 hazard models for example. We use ESHM13 to estimate the hazard for facilities in Europe, and EMME14 for estimating the hazard in the Middle East. As these areas border each other, it is possible that an earthquake would disrupt facilities in both Europe and the Middle East. Our approach ignores this possibility, and thus potentially overestimates the total number of earthquakes (since multiple hazard models may model the same seismic sources), while underestimating the impact of some of these earthquakes. This problem is an inherent disadvantage of combining multiple hazard models.

1) Results: We first consider the number of disrupted facilities. As can be seen in Fig. 7, the mean probability that at

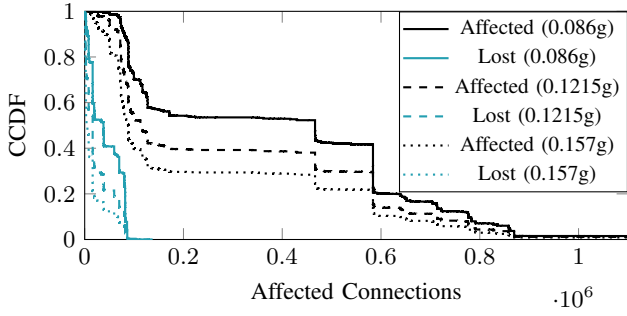


Fig. 9. The complementary CDF of the maximum number of connections that are affected by a single earthquake within a period of 50 years.

least one facility will be disrupted within 50 years is nearly 1. Worryingly, there are many events that would disrupt multiple facilities at once. There is a 10% probability that at least 20 facilities will be disrupted by a single earthquake. Given the level of facility sharing between IXPs, this could have a significant impact on the Internet.

Contrary to our expectations, the number of disrupted IXPs is often lower than the number of disrupted facilities. Furthermore, the worst-case number of simultaneously disrupted facilities is quite a bit lower than the worst-case number of simultaneously disrupted IXPs: 72 facilities compared to 46 IXPs. This shows that a number of IXPs are distributed over facilities that can be struck by the same earthquake.

For comparison, we also consider higher PGA thresholds (Fig. 7). While there is a clear decrease in earthquake impact if we increase the threshold to 0.157g (roughly corresponding to a macroseismic intensity of 7), the probability of simultaneous facility disruption is still quite high: There is a 3.7% probability that at least 20 facilities will simultaneously experience this level of shaking. Nevertheless, these results show that the choice of PGA threshold greatly influences our results, and that, since we chose a more pessimistic threshold, we are potentially overestimating the impact of earthquakes on IXPs.

Our results raises the question if IXPs spread their facilities over a large enough area. We compute the probability of full disruption of each of the 828 IXPs whose facilities are located in the area covered by the hazard models. It seems like distributing IXPs over multiple facilities helps: the median probability that an IXP is fully disrupted at least once in a 50-year period is 0.0118, while the median probability that an IXP with at least two facilities is fully disrupted is 0.00220. As can be seen in Fig. 8, IXPs with more facilities tend to have a lower probability of experiencing full disruption. These results should also translate to individual ASes. By peering with the same neighbors at multiple locations, an AS can significantly reduce the risk of earthquakes to its connectivity.

2) *Impact on Connectivity*: To get a better idea of the impact of these events, we again consider the ASes hosted at each IXP. We assume that, within each IXP, every AS peers with every other AS. We then define a unique (potential) *connection* for every pair of ASes that share at least one IXP. While this is an overestimate of the actual peering

density at each IXP, the number of connections should be roughly proportional to the actual number of peering links. Furthermore, the loss in connections due to IXP disruption is equivalent to the loss in available peering links at IXPs.

We assign two impact metrics to each event: (1) the number of affected connections, and (2) the number of lost connections. If two ASes share a disrupted facility, we mark their connection as affected. If the two ASes share no other undisrupted facility, the connection has no remaining backup and we mark it as lost. Note that this does not mean that these two ASes are completely disconnected from each other (packets can potentially still be routed through other ASes or through direct peering outside an IXP), but it does mean that these two ASes can not exchange packets directly at any remaining IXP. In this manner, the metric is a good indicator of impact on the IXP ecosystem.

