A Holistic Framework for Bridging Physical Threats to User QoE

USC/ISI Technical Report ISI-TR-687, December 2013[*]

Xue Cai
USC/ISI, Marina del Rey, CA
xuecai@isi.edu

John Heidemann
USC/ISI, Marina del Rey, CA
johnh@isi.edu

Walter Willinger
Niksun, Inc., Princeton, NJ
wwillinger@niksun.com

ABSTRACT

Submarine cable cuts have become increasingly common, with five incidents breaking more than ten cables in the last three years. Today, around 300 cables carry the majority of international Internet traffic, so a single cable cut can affect millions of users, and repairs to any cut are expensive and time consuming. Prior work has either measured the impact following incidents, or predicted the results of network changes to relatively abstract Internet topological models. In this paper, we develop a new approach to model cable cuts. Our approach differs by following *problems* drawn from real-world occurrences all the way to their impact on *end-users*. Because our approach spans many layers, no single organization can provide all the data needed to apply the model. We therefore perform *what-if* analysis to study a range of possibilities. With this approach we evaluate four incidents in 2012 and 2013; our analysis suggests general rules that assess the degree of a country's vulnerability to a cut.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Modeling techniques; C.2.1 [**Network Architecture and Design**]: Network topology; C.2.5 [**Local and Wide-Area Networks**]: Internet

General Terms

Measurement

Keywords

threat, model, submarine cable cut, user, quality-of-experience (QoE), web service, streaming video

1. INTRODUCTION

The Internet is of great importance today, and as critical infrastructure, the impact of various *threats* it faces needs to be carefully studied and understood. Threat models are built to provide understanding about how specific threats change the Internet, how the network reacts to such threats, and how the various threats affect end users.

In this paper, we focus on *submarine cable cuts*, a specific class of threats that impact the physical infrastructure of the Internet. Understanding the impact of submarine cable cuts is essential for three reasons. First, judging form the many reported real-world incidents [5–7, 11, 19, 27], they occur rather frequently and can have considerable impact. Second, the *majority* of international traffic travels over fewer than 300 submarine cables around the globe (Figure 1 shows cables in 2013). A single cut can profoundly affect millions of users and businesses. Third, recovery from a cut can be slow, with typical repair times as long as several weeks.

Figure 2 illustrates the challenge of submarine cable cuts. In this example, four landing stations are connected via a submarine cable system (SCS 1). Each landing station connects to users via some terrestrial networks. Various online services are replicated in facilities connected to landing stations, with differing deployments depending on user distribution and cost. A cable cut has broken the cable segment between stations 1 and 2.

One may simply treat this problem as a path-finding problem, focusing on graph properties. Since after the cut the graph is still connected, this naive model implies there are minimal results from the cut. Since all stations are still reachable, the only harm is increased path length between stations 1 and 2.

However, this naive model does not reflect important aspects of Internet operation: *users* and *services*, as

[*] Partial support for this research is sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, HSARPA, Cyber Security Division, BAA 11-01-RIKA and Air Force Research Laboratory, Information Directorate under agreement number FA8750-12-2-0344 ("LACREND", John Heidemann and Xue Cai), and via SPAWAR Systems Center Pacific under Contract No. N66001-13-C-3001 ("Retro-future", John Heidemann). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views contained herein are those of the authors and do not necessarily represent those of DHS or the U.S. Government.

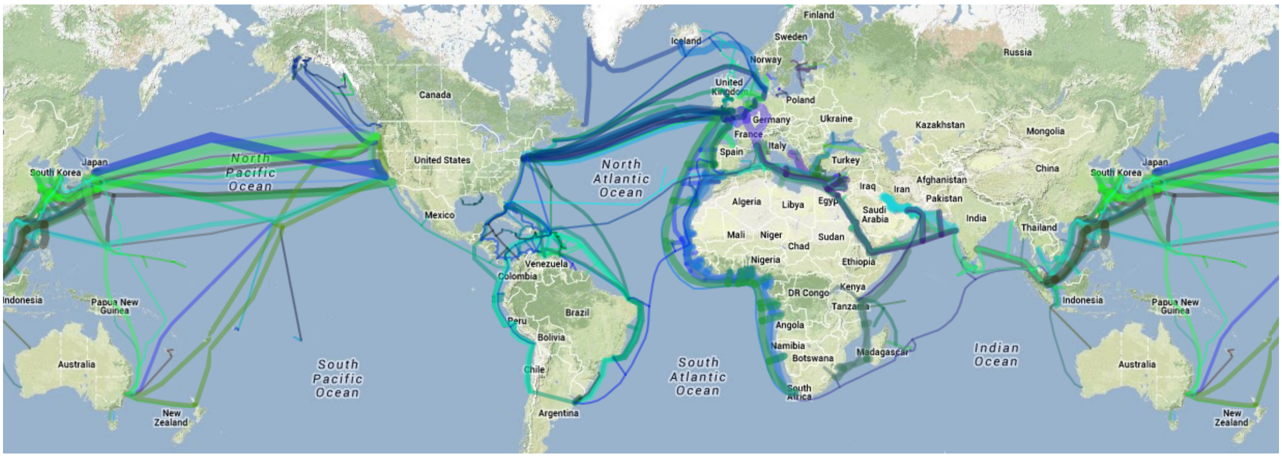

Figure 1: Global Submarine Cable Map in 2013 [29]

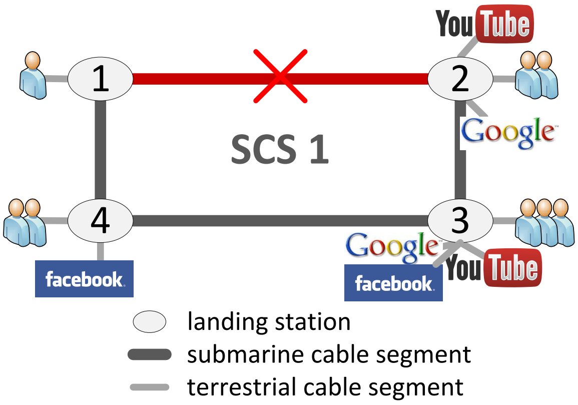

Figure 2: The problem to solve.

well as the *diverse mechanisms* each Internet layer uses to establish, protect, and restore data channels. Users interact with application-layer services, while the cut happens at the physical layer. To understand the impact of a cut on users we must consider and model the cascade of interactions from the physical all the way to the application layer.

Capturing this cascade is challenging, since each layer has diverse mechanisms that support communications. A user's access to a service depends not only on connectivity of a physical medium, but also on virtual data channels that must be provisioned at intermediate layers. The different mechanisms at each layer require separate models to capture their unique functionalities for fault-tolerance and recovery.

Rich Internet connectivity means that while completely disconnected users may be rare, unacceptable performance is a more common user experience. We must therefore also evaluate *user-perceived qualities* as measured by Quality-of-Experience (QoE). Network changes affect QoE in service-specific ways. For example, a user browsing web pages cares about the page-loading time, whereas a user watching videos concerns factors such as how long it takes to start the video, how often the video re-buffers, and what the video quality is.

Prior threat models are often influenced primarily by *data availability*. For example, the naive model we mentioned earlier follows directly from knowledge of the topology, but omits layer interactions and end-user impacts. Data-driven approaches can misplace risk by emphasizing threats that are unlikely and defining harms that are abstract from real-world users. For example, wide availability of data about AS topologies encourages threat models involving node and link removals in AS graphs. Since AS graphs represent business relationships, this graph manipulation has at best limited relationship to real-world events [3, 15].

The first contribution of our paper is to *frame the modeling need as spanning real-world threats at lower layers to end-user harm* (§ 3). To address this challenge we bring together a number of existing models of Internet components and show how they can fit together to identify the *essential mechanisms* at intermediate layers that change threat outcomes on end-users.

Even with carefully selected models, no single organization is likely to have all the data needed to populate models from the physical layer to users. Our second contribution is to show how *to apply what-if modeling to network threats* (§ 4.4). The ability to explore a range of possibilities allows one to make qualitative claims about possible outcomes in the face of incomplete data. As one example, we use Quality-of-Experience models (§ 3.6) to study a range of possible current and future outcomes to users that might result from a submarine cable cut.

Our third contribution is to illustrate our approach by *exploring four real-world incidents* of submarine cable cuts in 2012 and 2013 (§ 4). Using our models and what-if analysis, we provide general rules that help assess what makes some countries more vulnerable to disruption (§5): service-self-sufficiency and diversified connectivity. Our models allow countries to evaluate their vulnerabilities to these risks and explore possible mitigating strategies.

Finally, although we focus on submarine cable cuts, many parts of our model also apply to other disruption threats.

2. RELATED WORK

Four areas relate to our work: other models that either predict the impact of submarine cable cuts or post-facto evaluation after cuts, models of other threats to Internet infrastructure, and non-threat models of parts of the Internet.

Models of submarine cable cuts Omer et al. provide a model to assess the impact of submarine cable cuts [25]. They construct a physical cable topology in which nodes are continents and edges are aggregations of inter-connecting submarine cables based on the public map [29]. They then hypothesize threats by removing nodes or edges and assessed impact by computing the amount of traffic could be delivered between continents after the threat. Unlike their work, we analyze real-world cuts, and we relate the impact to end-users. In addition, we pay attention to the diverse data-transmission mechanisms on layers which are not present in their work.

Measurements of submarine cable cuts Many researchers have measured the consequences of submarine cable cuts [11,26,28]. In contrast, our model can provide implications before a cut happens. Nevertheless, these measurements are valuable as ground truth to validate and correct our model.

Models of other threats Because of the availability of data of AS topologies, past threat models typically build on the AS graph, modeling threats as removals of nodes or edges from it. Albert et al. [2] first analyzed errors (accidental removal of nodes) and attacks (intentional removal of nodes), assessing impact as network-diameter increase and fragmentation. Dolev et al. [13] builds on this model, but with the consideration of the network-layer transmission mechanism. They note that connectivity between ASes does not imply reachability—a valid AS path must be *valley-free* [18]. Wu et al. [43] further enriches the model and assessed impact as the reachability changes between all AS pairs. Different from their work, we start with threats drawn from real-world incidents and assess impact not on the network layer but on end-users.

