Solar Superstorms: Planning for an Internet Apocalypse

Sangeetha Abdu Jyothi

University of California, Irvine and VMware Research

sangeetha.aj@uci.edu

ABSTRACT

Black swan events are hard-to-predict rare events that can significantly alter the course of our lives. The Internet has played a key role in helping us deal with the coronavirus pandemic, a recent black swan event. However, Internet researchers and operators are mostly blind to another black swan event that poses a direct threat to Internet infrastructure. In this paper, we investigate the impact of solar superstorms that can potentially cause large-scale Internet outages covering the entire globe and lasting several months. We discuss the challenges posed by such activity and currently available mitigation techniques. Using real-world datasets, we analyze the robustness of the current Internet infrastructure and show that submarine cables are at greater risk of failure compared to land cables. Moreover, the US has a higher risk for disconnection compared to Asia. Finally, we lay out steps for improving the Internet's resiliency.

CCS CONCEPTS

• **Networks** → **Network reliability**; **Network structure**;

KEYWORDS

Internet Resilience, Internet Topology, Solar storms

ACM Reference Format:

Sangeetha Abdu Jyothi. 2021. Solar Superstorms: Planning for an Internet Apocalypse. In *ACM SIGCOMM 2021 Conference (SIGCOMM '21)*, August 23–27, 2021, Virtual Event, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3452296.3472916>

1 INTRODUCTION

What will happen if there is a global Internet collapse? A disruption lasting even a few minutes can lead to huge losses for service providers and damages in cyber-physical systems. The economic impact of an Internet disruption for a day in the US is estimated to be over \$7 billion [1]. What if the network remains non-functional for days or even months? This is the worst-case scenario, which, fortunately, we have never encountered in recent history. Threats to the Internet range from man-made cyber attacks to natural disasters such as earthquakes. The Internet is also affected by black swan events such as the Covid-19, which profoundly alter human lives and, in turn, our Internet usage. However, the influence of these indirect threats on the Internet is only secondary, with the worst-case impact often limited to reduced speeds.

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License. *SIGCOMM '21*, August 23–27, 2021, Virtual Event, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8383-7/21/08.
<https://doi.org/10.1145/3452296.3472916>

One of the greatest dangers facing the Internet with the potential for global impact is a powerful solar superstorm. Although humans are protected from these storms by the earth's magnetic field and atmosphere, they can cause significant damage to man-made infrastructure. The scientific community is generally aware of this threat with modeling efforts and precautionary measures being taken, particularly in the context of power grids [41, 43]. However, the networking community has largely overlooked this risk during the design of the network topology and geo-distributed systems such as DNS and data centers.

A Coronal Mass Ejection (CME), popularly known as solar storm, is a directional ejection of a large mass of highly magnetized particles from the sun. When the earth is in the direct path of a CME, these magnetized and charged solar particles will interact with the earth's magnetic field and produce several effects. In addition to spectacular auroral displays, they produce Geomagnetically Induced Currents (GIC) on the earth's surface through electromagnetic induction. Based on the strength of the CME, in extreme cases, GIC has the potential to enter and damage long-distance cables that constitute the backbone of the Internet.

The largest solar events on record occurred in 1859 and 1921, long before the advent of modern technology. They triggered extensive power outages and caused significant damage to the communication network of the day, the telegraph network. The probability of occurrence of extreme space weather events that directly impact the earth is estimated to be 1.6% to 12% per decade [42, 65]. More importantly, the sun was in a period of low activity in the past three decades [61] from which it is slowly emerging. Since this low phase of solar activity coincided with the rapid growth of technology on the earth, we have a limited understanding of whether the current infrastructure is resilient against powerful CMEs.

In this paper, we analyze the threat posed by solar superstorms to the Internet infrastructure and the steps to be taken to mitigate its effects. First, we ask a key question: is the threat significant, and should we factor this in Internet topology design and infrastructure deployment (§ 2)? Second, we study the impact of solar storms on key building blocks of the Internet infrastructure — long-haul land and submarine cables (§ 3). Third, using real-world datasets and a wide range of failure models, we quantify the impact of solar superstorms on the Internet infrastructure (§ 4). Finally, we lay out steps to manage the perils associated with solar superstorms (§ 5).

Long-distance fiber cables and communication satellites are susceptible to damage from solar storms through induced currents and direct exposure, respectively (§ 3). In cables, the optical fiber itself is immune to GIC. However, long-haul cables have repeaters to boost the optical signals spaced at intervals of 50 – 150 km which are powered using a conductor. These repeaters are vulnerable to GIC-induced failures, which can lead to the cable being unusable. GPS and communication satellites which are directly exposed to solar storms will suffer from lost connectivity during the event,

potential damage to electronic components, and in the worst case, orbital decay and reentry to earth (particularly in low earth orbit satellites such as StarLink [14]).

In order to study the impact of CMEs on terrestrial networks, we use a comprehensive set of Internet topology datasets, including submarine and land cables, DNS root servers, IXPs, Internet routers, etc. Since accurate modeling of repeater failures is not available, we employ a broad range of failure models derived based on GIC characteristics.

Our experiments provide several interesting insights regarding the Internet topology and its vulnerabilities. First, the topology is skewed with respect to Internet user distribution. There is a higher concentration of infrastructure elements on higher latitudes that are more vulnerable to solar superstorms. Second, submarine cables are more vulnerable than land cables, primarily due to their larger lengths. Third, different regions will be impacted differently. The US is highly susceptible to disconnection from Europe. Europe is in a vulnerable location but is more resilient due to the presence of a larger number of shorter cables. Asia has relatively high resilience with Singapore acting as a hub with connections to several countries. Finally, we analyze the impact on various Internet systems. DNS root servers are less vulnerable since they are highly geo-distributed. Google data centers have better resilience than Facebook's. A large fraction of Autonomous Systems have a presence in the higher latitudes, but a majority of them are geographically restricted to a smaller area.

Although the highest priority system for recovery during a solar event will be the power grid, the Internet is also a critical infrastructure necessary for disaster management. While this paper focuses on the vulnerabilities of the Internet infrastructure alone, a discussion on the interdependence with power grids and associated challenges are presented in § 5.

In summary, we make the following contributions:

- We present the first study that analyzes threats to the Internet infrastructure posed by a high-risk event: solar superstorms.
- We identify several vulnerabilities in the design of current Internet topology and associated geo-distributed infrastructure such as DNS and Autonomous Systems.
- We show that the Internet infrastructure distribution is skewed with respect to the user population. Internet infrastructure components are concentrated in higher latitudes that are susceptible to solar events.
- We investigate the impact of Geomagnetically Induced Currents on various infrastructure components and show that submarine cables are at the highest risk of damage.
- We demonstrate that the potential impact of solar superstorms on different regions varies widely. The US is highly vulnerable to disconnection compared to Asian countries.
- We discuss several open questions on improving Internet resiliency, including how to factor in solar superstorms during the design of Internet topology and other Internet sub-systems.

This paper does not raise any ethical concerns.

2 MOTIVATION: A REAL THREAT

In this section, we present a discussion on threats posed by solar activity and the likelihood of extreme solar events that can affect the earth.

2.1 Solar flares and CMEs

Solar activity waxes and wanes in cycles, with a period length of approximately 11 years [23]. During solar maxima, there is an increase in the frequency of two solar phenomena, solar flares and Coronal Mass Ejections (CMEs), both caused by contortions in the sun's magnetic fields [35].

Solar flares involve large amounts of emitted energy as electromagnetic radiation. Although flares can reach earth in 8 minutes, they affect only the upper layers of the atmosphere, particularly the ionosphere, causing disruptions to satellite communication and GPS. Solar flares do not pose any threat to terrestrial communication or other infrastructure.

A *Coronal Mass Ejection* (CME) involves the emission of electrically charged solar matter and accompanying magnetic field into space. It is typically highly directional. This cloud of magnetized particles can take 13 hours to five days to travel to the earth. They cannot penetrate the atmosphere and affect humans directly. However, they will interact with the earth's magnetic field and induce strong electric currents on the earth's surface that can disrupt and even destroy various human technologies. This will occur only if the earth happens to be on the path of a CME.

