RiskRoute: A Framework for Mitigating Network Outage Threats

Brian Eriksson
Technicolor Palo Alto
735 Emerson St.
Palo Alto, CA 94301
brian.eriksson@technicolor.com

Ramakrishnan Durairajan
University of Wisconsin-Madison
1210 W Dayton St
Madison, WI 53706
rkrish@cs.wisc.edu

Paul Barford
University of Wisconsin-Madison
1210 W Dayton St
Madison, WI 53706
pb@cs.wisc.edu

ABSTRACT

A comprehensive understanding of outage threats is critical for robust network design and operation, and evaluating cost trade-offs for recovery planning. In this paper, we describe a study of network infrastructure events due to outage events and a framework for mitigating these risks through backup routing and additional provisioning. We evaluate risk via the concept of *bit-risk miles*, the geographically-scaled outage risk of traffic in a network. Our focus on bit-risk miles allows for first-of-its-kind analysis of the tradeoffs of shortest path routing and risk-averse routing. We leverage the concept of bit-risk miles to present *RiskRoute*, a flexible routing framework that allows for backup routes to be configured to respond to both historical and immediately forecasted outage threats. Specifically, RiskRoute is an optimization framework that minimizes bit-risk miles between arbitrary points in a network. RiskRoute also reveals the best locations for provisioning additional network infrastructure in the form of new PoP-to-PoP links for single-network domains, and the best new peering relationships for multi-network domains. To assess and evaluate RiskRoute, we assemble diverse data sets including (*i*) - detailed topological maps and peering relationships of Internet Service Providers (ISPs) in the US, and (*ii*) - historical information on different types of natural disasters which threaten physical infrastructure. Our analysis reveals the providers that have the highest risk to disaster-based outage events. We also provide provisioning recommendations for network operators that can in some cases significantly lower bit-risk miles for their infrastructures.

Categories and Subject Descriptors

C.2.3 [**Computer Systems Organization**]: Computer-Communication Networks - Network Operations

Keywords

Routing, Network Robustness, Network Outages

CoNEXT'13, December 9-12, 2013, Santa Barbara, California, USA.

ACM 978-1-4503-2101-3/13/12 ...$15.00.
http://dx.doi.org/10.1145/2535372.2535385.

1. INTRODUCTION

Outages and loss of connectivity can have a significant impact on businesses and users who depend on communications infrastructure. In recognition of this, high availability is one of the most, if not *the* most, important operational objective of Internet Service Providers. Indeed, five-nines or 99.999% availability[1] is routinely reflected in service level agreements from ISPs and cloud computing providers.

This requirement presents significant design and engineering challenges as network outages can have a number of causes, ranging from accidents (*e.g.,* the Baltimore Howard Street Tunnel fire [1] or Mediterranean Cable Cuts [2]), misconfigurations (*e.g.,* Pakistani Youtube routing [3]), terrorism (*e.g.,* the World Trade Center attack [4] or a potential Electromagnetic Pulse Attack [5]), or censorship (*e.g.,* response to the 2011 Egyptian uprising [6]). One cause of Internet outages is natural weather-related phenomena, such as hurricanes [7], earthquakes [8], or even small-scale thunderstorms (*e.g.,* a recent measurements study [9] showed that thunderstorms result in a three-fold increase in Internet outages). Unique to a majority of weather-related outages is *(i)* - the ability to assess short-term potential threats before they occur, specifically via weather forecast analysis, and *(ii)* - the ability to assess long-term threats to specific geographic areas based on historical records of events and outages.

Our study of outage risk is informed by recent focus on *bit-miles* of Internet traffic. The traffic exchange policy of Level 3 [10] defines bit-miles as "the number of air miles a party's Internet network carries Internet traffic from the source or destination". Extending this concept, we present a detailed characterization of network outage risk via the introduction of *Bit-Risk Miles*.

Definition 1. The **bit-risk miles** of network traffic is the geographic distance traveled by the traffic plus the expected outage risk encountered along the specified routing path.