Models of the Internet Much prior work model how different parts of the Internet work without explicitly considering threats. These models provide useful input to our work. In particular, Feamster et al. [16] model how BGP selects paths for traffic flows. Mok et al. [32] model how flow condition affects video streaming qualities, while Zhang et al. [44] model video telephony. Researchers in [10, 12, 24, 32] model how service qualities affect user QoE. We incorporate some of the models above to build our multi-layer threat model.

3. MODELING CABLE CUTS

To understand the impact of cable cuts on the real world, we follow the problem from real-world threats to user-relevant harms, bridging them with our holistic model.

3.1 Model Overview

Our approach to model cable cuts is *problem-driven*, which contrasts with models that are built around the constraints of available data. We identify the threat and harms, then identify what role each takes in the Internet and determine how they relate. This approach is challenging, because the relationship of network components is not always obvious, and because components are often “black boxes” where obtaining data can be difficult. This section presents an overview about how we follow this approach to model cable cut impact.

Step 1: Selecting threat and harm We model submarine cable cuts because of their frequent occurrence, traffic importance, and long repair time. We choose to model harms as degraded QoE because it is what users care about.

Cable cuts happen to cable segments at the physical layer, while users access services at the application layer. Thus, to assess how cable cuts affect we must bridge these layers modeling intermediate layers. Before we discuss this process, we first provide background to bring out the basic idea.

Background of the Internet Logically, the Internet is structured as layers. Between two adjacent layers, *the lower layer provides a communication channel for the upper layer* and thus directly affects its communication quality.

The QoE of users is directly shaped by the quality of the communication on the *application layer*.

The communication further relies on lower *transport layer* to transmit its traffic and *network layer* to find a path for the traffic. The path then relies on lower *link layer* to establish channels that support the links composing the path.

Eventually, the link layer needs a physical medium (such as cable segment) to support its virtual channels. As a result, *any damage made to a physical medium will cascade up the stack*.

Step 2: Bridging threat to harm Our approach is to tie the threat to harms by *successively modeling how changes of lower-layer communication channels affect*

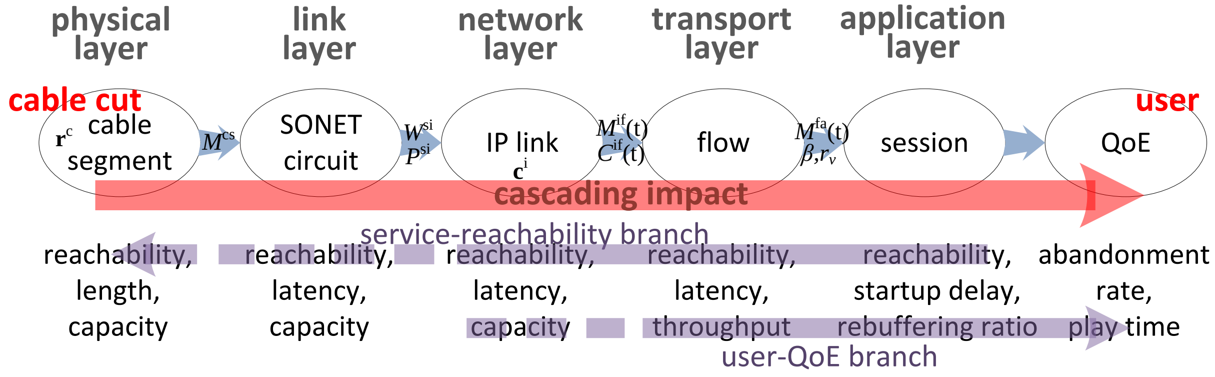
physical layer
link layer
network layer
transport layer
application layer
cable cut
r^c cable segment
M^cs
SONET circuit
W^si P^si
IP link c^i
M^if(t) C^if(t)
flow
M^fa(t) β,r_v
session
user
QoE
cascading impact
service-reachability branch
reachability, length, capacity
reachability, latency, capacity
reachability, latency, capacity
reachability, latency, throughput
reachability, startup delay, rebuffering ratio
abandonment rate, play time
user-QoE branch

Figure 3: The general picture of the model.

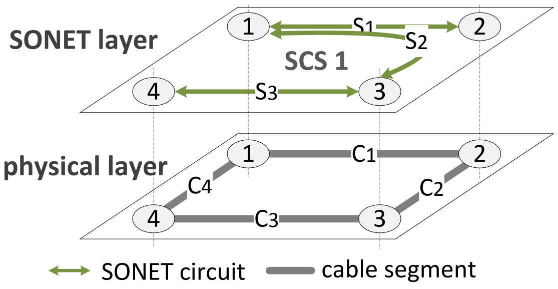

Figure 4: SONET circuits rely on cable segments as physical medium, but have to be provisioned to transmit data.

upper-layer communication quality, from threat layer to harm layer.

There are three benefits to this approach. First, with mostly independent models at each layer, each component can be treated more or less in isolation, making each layer simpler and easier to interpret. Second, we can explore and validate components separately to increase our overall confidence in the approach. Third, the approach provides a framework to identify the different components' relationships so as to capture a more holistic view of the problem.

Figure 3 shows the general picture of our model. The long solid red arrow represents the cascading impact from cable cuts to users. To model this impact, we break the model into five sub-models (five short arrows) that each addresses a direct impact.

To summarize, the problems that each sub-model is going to address in later sections are:

1. how does a cable cut break SONET circuits (Section 3.2)?

2. how does the breakage of SONET circuits break or impair upper-layer IP links (Section 3.3)?

3. how does the change of IP link condition affects flows that traverse through the link (Section 3.4)?

4. how does the application system adjust its session qualities to adapt to the new flow condition (Section 3.5)?

5. how does the change of session qualities affect user perceived QoE (Section 3.6)?

Modeling across multiple layers is a challenging task. To keep model manageable, we avoid details that can be captured in existing layers, such as WDM whose ring protection mechanism is captured adequately in our SONET model. We also do not model transient effects brought by mechanisms such as fast re-routing in MPLS, but instead focus on impact that lasts for at least days.

3.2 From Cable Cut to SONET Circuits

We first tie physical damage to the SONET link layer.

A *SONET circuit* is a virtual circuit between two SONET devices. Logically, a SONET circuit corresponds to physical and link layers in the OSI model [22]. Most submarine cable systems use SONET, connecting devices that are physically located at landing stations near the coast.

We model SONET circuits because they are statically provisioned and do not necessarily exist between all landing station pairs. Thus absence of a logically provisioned SONET circuit can leave a physical connection useless. For each SONET circuit, we evaluate reachability based on if a logical circuit is provisioned. In principle, circuits have latency and capacity, but we model those as part of the IP link described later.

Reachability A cable cut breaks one or more cable segments and thus changes the physical topology of a SONET submarine cable system. Because SONET circuits are virtual circuits which need to be provisioned, we can not simply examine reachability by finding paths in the changed topology.

Figure 4 gives an example. Four cable segments ($C_{1,2,3,4}$) connecting four landing stations compose the physical

sub-model	source
modeling cable cut breaking SONET circuits	this paper
modeling SONET circuits affecting IP links	this paper
modeling IP links affecting flows	this paper and [16]
modeling flows affecting sessions	video streaming† [32], video telephony [44], gaming [9]
modeling sessions affecting QoE	video streaming† [12,24,32], VoIP [10], gaming [9]

Table 1: Sources of Sub-models. Daggers: sub-models used in this paper.

topology of cable system SCS 1. However, only three station pairs can communicate with each other through the provisioned SONET circuits ($S_{1,2,3}$). Pairs without SONET circuits in between will not be able to communicate even if they are physically connected (for example, station 1 and 4), because SONET circuits are assigned statically and human intervention is often required to reconfigure them.

To model how a cable cut breaks circuits between station pairs, we define this static mapping between SONET circuits and cable segments as matrix M^{cs}. Each matrix element $a_{ij} = 1$ (otherwise 0) if and only if SONET circuit i traverses through cable segment j. In the example shown in Figure 4,

$$M^{cs} = \begin{array}{c} \\ S_1 \\ S_2 \\ S_3 \end{array} \begin{array}{c} \begin{array}{cccc} C_1 & C_2 & C_3 & C_4 \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \end{array}$$

When a cable cut happens, the SONET circuits that traverse the broken segments will be affected. The consequence is straightforward, all these circuits will break.

We translate this impact to the equation shown below.

$$\mathbf{r^s} = M^{cs} \cdot_{\wedge} \mathbf{r^c} \tag{1}$$

$\mathbf{r^c} = [\ r_1\ r_2\ \cdots\ ]^T$ is the column vector denoting the reachability status of all cable segments (*true* if reachable, *false* if broken), while $\mathbf{r^s}$ represents all SONET circuits.

The operator $\cdot_{\wedge}$ captures the fact that *a SONET circuit is reachable if and only if all cable segments it traverses through are reachable.* It is slightly different than the matrix multiplication (instead of sum, it computes the conjunction). It is defined as the following equation: $(A \cdot_{\wedge} B)_{ij} = \bigwedge_{k=1}^{m} a_{ik} b_{kj}$ where $(A \cdot_{\wedge} B)_{ij}$ is the element in i^{th} row and j^{th} column.

3.3 From SONET Circuits to IP links

We next model how SONET circuits affect IP links. An *IP link* is a virtual channel between two adjacent devices identified by IP addresses at the network layer. We use hop-by-hop IP links to model routing, and *IP paths* to refer to a series of IP links over an inter-networks.

We model IP links explicitly to allow for SONET circuit diversity. If SONET ring protection mechanism is used (two circuits supporting one link), the threat impact might be contained at the IP link and thus no impact on users.