2.2 Past CME events

The first recorded CME with a major impact on the earth is the Carrington event (Sep 1, 1859) [21]. This cataclysmic CME reached the earth in just 17.6 hours owing to its very high speed. The communication network of the day, the telegraph network, suffered from equipment fires, and several operators experienced electric shocks. This caused large-scale telegraph outages in North America and Europe. Even when power was disconnected, telegraph messages could be sent with the current generated by the CME. A recent study [50] which analyzed the risks posed by a Carrington-scale event to the US power grid today found that 20 – 40 million people could be without power for up to 2 years, and the total economic cost will be 0.6 – 2.6 trillion USD. To the best of our knowledge, there are no existing studies on the risks posed to the Internet infrastructure by such an event.

The largest geomagnetic storm of the 20th century, which occurred in May 1921, named the New York Railroad superstorm based on its impact on NY telegraph and railroad systems, also caused widespread damage across the globe [47]. Note that the strongest CME of the past century happened before widespread electrification. A smaller-scale CME had caused the collapse of the power grid in the entire province of Quebec, Canada, and over 200 grid problems at various locations in the US in 1989 [59]. However, this was only a moderate-scale CME.

A CME of Carrington-scale missed the earth by merely a week in July 2012 [62]. Fortunately, given the highly directional nature of CMEs, they can cause significant damage only when the earth is in the direct path.

2.3 Can we predict the next large event?

Similar to other natural phenomena such as earthquakes and the collapse of a star into a black hole, solar activity is extremely difficult to predict. To make matters worse, unlike earthquakes, we have limited data on intense solar phenomena that impact the earth because they are rarer and more difficult to study. Although it is impossible to forecast the exact occurrence of a catastrophic solar event and prediction of such events continues to be one of the hardest challenges in astrophysics, scientists have developed several models based on past observations.

Frequency estimates based on limited data for a direct impact currently range from 2.6 to 5.2 per century [16, 17, 51, 52]. There are also several studies assessing the probability of occurrence of a Carrington-scale event. Current estimates range from 1.6% [42] to 12% [65] probability of occurrence per decade for a large-scale event (note that the probability of occurrence per decade of a once-in-a-100-years event is 9%, assuming a Bernoulli distribution where events are independent). Today, there are several models with various knobs to capture the behavior of solar cycles. But the sensitivity of these knobs and the actual behavior of the sun remains largely elusive, with no clear winner across models. However, there is another factor that increases the risk of solar storms in the near term (next couple of decades).

The frequency of CMEs is not uniform across solar maxima. In addition to the 11-year cycle, solar activity also goes through a longer-term cycle in approximately 80 – 100 years called the Gleissberg cycle [33, 61]. This cycle causes the frequency of high-impact events like CMEs to vary by a factor of 4 across solar maxima [51]. The most recent solar cycles, cycle 23 (1996-2008) and cycle 24 (2008-2020), are a part of an extended minimum in the current Gleissberg cycle [30, 31]. In other words, ***modern technological advancement coincided with a period of weak solar activity and the sun is expected to become more active in the near future.*** Hence, the current Internet infrastructure has *not* been stress-tested by strong solar events.

Early predictions for the current solar cycle, which began in 2020, ranged from weak [19, 71] (a part of the current Gleissberg minimum) to moderately strong [28, 46, 67]. However, a recent study from November 2020 [53] suggested that this cycle has the potential to be one of the strongest on record. Recent estimates for the number of sunspots at the peak of this cycle are between 210 and 260 (a very high value) [37, 53]. In contrast, the previous cycle that ended in 2019 had a peak sunspot number of 116. Since CMEs often originate in magnetically active regions near sunspots, a larger number of sunspots will increase the probability of a powerful CME. If this estimate [53] proves accurate, it will also significantly increase the probability of a large-scale event in this decade. The actual strength of this cycle will be evident only later in the decade as the solar cycle progresses.

In the 20th century, the Gleissberg minimum point was in 1910 [31] and the largest CME of the century occurred a decade later in 1921 [47]. The past 2 – 3 solar cycles, which coincided with the birth and growth of the Internet were very weak. Given that a strong solar cycle that can produce a Carrington-scale event can occur in the next couple of decades, we need to prepare our infrastructure now for a potential catastrophic event.

3 IMPACT ON NETWORKS

Having established that solar superstorms are a real threat with a significant probability of occurrence in the near- and long-term, in this section, we discuss its impact on networks. We provide a brief overview of how CMEs produce geomagnetically induced currents on the earth's surface and how they affect Internet cables. We also briefly mention the effects on satellite communication. However, the focus of this paper is the impact on terrestrial communication networks, which carry the majority of the Internet traffic.

3.1 Geomagnetically Induced Current

CMEs produce variations in the earth's magnetic field, which in turn induce geoelectric fields on the earth's conducting surface (i.e., land and ocean floor). These spatiotemporally varying electric fields are responsible for the generation of Geomagnetically Induced Currents (GIC) [32, 64], as high as 100-130 Amps [58], that can flow through any extended ground-based conductive systems such as power grids, networking cables, etc. This electromagnetically induced current enters/exits long-distance conductors from grounded neutral, causing destruction of electrical equipment such as transformers/repeaters and, in turn, large-scale power outages/Internet outages spanning many states or even countries. The amplitude of GIC depends on a variety of factors, such as the time derivative of the geomagnetic field and the resistivity of the earth's crust and upper mantle.

Several factors influence the strength of GIC. (i) *Conductor length*: GIC is primarily induced in "long conductors" since the current is proportional to the area of the loop formed by the two grounds and the cable [54]. Hence, power grids [41, 43], oil and gas pipelines [72], networking cables, etc. are most vulnerable. Geographically localized infrastructure such as data centers can be protected using Transient Voltage Surge Suppression (TVSS) devices which are relatively inexpensive (~\$1000s). (ii) *Latitude*: Higher latitudes are at a significantly higher risk [63, 68, 69] (similar to other solar effects such as auroras). During the medium-scale 1989 event, the magnitude of the induced electric field dropped by an order of magnitude below 40°N [63]. During the Carrington event, estimates show that strong fields extended as low as 20°N [63] (limited measurements available from 1859). Recent studies show that GIC of small magnitudes can occur at the equator [22, 68, 75]. But the strength of variations in the field in equatorial regions is significantly lower than that in higher latitudes [22, 68, 75]. (iii) *Geographic spread*: Since GIC is caused by changes to the earth's magnetic field, it affects wide areas and is not restricted to the portion of the earth facing the sun. (iv) *Orientation of conductor*: Since CME-induced fluctuations do not have a directional preference (e.g., North-South vs. East-West), conductors along different orientations on earth are at equal risk [32].

Note that seawater has high conductivity [26]. The presence of highly conductive seawater over more resistive rocks increases the total conductance of the surface layer [27]. Hence, the ocean does not reduce the impact of GIC but increases it. For example, a study that modeled the geoelectric fields and potential GIC impact around New Zealand reported conductance in the range 1-500 S on land and 100-24,000 S in the ocean surrounding New Zealand [27]. A higher conductance implies that a higher GIC could be induced.

3.2 Impact on Long-Distance Cables

3.2.1 Understanding the vulnerabilities. Long-distance land and submarine cables carry signals in optical fibers. The fiber in the cable is immune to GIC, unlike the previous generation of coaxial cables, since it carries light and not electric current. However, long-haul cables that stretch hundreds or thousands of kilometers also have an accompanying conductor that connects repeaters in series along the length of cables called the power feeding line [4]. This conductor is susceptible to GIC [5, 49].

