Can LEO Satellites Enhance the Resilience of Internet to Multi-Hazard Risks?

Aleksandr Stevens, Blaise Iradukunda, Brad Bailey,
Ramakrishnan Durairajan
University of Oregon

Abstract. Climate change-induced and naturally-occurring multi-hazard risks (e.g., Cascadia megathrust earthquake followed by tsunamis in the U.S. Pacific Northwest or PNW) threaten humanity and society, in general, and critical Internet infrastructures, in particular. While mitigating the impacts of these hazards, in isolation, on terrestrial infrastructures has been the focus of prior efforts, we lack an in-depth understanding of infrastructure hardening efforts using *non-terrestrial deployments* such as low earth orbit or LEO satellites in the face of multi-hazard risks.

The main goal of this work is to evaluate whether LEO satellites can bolster the resilience of Internet infrastructure in the Pacific Northwest (PNW) against multi-hazard risks[1]. To this end, we have developed a first-of-its-kind simulator called MAZE[2] to understand the impacts that multi-hazard risks, each of which combined or in isolation, pose to wired and wireless infrastructures in the PNW. Using MAZE, we address two key challenges faced by first responders today: (1) navigating the cost vs. performance trade-offs in the hybrid routing of traffic between terrestrial and non-terrestrial networks during disasters, and (2) comparing the efficacy of using LEO satellites against a terrestrial risk-aware routing strategy (ShakeNet) and a global satellite network (BGAN) for emergency communication during multi-hazard risks. Our assessments show that LEO satellites offer two orders of magnitude latency improvement and 100s of thousands of dollars in saving, all while maintaining network connectivity in the face of multi-hazard risks. To demonstrate the practicality and versatility of MAZE, we perform two case studies including testing a traffic prioritization scheme for LEO satellites and assessing the impacts of cascading risk on network infrastructures along the U.S. west coast.

1 Introduction

The Internet plays a central role in our daily lives. However, since its inception, the Internet has grown increasingly exposed to small- and large-scale climate change [13, 16] and naturally occurring risks [17, 29, 35, 15, 11, 45]. Such risks can have significant consequences including economic loss and loss of connectivity for large sections of users/businesses for extended periods.

[1] Although we use the PNW as a demonstrative case in this work, we note that the solution can be applied to various geographic regions, at different granularities (e.g., city vs. state), and for a range of single- or multi-hazard risk scenarios.

[2] Source code of the MAZE simulator: `https://gitlab.com/onrg/maze`.

To illustrate, consider a multi-hazard risk such as a Cascadia megathrust (**M**9) earthquake followed by a tsunami in the Pacific Northwest (PNW). Since PNW is the Internet hub for several national and international networks, and hyperscale cloud providers, it is estimated that the state of Oregon alone, for example, could see an economic loss of several billion dollars due to damages to its Internet infrastructure [6, 35]. Importantly, several critical monitoring and alerting systems (e.g., ShakeAlert [24] and ALERTWildfire [47]) depend on a resilient Internet to both detect events and disseminate alerts to first responders and the public. Hence, multi-hazard risks can severely impact our ability to use the Internet infrastructure to continuously monitor and alert on those events.

State-of-the-art from industry and academia alike that seek to mitigate the impacts of such risks are found wanting. For one, prior efforts focus only on isolated natural hazards on terrestrial network infrastructures [17, 35, 15, 11, 45, 16]; none of them have investigated multi-hazard risks to the best of our knowledge. Second, emergency first-responder networks (e.g., AT&T's FirstNet [21] and Verizon's Frontline [51]) rely on terrestrial infrastructures such as cell towers for disaster mitigation. However, as shown recently in [35, 11, 16], terrestrial infrastructures are susceptible to major failures in the face of natural hazards. Third, risk-aware mitigation strategies (e.g., ShakeNet [35] and RiskRoute [17]) seek to harden network infrastructures by establishing geographically longer terrestrial backup routes that are less susceptible to risks. The success of such strategies relies critically on two assumptions: (1) backup routes do not suffer any damage during disasters, and (2) traffic can be routed *around* major damaged areas using those backup routes. However, in the face of multi-hazard risks, these assumptions are unlikely to hold. For example, what if primary routes are affected by an earthquake and backup routes are affected by landslides? Finally, as we show in our evaluation, existing hardening efforts that leverage non-terrestrial satellite deployments such as Broadband Global Access Network (BGAN) [20] are costly and offer sub-optimal performance during hazards.

Assessing the impacts of multi-hazard risks to Internet infrastructure is fraught with two key challenges. First is the cascading nature of the problem. For example, a Cascadia megathrust earthquake can result in several geographically-distant follow-up events such as tsunamis on the coast of Japan, landslides throughout the U.S. west coast, etc. Second is the lack of a simulation capability to quantify the benefits of hardening efforts that seek to mitigate the impacts of multi-hazard risks using non-terrestrial or hybrid deployments. For example, based on our interactions with first responders, it is unclear how to navigate the cost vs. performance trade-offs in the hybrid routing of traffic between terrestrial and non-terrestrial networks during disasters.

To address these challenges, the main goal of this work is to evaluate whether LEO satellites can bolster the resilience of Internet infrastructure in the Pacific Northwest (PNW) against multi-hazard risks. To this end, we have developed a first-of-its-kind simulator called MAZE to understand the impacts that multi-hazards, each of which combined or in isolation, pose to wired and wireless infrastructures in a geographic region of interest. MAZE can be used by first

responders and federal agencies to compare and contrast the benefits of various emergency communication and disaster mitigation strategies in a practical and repeatable manner.

At the core of MAZE are the following novel aspects. First, MAZE allows simulations of cost-effective data transfer between different points in a network using a wide variety of backup network routing strategies (e.g., LEO, ShakeNet, shortest-path, etc.). This enables the navigation of performance (e.g., latency) vs. cost trade-offs in using terrestrial and/or non-terrestrial routing strategies. Second, MAZE also enables the simulation of hybrid routes where certain network hops are done via terrestrial fiber and other network hops are performed via LEO constellations. This enables the simulation of complex, hybrid routing scenarios where a subset of terrestrial infrastructure is still functional for routing but the backup paths for damaged infrastructure could be established with LEO satellites. Finally, MAZE is built to be easily extensible to allow the simulation of new backup routing strategies. This allows easy evaluation of other futuristic multi-hazard risk scenarios and novel emergency communication strategies to mitigate them.

Using MAZE, we compare the efficacy of LEO satellites for infrastructure hardening against two baselines: a terrestrial risk-aware routing strategy (i.e., ShakeNet) and a global satellite network (i.e., BGAN). Our assessments show that LEO satellite-based hardening strategies offer two orders of magnitude latency improvement and 100s of thousands of dollars in saving, all while maintaining network connectivity in the face of multi-hazard risks. In addition, we analyze the percentage of emergency responders that can be serviced based on different budgetary restrictions. We find that with only 0.006% of current budget for OR and 0.018% of current budget for WA, these states can benefit from LEO-based backup communication to provide 100% service to first responders. To showcase both the practicality and extensibility of MAZE, we apply it to two case studies. The first case study seeks to test the feasibility/practicality of a traffic prioritization scheme for LEO satellites, whereas the second one aims to assess the impacts of a cascading risk on network infrastructures across three states in the U.S. west coast.

Contributions and Roadmap. This work makes the following key contributions.

- A first-of-its-kind empirical illustration of multi-hazard risks to Internet infrastructures in PNW (§ 2).
- MAZE, which is novel yet practical discrete-event simulator to compare and contrast the benefits of various emergency communication and disaster mitigation strategies in a repeatable manner (§ 3).
- Evaluation of MAZE with realistic isolated and multi-hazard risk scenarios (§ 4).
- Two case studies that demonstrate how MAZE can be used to transition a research prototype on traffic prioritization using LEO satellite into practice (§ 5.1), and enrich the resiliency analysis of researchers with practical issues faced by first responders during multi-hazard disaster scenarios (§ 5.2).

2 Background, Motivation and Related Work

2.1 Motivation

Multi-hazard risks are on the rise and are increasingly affecting terrestrial critical network infrastructures. In recent years, there has been a steady increase in climate change-induced as well as naturally occurring risks in the U.S. that have affected substantial regions, such as the rise in wildfires in the west [7] and tornadoes in the south [23]. One concerning risk for the Pacific Northwest (PNW) is the possibility of a megathrust earthquake. PNW lies in the Cascadia Subduction Zone [6], a region where a megathrust earthquake that will cause major infrastructure damage is expected to hit in the near future. Areas in PNW that experience very strong shaking on the Modified Mercalli Intensity (MMI) scale during the megathrust earthquake are likely to see significant damage to cell towers, IXPs, fiber lines, and other essential communication infrastructure [35]. A megathrust earthquake is also expected to cause a series of cascading disasters, including a tsunami and landslides down coastal Oregon [10]. In addition to the threat of a future earthquake, wildfires, a yearly occurrence in the PNW, have been shown to damage essential communication infrastructure such as cell towers [11].