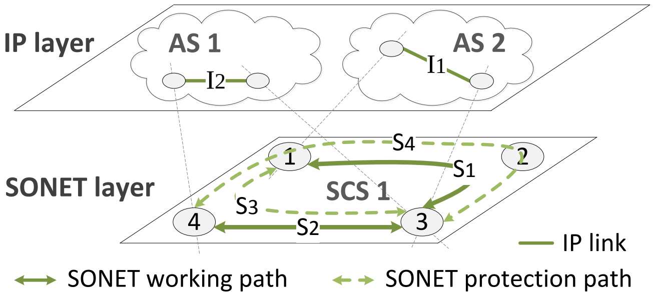

Figure 5: SONET systems with ring protection mechanism use two circuits (working and protection path) to support an IP link.

Two properties of an IP link are important to us: *reachability* and *latency*, because they are sufficient enough to capture the threat impact that will further propagate to users. Latency is the data propagation time, ignoring queuing delay, over the IP link. Normally it is fixed, but after a cable cut it may take a different value if a different SONET circuit is selected.

In principle, the cable cut could also change an IP link's capacity if multiple link-layer channels are supporting it via link aggregation. For simplicity, we model this situation using multiple IP links, each supported by a single link-layer channel at one time (working and protection SONET circuits do not work simultaneously).

Reachability A SONET circuit is supported by a *series* of cable segments, whereas an IP link is supported by one or two *parallel* SONET circuits. Due to this difference, the way the impact propagates is slightly different. The primary circuit is the *working path*, while the secondary one is the *protection path*. This protection scheme is known as the Multiplex Section-Shared Protection Ring (MS-SPring) or just "ring protection mechanism" [30]. As an example, in Figure 5, the working path S_1 and protection path S_3 together support IP link I_1. The protection path is optional.

Since the working and protection path are in parallel (rather than in series as cable segments), an IP link is reachable as long as *at least one* SONET circuit is reachable. Thus, unlike with SONET circuits, any active SONET circuit supports the IP link. We model this effect using the following equation:

$$\mathbf{r^i} = (W^{si} \cdot \mathbf{r^s}) \vee (P^{si} \cdot \mathbf{r^s}) \qquad (2)$$

Where, analogous to $\mathbf{r^s}$, $\mathbf{r^i}$ is the column vector denoting the reachability of all IP links. W^{si} and P^{si} are matrices mapping IP links to their working and protection paths, respectively. In the example shown in Figure 5,

$$W^{si} = \begin{array}{c} \\ I_1 \\ I_2 \end{array} \begin{array}{c} \begin{array}{cccc} S_1 & S_2 & S_3 & S_4 \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \end{array}, P^{si} = \begin{array}{c} \\ I_1 \\ I_2 \end{array} \begin{array}{c} \begin{array}{cccc} S_1 & S_2 & S_3 & S_4 \end{array} \\ \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

Because SONET circuits can be sold to different ISPs, the IP links they support can reside in different ISPs' networks. We show ISPs as different Autonomous Systems (ASes) in Figure 5.

Latency In cases where an IP link is still reachable after the cable cut, its latency may increase if protection path has higher latency than the working path. For long-haul IP links in modern networks, propagation delay is the major component of IP link latency. We assume *capacity* of both circuits is the same, as is typical in practice [30].

We thus model the impact on latency by the following equation:

$$\mathbf{l^i} = \begin{cases} W^{si} \cdot \mathbf{l^s} & \textit{if working path functions} \\ P^{si} \cdot \mathbf{l^s} & \textit{otherwise} \end{cases} \qquad (3)$$

where $\mathbf{l^i}$ is the column vector denoting the latency of all IP links, while $\mathbf{l^s}$ denotes latency of SONET circuits. An element in $\mathbf{l^s}$ is a finite number unless its corresponding circuit is broken. If broken, the value of the element equals to ∞.

Note that we have ignored the queuing delay might induced by the cable cut. A cable cut might causes some core IP links congested and thus increases their queuing delay. However, we believe in such cases, congestion-reactive traffic will back off [17], and routers will drop packets in the queue as there is no benefit to keep a queue when the link is heavily congested.

3.4 From IP Links to Transport-layer Flows

IP links connect devices; we next consider flows that represent traffic over an IP path (a series of IP links). We use the traditional network definition of a transport-layer flow: a series of packets sent between two network endpoints identified by two IP addresses, two port numbers, and the protocol.

We choose to model flows for three reasons: they add multi-hop, routing, and congestion control, bridging IP links to applications. Our goal is to capture properties: *reachability*, *latency*, and *throughput*. Reachability is affected by multi-hop communication and routing that can find paths around failed links. Latency is affected by path changes that increase path length. Finally, completing flows can trigger congestion control and change effective throughput, an important factor in application quality.

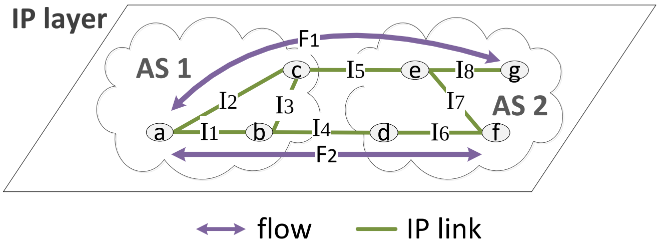

Figure 6: Traffic flows between two endpoints rely on the network layer to find a path composed of IP links. The path must comply with policies configured in routers.

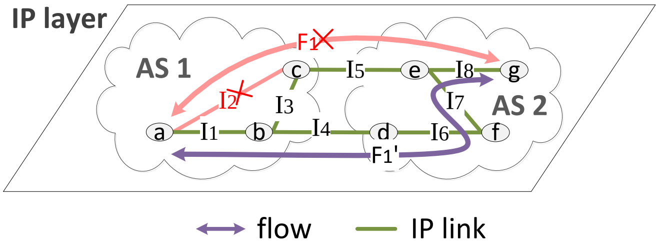

Figure 7: Flows are dynamically routed based on current IP link state for robustness.

Reachability We model flow reachability over IP links in the same manner as SONET circuit reachability over cable segments (Section 3.2). The way IP link reachability affect flow reachability However, unlike statically provisioned working and protection SONET circuits, the network layer uses dynamic routing to select from several possible paths. For example, in Figure 7, the flow between a and g may take two different IP paths. We thus model the impact on flow reachability with:

$$\mathbf{r^f}(t) = M^{if}(t) \cdot_{\wedge} \mathbf{r^i} \qquad (4)$$

This equation is similar as Equation 1, but the mapping is varies over time t as $M^{if}(t)$, unlike the the static M^{cs} in Equation 1.

In the example shown in Figure 6,

$$M^{if}(t) = \begin{array}{c} \\ F_1 \\ F_2 \end{array} \begin{array}{c} \begin{array}{cccccccc} I_1 & I_2 & I_3 & I_4 & I_5 & I_6 & I_7 & I_8 \end{array} \\ \begin{pmatrix} 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \end{pmatrix} \end{array}$$

Note that after a cable cut, $M^{if}(t)$ is likely to change. $M^{if}(t)$ is collaboratively decided by distributed routers running Border Gateway Protocol (BGP) and Interior Gateway Protocol (IGP) based on current IP reachability. Feamster et al. have proposed an algorithm to compute this matrix by emulating the route selection process of each ingress router for each destination prefix [16]. In principle, $M^{if}(t)$ can also capture load balancing and traffic engineering, but modeling these factors is future work.

Latency Flow latency is the sum of each IP links latency on its path. We capture the latency of each flow $\mathbf{l^f}(t)$ as:

$$\mathbf{l^f}(t) = M^{if}(t) \cdot \mathbf{l^i} \tag{5}$$

Throughput Throughput is affected both by IP link capacity and traffic and congestion control on each link. We assume most traffic is congestion-reactive [17], and therefore over medium-timescales each flow will converge on a fair share at its bottleneck. We thus model flow throughput $\mathbf{c^f}(t)$ as:

$$\mathbf{c^f}(t) = C^{if}(t) \cdot_{min} \mathbf{c^i} \tag{6}$$

where $\mathbf{c^i}$ denotes capacity of IP links. $C^{if}(t)$ is a matrix denoting the fraction of capacity occupied by each flow on each IP link. It is derived from $M^{if}(t)$ by computing the multiplicative inverse of number of flows on each IP link. Here operator $\cdot_{min}$ computes the minimum value over the vector. That is: $(A \cdot_{min} B)_{ij} = Min_{k=1}^{m} a_{ik} b_{kj}$ where $(A \cdot_{min} B)_{ij}$ is the element in i^{th} row and j^{th} column.

3.5 From Flows to Sessions

Applications often use one or more flows to realize complex network services; we call this exchange of the information a *session*. (Our sessions are somewhat more general than the OSI session layer, and are implemented in applications and libraries.)

We model sessions as a bridge between flows and application QoE. This bridge allows us to identify metrics that are application-specific but lower-level than users might care about. These metrics are useful because they are common to several different models of QoE, and because they identify measurable things in the network that we can verify. We expect each application to require distinct session information. We draw on prior work in modeling multiple applications, focusing on video streaming using a model developed by Mok et al. [32]. We focus on one generic session property, *reachability*, and three *application-specific properties of video streaming*, which later fit into the session-QoE model. Other applications that could be used within this framework include video telephony, VOIP, gaming, and newly emerged cloud applications.

Reachability We consider the reachability of a session equals to the one of the transmission flow. A video streaming session may initiate one or a series of transmission flows to transfer video segments to the user [20]. In case where multiple flows are used, we consider them as one flow but with changing endpoints (servers). There are also other flows involved in a streaming session, such as DNS queries that map the service to servers. But because these flows have much less influence over the session quality, we thus do not model them.

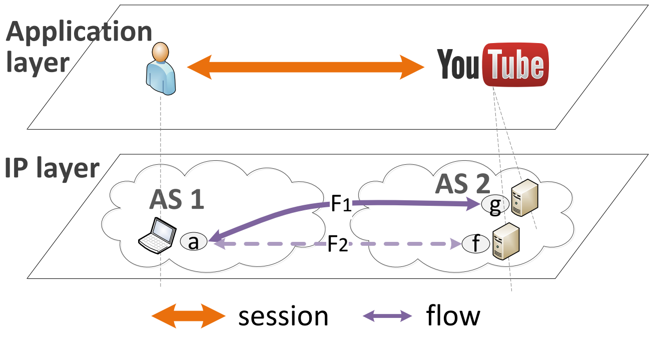

Figure 8: A session between an user and a service relies on one or multiple flows between the user client and server(s).