Power Feeding Equipments (PFEs), located in landing stations at the ends of the cable, power the repeaters which are connected in series via the power feeding line with a 1.1 Ampere regulated current. The resistance of the power feeding line is approximately $0.8 \Omega/\text{km}$. However, the actual voltage requirement is influenced by several factors, including earth potential difference at either end of the cable, the number of spare repeaters in the cable, etc. Considering these factors, a 10 Gbps 96-wave 9000 km long cable typically requires a power feeding voltage of about 11 kV and approximately 130 repeaters connected in series [76]. In practical deployments, inter-repeater distance vary from 50 to 150 km [48, 66].

Note that the repeaters are designed to operate at $\sim 1\text{A}$ current [73, 76]. However, as discussed in § 3.1, GIC during strong solar events can be as high as 100 – 130 Amperes. This is $\sim 100\times$ more than the operational range of these repeaters. Thus, in the event of a solar superstorm, *repeaters are susceptible to damage from GIC*. Moreover, even a single repeater failure can leave all parallel fibers in the cable unusable due to weak signal strength or disruption of power.

3.2.2 Recovery challenges. The specified lifetime of repeaters in submarine cables is 25 years [24]. Once deployed under the ocean, they are typically highly resilient unless the cable is damaged by human interference. This is a design requirement since the replacement/repair of repeaters or parts of the cable is expensive. Commonly, underwater cable damages are localized, and typical causes are fishing vessels, ship anchors, or earthquakes. When damage occurs, the location of the damage is first identified using tests from the landing sites, and then a cable ship is sent to the location for repair. This repair process can take days to weeks for a single point of damage on the cable.

The current deployment of submarine cables has never been stress-tested under a strong solar event. Due to the lack of real-world data on GIC effects on repeaters, the potential extent of damage (the number of repeaters that could be destroyed) and the time required to repair significant portions of a cable are unknown. However, the extent of damage is not dependent on the distance between the repeaters. It depends on the distance between the ground connections. GIC enters and exits the power feeding line at the points where the conductor is grounded, even when the cable is not powered. The potential difference between these earth connections determines the strength of GIC entering the cable. A short cable ($<50 \text{ km}$) without any repeaters does not require a ground connection. However, longer cables are grounded at several intermediate points, at intervals of 100s to 1000s of kilometers. For example, Google's Equiano cable connecting South Africa to Portugal has nine branching units [9] which are connected to the sea earth.

During the moderate storm of March 1989 [59] (one-tenth the strength of 1921 storm), considerable potential variations were measured on the only long-distance submarine cable of the time, the AT&T cable between New Jersey, US and the UK [56]. Since the variations were of moderate magnitude, the PFE at landing stations could control the impact. However, the cost for designing and deploying PFEs that are able to handle the highest possible variations is exorbitant and hence, not practical [45]. A TVSS device at the landing point cannot protect repeaters along the length of the cable.

Modeling the impact of GIC on repeaters and cables is an extremely difficult task. During a CME, induced GIC at a location will depend on the strength of the induced electric field, its direction, ground/crust composition, and a variety of other factors. Moreover, while long-distance cable endpoints are well-known, we have limited information about intermediate grounding points and the extent of GIC possible between them based on terrestrial characteristics. In short, a detailed analysis of the impact of CMEs on long-distance cables is hindered by the limited availability of data as well as the lack of tools and techniques that facilitate analysis at the intersection of astrophysics, electrical engineering, and optical engineering.

In this paper, we take the first step towards bridging the gap in our understanding of the impact of CMEs on terrestrial communication infrastructure using an extensive set of infrastructure datasets and a broad range of failure models. To the best of our knowledge, this is the first effort to bring together various pieces required to understand this threat in the context of the Internet. Note that more sophisticated models, particularly for repeater failures, can be plugged into our analyses in the future when they become available. This paper uses a wide range of probabilistic models to study the space since accurate modeling is not available. Real-world datasets and models are used in all contexts where feasible and available.

3.3 Impact on Satellites

Communication satellites are among the severely affected systems. The damages are not caused by GIC but due to direct exposure to highly charged particles in CMEs. These particles do not reach the earth's surface since they are mostly blocked by the atmosphere. Threats to communication satellites include damage to electronic components and extra drag on the satellite, particularly in low-earth orbit systems such as Starlink [14], that can cause orbital decay and uncontrolled reentry to earth [39].

Thus, *both surface-based and satellite-based communication systems are under high risk of collapse if a Carrington-scale event occurs again*.

4 ANALYZING THE IMPACT

To better understand the risk, we analyze various infrastructure components that constitute the current Internet topology. In addition to the threats posed to physical infrastructure, we also evaluate the implications for software building blocks of the Internet. All datasets except the private ITU land cables and the code used for experiments in this paper are available at <https://github.com/NetSAIL-UCI/Internet-Resilience>.

Figure 1: IXPs, long-distance transmission links on land, and submarine cables [10].

Figure 2: Public data centers and colocation centers [2].

4.1 Datasets

In this section, we give a detailed overview of datasets used in the physical infrastructure analysis.

4.1.1 Submarine cable map. The submarine cable map [15] consists of 470 cables which interconnect 1241 landing points (nodes) across the globe with latitude/longitude information of the landing points. Each cable has multiple branches and thus interconnects several cities. The longest cable has a length of 39,000 km.

4.1.2 Intertubes dataset. We rely on the Intertubes [29] dataset with the US long-haul fiber endpoints for analysis of land transmission lines in the US.

4.1.3 ITU dataset. The TIES version of the ITU transmission map consists of land and submarine communication network information. This map is built using data from several sources ranging from direct information obtained from operators to secondary maps built by organizations for larger regions or countries. Since submarine cables are considered as a separate dataset, we use only land cables from ITU in our analysis. Moreover, while the original dataset contains fiber and microwave links, we restrict our analysis to fiber links. The ITU dataset contains both long-haul and short-haul links. Note that we have not removed the Intertubes dataset from the ITU dataset to prevent unnecessary network partitioning in the bigger dataset. The ITU dataset contains 11,737 network links from across the globe that interconnect 11,314 nodes. The

exact latitude and longitude information are not available for this dataset. But nodes labeled with location names and the links are available. A snapshot of submarine cables, long-haul fiber, and IXP locations from the ITU website [10] is shown in Figure 1.

4.1.4 CAIDA dataset. We use the CAIDA Internet Topology Data Kit to analyze the distribution of routers across the globe as well as the extent of the spread of Autonomous Systems (AS). This dataset contains latitude and longitude information of 46,032,818 Internet routers. This dataset also contains router to AS mapping for 46,014,869 routers across 61,448 Autonomous Systems (ASes), which allows us to study the geographic spread of these ASes. We use the topology derived using MIDAR and iffinder. This topology has the highest confidence aliases with very low false positives.

4.1.5 DNS root servers. DNS root server locations of 1076 instances across 13 root servers are obtained from the root server directory [3].

4.1.6 IXP dataset. The IXP dataset obtained from the PCH Internet Exchange directory [13] contains 1026 locations across the globe, including their coordinate information.

4.1.7 Data center map. A visual snapshot of data centers across the globe was obtained from Data Center Map [2] (Figure 2). We also analyze data center locations of large content providers such as Google [6], Facebook [38], etc.

4.1.8 Population dataset. To estimate the fraction of population affected at each latitude, we use the gridded population data (with population per cell of length and breadth 1°) from NASA Socioeconomic Data and Applications Center [70].

4.2 Infrastructure Distribution

Due to the location dependence of solar storm impact, understanding the physical infrastructure distribution is the first step towards understanding infrastructure vulnerabilities.

4.2.1 Methodology. As discussed in § 3, long fiber cables with conductors for powering repeaters are at a significantly high risk. There are two key factors that affect a cable's risk of damage under GIC: *location of its end points* and *its length*.

First, locations above $40^\circ N$ and below $40^\circ S$ are more affected by solar superstorms (we use the conservative threshold of 40° stated in [63]). Various studies consider different thresholds in the range $40 \pm 10^\circ$. Hence, we evaluate the impact of solar superstorms on the Internet by first analyzing the distribution of network topology and related systems such as DNS, IXPs, etc. across various latitudes. Second, longer cables are at a higher risk (§ 3.2). Hence, we analyze the link lengths across various datasets to understand the risks faced by both land and submarine cables.