We note that a large majority of connections have a backup (Fig. 9). This shows the power of peering at multiple IXPs. Even if some facilities are disrupted by an earthquake, there is often another facility available that serves as a suitable backup.

That being said, the number of lost connections is still very high, even at higher probabilities (and at higher PGA thresholds). At best, this means that in case of a strong earthquake, a large number of BGP routes will need to be rerouted. At worst, ASes will be completely disconnected from the rest of the Internet.

B. Increasing Redundancy - A Novel Metric

As we discussed in the previous section, operators can reduce the impact of earthquakes on their ASes by peering at multiple facilities. However, selecting a new facility is not trivial. Clearly, peering at a facility with low probability of exceeding damaging levels of shaking helps reduce the risk of earthquakes. The results from Section IV-B suggest this factor is already taken into account: facilities that host more ASes tend to have a lower probability of exceeding damaging levels of shaking. However, only peering at low-risk facilities might not always be possible or cost-efficient, and, although less frequently, even a low-risk facility can be struck by an earthquake. Thus, to effectively reduce the risk of earthquakes, an operator would need to consider both the probability that its facilities will be disrupted by the same earthquake, as well as the redundancy of connections at each of its facilities.

We propose a novel metric for evaluating sets of peering locations with respect to earthquake risk. Our metric can be applied to IXP facilities, as well as to private peering. The aim of the metric is to ensure the probability that any of a selection of important connections is disconnected by an earthquake remains below a pre-selected threshold.

Definition 1 (Earthquake-Resistant Peering Metric): Suppose we are given a set of weights w_i for all ASes, a set of potential facilities F , the cost of peering at each facility $f \in F$, $c(f)$, and a threshold, $t \in [0, 1]$. Let $h_i \subseteq F$ be the subset of all facilities hosting AS i .

Given a selection of facilities $s \subseteq F$, the mean probability that the connection with AS i will be disrupted due to an

earthquake in an arbitrary time-frame is equivalent to the mean probability that facilities $h_i \cap s$ will simultaneously be disrupted due to an earthquake in that time-frame. We denote this probability by $p(h_i \cap s)$, and compute it using Equation 2.

We define the value of a selection of facilities $s \subseteq F$ as

$$\sum_i w_i I_i(s) - \sum_{f \in s} c(f) \quad (7)$$

where

$$I_i(s) = \begin{cases} 1 & \text{if } h_i \cap s \neq \emptyset \text{ and } p(h_i \cap s) \leq t \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

Note that one can easily extend this metric to require connectivity with only one out of a set of ASes, or to set individual thresholds per AS.

1) *Evaluation:* To evaluate our metric, we set a threshold of 0.01 in 50 years, and extract all ASes with at least one connection with a disruption probability above this threshold. We filter out any facilities outside of our hazard models, and any ASes peering at one of these facilities. Our goal will be to increase the resilience of the remaining 4,594 ASes against earthquakes, by connecting each AS to one additional facility.

For the purpose of this experiment, we consider each combination of IXP and facility (hosting the IXP) to be a unique facility. For each AS we aim to protect, we set the cost of each facility to 0, the weight of each of its current peers to 1, and the weight of all other ASes to 0. That is, our goal is to find the IXP-facility pair that protects as many of the current connections as possible.

Out of the 4,594 ASes, we find a new facility for 4,420. For the other 174 ASes, there is no possible facility that would reduce the disconnection probability of any of its peers to below our 0.01 threshold. The mean number of connections that were previously unsafe that can be protected by adding a single facility is 31.8%. But, the mean distance between the closest old facility and this new facility is 2,579km.

If we restrict ourselves to the countries each AS currently peers at, we find a solution for 3,721 ASes. The average distance to the new facility is now 569km, and the facility protects an average of 26.6% of previously unsafe connections.

For 2,280 ASes (almost 50%), we can even find a new facility within 100km of their old facilities. These facilities protect an average of 20.2% of previously unsafe connections, while their average distance to the old facilities is only 24km.