Extent of infrastructure damage in the PNW. To empirically illustrate the problem, we quantify the likely extent of infrastructure damage[3] in Oregon (OR) and Washington (WA) due to multi-hazard risk—Cascadia M9 earthquake followed by a tsunami in the U.S. West Coast. To this end, we use the USGS national seismic hazard maps from 2014 that are integrated into ShakeNet [35] and tsunami flooding models developed by the Washington Geological Survey (WGS) along with the National Oceanic and Atmospheric Administration (NOAA) [41] to predict how much infrastructure is susceptible to a multi-hazard disaster. We perform this by loading the damage models into ArcGIS [1] and performing a layer *merge*, creating a compound output layer which contains the total combined area containing susceptible infrastructure. We then performed an *Intersect* with the ShakeNet infrastructure layers which gave us a count of different types of affected infrastructure in the multi-hazard damage zone. For context, the seismic hazard resulting from earthquakes is commonly measured by indicating the expected frequency of shaking, expressed either as "return periods" corresponding to specific timeframes (e.g., every 50 years) or as the probability of surpassing a certain threshold (e.g., 2% or 10% probability of exceedance) within a defined interval [35].

Total counts of affected node infrastructures and surrounding fiber infrastructure (in km) is shown in Table 1. We note that the damage counts for infrastructure in major metropolitan areas such as Points-of-Presence (PoPs), Data Centers, and Colocation Facilities largely stayed the same between an isolated earthquake incident and a multi-hazard disaster scenario. This follows as most

[3] In this context, the "impact" of infrastructure damage is characterized by complete failures resulting in the absence of any service.

Fiber (km)	Cell Towers	PoPs	Data Centers	Colos
32,057	204,585	422	32	59

Table 1: Count of infrastructure that are prone to damage during a Multi-Hazard Cascadia Earthquake+Tsunami scenario for expected PGAs with 10% probability of exceedance in the next 50 years [35].

of this infrastructure is inland where tsunami damage is unlikely to reach (i.e., Portland), or in areas that are near the water but were already assumed to have major earthquake damage due to violent shaking (i.e., Seattle). Counts for cell towers and length of fiber effected for the multi-hazard scenario increased compared to the standalone earthquake scenario as fiber and cell towers on coastal Oregon, along with submarine fiber, would be additionally affected given the added tsunami and flooding.

Population affected by infrastructure damage in the PNW. In addition to estimates of susceptible infrastructure, we also estimated the total population that may be affected by damage to communication infrastructure. By using the United States Census Bureaus 2021 population census by county [9] and seeing which counties overlapped with areas with expected infrastructure damage, we were able to get an estimate of the total population that will be impacted by damage to communication infrastructure. These estimates are shown in Table 2. In Washington, 89.2% of the state population is expected to be impacted by infrastructure damage during a multi-hazard earthquake scenario. Similarly, in Oregon, 91.4% of the population is expected to fall within zones experiencing high chance of infrastructure damage.

	Total Population	Emergency Responders
OR	3,879,430	6,681
WA	6,901,149	18,189

Table 2: Count of people in areas prone to damage during a Multi-Hazard Cascadia Earthquake+Tsunami scenario split into total population and emergency responders.

We also applied these population percentages to data from the Bureau of Labor Management [39, 38, 40] regarding total number of emergency responders employed by state to get the number of emergency responders whose communication channels are likely to be impacted by infrastructure damage. This quantification regarding the amount of infrastructure damage and expected effected population makes it clear that alternative communication strategies need to evaluated in the event of a multi-hazard disaster scenario in the PNW. This evaluation of communication strategies is especially prudent with regards to emergency responders as they will need reliable communication in the event of a natural disaster in order to service the community and save lives.

2.2 State-of-the-art and their Limitations

Existing emergency responder networks cannot function in the face of multi-hazard risks. The majority of past work on emergency communications has to do with establishing dedicated communication networks to avoid network congestion and to ensure that first responders can communicate with each other in the face of natural disasters. For example, there are several fully-functional emergency communication networks in the U.S. including AT&T's FirstNet [21] and Verizon's Frontline [51]. However, these emergency responder networks, e.g., use the Long-Term Evolution (LTE) standard and, rely on terrestrial cell towers. These cell towers, however, are likely to be non-functional following major natural disasters such as a megathrust earthquake [35] and wildfires [11], let alone multi-hazard risks.

Risk-mitigation strategies consider only isolated risks and cannot effectively address multi-hazard risks. Risk-aware terrestrial routing and mitigation strategies such as ShakeNet [35] and RiskRoute [17] seek to harden network infrastructures against natural disasters by establishing geographically longer terrestrial backup routes that are less susceptible to risks. However, such strategies rely heavily on the assumption that certain areas, with backup routes, suffer little-to-no damage for routes to remain functional. However, this assumption is unlikely to hold during multi-hazard risks e.g., megathrust earthquakes in PNW followed by tsunamis or landslides in the U.S. west coast. Consequently, these strategies only account for routing *around* major damaged areas and fail to account for how communication *within* damaged areas can be established.

Existing GEO satellite-based efforts are either costly or offer suboptimal performance. Some of the aforementioned emergency communication networks extend beyond LTE and offer satellite communications using GEO satellites. For example, FirstNet currently offers satellite-based communication called Broadband Global Area Network (BGAN) [20] with upload speeds of up to 492Kbps for thousands of dollars [18]. Interest has also been expressed by FirstNet into possibly integrating LEO satellites into their network for emergency communications as a newer alternative to BGAN [28]. However, FirstNet has yet to formally announce plans to make use of LEO satellite constellations as the technology is still very new and evaluations of how these networks would fair in disaster scenarios compared to current systems such as BGAN is an open problem.

Recent work uses satellite constellations for post-disaster communications [34]. However, this work focuses primarily on mixing LEO satellites with the BeiDou Navigation Satellite (BDS) system [25]. BDS system, which is independently developed by China, is composed of GEO and MEO (Medium Earth Orbit) satellite constellations. Nevertheless, deployment of proprietary systems like BDS motivates our work on how current commercial LEO satellite options (e.g., Starlink [49], Kuiper [46], and Telesat Lightspeed [50]) would fair in the face of multi-hazard risks in the U.S.

Research has also been done by NASA into Delay/Disruption Tolerant Networking (DTN), which focuses on running space packet protocols that are re-

silient to the specific high latency and high fault communication channels between terrestrial (e.g., ground stations) and non-terrestrial (e.g., GEO satellite) communication infrastructures [12]. While this might apply to disaster scenarios, to the best of our knowledge, their current focus is limited to space missions. Furthermore, as we show in § 4, with the use of LEO satellites, latency stays low enough to allow continued use of conventional routing protocols through the satellite constellations, rendering DTN unnecessary in this case.

2.3 Opportunity

Low-cost low earth orbit (LEO) satellite-based communication can harden network infrastructures against terrestrial multi-hazard risks. Because of the rise in multi-hazard risks that threaten terrestrial communication infrastructure, making them non-functional, and the aforementioned shortcomings of prior efforts, it is important to evaluate alternative means of communication. One such alternative is LEO satellite-based communication which is attractive for disaster communications due to three key reasons. First, except for satellite ground stations, LEO satellites stay beyond the reach of terrestrial multi-hazard risks.[4] Second, an industry push in the deployment of commercial LEO satellite networks has resulted in relatively cheap hardware for connecting users to these satellite constellations. For example, a typical Starlink Satellite Kit is only USD 599 [48]. This means that scaling out to fit the bandwidth and network quality needs for a variety of disaster scenarios can be done cost-effectively. Third, because LEO satellites are much closer to earth than the older geosynchronous equatorial orbit (GEO) satellites that are currently in use by first responders, latency is significantly reduced during emergency communication scenarios.

2.4 The Key Challenge

Need: a capability to understand the practicality of LEO satellites for multi-hazard risk mitigation. While the above benefits are compelling, we lack a capability to assess and evaluate the practicality of how LEO satellites-based disaster communications compare and contrast against other mitigation efforts both in terms of network performance and cost. In light of this challenge, we have four options. (1) *Analytic modeling* can be used to examine idealized behaviors of the disasters and mitigation strategies but it lacks realism. For example, it is unclear how to mathematically capture disaster scenarios and their impacts in practical deployment settings. (2) *Testbed-based evaluation* can be used to test details (e.g., implementation of mitigation strategies) but has scalability issues. Practically speaking, it is also unclear how to create disaster scenarios in testbeds. (3) *In-situ evaluation*, can be used to test more complete

[4] Similar to other terrestrial infrastructures (e.g., fiber-optic cables), ground stations are susceptible to availability and resiliency issues resulting from multi-hazard risks.

implementations of mitigation strategies with high practicality but it lacks repeatability. (4) *Simulations* offer a promising path to test mitigation strategies and disaster scenarios in a practical and repeatable manner. Due to this reason, we take the fourth option and build a simulator to test and assess the LEO satellite-based disaster communications for multi-hazard risks.