4.2.2 Evaluation. In Figure 3, we plot the probability density function of submarine endpoints and world human population (each with densities calculated over 2° intervals). Although a significant fraction of the population in the northern hemisphere is below the 40^{th} parallel N, there is a higher concentration of submarine endpoints at higher latitudes.

Figure 3: PDF of population and submarine cable end points with respect to latitude.

While smaller countries have limited flexibility regarding the location of their submarine endpoints due to their tight geographical boundaries, we observe that there is room for improvement, particularly in larger nations like the US. While submarine cables between the US and Asia are more uniformly distributed along the west coast from Seattle to Southern California, there is a higher concentration of cables between the North East and the northern parts of Europe. There is only a single cable connecting Florida with Portugal and Spain in southern Europe (below $40^\circ N$).

We analyze the location distribution of other components in the Internet ecosystem (data centers, DNS root servers, IXPs) and observe a similar pattern of higher density at higher latitudes. The distribution of public data centers and colocation centers are shown in Figure 2 which follows the same pattern. In Figure 4, we show that 31% of submarine endpoints, 40% of Intertubes endpoints, 43% of IXPs, 38% of Internet routers, and 39% of DNS root servers are located above 40° . Moreover, another 14% of submarine endpoints have a direct link to these nodes, putting these locations at risk of GIC induced currents as well. However, only 16% of the world population is in this region.^{1 2}

We also evaluate the average cable length in the US long-haul fiber network, the global ITU land fiber network, and the global submarine cable network. The US long-haul fiber dataset [29] only provides approximate node locations and link information. Since these cables are known to be located adjacent to the US road system [29], we estimate the link length as the driving distance between the endpoints using Google maps API. From the publicly available

¹Note that since the Internet user population in developing countries grew rapidly in the past two decades and the Internet infrastructure deployment has not advanced at the same pace, observations from past work such as the Internet infrastructure being located predominantly where the users reside [44] are not valid today.

²While the Internet user distribution is not the same as the population distribution, these are very similar, and our conclusion regarding the distribution of Internet infrastructure in relation to users holds. The difference in percentage points between the population of a continent and Internet users in a continent as a fraction of the world is at most 5.5% [8]. For e.g., Asia has 55% of the world's population and 52% of the world's Internet users (difference of 3% in Asia, the largest difference of 5.5% in Africa). While Internet user statistics based on latitude are not available, using the highest rate of Internet penetration above 40° and recent data on total Internet users [8], an upper bound on the percentage of Internet users in this region is 24% of the world population. In short, the large skew of Internet resources remains true even when we restrict the comparison to Internet users and not the total population.

(a) Long-Distance Cable endpoints

(b) Other infrastructure

Figure 4: Distribution of network elements and population as percentage above latitude thresholds. One-hop endpoints are submarine endpoints within a direct connection to points above the threshold.

submarine dataset [15], we use 441 out of 470 cables for which length information is available.

In Figure 5, we observe that cable lengths are an order of magnitude higher in the submarine network (775 km median, 28000 km 99th percentile, and 39000 km maximum). A large fraction of land cables are not vulnerable to GIC since they are shorter than 150 km and hence, do not need repeaters. Due to the relatively large link lengths and presence of repeaters, submarine cables are more vulnerable to failures. They are also more difficult to repair [25].

4.3 Infrastructure Resilience

The impact of a solar event extends well beyond the event based on the extent of damage caused and the time needed for recovery. In this section, we report results on preliminary experiments characterizing the vulnerability of long-haul networks.

4.3.1 Methodology. We evaluate the resilience of long-distance cables using a broad range of repeater failure models. In practical deployments, inter-repeater distance range from 50 to 150km [48, 66].

Figure 5: Cable lengths of submarine cables (across the globe), long-haul fiber on land (US only) and ITU land cables (across the globe).

We evaluate the impact of repeater failure on connectivity at three values of inter-repeater distances: 50 km, 100 km, and 150 km on the US land network, the ITU global land network, and the global submarine network. At a repeater distance of 150 km, 258 out of 542 US land cables, 8443 out of 11,737 ITU global land cables, and 82 out of 441 submarine cables do not need a repeater. At 150km, the average number of repeaters per cable is 0.63, 1.7, and 22.3 in the ITU land network, the US land network, and the submarine network, respectively.

The extent of failure of repeaters will vary based on their location, their type, the strength of the solar storm, etc. Since the actual probability of failure of repeaters is not known, we conduct tests using a broad range of models. We assume repeaters are located at constant intervals and have the same probability of failure on each cable. If at least one repeater fails, we mark the cable as dead. In both land and submarine networks, we assume a node is unreachable when all its connected links have failed.

4.3.2 Uniform repeater failure probability. In the uniform failure probability analysis, all repeaters have the same probability of failure. This helps us identify vulnerabilities across networks at a coarse granularity.

In Figures 6 and 7, we compare the extent of failure of links and nodes in the submarine cable network, the US long-haul network, and the ITU global fiber network at various probabilities of repeater failures. For each value of the probability of failure, we repeat the experiment 10 times for each network and plot the mean and the standard deviation. We observe that even a 1% failure rate for repeaters at an inter-repeater distance of 150 km can lead to the failure of 14.9% submarine cables leaving 11.7% submarine endpoints unreachable. At the same probability, only 1.7% of US long haul fiber cables and 0.6% of ITU fiber cables fail, with a very small fraction of nodes unreachable (0.07% and 0.1% respectively). During catastrophic events with a large probability of repeater failure, at an inter-repeater distance of 150 km, nearly 80% of undersea cables will be affected, leaving an equal fraction of endpoints unreachable, whereas 52% of cables and 17% of nodes in the US land network are affected.

4.3.3 Non-uniform repeater failure probability. In practice, the magnitude and the reach of GIC are dependent on the location of the infrastructure. Higher latitudes are at a greater risk (§ 3.1). The risk is lower but not zero at lower latitudes as well. However, the strength of the induced currents will be much lower at lower latitudes. Hence, we conduct repeater failure analysis using two realistic failure models with latitude-dependent probabilities of failure.

Repeaters in a cable are assigned a probability of failure based on the highest latitude (L) endpoint of the cable. The three levels of failure are demarcated by latitudes 40° and 60° ($L > 60$, $40 < L < 60$, and $L < 40$). We estimate loss of connectivity under two states, S_1 and S_2 , with high ($[1, 0.1, 0.01]$) and low ($[0.1, 0.01, 0.001]$) repeater failure probabilities across three latitude levels.

In Figure 8, we show cable and node failures with non-uniform repeater failure probabilities across latitudes. We observe that link and node failures are an order of magnitude higher in the submarine network under both high and low failure rate states. 43% cables lose connectivity in S_1 with failure probabilities $[1, 0.1, 0.01]$ compared to 80% under uniform failure probability of 1. Even under the low failure scenario (S_2), around 10% of nodes and cables are vulnerable in the submarine network. However, these probabilities are negligible in the US land network. This analysis was not conducted for the ITU land network because the exact latitude and longitude of locations are not available. However, given that the ITU network performs better than the US land network under all scenarios in the uniform failure probability analysis, we can consider the US land network performance as an upper bound for the ITU land network performance.

In short, submarine cables are at a significantly higher risk than fiber cable networks on land.

4.3.4 Connectivity analysis at country-scale. The cable- and node-level analysis presented in the previous section does not give us a comprehensive view of how various countries are affected. To fill this gap in our understanding, next, we analyze the connectivity at the scale of countries across various uniform and non-uniform probabilities of repeater failures. We investigate city-, region- and country-level connectivity. Key observations based on the more realistic non-uniform probability scenarios, S_1 (high failure) and S_2 (low failure), are presented below at the scale of countries.