Prior work shows that there is typically only one TCP flow at a time [20], we therefore consider the reachability of a session equals to the one of the single flow at a given time.

Most video streaming services employ Content Delivery Networks (CDNs) that distribute content around the Internet in caches to reduce latency to the user, and bandwidth costs to the provider. Different CDN caches provide redundancy to the session, just as backup SONET circuits do to the IP link. Reachability of to any cache allows the service to proceed.

For example, in Figure 8, the session the user and video streaming service *YouTube* depends on the underlying flow F_1. But if anything goes wrong with F_1, the user can still access the service by F_2 which goes to another cache providing the same content.

We model session reachability ($\mathbf{r^a}(t)$) as the following equation (We use superscript $^\mathbf{a}$ to denote "application". In principle, we should use $^\mathbf{s}$ to denote "session", however, $^\mathbf{s}$ is already used for "SONET circuit").

$$\mathbf{r^a}(t) = M^{fa}(t) \cdot \mathbf{r^f}(t) \tag{7}$$

where $M^{fa}(t)$ is the dynamic mapping between sessions and flows (or services and servers). Note that unlike Equation 2 where redundancy is expressed by disjunction, redundancy here is expressed by dynamic. $M^{fa}(t)$ will automatically change after a failure to restore the session reachability.

Application-specific properties of video streaming We also draw on several properties specific to video streaming: *video bitrate*, *startup delay*, and *rebuffering ratio*. These properties have been developed in models specific to video evaluation [12, 24, 32]; we adapt them to our model and flow rate c^f.

Video bitrate r_v is often congestion adaptive, picking certain bitrate tiers in modern streaming players [20]; we model this mechanism by picking the largest r_v just less than c^f. Video startup delay (d_s) is

$$d_s = \frac{\beta \times r_v}{c^f} \tag{8}$$

a function of video buffer size β, video bitrate, and flow rate (from Mok et al. [32]). Rebuffering ratio is also important. We model it as

$$r_b \approx \frac{r_v}{c^f} - 1 \tag{9}$$

aggregated from two separate models in [32] (they define separate rebuffering time and frequency).

We define these network-level metrics of video performance in our session model because they serve as input to several different QoE models (described in the next section), and because they help identify how specific network phenomena cause problems to the user experience.

Other applications Video telephony is also an important application on the Internet. Well-known specific systems include Skype video calls, Google hangout, and iChat. Zhang et al. [44] proposed models to predict three session qualities (sending rate, video rate, and frame rate) of Skype.

Compared with video streaming, video telephony is more sensitive to real-time condition such as latency. However, as the models in [44] show, throughput is still the most important factor.

3.6 From Sessions to QoE

We can now bring our model to the user by modeling their Quality of Experience in specific applications. We can reflect reachability as completely unacceptable QoE ($-\infty$), but more QoE is interesting when it reflects more subtle differences. As a concrete application where we can estimate QoE, we continue to focus on video streaming.

Application-specific QoE of video streaming We survey four QoE models of video streaming developed in three prior papers [12, 24, 32]. These models are induced by analyzing data sets of varying sizes. The model in [32] is based on lab experiments including 270 views from 10 viewers; while the two in [24] are based on a much larger set from Akamai (23 million views from 6.7 million viewers). The model in [12] is drawn from the largest and most diverse dataset, including 300 million views from 100 million viewers in a week, from various content providers. We thus first choose the model in [12]. In addition, we also incorporate one model in [24] as another branch to complete our model. These two models focus on two different aspects of QoE (play time and abandonment rate) and we think they are both useful.

The model in [12] uses *decreased video play time* (Δ_{QoE_P}) to indicate user QoE (less play times indicates worse experience), and studies how it is shaped by rebuffering ratio (r_b). The model can be formalized by the following equation:

$$\Delta_{QoE_P} = \begin{cases} -1 \; minute/\% * r_b & video\ on\ demand \\ -3 \; minute/\% * r_b & live\ video \end{cases} \tag{10}$$

which means users watch 1 or 3 minutes less every 1% more rebuffering for two types of video. Note that this model has a range where it is applicable. It only applies to rebuffering ratio less than 10%. Beyond that, QoE is too bad to be applicable.

Focusing on another aspect of user experience, the model in [24] uses *negative video abandonment rate* (QoE_A) to indicate user QoE and studies the causality between it and the startup delay (d_s).

$$QoE_A = -(d_s - 2) * 5.8\% \tag{11}$$

The above equation means that users start to abandon videos after 2 seconds of startup delay, and abandonment rate raises by 5.8% every one more second delay. This model also has its application range. The authors did not discuss it directly, but by examining their regression graph (Figure 10 in [24]), we conclude that this model only applies to startup delay less than 10 seconds.

Other applications Although we focus on video streaming, QoE models exist for other applications such as Internet telephony. Chen et al. [10] proposed a model to predict Skype voice call QoE from session quality. Their model shows that Skype sending rate and the jitter of sending rate are the two most important factors that ensure a good quality of experience for users.

Online gaming is another application area where QoE can be modeled. Chang et al. [9] has proposed a model for QoE for on-line gaming, showing that both display frame rate and frame distortion are critical to user experience. Such models could fit in our framework.

3.7 Data needed for the model

So far, we have completed our model by incorporating prior models and developing ones that are needed.

This model helps one to predict what users will be affected by the cable cut and how the cut affects their video streaming experience. However, to conduct any useful prediction, one needs to collect real-world data as input and parameters to the model.

Table 2 lists the data needed. As we can see from the table, almost all data are *proprietary* and thus hard to obtain. However, we can still gather some through measurement and online documents (see citations in the table).

Some data are not only proprietary, but also *dynamic* (notations with (t)), which means they may be hard to obtain even for service providers. For example, the

layer	nota-	meaning	value/	prop-	in equa-
	tion		source	rietary	tion
PHY	$\mathbf{r}^{\mathbf{c}}$	reachability status of cable segments	cable owners	y	1
PHY to SONET	M^{cs}	mapping from SONET circuits to cable segments	cable owners	y	1
SONET	$\mathbf{l}^{\mathbf{s}}$	latency of SONET circuits	cable owners	y	3
SONET to IP	W^{si}	mapping from IP links to SONET working paths	ISPs	y	2, 3
	P^{si}	mapping from IP links to SONET protection paths	ISPs	y	2, 3
IP	$M^{if}(t)$	mapping from flows to IP links at time t	ISPs	y	4, 5
	$C^{if}(t)$	mapping of IP link bandwidth to flows at time t	ISPs	y	6
	$\mathbf{c}^{\mathbf{i}}$	capacity of IP links	ISPs	y	6
IP to APP	$M^{fa}(t)$	mapping from sessions to flows at time t	app providers	y	7
APP	β	video buffer size t	$0.5 \sim 5$ min [1]	y	8
	r_v	video bitrate t	$0.35 \sim 3.8$ Mb/s [1]	y	8, 9

Table 2: Data needed for the model.

mapping from flows to IP links ($M^{if}(t)$) is dynamic. It is governed by complex routing protocols that reside in distributed routers according to the current IP link state. Often, even internal operators find it hard to predict routes for flows.

We see two approaches to obtain dynamic data. First, one can log the information for a period of time long enough to predict future behavior. Most network traffic has strong diurnal and weekly periodicity that allows trend identification. Alternatively, one can build models that infer dynamic behavior from slower changing information, such as prior work in routing [16] and traffic matrix estimation [31, 45].

In some cases, data may be unavailable, either to researchers or operators. Although missing data makes *specific* outcomes difficult to predict, modeling makes it relatively easy to quickly study a range of parameters. Such a study can suggest if negative outcomes are likely or unlikely over different possibilities.

4. CASE STUDIES

After constructing the model in Section 3, we next apply our model to understand real-world incidents (see Table 3). Specifically, we characterize specific aspects of networks that make countries more or less vulnerable to threats. In addition, we also explore how one can apply the model when facing incomplete data. We find that service self-sufficiency (hosting services near users), and geographic diversity of circuits both help insulate a country from outages.

In this section we focus on Bangladesh and its 2012 cut (the first incident in Table 3). We explored this incident concurrent with developing our model. Subsequently, we applied our model to the three other cases listed in Table 3. Although each scenario requires new parameters, our model is effective at evaluation of these additional cases, suggesting it generalizes and is not overfit to a single occurrence. Due to space constraints, detailed discussion of these cuts is in Appendix B.

The SeaMeWe-4 cable cut happened in 2012 (Section 4.1) had a significant impact on Bangladesh; to understand its cause we first apply our model for an explanation (Section 4.3). Besides the explanation, we would also like to quantify the impact, especially on user QoE, beyond what has been reported in public news (Section 4.4). One step further, countries which suffer from cable cuts would also like to know how to mitigate the impact. We therefore present a method to help countries address this issue (Appendix A).

The process to apply our model has been discussed in Section 3 and shown in Figure 3. However, to address incomplete data, we have slightly changed the course. We briefly describe this modified process (Section 4.2) and apply it to one of our examples for illustration. Finally, we would like to share what we have learned about addressing incomplete data in a generic scenario (Section 4.5).

4.1 Incident Overview

SeaMeWe-4 submarine cable system connects 17 landing stations from South East Asia via Middle East to Western Europe (Figure 9), in a bus-like topology [36]. While submarine cables are often rings, geography forces a linear topology here. The system is managed by a consortium composed of 16 telecommunication companies and spans about 20,000 km, supporting communication at 1.28 Tb/s [42].

We analyze the 6 June 2012 cable cut occurring 60 km outside Singapore that disconnected it from other stations [5].