Fig. 10 shows the effect of peering at *all* of these facilities on the number of lost connections during an earthquake. Since we chose to protect currently existing connections, we only consider these original connections. We can see that connecting to additional facilities did indeed protect many connections against earthquakes. Interestingly, while peering at additional facilities within the same country increased the resilience of both individual ASes and the Internet as a whole against earthquakes, restricting facilities to a distance of 100km of old facilities greatly reduced the benefit to the overall resilience of the Internet.

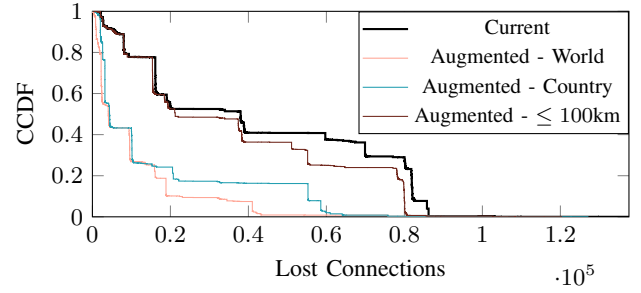


Fig. 10. The complementary CDF of the maximum number of original connections that are lost due to a single earthquake within a period of 50 years, before and after spreading ASes over more facilities.

VI. DISCUSSION

Our analysis is a best-effort analysis of the risk of earthquakes to global IXP infrastructure. The maps included in this paper are not meant to be used to support any important decision involving human life, capital and movable and immovable properties. Due to the scale of our analysis, and our selection of hazard models, a number of concessions were made. The ground motion systems of our hazard models cannot account for the intra-event spatial correlation of ground motions. Furthermore, due to a lack of data, we assume the conditions of each site are equivalent (to reference rock). These conditions affect the level of shaking, and it is possible that some facility locations have been purposely placed in areas that are less susceptible to earthquakes. In addition, since we lack data on the characteristics of each facility as well, we say a facility is disrupted if it experiences potentially damaging levels of shaking. When building characteristics are known, one can use fragility curves to estimate a probability of damage instead.

VII. CONCLUSION

We have conducted the first global study of the risk of earthquakes to Internet infrastructure. We find that a large number of IXP facilities are at risk of earthquakes. On the positive side, IXP facilities that host a large number of ASes tend to be located in less earthquake-prone areas.

We confirm the effectiveness of spreading out over multiple facilities: IXPs with more facilities tend to have a greatly reduced probability that all their facilities are disrupted simultaneously. However, we find that not all ASes spread out over IXPs sufficiently. To this end, we have proposed a novel metric for selecting new peering locations, which takes into account earthquake hazard at both current and new facilities, and the probability of combined facility failures. We have demonstrated the effectiveness of our metric in reducing the number of lost peering connections by finding a selection of new facilities for ASes that are currently at risk of earthquakes.

APPENDIX

The hazard models for Europe, Australia, the Indian sub-continent, the Middle East, and Papua New Guinea include initial configuration files. For these models, we kept the

calculation and site parameters. For all other hazard models, we set the site attributes to the same reference values used in the ESHM13 (corresponding to a reference rock condition matching Eurocode 8 Type A). And we set `rupture_mesh_spacing` to 5^8 , `width_of_mfd_bin` to 0.1, and `area_source_discretization` to 10. To prevent very high, potentially unrealistic estimates of the level of shaking, the tail-end of the ground motion distribution is usually cut off; We apply a truncation level of 3.

One of the more important parameters is the maximum distance between ruptures and sites at which the OpenQuake engine still considers the rupture when computing the hazard at a site. We set this distance to an, in our eyes, conservative level of 800km. For Canada, we indicate a maximum distance per tectonic region type, as described in [34].