US: Under low failures (S_2), on the West coast, while most cables connected to Oregon fail, connectivity from California to Hawaii, Japan, Hong Kong, Mexico, Costa Rica, etc. are unaffected. Under similar conditions on the East coast, connectivity between the North East (and Canada) to Europe fail completely with a probability of 0.8. With a probability of 0.2, connectivity of all but one cable is lost. Connections from Florida to Brazil, the Bahamas, etc. are not affected under the low failure scenario. Under high failures (S_1), on the West coast, all long-distance connectivity is lost except for one cable interconnecting Southern California with Indonesia/Hawaii/Micronesia/ Philippines. On the East coast, only long-distance cables from Florida to the Caribbean, Virgin Islands, Columbia, etc. remain operational. US-Europe connectivity is lost with a probability of 1.0.

While Hawaii loses its connectivity to Australia, it remains connected to the continental US and Asia even under high failures (S_1).

Figure 6: The impact of repeater failure on cables under uniform probability of failure of repeaters. A cable fails if at least one of its repeaters fails. Error bars denote standard deviation computed over 10 trials.

Figure 7: The impact of repeater failure on nodes under uniform probability of failure of repeaters. A node is unreachable when all cables connecting to it fail. Error bars denote standard deviation computed over 10 trials.

Figure 8: Cable and nodes failures under two states of non-uniform repeater failures (S_1 , S_2). In each state, repeaters in a cable are assigned a failure probability based on the highest latitude (L) endpoint of the cable. Three levels of failure are $L > 60$, $40 < L < 60$, and $L < 40$. Assigned failure probability per repeater in S_1 is [1, 0.1, 0.01] and in S_2 is [0.1, 0.01, 0.001] across the three levels respectively.

Alaska, on the other hand, loses all its long-distance connectivity except its link to British Columbia in Canada.

China: With low failures (S_2), about 56% of connections are unaffected. However, the densely populated city of Shanghai loses all its long-distance connectivity even under this scenario. This is because all cables connecting to Shanghai are at least 28,000 km and interconnect multiple cities. Under high failures (S_1), China loses all its long-distance cables except one (connecting to Japan, Philippines, Singapore, and Malaysia).

India: The majority of cables connecting to India are unaffected, and none of the cities are disconnected at low failure probabilities (S_2). Even under the high-failure scenario (S_1), some international connectivity remains (e.g., India to Singapore, Middle East, etc.). Unlike in China, the key cities of Mumbai and Chennai do not lose connectivity even with high failures.

Singapore: Even under high failures, several cables connecting to Singapore remain unaffected. Reachable destinations under S_1 include Chennai (India), Perth (Australia), Jakarta (Indonesia), etc.

UK: While the UK loses most of its long-distance cables under the high failure scenario, its connectivity to neighboring European locations such as France, Norway, etc. remains. However, connectivity to North America is lost.

Figure 9: (a) An AS has a reach above a particular latitude if it has at least one router with latitude coordinates above the threshold. ASes above 40° are more vulnerable. (b) The spread of an AS, measured using the location coordinates of routers in the AS, is the difference between the highest and the lowest latitudes of its component routers. Note that 1° latitude spread approximately equals to 111 km.

South Africa: Even under the high-failure scenario (S_1), although it loses some capacity, South Africa continues to retain its connectivity to both the Eastern and the Western Coast of Africa. The cable interconnecting South Africa and Portugal (and other intermediate locations including Nigeria), as well as South Africa and Somalia (and other intermediate locations including Mozambique and Madagascar), are unaffected.

Australia and New Zealand: With high failures (S_1), New Zealand loses all its long-distance connectivity except to Australia. Similarly, while Australia retains most of its connectivity to nearby islands in addition to domestic connections, the longest unaffected cable interconnects it with Jakarta (Indonesia) and Singapore.

Brazil: Interestingly, even under high failures (S_1), Brazil will retain its connectivity to Europe in addition to other parts of South America such as Argentina. However, it will lose its connectivity to North America. It is interesting to note that the US loses its

connectivity to Europe under this failure scenario, but Brazil does not. This is because the Ellalink cable connecting Brazil to Portugal is 6,200 km, while the cable connecting Florida to Portugal is much longer at 9,833 km.

Across several countries, there are short cables interconnecting nearby nodes. For example, cables interconnecting various islands in Hawaii or that connecting two landing stations in Rhode Island. Across both high- and low-latitude locations on all continents, such cables are unaffected even under high repeater failure rates.

4.4 Systems Resilience

In the previous section, we analyzed the impact on the physical infrastructure. The impact on Internet users will, however, depend on systems that run on this physical infrastructure. In this section, we discuss our investigation on Autonomous Systems (ASes), hyperscale data centers, and DNS servers.

4.4.1 Implications for Autonomous System Connectivity.

Ideally, we want to understand the ASes that will be affected when submarine cables fail. However, this will require AS to cable mapping, which is currently unavailable. Fortunately, CAIDA provides a dataset with Internet router locations and router to AS mappings. This dataset contains router to AS mapping for 46,014,869 routers across 61,448 ASes. We use this dataset to understand the impact on ASes. Note that there exist several known errors in the location mapping employed by the CAIDA dataset [34]. Hence, the evaluations in this paper are affected by these errors as well.

First, the impact on an AS depends on its presence in the vulnerable latitude region (above 40°). Hence, we measure the percentage of ASes that have at least one router instance in this region. In Figure 9(a), we observe that 57% of ASes have a presence in the vulnerable regions.

Second, an AS has a higher probability of being affected if it has a large spread. We measure the spread of an AS using the location coordinates of its routers. The evaluation is restricted to latitude spread since the estimation of longitude boundaries is more difficult and error-prone with the given dataset. An AS' latitude spread is measured as the difference between the highest and the lowest latitudes of its component routers. In Figure 9(b), we observe that 50% of ASes have a latitude spread less than 1.723° and 90% of ASes have a latitude spread less than 18.263° (1° latitude spread $\approx$ 111 km). This shows that a vast majority of ASes are geographically localized and do not have a large spread.

A geographically restricted AS is less likely to be directly impacted by solar superstorms. A large spread for an AS may not always mean that it has long links. However, with a large spread, it is likely that an AS will be directly impacted or indirectly impacted within one hop by a severe event.

Overall, a large number of ASes with low spread may help in limiting the impact. However, a significant fraction of ASes have some presence in the more vulnerable regions.

4.4.2 Implications for Hyperscale data centers.

Based on the characteristics of solar superstorms-based failures at the level of nodes, cities, and countries in the previous section, next we discuss the implications for large-scale data centers that serve a large fraction of Internet services and content today. In particular, we

consider public information on data center locations of Google [6] and Facebook [38].

While the majority of Google data centers [6] are located in the US, they are spread across latitudes and longitudes. In the event of high failures, data centers in South Carolina and Georgia as well as Las Vegas are more likely to be located close to active long-distance cables. In Asia, one of Google's data centers is located in Singapore, a location less likely to lose connectivity during solar superstorms. European Google data centers are located in countries with short interconnecting submarine cables to the rest of Europe. While this area is highly susceptible to power failures, Internet failures due to repeater damage are less probable due to the shorter length of cables in this region.

Facebook's data centers [38] are predominantly located in the northern parts of the Northern hemisphere. Facebook does not operate any hyperscale data centers in Africa or South America, unlike Google. Owing to the limited geographic spread of data centers, Facebook will have less resilience in the event of solar superstorms.

4.4.3 Implications for DNS servers. DNS root servers are highly geographically distributed. Although the distribution is not proportional to Internet users (Africa with more Internet users than North America has nearly half the number of DNS servers), DNS root servers are widely present on all continents. Hence, DNS root servers are resilient in the face of solar superstorms.

While location information of all DNS root servers is publicly available (including all anycast server locations), data on more than 1500 Top-Level Domain servers and other DNS root zone servers are limited. While the IP address of these servers can be obtained from DNS records, it is difficult to identify their locations, particularly because they typically employ anycast.