3 Design and Implementation of MAZE Simulator

3.1 Overview of MAZE Simulator

Motivated by the above need, we design MAZE: a first-of-its-kind simulator to understand the impacts that multi-hazard risks, each of which combined or in isolation, pose to wired and wireless infrastructures in a geographic area of interest. MAZE seeks to empower first responders and federal agencies with a *practical* tool to compare and contrast the benefits of various emergency communication and disaster mitigation strategies in a *repeatable* manner.

Novelty of MAZE. MAZE is designed with three novel aspects in mind. First, to empower first responders and disaster response agencies, MAZE allows simulations of cost-effective data transfer between different points in a network using a wide variety of backup routing and mitigation strategies (e.g., LEO, ShakeNet, shortest-path, etc.). This enables the navigation of performance (e.g., latency) vs. cost trade-offs in using terrestrial and/or non-terrestrial routing strategies. Second, MAZE also enables the simulation of hybrid routes where certain network hops are done via terrestrial fiber and other network hops are performed via LEO constellations. This enables the simulation of complex, hybrid routing scenarios where a subset of terrestrial infrastructure is still functional for routing but the backup paths for damaged infrastructure could be established with LEO satellites. Third, MAZE is built to be easily extensible to allow simulation of new backup routing strategies. This allows easy evaluation of other futuristic multi-hazard risk scenarios and novel emergency communication strategies to mitigate them.

Implementation of MAZE. MAZE consists of capabilities to (1) create multi-hazard risk scenarios (§ 3.2), (2) identify all terrestrial and non-terrestrial network routes (henceforth, *hybrid routes*) (§ 3.3), and (3) compare and contrast cost vs. performance trade-offs to provide decision support for first responders e.g., to pick the most optimal hybrid route (§ 3.4). At its core, MAZE builds on top of a geographic information system called ArcGIS [1]. ArcGIS provides robust visualization and geo-analytic capabilities atop an object-relational database. The database is purposefully built for datasets with geo-anchored features (e.g., network infrastructure node with <latitude, longitude> point features, census blocks as polygon features, etc.).

Scope of this work. We limit the geographic scope of MAZE to the U.S. Pacific Northwest (PNW) and west coast, and three disaster scenarios (i.e., earthquake, wildfire, and earthquake followed by a tsunami; the former two are isolated disasters whereas the latter is a multi-hazard risk). We also consider a cascading

risk scenario: an event that starts as an isolated risk but ends up as a multi-hazard risk These scenarios are not prescriptive but are representative of first responders' critical needs in PNW. Without loss of generality, MAZE can be used for different multi-hazard risks, geographic locations of interest, and at different granularities (e.g., city vs. state).

3.2 Creating Multi-hazard Risk Scenarios

MAZE offers an interface to either create realistic disaster scenarios manually or plug in existing risk models from the Earth Sciences and Hazards communities. In this work, we consider models for three disaster scenarios, each of which are in formats (e.g., KML) supported by ArcGIS.

First is a wildfire scenario, which we obtain from the Northwest Interagency Coordination Center's historical wildfire map [37]. The goal is to demonstrate how MAZE can be used to assess the efficacy of different emergency communication strategies in the context of past wildfires in the PNW. As a possible candidate model, we chose 2020 Lionshead and Beachie Creek wildfires due to their sheer size, as well as their proximity to the Salem and Portland areas, both of which are highly populated. These wildfire scenarios are shown in Figure 1. Another reason for choosing this scenario is the fact that wildfires have been shown to damage cell tower infrastructure [11]. This means that primary emergency communication networks like FirstNet [21] are susceptible to failures and calls for alternative strategies.

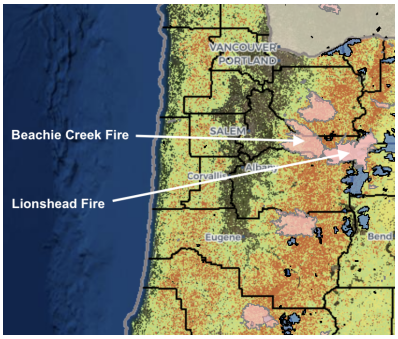

Fig. 1: PNW Beachie Creek and Lionshead Wildfires Disaster Scenario constructed using Northwest Interagency Coordination Center's historical wildfire map [37].

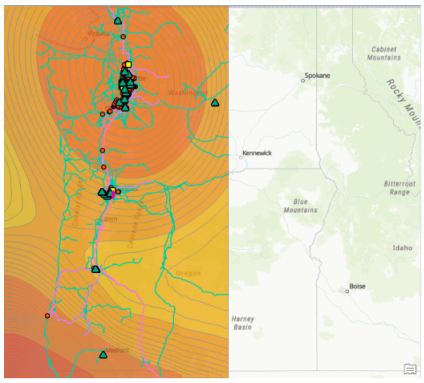

Fig. 2: PNW Megathrust Earthquake Disaster Scenario constructed using ShakeNet [35].

To create a model for an earthquake disaster scenario and estimate where the majority of infrastructure damage will occur, we use earthquake shaking

models from the ShakeNet [35] framework and integrate them into MAZE. Note that this does *not* mean MAZE will identify the same backup routes for affected areas, similar to Mayer et al. [35]. This simply means that any area expected to experience very strong shaking is likely to suffer significant infrastructure damage and is therefore a viable area to evaluate alternate communication strategies within. For example, the inner ring engulfing Seattle in Figure 2 represents the area expected to experience very strong shaking in the PNW. This inner ring is the damaged area we use in our simulations for the earthquake disaster scenario.

Third, from our conversation with hazards scientists and first responders, we create a multi-hazard risk scenario by augmenting the geographical areas from two different models. First are the areas from Mayer et al. [35] that are prone to very strong shaking due to the Cascadia M9 earthquake. Second are the areas susceptible to tsunamis and landslides along the Oregon coast following the megathrust earthquake [10].

3.3 Identifying Backup Routes

Given a risk model (such as the ones described above) and the resulting geographic locations of damage, the next component in MAZE seeks to identify all possible backup routes between a source-destination pair for emergency disaster response. To illustrate, consider the scenario in Figure 3 where first responders from Salem (indicated as point A) are trying to coordinate with a disaster response agency in Harney Basin (indicated as point B). For this scenario, the first responders are interested in identifying all types of infrastructures to establish backup communication paths between the two points. Here, an infrastructure type can be terrestrial (e.g., routes that use fiber-optic cable), non-terrestrial (e.g., routes that use LEO satellites), or hybrid (e.g., routes that use a combination of terrestrial and non-terrestrial infrastructures).

MAZE performs this identification in three steps. Step 1 is to *fuse* terrestrial (e.g., fiber routes, cell towers, etc.) and non-terrestrial (e.g., satellite ground stations, locations of LEO satellites, etc.) network infrastructures from ShakeNet framework [35] atop the risk models described above to obtain functioning vs. non-functioning infrastructures. We use the *overlap* tool in ArcGIS to fuse infrastructure datasets with the models. Here, we define functioning infrastructures as those that are intact and are not impacted by the disaster considered in the risk model. By doing so, we implicitly consider "reliability" as a key design metric in MAZE. We encode infrastructure information as a path graph $P_{A,B}$ between two points A and B. $P_{A,B}$ is represented by a series of network segments $(x_1, x_2), (x_2, x_3), \ldots, (x_{n-1}, x_n)$. Each vertex x has two parameters: (1) a "type" parameter that stores hybrid, terrestrial, or non-terrestrial and describes the type of vertex; and (2) the geographic location of the vertex as latitude and longitude pairs. Note that x_1's location is the same as point A's location and x_n's location is the same as point B's.