REFERENCES

- [1] M. Kazama and T. Noda, "Damage statistics (summary of the 2011 off the pacific coast of tohoku earthquake damage)," *Soils and Foundations*, vol. 52, no. 5, pp. 780–792, 2012, special Issue on Geotechnical Aspects of the 2011 off the Pacific Coast of Tohoku Earthquake.
- [2] Z. E. Khaled and H. McHeick, "Case studies of communications systems during harsh environments: A review of approaches, weaknesses, and limitations to improve quality of service," *International Journal of Distributed Sensor Networks*, vol. 15, no. 2, 2019.
- [3] S. Giovinazzi, A. Austin, R. Ruiter, C. Foster, M. Nayyerloo, N.-K. Nair, and L. Wotherspoon, "Resilience and fragility of the telecommunication network to seismic events," *Bulletin of the New Zealand Society for Earthquake Engineering*, vol. 50, no. 2, pp. 318–328, Jun. 2017.
- [4] M. Di Bartolomeo, G. Di Battista, R. di Lallo, and C. Squarcella, "Is it really worth to peer at ixps? a comparative study," in *2015 IEEE Symposium on Computers and Communication (ISCC)*, 2015, pp. 421–426.
- [5] J. Oostenbrink and F. Kuipers, "Computing the impact of disasters on networks," *SIGMETRICS Perform. Eval. Rev.*, vol. 45, no. 2, p. 107–110, Oct. 2017.
- [6] J. Tapolcai, B. Vass, Z. Heszberger, J. B  r  , D. Hay, F. A. Kuipers, and L. R  nyai, "A tractable stochastic model of correlated link failures caused by disasters," in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, 2018, pp. 2105–2113.
- [7] X. Wang, M. Chen, and S. Lu, "Modeling geographically correlated failures to assess network vulnerability," *IEEE Transactions on Communications*, vol. 66, no. 12, pp. 6317–6328, 2018.
- [8] A. Valentini, B. Vass, J. Oostenbrink, L. Cs  k, F. Kuipers, B. Pace, D. Hay, and J. Tapolcai, "Network resiliency against earthquakes," in *2019 11th International Workshop on Resilient Networks Design and Modeling (RNDM)*, 2019, pp. 1–7.
- [9] H. Talebayan, K. Leelardcharoen, L. Due  as-Osorio, B. J. Goodno, and J. I. Craig, "Congestion and observability across interdependent power and telecommunication networks under seismic hazard," *Earthquake Spectra*, vol. 37, no. 4, pp. 2892–2919, 2021.
- [10] S. A. Jyothi, "Solar superstorms: Planning for an internet apocalypse," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, ser. SIGCOMM '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 692–704.
- [11] S. Anderson, C. Barford, and P. Barford, "Five alarms: Assessing the vulnerability of us cellular communication infrastructure to wildfires," in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 162–175.
- [12] B. Eriksson, R. Durairajan, and P. Barford, "Riskroute: A framework for mitigating network outage threats," in *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 405–416.
- [13] R. Durairajan, S. Ghosh, X. Tang, P. Barford, and B. Eriksson, "Internet atlas: A geographic database of the internet," in *Proceedings of the 5th ACM Workshop on HotPlanet*, ser. HotPlanet '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 15–20.
- [14] R. Durairajan, C. Barford, and P. Barford, "Lights out: Climate change risk to internet infrastructure," in *Proceedings of the Applied Networking Research Workshop*, ser. ANRW '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 9–15.
- [15] J. Mayer, V. Sahakian, E. Hoof, D. Toomey, and R. Durairajan, "On the resilience of internet infrastructures in pacific northwest to earthquakes," in *Passive and Active Measurement: 22nd International Conference*, 2021, pp. 247–265.
- [16] The CAIDA UCSD IXPs Dataset, 2021-07. CAIDA. [Online]. Available: <https://www.caida.org/catalog/datasets/ixps>
- [17] GeoNames. [Online]. Available: <http://www.geonames.org/>