4.4.4 Summary.

- The distribution of Internet infrastructure is skewed when compared to the distribution of Internet users. The infrastructure concentration in higher latitudes poses a greater risk during solar storms.
- The investigation on the impact of GIC shows that submarine cables have a higher risk of failure compared to land cables. This is primarily due to their large lengths.
- The US is one of the most vulnerable locations with a high risk of disconnection from Europe during extreme solar events. Intra-continental connections in Europe are at a lower risk due to the presence of a large number of shorter land and submarine cables interconnecting the continent.
- In Asia, Singapore will retain good connectivity to neighboring countries even under severe storms. Chinese cities are more prone to lose connectivity than Indian cities because they connect to much longer cables. The cables running along the eastern and western coasts of Africa are less prone to failures.
- The cable between Brazil and Europe has less probability of being affected compared to cables connecting the US and Europe.
- Australia, New Zealand, and other island countries in the region will lose most of their long-distance connections. But local connectivity, as well as connections to Singapore, are less vulnerable.

- A large fraction of ASes have a presence in vulnerable regions, however, the vast majority have a small spread. The extent of impact on an AS and its customers will depend on a combination of these factors.
- Google data centers have a better spread, particularly in Asia and South America. Facebook is more vulnerable.
- DNS root servers are highly distributed and hence not vulnerable.

5 PLANNING FOR THE FUTURE

Although we have sentinel spacecraft that can issue early warnings of CMEs providing at least 13 hours of lead time, our defenses against GIC are limited. Hence, we need to prepare the infrastructure for an eventual catastrophe to facilitate efficient disaster management. Towards this goal, we outline several directions.

5.1 Internet Infrastructure Design

As shown in our analysis, the current Internet infrastructure is heavily concentrated in higher latitudes that are at a greater risk for GIC. We need to factor in this threat during infrastructure expansion.

With the increased melting of Arctic ice, there are ongoing efforts to lay cables through the Arctic [11, 12]. While this is helpful for improving latency, these cables are prone to higher risk. During topology design, we need to increase capacity in lower latitudes for improved resiliency during solar storms (although latency is higher). Moreover, since links from the US and Canada to Europe and Asia are highly vulnerable, adding more links to Central and South America can help in maintaining global connectivity. South America is more likely to maintain connectivity to Europe and Africa.

At submarine cable landing points, particularly in the low latitudes, it is important to have mechanisms for electrically isolating cables connecting to higher latitudes from the rest, to prevent cascading failures.

A higher density of data centers and IXPs in northern parts of Europe and America also poses a threat to Internet services. Data center and application service providers should be cognizant of solar threats during new deployments. We need to develop standardized tests for measuring end-to-end resiliency of applications under such extreme events. Specifically, systematic modeling of potential disruptions to the Internet, from the physical layer to various applications, through collaborations between astrophysicists, electrical engineers, and networking researchers is critical for improving Internet resiliency.

This paper focused on terrestrial infrastructure only. With increasing deployments of low earth orbit satellites [14], it is also important to study the impact of solar superstorms on satellite Internet constellations that are directly exposed to powerful CMEs.

5.2 How to Use the Lead Time?

A CME that originates in the sun reaches the earth at least 13 hours, typically 1-3 days, later. This provides the infrastructure operators an opportunity to devise a shutdown strategy that minimizes connectivity loss during and after impact. In power grids, since GIC is superimposed with generated current, a key strategy involves reducing or completely shutting down power generation.

In the Internet infrastructure, the shutdown strategy needs to focus on two aspects: (i) protecting the equipment during a solar event, and (ii) ensuring the continuation of services after the event anticipating partial damage to infrastructure (damaged submarine cables, satellites, etc.). Similar to power grids, powering off is the easiest solution for equipment damage prevention. However, note that this only provides limited protection since GIC can flow through a powered-off cable. Since the peak current is reduced slightly by powering off, this can help only when the threat is moderate.

Planning for post-impact connectivity is a much harder problem, especially with limited modeling available for the extent of cable damage. Search engines, financial services, etc. should geodistribute critical data and functionalities so that each partition (potentially disconnected landmasses such as N. America, Eurasia, Australia, etc.) can function independently. Also, service and content providers should pre-provision high priority service for critical applications such as 911, hospitals, fire departments, etc.

5.3 Piecing Together a Partitioned Internet

We need to rethink the network environment in the event of a partial or complete disconnection [57]. This includes designing ad-hoc network connectivity mechanisms (e.g., Project Loon [36]) and peer-to-peer applications that can bootstrap connectivity locally. User-powered mesh networks [7] that proved valuable during other natural disasters such as earthquakes can also help during solar superstorms. However, unlike other localized disasters, wide-area connectivity disruption is a unique challenge associated with solar storms. To tackle this problem, we need to examine alternatives such as backup interdomain protocols that allow multiple paths and more resilient Internet architectures (e.g., SCION [18]). More broadly, designing and installing in advance a seamless protocol that can piece together all available modes of communication, including cables, satellites, and wireless, across multiple administrative domains is critical for fast recovery.

5.4 Devising New Resilience Tests for Internet Systems

Current best practices on fault tolerance and resilience evaluation in software systems revolve around failure models that consider a limited number of failures within and across locations. Large-scale infrastructure failures spanning broad swaths of the Internet are absent in the literature. This was primarily due to the fortuitous absence of such a catastrophic event in the past two decades when the Internet infrastructure grew rapidly. Hence, our understanding of the impact of a solar superstorm on Internet sub-systems (e.g., Autonomous Systems) and Internet-based systems (e.g., cloud services, Voice over IP) is very limited. We need to devise standard practices in resilience testing involving large-scale failures.

5.5 Interdependence with Power Grids

The economic impact of Internet disruption for a day in the US is estimated to be \$7 billion [1], while the same due to electric grid failure is estimated to be more than \$40 billion [60]. The power grid is the most critical infrastructure which almost all other systems rely on. The impact of solar superstorms on power grids has been an active area of research over several decades [20, 40, 41, 43, 55,

74]. Today, there is a tight interdependence between power grids and networks. Smart grids rely on either their own private WAN networks or the public Internet for their functioning. As a result, failures of power grids and the Internet and other communication networks are more tightly coupled.

Although power grids and the Internet have varying levels of impact on society, there are several differentiating features in the nature of failures and hence, the recovery mechanisms required. First, power grids and the Internet differ in their structure and organization. For example, in the US, there are three regional power grids. If the power grid in New York fails, it will not cause any significant effects in California. The Northeast will be without power, but it is not possible to transfer electricity from California to New York and cause any power overload in California. On the other hand, when all submarine cables connecting to NY fail, there will be significant shifts in BGP paths and potential overload in Internet cables in California. In short, the Internet is more global compared to power grids, and even regional failures can result in significant consequences for the broader Internet.

Second, the length of power cables will closely follow the ITU land dataset. Both power lines and Internet cables on land are typically in close proximity to road and rail networks. Hence, longer submarine cables may be susceptible to higher risks. The key roadblock associated with replacing transformers in power grids is the manufacturing time for new equipment. For submarine cables, in addition to equipment availability, access to failure location for repair is also a challenge.

Modeling of the interdependence of power grids and the Internet, potential cascading failures between the two systems due to their coupling, and design of recovery mechanisms for the mutually dependent systems are problems that need to be tackled for increasing the robustness of our critical infrastructure.

6 CONCLUSION

In this paper, we show that a powerful solar superstorm has the potential to cause massive disruption of the Internet. Astrophysicists estimate the likelihood of a solar storm of sufficient strength to cause catastrophic disruption occurring within the next decade to be 1.6 – 12%. Paying attention to this threat and planning defenses against it, like our preliminary effort in this paper, is critical for the long-term resilience of the Internet. Several challenges remain open in this space. How can we model infrastructure failures more accurately? How do we factor in solar threat during Internet infrastructure and systems design? How can we help operators in making disaster preparation and recovery plans? We anticipate that this paper will provide an initial impetus towards answering these important questions.

ACKNOWLEDGMENTS

I sincerely thank the anonymous reviewers, the shadow shepherd, and the shepherd, Adrian Perrig, for their helpful comments and feedback. I would also like to thank Kishor Kumar Kalathiparambil for verifying the correctness of astrophysics discussions in this paper, and Aditya Akella and Sujata Banerjee for their feedback on an early version of the paper.