Step 2 is to *enumerate* all possible backup paths that could be established atop the functioning infrastructures that could potentially be used to route traffic between any two points (e.g., A and B). These enumerated paths are also

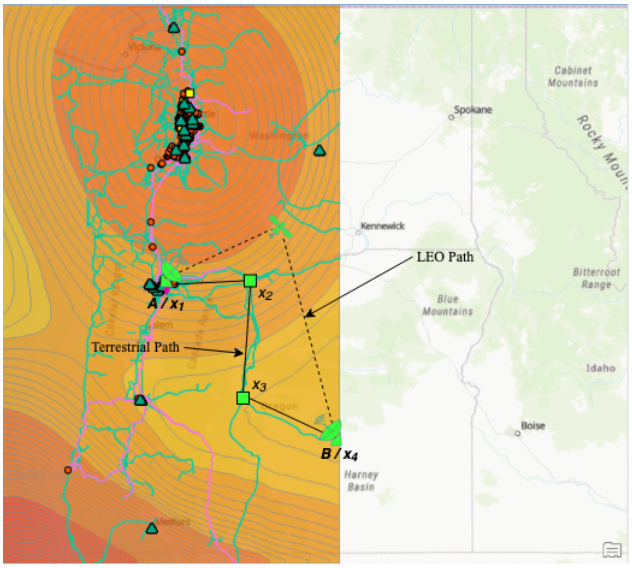

Fig. 3: PNW hybrid route example.

annotated with performance metrics[5] (e.g., latency values) from different sources (e.g., speed-of-light based round-trip times or RTT for fiber-optic paths between terrestrial points, estimates of ground station-to-LEO satellite and inter-satellite links from Chaudhry et al. [14]). Concretely, a network segment (x_1, x_2) will be annotated differently depending on the type of the vertices. For example, if x_1 and x_2 are of different types (i.e., ground station to terrestrial, or terrestrial to ground station), or they are both of terrestrial type, then they will be annotated speed-of-light RTT estimates. Otherwise, if x_1 and x_2 are both ground stations, the latency estimates will be calculated as if it was routed through the LEO satellite constellation using [14]. An example for both these annotations can be seen in Figure 3 where the path can be established through a LEO constellation between two ground stations (x_1, x_4) or through terrestrial infrastructure segments $(x_1, x_2), (x_2, x_3), (x_3, x_4)$ sequentially. Nevertheless, the latency of a specific network segment (x_1, x_2) performed at time t will be notated as $RTT_t(x_1, x_2)$ and will be in milliseconds. The final RTT of the whole path at time t, $RTT_t(P_{A,B})$, will be defined as the sum of the RTTs of the individual network segments, or

$$RTT_t(P_{A,B}) \;=\; RTT_t(x_1, x_2) \;+\; RTT_t(x_2, x_3) \;+\; \ldots + \; RTT_t(x_{n-1}, x_n)$$

To make the performance values amenable for the simulations, we generate $T/\Delta t$ RTT values where T is the total time of a simulation in milliseconds and

[5] In this work, we consider latency as the key performance metric because it translates directly to response times of first responders during a disaster. In future work, we plan to consider other metrics such as path congestion, throughput, among others.

Δt is the time increment in milliseconds. So, a simulation would produce the following list:

$$RTT_0(P_{A,B}),\ RTT_{1*\Delta t}(P_{A,B}),$$
$$RTT_{2*\Delta t}(P_{A,B}),\ \ldots,\ RTT_T(P_{A,B})$$

Note that, as shown by Lai et al. [33] and subsequently confirmed by Izhikevich et al. [27], the final RTTs obtained using the equation above are an *underestimate* of the real-world RTTs. We note that the difference arises due to both system-level overheads (e.g., packet processing) and operational factors (e.g., congestion arising from an increased demand surge following a disaster). This difference is discussed further in § 6.2.

Step 3 is to *delegate* the simulation of routing to appropriate sub-simulators based on the type of infrastructure segments in backup paths. For non-terrestrial segments, we use Hypatia framework [31]. In essence, Hypatia works by generating the position of the satellites in the LEO constellation over time T at time increments Δt. We also provide Hypatia with parameters such as satellite altitude, number of orbits, etc., which correspond to specific commercial LEO constellation configurations such as Starlink [49], Kuiper [46], and Lightspeed [50]; see source code here [4] for all configurable parameters, including satellites per orbit, ISL link capacity, among others. Similarly, for simulating routing atop the terrestrial segments, we use the *route planner* capability in ShakeNet framework [35]. At the end of the simulation, all functional backup non-terrestrial and terrestrial paths will be produced by Hypatia and ShakeNet respectively.

3.4 Navigating Cost vs. Performance Trade-offs

Network performance is important, however, the cost is arguably an equally important factor as it defines whether or not a solution is likely to be implemented by different government departments that are on a strict budget. In light of this, one of the key challenges faced by first responders today is the lack of decision support to identify the most performant yet cost-effective backup path during a disaster scenario. Another challenge is to compare and contrast the cost vs. performance trade-offs of different emergency communication strategies (e.g., Starlink vs. FirstNet vs. ShakeNet) for the backup paths.

To tackle these challenges, we (1) frame the issues as linear path optimization formulations (e.g., minimize latency, minimize cost), which are omitted due to space reasons, (2) input all backup paths with annotated performance values as well as the cost of operation, and (3) offer decision support to first responders e.g., to choose *the* most performant and cost-effective backup path. To calculate cost of operations in (2), for each strategy, we first calculate the cost for establishing a backup path i as follows:

$$Cost_{path}(i) = t + \left\lceil \frac{a}{b} \right\rceil \times p$$

where a is the average amount of upload bandwidth each user needs, b is the max. upload bandwidth supported, t is the capital equipment cost, and p is the

cost of a network plan. We then calculate the total cost of operation using the individual backup paths. Concrete, for a specific disaster scenario (s), we use the total number of users u that need to be supported on each one of the backup paths on the network as follows.

$$Cost_{network}(s) = u \times t + \left\lceil \frac{u \times a}{b} \right\rceil \times p$$

We then limit the possible network paths to only utilize this subset of hardware and bandwidth allocated, allowing us to estimate the total cost of operation based on these maximum hardware and available bandwidth restrictions. Note that we chose to use upload bandwidth over download bandwidth as our variable because upload bandwidth is conventionally lower than download bandwidth. For bi-directional communication, available upload bandwidth would therefore be the limiting factor.

4 Evaluation of MAZE

For each disaster scenario (in § 3.2), we seek to evaluate the efficacy of LEO constellation's performance and cost against two baselines: (1) risk-aware terrestrial routing strategy using ShakeNet [35], and (2) GEO satellite-based emergency communication strategy i.e., BGAN [20]. This first baseline will shed light on how well LEO satellites fare in comparison with geographically longer yet risk-aware terrestrial-only routing strategy. The second baseline will tell us how much better, if at all, LEO satellites are compared to the current standard satellite networks used by emergency responders today. We also employ MAZE to analyze the percentage of emergency responders that can be serviced based on different budgetary restrictions imposed by a particular state.

4.1 Performance Comparison of Emergency Communication Strategies

Using MAZE, we simulated the three disaster scenarios for 100 seconds in the Pacific Northwest. The network paths for the three scenarios tested were between points in the damaged areas with the largest geographic separation and population (i.e., Seattle and Portland). We obtained latencies for the MAZE-selected backup path. We used Hypatia (for Starlink LEO satellite-based backup path) and ShakeNet (for terrestrial backup path) to obtain latencies for the other baselines. We also compared these against the BGAN latencies reported by Inmarsat [3].

Each of the strategies along with the simulated performance is shown in Figure 4. For wildfires, average RTTs increased by 2.153 milliseconds when routing through the Starlink constellation instead of the terrestrial ShakeNet path. This is due to ShakeNet's selection of a geographically closer yet more efficient backup route via Yakima instead of the direct route between Portland and Seattle. However, for earthquake and multi-hazard disasters, Starlink outperforms

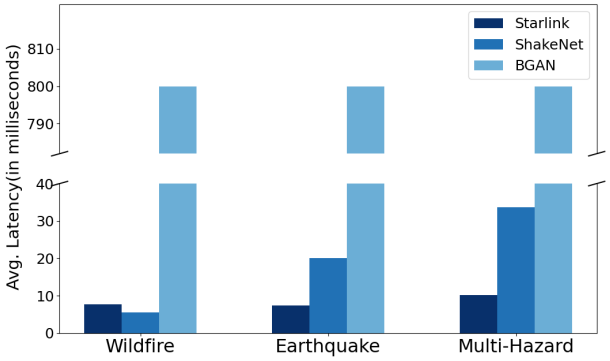

Fig. 4: Average latency comparison between Starlink, ShakeNet Fiber, and BGAN networks for Wildfire, Earthquake, and Multi-Hazard disasters in PNW.

the terrestrial ShakeNet solution, offering a latency decrease of 2.7x and 3.3x respectively. This shows that as damage area increases, LEO satellites provide a more consistent latency when compared against risk-aware terrestrial routing solutions like ShakeNet. This is likely because as the damaged area increases, which is typical of multi-hazard risks, solutions like ShakeNet require very long terrestrial routes to be used to compensate for routing around large damage areas, dramatically increasing latency.