A naive model of reachability might suggest that this cut would affect Internet users in Singapore. However, public reports suggest that Singapore users were *barely affected*, while Bangladeshi users experienced significant problems [5]. Press reports suggest about eight million Bangladesh netizens suffered very slow connections af-

incident	victim country	cables (total)	cables (cut)	self- sufficiency	geo- diversity	geo- weakness	capacity drop
SeaMeWe-4'12 [5]	Bangladesh	1	1	4%	low	eastbound to Singapore	67%*
SeaMeWe-4'13 [34]	Pakistan	4	2	0%	medium	westbound to Europe	60%†
IMEWE'12 [27]	Lebanon	1	1	8%	low	westbound to France	100%†
TEAMS'12 [6]	Kenya	3	1	4%	medium		20%†

Table 3: Four real-world incidents we have studied. Asterisks (*): estimated, daggers (†): reported.

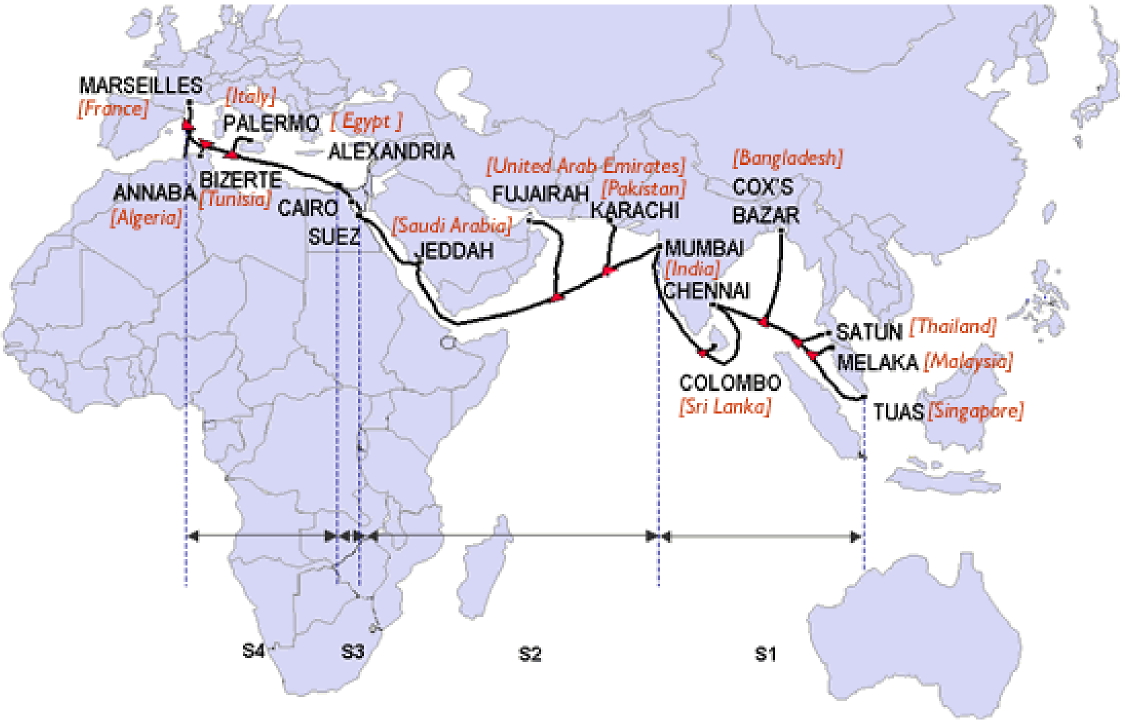

Figure 9: Physical topology of SeaMeWe-4 [35].

ter the cut. We did not find news for Thailand and Malaysia which likely implies that these two countries were also barely affected.

4.2 Applying the Model

We next briefly describe the applying process. The ideal process is shown as the solid red line in Figure 3, however, in order to address data incompleteness, we instead follow two branches of it visualized by dashed purple lines. (The dashed parts of the two branches are where data is unavailable. To address this problem, we perform what-if analysis on these parts.) The *service-reachability* branch reverses the bottom-up order of the ideal process, starts from the session reachability, and works down the stack to the cable segment reachability. The *user-QoE* branch keeps the bottom-up order, starts from the IP link capacity, and ends at video play time, based on the prerequisite that session reachability is satisfied. We next describe the two branches in details.

Service reachability branch integrates Equation 7, 4, 2 and 1 to obtain the cross-layer impact on service reachability by following a top-down direction. Data needed in this branch are obtainable through measurement or educated inference, therefore this path allows us to quantify service reachability.

To analyze service reachability, we first apply Equation 7 to map service reachability to flow reachability. We obtain the data needed in this step ($M^{fa}(t)$) from distributed DNS queries and BGP routing tables. Distributed DNS queries can discover more servers than from a single vantage point, and thus decrease the possibility of underestimating service reachability. We therefore query server addresses from all major Singaporean and Bangladeshi ISPs. We limit our scope to Singaporean and Bangladeshi users only to make the problem more manageable. More specifically, we only consider sessions between users and top services within these two countries respectively. We identify top services as the top 25 websites for each country provided by Alexa [4], and users by IP addresses announced by major Singaporean and Bangladeshi ISPs.

One layer down, we then map flow reachability to IP link reachability by applying Equation 4. We obtain the data needed here ($M^{if}(t)$) from traceroutes targeted either to servers or users supplemented by AS path inference.

Down to the bottom of the stack, we analyze cut effects on IP link reachability by applying Equation 2 and Equation 1. The data needed in these two steps (M^{cs}, W^{si}, P^{si}) are proprietary, and also hidden under the IP layer within each network's boundaries. Only the ISPs have this data and can measure it, so instead we make educated estimates based on ISPs' public documents. We use published IP [38] and physical topologies [39]; some ISPs make this data available to attract customers. We can compare these two topologies to infer what IP links depend on what cable segments.

One can follow the above service-reachability branch to analyze cut effects on service reachability, which provides a binary answer (either accessible or not). QoE depends on service reachability (QoE only makes sense when services are reachable), but contains much richer information about user satisfaction. We next describe the other branch that examines user QoE based on the assumption that services are reachable.

User QoE branch integrates Equation 6, 9 and 10. This analysis requires proprietary information including IP link capacity and flow traffic matrix ($\mathbf{c^i}$ and $C^{if}(t)$ in Equation 6). Even for ISPs where this information is known, the exact values change over time. We therefore study a *range* of values in the parameters space to understand the range of conditions where the network is robust or fragile. Thus our approach can both answer what-if questions, where one provides or speculates about specific parameters, and project beyond current usage to possible future scenarios.

We study link capacity ($\mathbf{c^i}$) and user traffic ($C^{if}(t)$). To simplify representation, we replace capacity of individual links with aggregate international capacity, and dynamic traffic with maximum flow count. This approximation is appropriate when most services are international (as we show they are for Bangladesh) and flows are congestion-reactive and therefore will converge on a fair share at the bottleneck international links. Figure 10 shows an example of the parameter space, with capacity shown against flow count.

4.3 Causes of Large Impact on Bangladesh

We next discuss the two weaknesses of Bangladesh's infrastructure that our analysis reveals: low service self-sufficiency and low geographic diversity of international circuits. These weaknesses result in harm to Bangladeshis.

We define *self-sufficiency* as a metric to quantify the degree a country depends on the outside world for Internet services. Specifically, it equals the number of top 25 websites that are hosted by any domestic servers. For Bangladesh, the self-sufficiency is very low (4%), meaning only one website is within its border. More specifically, among the 24 popular websites hosted abroad, 16 are foreign or global services (such as Google and Facebook) eight are local but hosted abroad (such as BanglaNews24 and BDJobs), suggesting an opportunity for new hosting services inside Bangladesh to improve self-sufficiency.

In contrast, Singapore is much more self-sufficient (52%). When considering only the top five websites, its self-sufficiency rises to 80%.

The low service self-sufficiency suggests that Bangladesh heavily depends on the outside world, and thus it is very important for Bangladesh to diversify its international outlets to cope with physical threats that disrupt regional connectivity.

However, at the time of the cut (mid-2012), SeaMeWe-4 was Bangladesh's only high-capacity international cable We confirmed this statement by searching through the complete list of submarine cables [42], and concluded that Bangladesh had no terrestrial connectivity at that time because a later Dec. 2012 terrestrial connection (via the *ITC* cable) appeared as major news [8]. Satellite or dialup links can provide service for some, but both are slow and do not support general traffic).

The low cable diversity is further intensified by the low circuit diversity. Most Bangladesh's international circuits are provisioned to the east connecting with Singapore and so were cut during the incident. The sudden disruption of eastbound circuits leads approximately to a 60-70% drop of Bangladesh's total international capacity. If Bangladesh had provisioned more backup circuits to the west connecting with global ISPs in Middle East or Europe, user traffic could shift to the west and threat impact would be much smaller.

In summary, the low geographic diversity of circuits, together with the low service self-sufficiency, has made Bangladesh vulnerable to cable cuts. Self-sufficiency will improve by either encouraging popular foreign services to deploy servers in-country, or if popularity of domestic services grows. geographic diversity of international connectivity is improved by adding circuits or cables to new destinations, as Bangladesh did in Dec. 2012 [8].

4.4 Impact on QoE in Different What-If Scenarios

This section studies the threat impact on user QoE for different possible scenarios. These different possibilities allow us to explore potential future situations, an approach that applies not only to cable cuts, but also other cases where international capacity supply or traffic demand changes (such as planned maintenance and flash crowds).

We explore a two-dimensional parameter space to study these what-if scenarios. We have briefly described the parameter space in § 4.2 which has international capacity and flow count as its two dimensions. Figure 10 shows how user QoE (represented by play time) varies in this space, as measured by:

$$\Delta_{QoE_P} = 100(1 - \frac{r_v y}{x})\ minute \qquad (12)$$

integrated from Equation 6, 9 and 10.

QoE models and Figure 10 simplify several aspects of Internet video. Rather than model adaptive video, we approximate r_v by fixing it at the basic bitrate of many services today ($r_v = 350$ kb/s). In addition, the QoE curve may vary by content type (for example, some animation can be encoded more efficiently); these differences are not reflected in current QoE models. Finally, we allocate bandwidth equally over all users in a country. In practices, their needs will vary. Future work could capture these effects by refining our model, perhaps with a multi-tier QoE equation.