REFERENCES

- [1] [n.d.]. Cost of Shutdown Tool. <https://netblocks.org/cost/>. (Accessed on 06/07/2021).
- [2] [n.d.]. Data Centers Map. <https://www.datacentermap.com/>.
- [3] [n.d.]. DNS root servers. <https://root-servers.org/>.
- [4] [n.d.]. Geomagnetic Effects on Communication Cables. <https://www.spaceweather.gc.ca/tech/se-cab-en.php/>.
- [5] [n.d.]. Geomagnetic Storms – Reducing the Threat to Critical Infrastructure in Canada. <http://www.solarstorms.org/CanadaPipelines.html>.
- [6] [n.d.]. Google Data Center Locations. <https://www.google.com/about/datacenters/inside/locations/>.
- [7] [n.d.]. GoTenna Mesh Networks. <https://gotennamesh.com/products/mesh>.
- [8] [n.d.]. Internet Statistics. <https://www.internetworldstats.com/stats.htm/>.
- [9] [n.d.]. Introducing Equiano, a subsea cable from Portugal to South Africa. <https://cloud.google.com/blog/products/infrastructure/introducing-equiano-a-subsea-cable-from-portugal-to-south-africa>.
- [10] [n.d.]. ITU interactive transmission map. <https://www.itu.int/itu-d/tnd-map-public/>.
- [11] [n.d.]. Major step towards a Europe-Asia Arctic cable link. <https://thebarentsobsewer.com/en/industry-and-energy/2019/06/mou-signed-set-arctic-telecom-cable-company>.
- [12] [n.d.]. Melting Arctic Ice Opens a New Fiber Optic Cable Route. <https://spectrum.ieee.org/tech-talk/telecom/internet/melting-sea-ice-opens-the-floodgate-for-a-new-fiber-optic-cable-route/>.
- [13] [n.d.]. PCH Internet Exchange Directory. <https://www.pch.net/ixp/dir>.
- [14] [n.d.]. Starlink. <https://www.starlink.com/>.
- [15] [n.d.]. Telegeography's Submarine Cable Map. <https://github.com/telegeography/www.submarinecablemap.com>.
- [16] L Barnard and Mike Lockwood. 2011. A survey of gradual solar energetic particle events. *Journal of Geophysical Research: Space Physics* 116, A5 (2011).
- [17] L Barnard, M Lockwood, MA Hapgood, Matt J Owens, Chris J Davis, and F Steinhilber. 2011. Predicting space climate change. *Geophysical research letters* 38, 16 (2011).
- [18] David Barrera, Laurent Chuat, Adrian Perrig, Raphael M. Reischuk, and Pawel Szalachowski. 2017. The SCION Internet Architecture. *Commun. ACM* 60, 6 (June 2017), 56–65. <https://doi.org/10.1145/3085591>
- [19] Prantika Bhowmik and Dibyendu Nandy. 2018. Prediction of the strength and timing of sunspot cycle 25 reveal decadal-scale space environmental conditions. *Nature communications* 9, 1 (2018), 1–10.
- [20] DH Boteler, RM Shier, T Watanabe, and RE Horita. 1989. Effects of geomagnetically induced currents in the BC Hydro 500 kV system. *IEEE Transactions on Power Delivery* 4, 1 (1989), 818–823.
- [21] Richard C Carrington. 1859. Description of a singular appearance seen in the Sun on September 1, 1859. *Monthly Notices of the Royal Astronomical Society* 20 (1859), 13–15.
- [22] BA Carter, E Yizengaw, R Pradipta, JM Weygand, M Piersanti, A Pulkkinen, MB Moldwin, R Norman, and K Zhang. 2016. Geomagnetically induced currents around the world during the 17 March 2015 storm. *Journal of Geophysical Research: Space Physics* 121, 10 (2016), 10–496.
- [23] SC Chapman, SW McIntosh, RJ Leamon, and NW Watkins. 2020. Quantifying the solar cycle modulation of extreme space weather. *Geophysical Research Letters* (2020), e2020GL087795.
- [24] Valerie Coffey. 2014. Sea change: The challenges facing submarine optical communications. *Optics and Photonics News* 25, 3 (2014), 26–33.
- [25] Valerie Coffey. 2014. Sea change: The challenges facing submarine optical communications. *Optics and Photonics News* 25, 3 (2014), 26–33.
- [26] Steven Constable. 2007. *Conductivity, Ocean Floor Measurements*. Springer Netherlands, Dordrecht, 71–73. https://doi.org/10.1007/978-1-4020-4423-6_30
- [27] T Divett, M Ingham, CD Beggan, GS Richardson, CJ Rodger, AWP Thomson, and M Dalzell. 2017. Modeling geoelectric fields and geomagnetically induced currents around New Zealand to explore GIC in the South Island's electrical transmission network. *Space Weather* 15, 10 (2017), 1396–1412.
- [28] ZL Du. 2020. The solar cycle: predicting the peak of solar cycle 25. *Astrophysics and Space Science* 365, 6 (2020), 1–5.
- [29] Ramakrishnan Durairajan, Paul Barford, Joel Sommers, and Walter Willinger. 2015. InterTubes: A study of the US long-haul fiber-optic infrastructure. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*. 565–578.
- [30] J Feynman and A Ruzmaikin. 2011. The Sun's strange behavior: Maunder minimum or Gleissberg cycle? *Solar physics* 272, 2 (2011), 351.
- [31] J Feynman and A Ruzmaikin. 2014. The Centennial Gleissberg Cycle and its association with extended minima. *Journal of Geophysical Research: Space Physics* 119, 8 (2014), 6027–6041.
- [32] Jennifer L. Gannon. [n.d.]. Geomagnetic Storms and Geomagnetically Induced Currents. <https://electricenergyonline.com/energy/magazine/966/article/Geomagnetic-Storms-and-Geomagnetically-Induced-Currents.htm>.
- [33] Adriana Garcia and Zadi Mouradian. 1998. The Gleissberg cycle of minima. *Solar Physics* 180, 1-2 (1998), 495–498.
- [34] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Han Zhang, Roya Ensafi, and Christos Papadopoulos. 2017. A look at router geolocation in public and commercial databases. In *Proceedings of the 2017 Internet Measurement Conference*. 463–469.
- [35] Max Gleber. [n.d.]. CME Week: The Difference Between Flares and CMEs. <https://www.nasa.gov/content/goddard/the-difference-between-flares-and-cmes>.
- [36] Google. [n.d.]. Project Loon. <https://loon.com>.
- [37] YB Han and ZQ Yin. 2019. A decline phase modeling for the prediction of solar cycle 25. *Solar Physics* 294, 8 (2019), 1–14.
- [38] Kim Hazelwood, Sarah Bird, David Brooks, Soumith Chintala, Utku Diril, Dmytro Dzhulgakov, Mohamed Fawzy, Bill Jia, Yangqing Jia, Aditya Kalro, et al. 2018. Applied machine learning at Facebook: A datacenter infrastructure perspective. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 620–629.
- [39] RB Horne, SA Glauert, NP Meredith, D Boscher, V Maget, D Heynderickx, and D Pitchford. 2013. Space weather impacts on satellites and forecasting the Earth's electron radiation belts with SPACECAST. *Space Weather* 11, 4 (2013), 169–186.
- [40] John Kappenman. 2010. *Geomagnetic storms and their impacts on the US power grid*. Citeser.
- [41] John G Kappenman. 2001. An introduction to power grid impacts and vulnerabilities from space weather. In *Space Storms and Space Weather Hazards*. Springer, 335–361.
- [42] Kathrin Kirchen, William Harbert, Jay Apt, and M Granger Morgan. 2020. A Solar-Centric Approach to Improving Estimates of Exposure Processes for Coronal Mass Ejections. *Risk Analysis* (2020).
- [43] Harold Kirkham, Yuri V Makarov, Jeffery E Dagle, John G DeSteele, Marcelo A Elizondo, and Ruisheng Diao. 2011. *Geomagnetic storms and long-term impacts on power systems*. Technical Report. Pacific Northwest National Lab.(PNNL), Richland, WA (United States).
- [44] Anukool Lakhina, John W Byers, Mark Crovella, and Ibrahim Matta. 2003. On the geographic location of Internet resources. *IEEE Journal on Selected Areas in Communications* 21, 6 (2003), 934–948.
- [45] Louis J Lanzerotti, Lester V Medford, Carol G MacLennan, and David J Thomson. 1995. Studies of large-scale Earth potentials across oceanic distances. *AT&T technical journal* 74, 3 (1995), 73–84.
- [46] FY Li, DF Kong, JL Xie, NB Xiang, and JC Xu. 2018. Solar cycle characteristics and their application in the prediction of cycle 25. *Journal of Atmospheric and Solar-Terrestrial Physics* 181 (2018), 110–115.
- [47] Jeffrey J Love, Hisashi Hayakawa, and Edward W Cliver. 2019. Intensity and impact of the New York Railroad superstorm of May 1921. *Space Weather* 17, 8 (2019), 1281–1292.
- [48] Adam Markow. [n.d.]. *Summary of Undersea Fiber Optic Network Technology and Systems*. http://hmorell.com/sub_cable/documents/Basics%20of%20Submarine%20System%20Installation%20and%20Operation.pdf.
- [49] James A Marusek. [n.d.]. *Solar storm threat analysis*. Impact, 2007.
- [50] Trevor Maynard, Neil Smith, and Sandra Gonzalez. 2013. Solar storm risk to the north american electric grid. *Lloyd's* 1 (2013), 11.
- [51] KG McCracken, GAM Dreschhoff, DF Smart, and MA Shea. 2004. A study of the frequency of occurrence of large-fluence solar proton events and the strength of the interplanetary magnetic field. *Solar Physics* 224, 1-2 (2004), 359–372.
- [52] KG McCracken, GAM Dreschhoff, EJ Zeller, DF Smart, and MA Shea. 2001. Solar cosmic ray events for the period 1561–1994: 1. Identification in polar ice, 1561–1950. *Journal of Geophysical Research: Space Physics* 106, A10 (2001), 21585–21598.
- [53] Scott W McIntosh, Sandra Chapman, Robert J Leamon, Ricky Egeland, and Nicholas W Watkins. 2020. Overlapping magnetic activity cycles and the sunspot number: forecasting sunspot cycle 25 amplitude. *Solar Physics* 295, 12 (2020), 1–14.
- [54] Dan McMurrow. 2011. *Impacts of Severe Space Weather on the Electric Grid*. <https://fas.org/irp/agency/dod/jason/spaceweather.pdf>.
- [55] AG McNish. 1940. The magnetic storm of March 24, 1940. *Terrestrial Magnetism and Atmospheric Electricity* 45, 3 (1940), 359–364.
- [56] LV Medford, LJ Lanzerotti, JS Kraus, and CG MacLennan. 1989. Transatlantic earth potential variations during the March 1989 magnetic storms. *Geophysical Research Letters* 16, 10 (1989), 1145–1148.
- [57] David Mendonça, Theresa Jefferson, and John Harrauld. 2007. Collaborative ad-hocacies and mix-and-match technologies in emergency management. *Commun. ACM* 50, 3 (2007), 44–49.
- [58] Satoko Nakamura, Yusuke Ebihara, Shigeru Fujita, Tada-nori Goto, N Yamada, S Watari, and Y Omura. 2018. Time domain simulation of geomagnetically induced current (GIC) flowing in 500-kV power grid in Japan including a three-dimensional ground inhomogeneity. *Space Weather* 16, 12 (2018), 1946–1959.
- [59] Sten Odenwald. [n.d.]. *The Day the Sun Brought Darkness*. https://www.nasa.gov/topics/earth/features/sun_darkness.html.
- [60] Edward J Oughton, Andrew Skelton, Richard B Horne, Alan WP Thomson, and Charles T Gaunt. 2017. Quantifying the daily economic impact of extreme space weather due to failure in electricity transmission infrastructure. *Space Weather*