Of the three strategies, BGAN, which is the current standard in satellite network communication for emergency responders in disaster scenarios, is 104x slower than Starlink and 144x slower than ShakeNet's risk-aware terrestrial routing in the wildfire disaster experiment performed. Similarly, for the megathrust earthquake disaster scenario, BGAN is 108x slower than Starlink and 40x slower than ShakeNet. For the multi-hazard scenario, BGAN is 79x slower than Starlink and 24x slower than ShakeNet.

4.2 Total Cost of Emergency Communication Strategies

Using MAZE, we next compare the cost of establishing backup paths using Starlink and BGAN networks for the wildfire, earthquake, and multi-hazard disaster scenarios. Specifically, we use the formula to calculate total costs (for scenarios) from § 3.4, and calculate the cost for supporting concurrent VoIP users on the Starlink and BGAN satellite networks:

$$Cost_{Starlink}(s) = u \times 599 + \left\lceil \frac{u \times 80}{9330} \right\rceil \times 110$$

and

$$Cost_{BGAN}(s) = u \times 4995 + \left\lceil \frac{u \times 80}{492} \right\rceil \times 2840$$

Based on our conversations with first responders, our calculations assume that the equipment is not shared by groups of users. Furthermore, the equations above assume users (e.g., first responders) are using the G.722 codec for VoIP which requires about 80 Kbps of upload bandwidth per caller [32]. We chose VoIP using G.722 as we could ensure that this was a communication method usable on both the Starlink and BGAN satellite networks. However, it is important to note that due to Starlink's much higher bandwidth support per user equipment and lower latency, more bandwidth-hungry and latency-sensitive services (e.g., HD video streaming) are viable on the Starlink network. We used the list price of $599 for Starlink user equipment with the network plan being $110/month [48]. The corresponding user equipment and plan prices for the BGAN network were $4,995 and $2,840/month respectively [2]. The available upload bandwidth per terminal and plan used was 9.33 Mbps for Starlink [19] and 492 Kbps for BGAN [18].

Table 3 shows total cost estimates for using Starlink- and BGAN-based emergency communication strategies in the modeled wildfire, earthquake, and multi-hazard disaster scenarios. We split costs into two parts in Table 3. First is the capital costs which represented by the $u \times t$ (i.e., Hardware) portion of the equation for user equipment. Second is the operational costs which is represented by $\left\lceil \frac{u \times a}{b} \right\rceil \times p$ (i.e., Network Plan column). For the wildfire scenario, we estimated up to 650 firefighters may have responded to the Lionshead and Beachie Creek wildfires [44]. Supporting 650 concurrent VoIP users on Starlink was 9x less expensive than supporting the same number of users on BGAN.

	Hardware	Network Plan	Total Cost
Starlink Wildfire	$389,350	$660	$390,010
BGAN Wildfire	$3,246,750	$301,040	$3,547,790
Starlink Earthquake	$2,755,400	$4,400	$2,759,800
BGAN Earthquake	$22,977,000	$2,124,320	$25,101,320
Starlink Multi-Hazard	$6,235,590	$9,900	$6,245,490
BGAN Multi-Hazard	$51,997,950	$4,808,120	$56,806,070

Table 3: Cost comparison between BGAN and Starlink for emergency responder use in Wildfire, Earthquake, and Multi-Hazard disaster scenarios.

In the case of the megathrust earthquake disaster scenario, it is estimated that ∼4,600 emergency responders will be concurrently using an emergency communication network [26]. To support these first responders, using Starlink is cost-effective i.e., also a 9x reduction in cost compared to BGAN. For the multi-hazard scenario where with an estimation of ∼10,410 emergency responders [8], Starlink is again cost-effective: 9x reduction in cost compared to BGAN.

4.3 Percentage of Serviceable Emergency Responders in PNW

The total number of emergency responders estimated to be within a zone that may experience any damage in a multi-hazard disaster scenario was estimated

to be 18,189 for Washington and 6,681 for Oregon in § 2. Given our equation for cost of using the Starlink constellation as a function of total concurrent users, we can calculate what percentage of emergency responders can be serviced based on different budgetary restrictions. These results can be seen in Table 4 separated by increments of 20%. The cost of serving the emergency responder population in Oregon (OR) ranges from $801,584 for 20% coverage to $4,008,299 for full coverage (100%). Similarly, in Washington (WA), the cost ranges from $2,182,083 for 20% coverage to $10,912,371 for complete coverage (100%).

	20%	40%	60%	80%	100%
OR	$801,584	$1,603,058	$2,404,642	$3,206,116	$4,008,299
WA	$2,182,083	$4,364,655	$6,547,227	$8,729,799	$10,912,371

Table 4: Cost of serving impacted emergency responder population by percentage in Oregon and Washington.

According to the National Association of State Budget Officers [42], the Oregon government spent a total of $66.8 billion in the 2021 fiscal year with Washington State spending a total of $60.5 billion. With these budget numbers in mind and the cost of servicing emergency responders in Table 4, we can deduce that to ensure 100% of emergency responders have reliable communication via Starlink in the event of a natural disaster, Oregon state would only have to allocate 0.006% of their yearly total expenditures and Washington State would only have to allocate 0.018% of their yearly total expenditures.

5 Case Studies

In this section, we present two case studies. These case studies demonstrate how extensible MAZE is (as mentioned in § 3.1), and how MAZE can be used to (1) transition a research prototype into practice (§ 5.1), and (2) enrich the multi-hazard resiliency analysis of researchers with practical issues faced by first responders (§ 5.2).

5.1 How to Prioritize Traffic Classes during Multi-Hazard Risks?

The first case study is obtained from *first responders* who were eager to assess the practicality of a traffic prioritization scheme proposed by Zhou et al. [52] in LEO satellites, and examine its performance under chaotic/varying network loads (which is very typical during multi-hazard risks) vs. a risk-aware routing strategy such as ShakeNet [35]. At its core, the scheme prioritizes various traffic classes (e.g., government vs. normal) in LEO satellite-based systems using a dynamic channel reservation algorithm [52]. In addition, the algorithm depends on three parameters—specifically, handover failures, new call blocking, and QoS decline—which are standard in essentially all LEO-based traffic prioritization

schemes. We note that the first responders were interested in this particular scheme because of its focus on a small number of user classes (i.e., "normal", "senior", and "government"). The primary challenge stems from the common hurdle faced during the "transition to practice": the lack of capabilities to effectively implement an idea from a research paper into practical application.

We implemented the algorithm described in [52] in Python and plugged into MAZE. Similar to Zhou et al., we modeled the rate of new call arrivals to the network as a Poisson distribution. Consequently, the term "network load" is just the Poisson parameter λ for that distribution. Additionally, we acknowledge that the three parameters employed in our simulations (below) are random and we were not able to determine how the parameters from Zhou et al. [52] are applicable in the context of natural disasters. Nevertheless, we note that these parameters can be adjusted by first responders in relevant risk scenarios.

We conducted a test of the scheme using the MAZE. The test involved a simple path in the PNW region, utilizing the Starlink constellation. The path spanned from "Seattle -> Portland GS -> San Francisco GS -> Los Angeles." In the path, GS refers to a ground station, and the connection between GS nodes was established through a LEO satellite from Starlink. The purpose of the test was to simulate a multi-hazard risk scenario (specifically, an earthquake followed by a tsunami in the PNW), as discussed in § 2. The simulation was run for ten seconds, and the highest priority was assigned to the government traffic class, prioritizing it over other classes.

During each simulation cycle, a list of RTT values (in milliseconds) was generated for each path segment, considering a specific network load value (λ). This process was repeated for various λ values and for each user class. The resulting lists of RTTs were then averaged over the entire duration of the simulation.

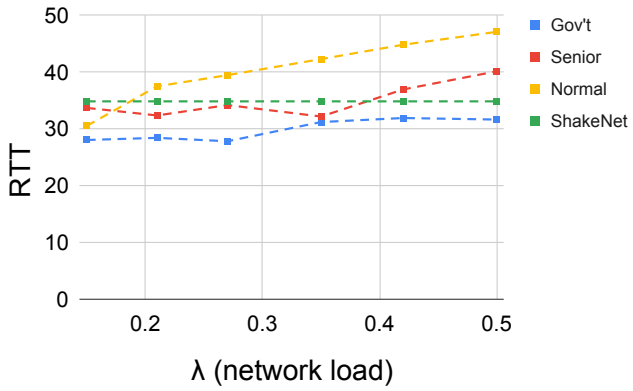

Fig. 5: Latencies experienced by different traffic classes under varying network loads.