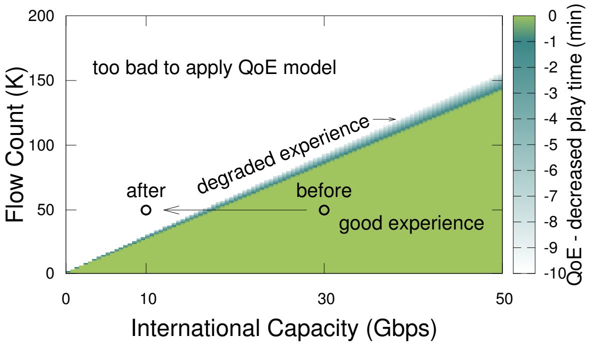

Figure 10: QoE - decreased play time (minute) in 2-D parameter space, $r_v = 350$ kb/s

As shown in the figure, there are three regions corresponding to different QoE in the parameter space. The bottom right green triangle region corresponds to good experience with zero decreased play time (0 in color box). The narrow middle strip (shown as shades of blue) just on top of it corresponds to degraded experience. Users in this region react by watching less time ranging from several seconds to 10 minutes (-1 ∼ -10). Note that this region is very narrow, indicating there is fairly little adaptively at these scales. The upper left white region is where the QoE model does not apply because the capacity is so small and flows are so many, in this region users are unlikely to use the service at all.

To answer what-if questions, we consider performance before and after a network change. As shown in the figure, this approach first picks a specific position in the space ("before" dot) representing the state of a given country before the cut happened, it then moves the initial state to another position ("after" dot) representing the state after the cut. The approach finally assesses the cut impact by comparing the QoE values at these two positions. In the example shown in Figure 10, the cut makes the QoE degrades from good to an intolerable level.

Theoretically, the before and after states can be at any positions in the parameter space. However, in practice, the before state typically appears in the good experience region and the direction and distance it moves only follows a finite set of ways. We next summarize major possible scenarios.

Steady demand (with decreased capacity) may be caused by a cable cut or simply a regular maintenance. In this scenario, users maintain their regular online activities despite the potential change of service qualities. Because the user traffic is unchanged, the direction the before state moves is fixed—it only moves horizontally to the left. The distance it moves reflects the capacity taken out of service by cable cuts or maintenance, in turn depending on the diversity of the country's connectivity.

The Bangladesh incident may match this scenario. We next conjecture about the possible impact on Bangladeshi users. Since the cable cut results in a 60-70% drop of Bangladesh's international capacity, the after state shifts to about 30% of the initial capacity (moving directly left). However, since the initial capacity (the before state) is unknown, we cannot exactly position the after state. Nevertheless, what-if analysis allows us to *bound regions* of before states that produce different outcomes. Figure 11(a) shows regions of before states where Bangladeshi users would feel (i) no impact (i.e., the after state is still within the good experience boundary) by inequality $100(1 - \frac{3r_v y}{x}) \geq 0$; (ii) moderate impact (i.e., the after state falls within the degraded experience boundary) by inequality $-10 \leq 100(1 - \frac{3r_v y}{x}) < 0$; and (iii) significant impact (i.e., the after state is outside the QoE model application range) by inequality $100(1 - \frac{3r_v y}{x}) < -10$. (Observe that these three inequalities are obtained by substituting x with $x/3$ in Equation 12 to reflect the capacity drop.)

Figure 11(a) also reveals interesting implications about connectivity planning—the trade-off between resilience and resources. The no-impact zone provides high resilience at the cost of over-provisioning capacity by a factor of 2. In contrast, the significant-impact zone, although vulnerable to cable cuts, can accommodate two-times more flows in normal situations. The amount of extra capacity to provide and how to plan is a crucial problem for many countries. Our models can guide countries in assessing such trade-offs, as we expand upon in Appendix A.

Decreased demand (with decreased capacity) could be caused by the same reasons as the first scenario, but reflects user defection (giving up) that often results from degraded service quality. With decreased demand, the after-state is both left and below the before state to reflect a decrease in number of flows as users defect, as shown in Figure 11(b).

Suppose Bangladesh falls into this scenario and half of the flows were withdrawn. We can again bound regions of before states by different outcomes they result in as shown in Figure 11(b). Defection extends the no-impact zone (compared with Figure 11(a)) for those users who remain, although the lower demand represents different cost of the cut.

This scenario matches intuition that when capacity decreases significantly, users wait (perhaps to come back later), or ISPs limit normal traffic so prioritized flows (such as control and business flows) can maintain their throughput. Note however that the price for having protecting prioritized flows is a potentially large number of significantly-impacted normal flows. Hence, the

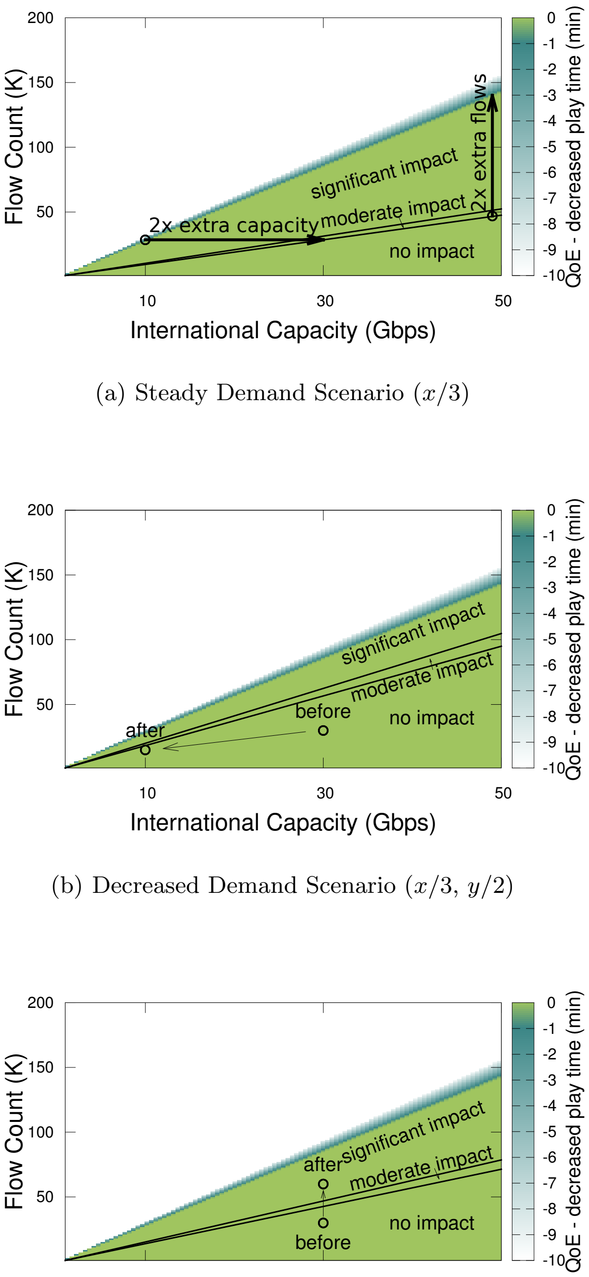

(a) Steady Demand Scenario ($x/3$)

(b) Decreased Demand Scenario ($x/3$, $y/2$)

(c) Increased Demand Scenario ($y \times 2$)

Figure 11: Different before states result in different outcomes.

fundamental solution is still overprovisioning with spare capacity.

Increased demand (with unchanged capacity) is typically not caused by cable cuts, since cable cuts usually reduce capacity. Instead, this scenario is often caused by flash crowds, such as global events that cause short-term traffic surges. The result is that the after state shifts *up* relative to before (see Figure 11(c)).

This scenario is not relevant to our cable cut for Bangladesh, but it shows the generality of our analytic approach. Similarly to the previous scenarios, we can bound regions of before states leading to different outcomes (Figure 11(c) shows the situation for doubled traffic). As we can see from the figure, to accommodate unexpected traffic surge (analogous to containing cable cut impact), over-provisioning of capacity is needed.

4.5 General Tactics to Address Incomplete Data

Our experiences suggest two general tactics to handle incomplete data: focusing on a subset of the problem to reduce the data needed and studying a range of possibilities. We next elaborate on each of these two tactics.

The first tactic generally applies to problems that require large amounts of hard-to-obtain data. To workaround missing data, we instead refocus on a subset of the problem that requires data that is more easily or readily available. In our specific case, the problem is to analyze the impact on all users in all countries given a cable cut. Solving this problem would need all related data about all users which is impossible to obtain. We instead focus only on a subset of all users. In the case at hand, this subset consists of users in Bangladesh because those users have reportedly been severely impacted by the cable cut. With this new focus, we only need data for network components that are concerned with Bangladesh and the users in that country. Before obtaining such data, we first identify the relevant components. To this end, we start with the users of interest who are represented in our model at the application layer and work our way down the stack. At each layer, we identify the components of interest. In short, we end up following the top-down service-reachability branch and not the ideal bottom-up process depicted in Figure 3.

The second approach addresses missing data by studying a range of possible values of the data. In our specific case, the required data to study user QoE is link capacity and user traffic. However, both data are proprietary and thus not available to us. To continue our analysis, we instead study a range of possible values of these data as we have demonstrated in Section 4.4. By performing the what-if analysis, we can not only study the current usage, but also predict possible future scenarios.

5. GUIDELINES TO UNDERSTAND AND MODEL THREATS

We next summarize the lessons we learned in framing, modeling, and analyzing network threats, and what they say about network design.