- 15, 1 (2017), 65–83.
- [61] Alexei N Peristykh and Paul E Damon. 2003. Persistence of the Gleissberg 88-year solar cycle over the last 12,000 years: Evidence from cosmogenic isotopes. *Journal of Geophysical Research: Space Physics* 108, A1 (2003), SSH–1.
 - [62] Tony Phillips. [n.d.]. *Near Miss: The Solar Superstorm of July 2012*. https://science.nasa.gov/science-news/science-at-nasa/2014/23jul_superstorm/.
 - [63] Antti Pulkkinen, Emanuel Bernabeu, Jan Eichner, Ciaran Beggan, and AWP Thomson. 2012. Generation of 100-year geomagnetically induced current scenarios. *Space Weather* 10, 4 (2012).
 - [64] Antti Pulkkinen, E Bernabeu, A Thomson, A Viljanen, R Pirjola, D Boteler, J Eichner, PJ Cilliers, D Welling, NP Savani, et al. 2017. Geomagnetically induced currents: Science, engineering, and applications readiness. *Space Weather* 15, 7 (2017), 828–856.
 - [65] Pete Riley. 2012. On the probability of occurrence of extreme space weather events. *Space Weather* 10, 2 (2012), 1–12.
 - [66] Patrice Le Roux, Mélanie Jaouen, and Ghislaine Vareille. [n.d.]. *A Long-Span Repeater for Regional Submarine Systems*. <https://www.suboptic.org/wp-content/uploads/2014/10/LeRoux-WeB1.pdf>.
 - [67] V Sarp, A Kilcik, Vasyi Yurchyshyn, JP Rozelot, and A Ozguc. 2018. Prediction of solar cycle 25: a non-linear approach. *Monthly Notices of the Royal Astronomical Society* 481, 3 (2018), 2981–2985.
 - [68] Roberta Tozzi, Igino Coco, Paola De Michelis, and Fabio Giannattasio. 2019. Latitudinal dependence of geomagnetically induced currents during geomagnetic storms. *Annals of Geophysics* (2019).
 - [69] Roberta Tozzi, Paola De Michelis, Igino Coco, and Fabio Giannattasio. 2019. A preliminary risk assessment of geomagnetically induced currents over the Italian territory. *Space Weather* 17, 1 (2019), 46–58.
 - [70] University Center for International Earth Science Information Network - CIESIN - Columbia. 2020. Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11. NASA Socioeconomic Data and Applications Center (SEDAC). <https://doi.org/10.7927/H4JW8BX5>
 - [71] Lisa A Upton and David H Hathaway. 2018. An updated Solar Cycle 25 prediction with AFT: the modern minimum. *Geophysical Research Letters* 45, 16 (2018), 8091–8095.
 - [72] A Viljanen, A Pulkkinen, R Pirjola, K Pajunpää, P Posio, and A Koistinen. 2006. Recordings of geomagnetically induced currents and a nowcasting service of the Finnish natural gas pipeline system. *Space Weather* 4, 10 (2006).
 - [73] VMasuo Suyama VMasato Nagayama VHaruki Watanabe and VHaruo Fujiwara VColin Anderson. 1999. WDM optical submarine network systems. *FUJITSU Sci. Tech. J* 35, 1 (1999), 34–45.
 - [74] Magnus Wik, A Viljanen, R Pirjola, A Pulkkinen, P Wintoft, and Henrik Lundstedt. 2008. Calculation of geomagnetically induced currents in the 400 kV power grid in southern Sweden. *Space weather* 6, 7 (2008).
 - [75] Yosuke Yamazaki and Michael J Kosch. 2015. The equatorial electrojet during geomagnetic storms and substorms. *Journal of Geophysical Research: Space Physics* 120, 3 (2015), 2276–2287.
 - [76] Kenichi Yoneyama, Hiroshi Sakuyama, and Akira Hagisawa. 2010. Construction technology for use in repeated transoceanic optical submarine cable systems. *NEC Technical Journal* 5, 1 (2010).