Figure 5 illustrates the average latencies experienced by each traffic class under different network load conditions. Two key observations can be made from this figure. Firstly, as per our configuration, the scheme proposed by Zhou et al. effectively prioritized government traffic over other classes. This can be deduced from the lower RTT values observed for government traffic compared to other classes. Secondly, and unsurprisingly, the risk-aware routing strategy exhibited overall stability, as indicated by consistent RTT values despite varying network loads. This is due to longer risk-aware routing chosen by ShakeNet. Overall, these results demonstrates how MAZE can serve as a decision-support tool for first responders, enabling them to assess different backup strategies and identify an appropriate approach that would perform well under diverse network load conditions.

5.2 What are the Impacts of a Cascading Risk to Network Infrastructures?

The second case study is more *academic* in nature and seeks to evaluate the consequences of cascading risks on network infrastructures along the U.S. West Coast, namely California (CA), Oregon (OR), and Washington (WA). Here, a cascading risk refers to an incident that initiates as a singular or isolated hazard (e.g., an earthquake like San Andreas or Cascadia) and has the potential to evolve into a multi-hazard situation (e.g., an earthquake followed by a tsunami).

In order to evaluate the effects of cascading risk on multiple states, we analyze the potential magnitude of infrastructure destruction by utilizing network assets such as long-haul fiber-optic cables, Starlink's satellite ground stations, and cell towers provided by ShakeNet. Figure 6 displays these infrastructure assets.

We also use the USGS national seismic hazard maps from ShakeNet, and tsunami flooding models developed by the Washington Geological Survey (WGS) along with the National Oceanic and Atmospheric Administration (NOAA) [41] to understand both the area of coverage of those disasters as well as how much infrastructure is susceptible to a cascading disaster. Seismic hazard includes both San Andreas earthquake as well as Cascadia earthquake. Figures 7 and 8 show areas in the U.S. west coast that are susceptible to damage from the two earthquakes with different magnitudes with 2% and 10% probability of exceedance, respectively, followed by tsunamis. Seismic hazard is commonly expressed by indicating the anticipated occurrence rate of shaking, either in terms of "return periods" relevant to specific timeframes (e.g., every 50 years) or as the probability of surpassing a certain threshold (e.g., 2% probabilty of exceedance or 10% probability of exceedance) within a defined interval [35].

Next, we load all these damage models along with infrastructure assets into MAZE and performing a layer *merge*, creating a compound output layer which contains the total combined area containing susceptible infrastructure. We then performed an *Intersect* with the infrastructure layers which gave us the extent of infrastructure damages due to a cascading risk.

We make several observations based on the data presented in Tables 5 and 6. Firstly, the magnitude of infrastructure damage escalates as the perceived inten-

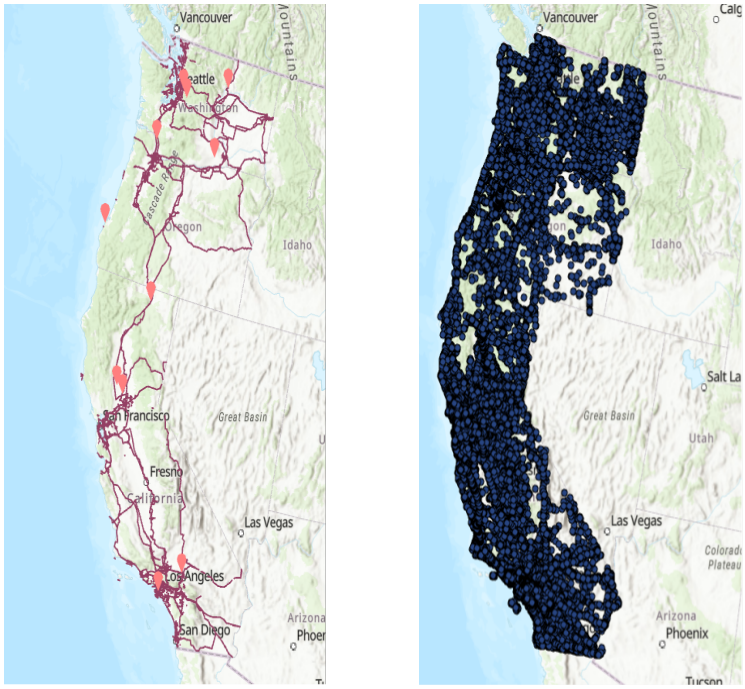

(a) Long-haul fiber optic cables and Starlink ground stations deployments.

(b) Cell towers deployments.

Fig. 6: Network infrastructures spanning across three states in the U.S. west coast.

	Cell towers		Fiber miles (km)		Ground stations	
Perceived shaking	10%	2%	10%	2%	10%	2%
Violent	516,474	237,294	11,285	11,024	5	6
Severe	462,040	597,941	9,106	15,171	1	4
Very strong	1,702	404,759	1,593	9,262	0	1

Table 5: Network infrastructures affected by earthquakes only scenario for expected PGAs with 10% and 2% probability of exceedances in the next 50 years [35].

sity of shaking experienced by an observer increases (e.g., violent, severe, very strong). Secondly, within each table, the extent of damage to each type of infrastructure asset resulting from earthquake shaking with a 10% probability of exceedance is generally lower than that with a 2% probability of exceedance. The only exception is violent shaking, where the damage from a 2% probability of exceedance is lower than that from a 10% probability. Third, when comparing the two tables, we observe a significant increase in the extent of damage to infrastructure assets due to cascading risks.

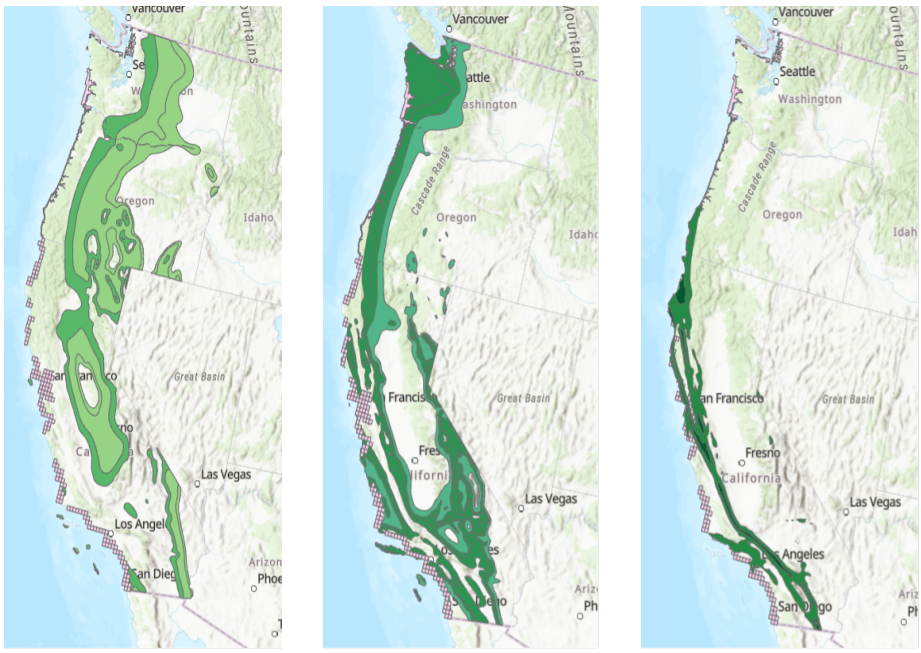

(a) M7 with PGA of 2% + tsunami

(b) M8 with PGA of 2% + tsunami

(c) M9 with PGA of 2% + tsunami

Fig. 7: Areas in the U.S. west coast that are susceptible to damage from earthquakes with different magnitudes that are expected with PGA of 2% followed by tsunamis.

	Cell towers		Fiber miles (km)		Ground stations	
Perceived shaking	10%	2%	10%	2%	10%	2%
Violent	689,316	531,974	21,306	20,675	4	6
Severe	598,491	737,067	14,017	25,198	1	4
Very strong	302,656	573,983	7,905	14,258	0	1

Table 6: Network infrastructures affected by earthquakes + tsunami scenario for expected PGAs with 10% and 2% probability of exceedances in the next 50 years [35].

In summary, this case study demonstrates the versatility of MAZE in assessing various risks across different regions of interest. It encompasses scenarios ranging from single-hazard to multi-hazard to cascading risks, highlighting the broad applicability of the tool.