First, *almost on every layer of the Internet, topological connectivity does not imply data reachability.* The topology could be a well-known AS and router topology on the network layer, or physical cable topology, or even application-layer client-server/peer-to-peer network topology. Prior work [13, 18, 43] has mainly focused on the connectivity/reachability issue of the AS topology. In this paper, we also showed that cable connectivity is not enough — a SONET circuit needs to be established to transmit data (§ 3.2 and 4.3). In fact, the data transmission on every layer needs to follow the layer's control protocol (such as SONET, Ethernet, MPLS, OSPF, BGP, TCP, RTCP and SIP) and it is these protocols that govern the data reachability, given the prerequisite that the two ends are topologically connected. Thus, to model data reachability, one must consider the behavior of these control protocols. We stress this principle because it can be easily forgotten and therefore cause modeling errors.

Second, *fault-recovery mechanisms reside on many layers, and new ones are frequently added.* Network routing is the best known recovery mechanism (§ 3.4). In this paper, we also identified the ring protection on SONET layer (§ 3.3) and server redundancy on the application layer (§ 3.5). Use of multiple servers to enhance reliability and performance is a relatively new mechanism used by content delivery networks (CDNs). Sometimes, when all of the lower layers fail to contain the threat impact, CDNs can mitigate the impact by delegating proper servers to serve users (In § 4.3, we show that Singaporean users are mainly served by local servers and thus were barely affected by the cut). Therefore, to correctly assess threat impacts, one must consider the fault-recovery mechanisms on all related layers, paying special attention to newly introduced ones.

Third, the *effects of threats on real users are strongly influenced by user behavior and network architectures.* This observation has implications on both modeling and network deployment. To model threats, it means that one must consider local users' preferences for services. One must also understand where modern CDNs deploy servers, since CDN nodes can make a "foreign" service local. To understand the impact of threats, one must identify common traffic sources and destinations instead of just picking arbitrary endpoints, as we have demonstrated (§ 4.2).

This observation also suggests that countries that wish to improve their network resilience can do so by improving self-sufficiency (encouraging use of local services and local replicas of global services), as well as by diversifying network connectivity. All of these "best practices" can reduce the impact of disruptions as shown by the examples in § 4.

Lastly, *reachability is the basis, but by far not enough to capture QoE for modern users.* Modern users' expectation has risen sharply along with the rapid development of the Internet. Years ago, being able to fetch a webpage is satisfying enough for many users, while nowadays, users abandon a webpage in seconds and regard a video that buffers frequently intolerable (see our discussion in Section 3.6). Therefore, to draw real-world attention, we need to shift the focus from reachability to QoE.

6. CONCLUSION

We have developed a holistic model that first relates low-layer physical threats with high-layer Quality-of-Experience for end users. Since no single organization has data that spans all these layers, we applied what-if analysis to understand possible outcomes in the face of gaps in specific data. We have applied our model to four incidents and identified low service self-sufficiency and low geographic diversity as two major vulnerabilities of developing countries. What-if analysis and our model can predict possible outcomes of future events, and the effects of mitigation strategies.

7. REFERENCES

[1] S. Akhshabi, A. C. Begen, and C. Dovrolis. An experimental evaluation of rate-adaptation algorithms in adaptive streaming over http. In *Proceedings of the second annual ACM conference on Multimedia systems*, pages 157–168. ACM, 2011.

[2] R. Albert, H. Jeong, and A.-L. Barabási. Error and attack tolerance in complex networks. *Nature*, 406:378–382, July 27 2000.

[3] D. Alderson, L. Li, W. Willinger, and J. C. Doyle. Understanding internet topology: principles, models, and validation. *Networking, IEEE/ACM Transactions on*, 13(6):1205–1218, 2005.

[4] Alexa. The top 500 sites in each country or territory. `http://www.alexa.com/topsites/countries`, June 2013.

[5] BBC. Bangladesh suffers Internet disruption after cut cable. `http://www.bbc.co.uk/news/technology-18366007`, June 2012.

[6] BBC. Ship's anchor slows down east african web connection. `http://www.bbc.co.uk/news/world-africa-17179544`, Feb. 2012.

[7] BBC. Egypt arrests as undersea Internet cable cut off Alexandria. `http://www.bbc.co.uk/news/world-middle-east-21963100`, Mar. 2013.

[8] bdnews24. Bangladesh connected with terrestrial cable. `http://biz-bd.bdnews24.com/details.php?id=237802&cid=4`, Dec. 2012.

[9] Y.-C. Chang, P.-H. Tseng, K.-T. Chen, and C.-L. Lei. Understanding the performance of thin-client gaming. In *Communications Quality and Reliability (CQR), 2011 IEEE International Workshop Technical Committee on*, pages 1–6. IEEE, 2011.

[10] K.-T. Chen, C.-Y. Huang, P. Huang, and C.-L. Lei. Quantifying Skype user satisfaction. In *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '06, pages 399–410, New York, NY, USA, 2006. ACM.

[11] K. Cho, C. Pelsser, R. Bush, and Y. 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*, page 2. ACM, 2011.

[12] F. Dobrian, V. Sekar, A. Awan, I. Stoica, D. Joseph, A. Ganjam, J. Zhan, and H. Zhang. Understanding the impact of video quality on user engagement. In *Proceedings of the ACM SIGCOMM 2011 conference*, SIGCOMM '11, pages 362–373, New York, NY, USA, 2011. ACM.

[13] D. Dolev, S. Jamin, O. Mokryn, and Y. Shavitt. Internet resiliency to attacks and failures under BGP policy routing. *Computer Networks*, 50(16):3183–3196, 2006.

[14] J. Doucette, W. D. Grover, et al. Capacity design studies of span-restorable mesh transport networks with shared-risk link group (srlg) effects. In *SPIE Opticomm*, 2002.

[15] J. C. Doyle, D. L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger. The robust yet fragile nature of the Internet. *PNAS*, 102(41):14497–14502, 2005.

[16] N. Feamster, J. Winick, and J. Rexford. A model of BGP routing for network engineering. In *SIGMETRICS*, pages 331–342. ACM, 2004.

[17] S. Floyd and K. Fall. Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM ToN*, 7(4):458–473, Aug. 1999.

[18] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM ToN*, 9(6):733–745, Dec. 2001.

[19] D. Greenlees and W. Arnold. Asia scrambles to restore communications after quake. `http://www.nytimes.com/2006/12/28/business/worldbusiness/28iht-connect.4042439.html?_r=1`, Dec. 2006.

[20] T.-Y. Huang, N. Handigol, B. Heller, N. McKeown, and R. Johari. Confused, timid, and unstable: picking a video streaming rate is hard. In *Proceedings of the 2012 ACM conference on Internet measurement conference*, IMC '12, pages 225–238, New York, NY, USA, 2012. ACM.

[21] H. H. Inu. Ipv6 deployment in bangladesh. http://meetings.apnic.net/__data/assets/pdf_file/0003/23664/BD-IPv6Bangladesh.pdf, 2011.

[22] ISO/IEC. ISO/IEC standard 7498-1. http://standards.iso.org/ittf/PubliclyAvailableStandards/s020269_ISO_IEC_7498-1_1994(E).zip, 1994.

[23] ISPAK. Internet Facts. http://www.ispak.pk, Apr. 2012.

[24] S. S. Krishnan and R. K. Sitaraman. Video stream quality impacts viewer behavior: inferring causality using quasi-experimental designs. In *Proceedings of the 2012 ACM conference on Internet measurement conference*, IMC '12, pages 211–224, New York, NY, USA, 2012. ACM.

[25] R. N. M. Omer and A. Mostashari. Measuring the resilience of the global Internet infrastructure system. In *Proc. of the IEEE International Systems Conference*, pages 156–162, Vancouver, Canada, 2009.

[26] D. Madory. East african internet resilience. http://www.renesys.com/2012/02/east-african-cable-breaks/, Feb. 2012.

[27] D. Madory. Lebanon Loses Lone Link. http://www.renesys.com/blog/2012/07/large-outage-in-lebanon.shtml, July 2012.

[28] D. Madory. Smw4 cut shakes up south asia. http://www.renesys.com/blog/2012/06/smw4-break-on-south-asia.shtml, June 2012.

[29] G. Mahlknecht. Greg's cable map. http://www.cablemap.info/, 2013.

[30] G. Maier, A. Pattavina, S. De Patre, and M. Martinelli. Optical network survivability: protection techniques in the WDM layer. *Photonic Network Communications*, 4(3-4):251–269, 2002.

[31] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot. Traffic matrix estimation: existing techniques and new directions. In *SIGCOMM '02*, pages 161–174, New York, NY, USA, 2002. ACM.

[32] R. K. Mok, E. W. Chan, and R. K. Chang. Measuring the quality of experience of HTTP video streaming. In *Integrated Network Management (IM), 2011 IFIP/IEEE International Symposium on*, pages 485–492. IEEE, 2011.

[33] C. C. of Kenya. Quarterly sector statistics report. http://www.cck.go.ke/resc/statistics/SECTOR_STATISTICS_REPORT_Q2_2010-11_x2x_x3x_x2x.pdf, Nov. 2010.

[34] Renesys. Intrigue surrounds smw4 cut. http://www.renesys.com/2013/03/intrigue-surrounds-smw4-cut/, Mar. 2013.

[35] SEA-ME-WE 4. Cable system configuration. http://www.seamewe4.com/inpages/cable_system.asp, Feb. 2013.

[36] SEA-ME-WE 4. Sea-me-we 4 potential customer material. http://www.seamewe4.com/pdfs/home/Customer_event/SMW4_Customer_event_with_Backhaul_slide.pdf, Feb. 2013.

[37] P. Sebos, J. Yates, G. Hjalmtysson, and A. Greenberg. Auto-discovery of shared risk link groups. In *Optical Fiber Communication Conference and Exhibit, 2001. OFC 2001*, volume 3, pages WDD3–WDD3. IEEE, 2001.

[38] SingTel. Coverage map. http://business.singtel.com/upload_hub/mnc/STiX_Factsheet_2010.pdf, 2010.

[39] SingTel. Our coverage. http://info.singtel.com/large-enterprise/products/global-connectivity/our-coverage, Feb. 2013.

[40] T. D. Star. Internet speed increases in one year. http://www.dailystar.com.lb/Business/Lebanon/2012/Nov-10/194587-internet-speed-increases-in-one-year.ashx, Nov. 2012.