- [18] OpenQuake Engine. Global Earthquake Model Foundation. [Online]. Available: <https://github.com/gem/oq-engine>
- [19] J. Woessner, D. Laurentiu, D. Giardini, H. Crowley, F. Cotton, G. Gr  nthal, G. Valensise, R. Arvidsson, R. Basili, M. B. Demircioglu, S. Hiemer, C. Meletti, R. Musson, A. Rovida, K. Sesetyan, and M. Stucchi, "The 2013 european seismic hazard model: key components and results," *Bulletin of Earthquake Engineering*, vol. 13, pp. 3553–3596, 2015.
- [20] GEM Hazard Models. Global Earthquake Model Foundation. [Online]. Available: <https://www.globalquakemodel.org/products>
- [21] T. Allen, J. Griffin, and D. Clark. (2019) The 2018 National Seismic Hazard Assessment for Australia: Model input files. Record 2018/032. Geoscience Australia. [Online]. Available: <http://dx.doi.org/10.11636/Record.2018.032>
- [22] N. Ackerley. Indian Subcontinent PSHA. [Online]. Available: <https://github.com/nackerley/indian-subcontinent-psha>
- [23] L. Danciu, K. Sesetyan, M. Demircioglu, M. Erdik, and D. Giardini, "Openquake input files of the seismicogenic source model of the 2014 earthquake model of the middle east (emme-project)," 2016.
- [24] S. Ullah, K. Abdrakhmatov, A. Sadykova, R. Ibragimov, A. Ishuk, D. Laurentiu, S. Parolai, D. Bindi, M. Wieland, and M. Pittore. (2015) Emca central asia seismic source model. v. 1.1. [Online]. Available: <https://doi.org/10.5880/GFZ.EWS.2015.002>
- [25] Papua New Guinea Seismic Hazard Assessment. Geoscience Australia. [Online]. Available: <https://github.com/GeoscienceAustralia/PNGSHA>
- [26] M. Stirling, G. McVerry, M. Gerstenberger, N. Litchfield, R. Van Disen, K. Berryman, P. Barnes, L. Wallace, P. Villamor, R. Langridge, G. Lamarche, S. Nodder, M. E. Reyners, B. Bradley, D. A. Rhoades, W. D. Smith, A. Nicol, J. Pettinga, K. J. Clark, and K. Jacobs, "National seismic hazard model for new zealand: 2010 update," *Bulletin of the Seismological Society of America*, vol. 102, no. 4, pp. 1514–1542, 2012.
- [27] Earthquake Hazards 201 - Technical Q&A. USGS. Accessed: 2021-07-20. [Online]. Available: <https://www.usgs.gov/natural-hazards/earthquake-hazards/science/earthquake-hazards-201-technical-qa>
- [28] E. H. Field, T. H. Jordan, and C. A. Cornell, "Opensha: A developing community-modeling environment for seismic hazard analysis," *Seismological Research Letters*, vol. 74, no. 4, pp. 406–419, 2003.
- [29] *The OpenQuake-engine User Manual. Global Earthquake Model (GEM) OpenQuake Manual for Engine version 3.11.2*, GEM, 2021.
- [30] K. S. Rukstales and M. D. Petersen. Data Release for 2018 Update of the U.S. National Seismic Hazard Model: U.S. Geological Survey data release. USGS. [Online]. Available: <https://doi.org/10.5066/P9WT5OVb>
- [31] M. Caprio, B. Tarigan, C. B. Worden, S. Wiemer, and D. J. Wald, "Ground motion to intensity conversion equations (gmices): A global relationship and evaluation of regional dependency," *Bulletin of the Seismological Society of America*, vol. 105, no. 3, pp. 1476–1490, 2015.
- [32] C. Worden, M. Gerstenberger, D. Rhoades, and D. Wald, "Probabilistic relationships between ground-motion parameters and modified mercalli intensity in california," *Bulletin of the Seismological Society of America*, vol. 102, no. 1, pp. 204–221, 2012.
- [33] ISO 3166, International Organization for Standardization Std. [Online]. Available: <https://www.iso.org/iso-3166-country-codes.html>
- [34] T. I. Allen, S. Halchuk, J. Adams, and G. A. Weatherill, "Forensic psha: Benchmarking canada's fifth generation seismic hazard model using the openquake-engine," *Earthquake Spectra*, vol. 36, no. 1_suppl, pp. 91–111, 2020.

⁸For Taiwan, we lowered this value to 1, as the model contains seismic sources with low magnitudes that cannot be represented properly with a mesh spacing of 5.