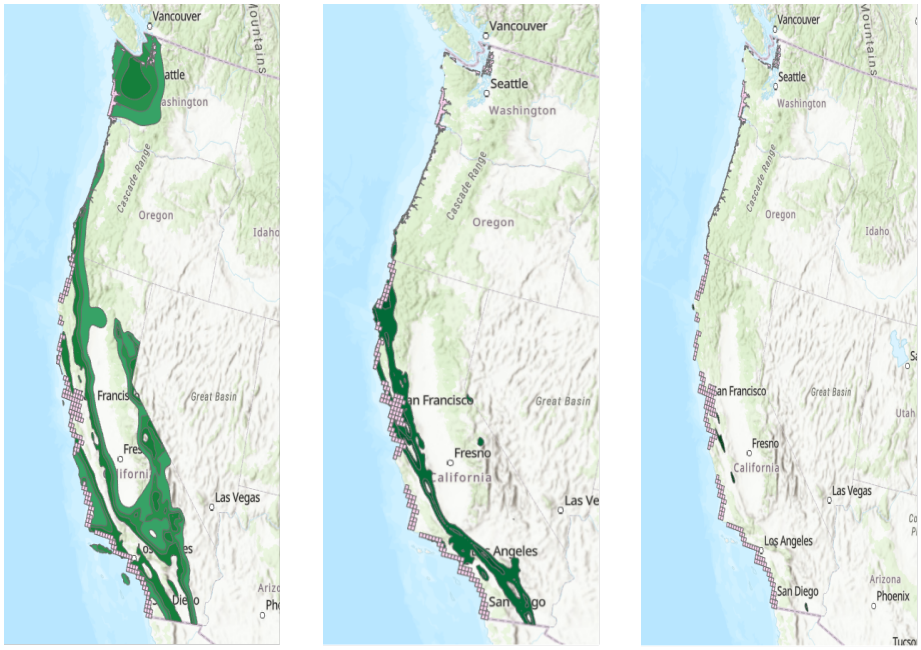

(a) M7 with PGA of 10% + tsunami

(b) M8 with PGA of 10% + tsunami

(c) M9 with PGA of 10% + tsunami

Fig. 8: Areas in the U.S. west coast that are susceptible to damage from earthquakes with different magnitudes that are expected with PGA of 10% followed by tsunamis.

6 Limitations and Discussion

6.1 Geographic Scope of MAZE and its Applicability

While the paper focuses on the U.S. PNW as a canonical example, we note that MAZE is a general-purpose tool. Concretely, MAZE (1) can be used to assess the efficacy of a wide-variety of backup routing and hazard mitigating strategies, and (2) is not limited to any particular geographic area. For (1), while one might contend that the LEO-based strategies used in this paper is hypothetical, we emphasize that our framework has the capability to incorporate actual performance measurements from Starlink or any LEO-based satellites between the selected sources and destination. In addition, evolving satellite connectivity landscape (e.g., StarLink mobility plans for first responders [5]) and how they fare during natural disasters could be studied using MAZE. For (2), without loss of generality, MAZE can be applied to any single- or multi-hazard risks, each occuring at different granularities (e.g., city vs. state vs. multiple states). We plan to study different combinations of (1) and (2) as part of future work.

6.2 Factors Affecting the Performance of Backup Paths

The need for reconfiguring backup paths over time as subsequent disasters arise is a key factor that will affect the performance on backup paths. Note that this work only considers the evaluation of backup paths for a one-off damage scenario. This is a rich area for future research. For example, MAZE can be extended as follows to evaluate disaster-aware dynamic reconfiguration of backup paths. First, add a new reliability metric to MAZE that keeps track of network load, packet drops, available paths, etc., over time as certain hops in backup paths degrade due to cascading disasters. Next, determine the operational threshold for the reliability metric, e.g., the extent of infrastructure damage that a path can withstand before dropping packets. Finally, apply the threshold on the metric to trigger subsequent re-routing via communication strategy of interest (e.g., LEO-based backup paths) during cascading disasters.

Another factor influencing the performance of backup paths is the gap between simulations and real-world operational constraints in LEO networks. This is especially important when evaluating the behavior of latency in LEO-based backup paths both before and after the disaster to account for factors such as simultaneous traffic congestion, infrastructure damage, and more. Furthermore, as pointed in § 3.3, the issue of RTT underestimates [33, 27] resulting from system-level and operational overheads aggravates this issue further. Addressing this issue in MAZE requires significant domain expertise in disaster modeling and consideration of real-world measurement data [27]. This is another ripe area for future research, given the recent uptick in measurement efforts that collect real-world RTTs of LEO networks [27, 22, 36, 30].

6.3 Non-terrestrial Risks to LEO Satellites

While LEO satellites could be used as "backup" communications infrastructure by first responders during terrestrial multi-hazard risks, as shown by Jyothi et al. [29], they are susceptible to another class of risk that is non-terrestrial in nature: Coronal Mass Ejections (CMEs). Potential impacts include disruption of communication due to damages resulting from radiation (e.g., charged particles from CMEs can disrupt communication), risk of Single Event Upsets (SEUs) (e.g., high-energy particles might flip one or more bits, impacting the sensitive electronics in LEO satellites), and power instabilities. Considering the limited knowledge we currently possess regarding the impact of these risks on LEO satellites as a whole, coupled with uncertainties surrounding the effectiveness of protective measures implemented by satellite operators, such as shielding sensitive components, we argue that LEO satellites should exclusively be utilized as backup communications infrastructure for terrestrial risks.

6.4 Lack of Community-wide Datasets

Due to an increase in the frequency and severity of natural disasters, it is crucial to have comprehensive dataset on the impact of such disasters on network

infrastructures. Community-wide datasets (1) serve as a central repository of valuable data, including historical records, real-time monitoring, and predictive models; and (2) facilitate informed decision-making, effective planning, and the implementation of appropriate mitigation strategies.

As discussed in § 2, previous efforts have examined the impacts of various risks on network infrastructure on a case-by-case basis (e.g., isolated events). However, as a measurement community, we currently lack comprehensive datasets that capture the patterns, trends, and risks posed by climate change and natural disasters to network infrastructures. The notable exception to this is the diligent work conducted by the Thunderping team [45, 43]. We believe that the availability of such datasets would enhance the realism and effectiveness of tools like MAZE.

The absence of datasets is not only a concern when it comes to understanding the impacts of natural hazards on network infrastructures but also extends to satellite providers and their associated costs. For instance, in addition to Inmarsat considered in MAZE, we know that there are several well-known satellite providers including HughesNet, ViaSat, among others. However, their pricing and service data are challenging to obtain, as they primarily provide consumer services with limited transparency regarding costs. This lack of pricing information hampers the assessment of the *long-term expenses* associated with using BGANs for communication during natural hazards. Additionally, the absence of data on the current hardware used by federal agencies for BGANs and the lifespan of these products makes it difficult to determine whether they offer any advantages over LEO-based services in the long run.

7 Summary and Future Work

Climate change-induced and naturally occurring multi-hazard risks are among the most significant threats to humanity and critical infrastructures alike. In this work, we seek to harden critical Internet infrastructures against multi-hazard risks. To this end, we develop a simple yet effective simulator called MAZE. Using MAZE, first responders and federal agencies can compare and contrast the benefits of various emergency communication, and can get better decision support on effective disaster mitigation strategies in a repeatable manner. We demonstrate the efficacy of MAZE by comparing LEO satellite-based emergency communication strategy against two baselines (i.e., ShakeNet and BGAN) in the face of different disaster scenarios. Our simulations show that LEO satellite-based hardening strategies offer two orders of magnitude latency improvement and 100s of thousands of dollars in saving, all while maintaining network connectivity during multi-hazard risks.

While one of the case studies demonstrates how MAZE can be used to transition a research prototype to the real world, three key challenges (listed below) remain in using LEO-based satellites for multi-hazard risk scenarios. We believe MAZE could be used to tackle each one of these challenges, which we plan to focus on as part of future work.

– First, MAZE could be extended to study the issue of scalability and network capacity of LEO satellite-based communication systems in the face of increasing demands during multi-hazard events.
– Second, MAZE could be used for assessing and validating the practicality and effectiveness of risk-aware routing protocols or mitigation strategies for LEO satellite-based communication during multi-hazard risks. This, of course, requires partnerships with industry, academic, and government stakeholders alike.
– Third, MAZE could be used to explore the policy and regulatory considerations associated with using LEO satellite networks for emergency communication. Potential opportunities include investigating spectrum allocation, licensing requirements, and coordination with government agencies to ensure compliance and seamless integration of LEO satellite systems for emergency communications.

Acknowledgements

We thank the anonymous reviewers and our shepherd, Nitinder Mohan, for their insightful feedback. This work is supported by the Internet Society (ISOC) Foundation. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of ISOC.