[41] T. D. Star. Lebanon experiences nationwide Internet blackout. http://www.dailystar.com.lb/Business/Lebanon/2012/Jul-02/179079-lebanon-experiences-nationwide-internet-bl ashx, July 2012.

[42] Submarine Telecoms Forum. Submarine Cable Almanac. http://www.subtelforum.com/articles/submarine-cable-almanac/, Feb. 2013.

[43] J. Wu, Y. Zhang, Z. M. Mao, and K. G. Shin. Internet routing resilience to failures: analysis and implications. In *ACM CoNEXT*, pages 25:1–25:12, New York, NY, USA, 2007. ACM.

[44] X. Zhang, Y. Xu, H. Hu, Y. Liu, Z. Guo, and Y. Wang. Profiling Skype video calls: Rate control and video quality. In *INFOCOM*, pages 621–629, 2012.

[45] Y. Zhang, M. Roughan, C. Lund, and D. Donoho. An information-theoretic approach to traffic matrix estimation. In *SIGCOMM '03*, pages 301–312, New York, NY, USA, 2003. ACM.

APPENDIX

A. IMPLICATIONS FOR CONNECTIVITY PLANNING

To help a country to plan its international connectivity, we present a method in this section.

A problem many countries, especially developing ones, face is that *in order to let a% of population to enjoy decent online experiences, how much international capacity should be provisioned?* More specifically, what submarine cable consortium should the country participate in? what circuits should be provisioned? and how much capacity for each circuit?

In addition to normal conditions, countries are also interested in whether the goal will be respected during abnormal situations (such as sudden capacity drop caused by cable cuts). This concern leads to more questions: *how much extra capacity needs to be provisioned? how to distribute the extra capacity among circuits?*

We next present a method to aid countries answering these questions based on the what-if study described in Section 4.4. The core idea is to *diversify the connectivity to limit the capacity changes brought by common threats, and therefore to achieve the maximal resilience using the minimal amount of resources.* Figure 12 helps to understand this idea. For a given traffic demand, the *minimal capacity* needed to achieve good QoE during normal condition resides on the boundary of the "good experience" zone. However, in order to be also resilient to threats, the country needs *extra capacity* to extend itself into the "no impact" zone. The shape of the good-experience zone is irrelevant to country connectivity and thus the minimal capacity is fixed for a given traffic demand. In contrast, the shape of the no-impact zone can be changed via careful connectivity planning, and the bigger the zone, the less extra capacity is required. This observation provides countries important implications—a good connectivity planning can save millions of dollars on infrastructure resources.

The no-impact zone is bounded by how much capacity and traffic varies during abnormal conditions. The larger the capacity drop, or the larger the traffic increases, the smaller the no-impact zone is. Hence, to extend the no-impact zone and in turn reduce the extra capacity needed, the country needs to *shrink the range of capacity drop and traffic increase.* We focus on how to shrink capacity drop in this paper and leave confining traffic increase for future work.

To shrink the capacity drop, the country needs to diversify its connectivity so a single threat only affects a limited number of circuits with a limited amount of capacity. A helpful concept here is the Shared Risk Group (SRG) which identifies resources (such as circuits) that are likely to be brought down by a common threat (for example, all Bangladesh's eastbound circuits to Singapore are in the same SRG). The total capacity of all resources within a SRG hence represents the capacity drop caused by the corresponding threat. The maximal capacity of all SRGs ultimately determines the boundaries of the no-impact zone.

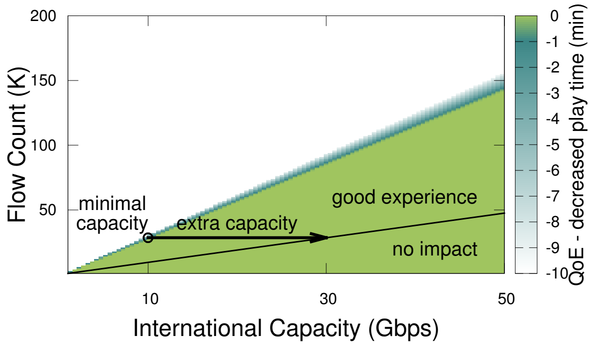

Figure 12: Capacity planning during normal and abnormal conditions.

There are many previous works on how to identify SRGs [14, 37]. For cable cuts specifically, the common SRGs are circuits going through the same cable conduits, the same straits, and the same landing stations. Therefore, countries need to pay special attention to make sure their circuits are diversified over different cables, straits, and destinations.

B. APPLYING TO OTHER INCIDENTS

In addition to Bangladesh, we also apply our model to three other cases as shown in Table 3. We examine if our model can explain other incidents and if low self-sufficiency and geo-diversity are still the major causes (Section B.1). We then discuss the different geographical regions each country is vulnerable to and provide recommendations (Section B.2). Finally, we hypothesize and compare the impact on user QoE in all incidents (Section B.3).

B.1 Generalizing Causes of Large Impact

We have concluded that low service self-sufficiency and low geographic diversity are the two major causes that make a country vulnerable to submarine cable cuts. We believe this conclusion also applies to many other incidents that impact different counties with different connectivity.

We examine three additional incidents to support our conclusion shown in Table 3. Our model provides a consistent explanation for all incidents. Except Pakistan, all significantly impacted countries (capacity drop $> 50\%$) show low self-sufficiency (ranging from 0% to 8%) and low geo-diversity. The reason why Pakistan (with a medium level of geo-diversity) also experienced large impact is because it had two cables out of service during the same period. The challenge Pakistan was facing was much bigger than Bangladesh and Lebanon, however, it still performed better than the other two countries (60% capacity drop compared with 67% and 100% drop).

The four incidents cover a wide range of geographical area (South Asia, Middle East, and Africa) and diverse connectivity, we therefore believe our model is generic enough to apply to many submarine cable cut incidents.

Although all incidents have shown that low geo-diversity is a major vulnerability, countries differ at the location where they are vulnerable to. We next discuss the weakness for each country.

B.2 Geographic Diversity of Different Countries

We next discuss the geographical regions each country is vulnerable to and how they could improve geo-diversity. Table 3 summarizes these geographical regions for each country.

We have learned that Bangladesh heavily relies on its eastbound circuits to Singapore for Internet access in Section 4.3. Thus, the region between Bangladesh and Singapore where SeaMeWe-4 traverses is Bangladesh's weakness. Any earthquake and ship anchors in this region pose threats to Bangladesh. As we have also mentioned in Section 4.3, Bangladesh could improve its geographic diversity by adding circuits or cables to new destinations, such as India, Middle East, and Europe.

Unlike Bangladesh, Pakistan's geographical weakness is westbound to Europe. We infer that westbound circuits through SeaMeWe-4 and all circuits through IMEWE (the other cable out of service) represent 60% of Pakistan's international capacity, and they were either broken or dis-functional during the incident. The heavy circuit provisioning in one direction makes Pakistan vulnerable to threats happening along the westbound routes to Europe, which almost all go through the Suez Canal, the biggest single point of failure of Pakistan's Internet access.

To improve geographic diversity, Pakistan could follow two ways. The first way is to provision more eastbound circuits to Asia. The second one is to establish circuits via different routes, such as through South Africa rather than through Egypt to reach Europe.

Lebanon has the lowest geographic diversity among all countries. We infer that it only provision westbound circuits to France for Internet access. Thus, a threat at any position of the cable route between Lebanon and France can bring the whole country down, which is what happened in the incident [41].

Lebanon could improve its geographic diversity by establishing Internet circuits to some countries other than France, and ideally in other directions. In this way, even if the westbound circuits are broken, Lebanon can still rely on eastbound circuits to avoid a complete Internet blackout.

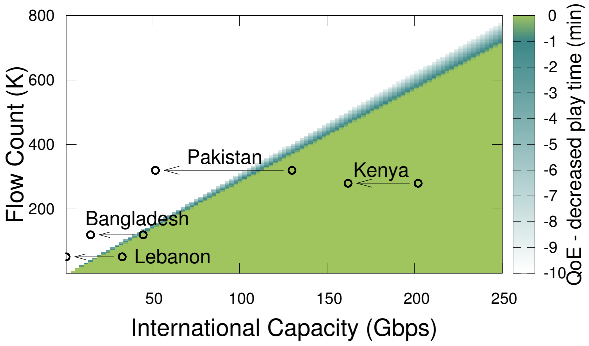

Figure 13: Estimated impact on user QoE in four incidents.

Compared with the previous three countries, we believe Kenya has better geo-diversity. Similar to Pakistan, Kenya also provisions circuits that go through Suez Canal to reach Europe. However, Kenya connects to at least two more destinations via two more routes for Internet access: northbound to United Arab Emirates and eastbound to Asia.

We have shown that geographical weakness is often caused by provisioning circuits heavily in one direction, to one destination, or via one route by examining four countries. This weakness can lead to significant drop of a country's international capacity. We next look at how this capacity drop affects user QoE.

B.3 Impact on User QoE

In this section, we analyze QoE (see Section 4.4) for all four incidents. We compare the potential impact on user QoE in different countries, and see how good connectivity planning could insulate users from cable cut impact. Figure 13 shows our estimation of the position of each country before and after the cable cuts. We gather the approximate international Internet capacity of each country from public web pages [21, 23, 33, 40]. These capacity numbers place each country on the capacity axis. We then position each country on the flow axis based on their number of Internet users. Finally, we decide how far the before state moves by estimated or reported capacity drop listed in Table 3.

From Figure 13, we can see that Kenya has done a good job to insulate its users from submarine cable cuts. It not only has provisioned abundant extra capacity, but it has good geographic diversity (one cable cut has only resulted in 20% capacity drop). Other countries are much closer to the edge of acceptable capacity. All exit the "good experience" region after a single cable cut. These countries do not have adequate extra capacity in intentional connectivity, visualized by the distance between the before state and the good experience zone boundary. Neither do they have good diversity of international capacity as shown by the relative length of the arrow (the shorter, the higher diversity).

From these additional cases, we conclude that a good connectivity planning could insulate users from impact brought by submarine cable cuts.