References

1. ArcGIS. `https://www.arcgis.com`.
2. Inmarsat BGAN Data Plans. `https://satellitephonestore.com/bgan-service`.
3. Inmarsat BGAN M2M. `https://www.inmarsat.com/en/solutions-services/enterprise/services/bgan-m2m.html`.
4. Satellite Configurations Used in MAZE. `https://gitlab.com/onrg/maze/-/raw/main/rtt_simulator/constellation_config.py`.
5. StarLink for Land Mobility. `https://www.starlink.com/business/mobility`.
6. The Oregon Resilience Plan – Cascadia: Oregon's Greatest Natural Threat. `https://www.oregon.gov/oem/Documents/01_ORP_Cascadia.pdf`.
7. *Fourth National Climate Assessment. Volume II, Impacts, risks, and adaptation in the United States. Report-in-brief.* U.S. Global Change Research Program, Washington, DC, 2018.
8. State Occupational Employment and Wage Estimates Oregon. `https://www.bls.gov/oes/current/oes_or.htm`, May 2021.
9. United states census bureau, county population: 2020-2021. `https://www.census.gov/data/tables/time-series/demo/popest/2020s-counties-total.html`, 2021.
10. Jonathan Allan, Joseph Zhang, Fletcher O'brien, and Laura Gabel. Columbia river tsunami modeling: Toward improved maritime planning response, 12 2018.
11. Scott Anderson, Carol Barford, and Paul Barford. Five alarms: Assessing the vulnerability of us cellular communication infrastructure to wildfires. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, page 162–175, New York, NY, USA, 2020. Association for Computing Machinery.

12. Scott Burleigh and Keith Scott. Interplanetary Overlay Network. https://www.inmarsatgov.com/firstnet/wp-content/uploads/2020/03/SATCOM-overview-firstnet.pdf, june 2020.
13. https://climate.nasa.gov/news/2926/can-climate-affect-earthquakes-or-are-the-connections-shaky/.
14. Aizaz U Chaudhry and Halim Yanikomeroglu. Optical wireless satellite networks versus optical fiber terrestrial networks: The latency perspective. In *30th Biennial Symposium on Communications*, pages 225–234. Springer, 2021.
15. Kenjiro Cho, Cristel Pelsser, Randy Bush, and Youngjoon Won. The japan earthquake: the impact on traffic and routing observed by a local isp. In *Proceedings of the Special Workshop on Internet and Disasters*, pages 1–8, 2011.
16. Ramakrishnan Durairajan, Carol Barford, and Paul Barford. Lights Out: Climate Change Risk to Internet Infrastructure. In *proceedings of Applied Networking Research Workshop*, 2018.
17. Brian Eriksson, Ramakrishnan Durairajan, and Paul Barford. Riskroute: A framework for mitigating network outage threats. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*, pages 405–416, 2013.
18. FirstNet. Satellite Solutions for FirstNet. https://www.inmarsatgov.com/firstnet/wp-content/uploads/2020/03/SATCOM-overview-firstnet.pdf, Mar 2020.
19. Josh Fomon. Here's How Fast Starlink Has Gotten Over the Past Year. https://www.ookla.com/articles/starlink-hughesnet-viasat-performance-q1-2022, June 2022.
20. A. Franchi, A. Howell, and J. Sengupta. Broadband mobile via satellite: inmarsat bgan. In *IEE Seminar on Broadband Satellite: The Critical Success Factors - Technology, Services and Markets (Ref. No. 2000/067)*, pages 23/1–23/7, 2000.
21. Jill C. Gallagher. The first responder network (firstnet) and next-generation communications for public safety : issues for congress, 2018.
22. Johan Garcia, Simon Sundberg, Giuseppe Caso, and Anna Brunstrom. Multi-timescale evaluation of starlink throughput. In *Proceedings of the 1st ACM Workshop on LEO Networking and Communication*, pages 31–36, 2023.
23. Vittorio A. Gensini and Harold E. Brooks. Spatial trends in united states tornado frequency. *NPJ climate and atmospheric science*, 1(1), 2018.
24. Douglas D Given, Richard M Allen, Annemarie S Baltay, Paul Bodin, Elizabeth S Cochran, Kenneth Creager, Robert M de Groot, Lind S Gee, Egill Hauksson, Thomas H Heaton, et al. Revised technical implementation plan for the shakealert system—an earthquake early warning system for the west coast of the united states. Technical report, US Geological Survey, 2018.
25. Chunhao Han. The beidou navigation satellite system. In *2014 XXXIth URSI General Assembly and Scientific Symposium (URSI GASS)*, pages 1–3, 2014.
26. King County Public Health. Division of emergency medical services 2021 annual report. https://kingcounty.gov/depts/health/emergency-medical-services/~/media/depts/health/emergency-medical-services/documents/reports/2021-Annual-Report.ashx, september 2021.
27. Liz Izhikevich, Manda Tran, Katherine Izhikevich, Gautam Akiwate, and Zakir Durumeric. Democratizing LEO Satellite Network Measurement. *arXiv preprint arXiv:2306.07469*, 2023.
28. Donny Jackson. FirstNet Authority seeks input on potential solutions to off-network challenges. *Urgent Communications*, 2021.

29. Sangeetha Abdu Jyothi. Solar superstorms: planning for an internet apocalypse. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 692–704, 2021.
30. Mohamed M Kassem, Aravindh Raman, Diego Perino, and Nishanth Sastry. A browser-side view of starlink connectivity. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 151–158, 2022.
31. Simon Kassing, Debopam Bhattacherjee, André Baptista Águas, Jens Eirik Saethre, and Ankit Singla. Exploring the "internet from space" with hypatia. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, page 214–229, New York, NY, USA, 2020. Association for Computing Machinery.
32. Connor Kimball. How much bandwidth is needed for voip. `https://www.avoxi.com/blog/how-much-bandwidth-is-needed-for-voip/`.
33. Zeqi Lai, Hewu Li, Yangtao Deng, Qian Wu, Jun Liu, Yuanjie Li, Jihao Li, Lixin Liu, Weisen Liu, and Jianping Wu. {StarryNet}: Empowering Researchers to Evaluate Futuristic Integrated Space and Terrestrial Networks. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 1309–1324, 2023.
34. Bozhong Li, Zifan Li, Hongxi Zhou, Xinpeng Chen, Yuanlong Peng, Peng Yu, Yun Wang, and Xue Feng. A system of power emergency communication system based bds and leo satellite. In *2021 Computing, Communications and IoT Applications (ComComAp)*, pages 286–291, 2021.
35. Juno Mayer, Valerie Sahakian, Emilie Hooft, Douglas Toomey, and Ramakrishnan Durairajan. On the resilience of internet infrastructures in pacific northwest to earthquakes, 2021.
36. François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. A first look at starlink performance. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 130–136, 2022.
37. Northwest Interagency Coordination Center. Northwest fire locations. `https://gacc.nifc.gov/nwcc/information/firemap.aspx`.
38. Bureau of Labor Management. Occupational Employment and Wage Statistics of EMTs, 2021.
39. Bureau of Labor Management. Occupational Employment and Wage Statistics of Firefighters, 2021.
40. Bureau of Labor Management. Occupational Employment and Wage Statistics of Paramedics, 2021.
41. Washington State Department of Natural Resources. Tsunami hazard maps, 2022.
42. National Association of State Budget Officers. State Expenditure Report. `https://www.nasbo.org/reports-data/state-expenditure-report`, 2021.
43. Ramakrishna Padmanabhan, Aaron Schulman, Dave Levin, and Neil Spring. Residential links under the weather. In *Proceedings of the ACM Special Interest Group on Data Communication*, pages 145–158. 2019.
44. Portland Fire & Rescue. FAQs. `https://www.portlandoregon.gov/fire/article/378460`, 2021.
45. Aaron Schulman and Neil Spring. Pingin' in the Rain. In *ACM IMC*, November 2011.
46. SkyBrokers. Amazon Kuiper Systems, LLC. urlhttps://sky-brokers.com/supplier/amazon-kuiper-systems-llc/, 2022.
47. Ken Smith, Graham Kent, David Slater, Gabe Plank, Mark Williams, M McCarthy, Frank Vernon, Neal Driscoll, Hans-Werner Braun, R Anderson, et al. Integrated multi-hazard regional networks: Earthquake warning/response, wildfire detection/response, and extreme weather tracking. *Applied Geology in California:*

Association of Environmental and Engineering Geologists (AEG) Special Publication, (26):599–612, 2016.
48. Starlink. Starlink kit. `https://www.starlink.com`.
49. Starlink. World's most advanced broadband internet system. `https://www.starlink.com/satellites`, 2021.
50. Telesat. Telesat Lightspeed LEO Network. `https://www.telesat.com/leo-satellites/`, Nov 2021.
51. Verizon. First responder benefits program. `https://www.verizon.com/business/solutions/public-sector/public-safety/programs/first-responder-benefits/`, 2020.
52. Jian Zhou, Xiaoguo Ye, Yong Pan, Fu Xiao, and Lijuan Sun. Dynamic channel reservation scheme based on priorities in leo satellite systems. *Journal of Systems Engineering and Electronics*, 26(1):1–9, 2015.