The bit-risk miles idea allows us to evaluate the trade-offs between routing using a geographically short path (*e.g.,* as might be selected due to Service Level Agreements (SLAs) that are focused on low latency) through infrastructure with high outage risk versus using a geographically longer path and directing traffic through infrastructure with low outage risk.

[1]Specifying downtime will not exceed more than about 26 seconds over a 30 day period.

Leveraging this concept of bit-risk miles, we present a generalized route optimization framework, which we call *RiskRoute*. Our focus is the geographic scope of networks (PoPs and links), which we collectively refer to as the physical infrastructure. The RiskRoute framework can be used to suggest routing changes that will minimize outage risk with respect to both historical and current outage forecasts. While we focus on natural disaster events as threats in this paper, our approach does not preclude other models for threats such as hardware fails or misconfigurations. The framework can be used to identify backup routes or provisioning changes in network infrastructure that will reduce bit-risk miles, and it can be used at both the intra- and inter-domain levels. RiskRoute does not, however, specify the mechanisms for deploying or implementing backup routes, nor does it consider global properties of BGP convergence in the case of inter-domain routing. Rather, our contribution is in the assessment of outage risk and development of a higher level framework for minimizing risk.

While significant prior work has been done on recovering from network outages after they have occurred (*e.g.,* [11, 12], etc.), to the best of our knowledge, no prior studies have examined the notion of proactively identifying routes that avoid areas of high network risk. Motivating this is the fact that preemptive avoidance of high risk areas is already a standard practice of network service providers. For example, the NTT, Level3, and Verizon networks all changed routing before Hurricane Sandy (*i.e.,* changed routing paths to avoid certain risky PoPs) presumingly to mitigate or avoid outages [13]. Unfortunately, these routing changes are currently done by-hand and are unlikely to be comprehensive with respect to the possible range of outage threats. Therefore, current best practices for preemptive routing changes are slow to execute and can only reasonably be expected on the largest of network links and for the most severe outage events. The RiskRoute framework enables both fine adjustments to routing tables in response to immediate threats as well as provisioning changes that can increase the general robustness of an infrastructure.

Our analysis explores the robustness of 23 networks in the continental United States using ground truth PoP location and connectivity information provided by ISPs. Our results point to clear disaster provisioning recommendations to best decrease the bit-risk miles of network traffic, including specific new links between currently unconnected PoPs and new peering relationships between ISPs. In summary, this is the first study to introduce the concept of bit-risk miles, to examine preventive routing based on probable future outage events, and to offer case studies of routing during real-world historical disasters and forecasts.

The remainder of this paper is structured as follows. Relevant prior studies are discussed in both Sections 2 and 3. The data sets used throughout the paper are described in Section 4. The bit-risk miles concept is introduced in Section 5, and exploited in the RiskRoute methodology in Section 6. Our experiments on real-world networks, ranging from small, local networks to nation-wide tier-1 ISPs are explored in Section 7, along with discussion about actionable lessons learned from these experiments. We summarize and describe future work in Section 8.

2. RELATED WORK

Outages in the Internet have been studied from the viewpoint of Internet reachability [14, 15] and detecting/localizing current network outages [11, 16, 17]. Analyzing the robustness of networks has been the focus of studies on network survivability [18], network resilience analysis [19, 20], and the development of resilient routing protocols [12, 21, 22, 23, 24]. Gorman considers the vulnerability of Internet infrastructure to malicious attacks in [25]. In contrast to this prior work on resilient routing, this study is the first to consider preventive routing and the first to use risk analysis to potentially avoid outages before they occur.

The experiments in this paper focus on Internet outage risk from natural disasters[2]. While the greater Internet has shown to be relatively robust to localized disaster events, they can still leave large sections of the population without network access for extended periods of time. This is a consistent theme in multiple studies of major global disasters, including the Fukushima earthquake and tsunami [8], the September 11th attack [4], and Hurricane Katrina [7]. Even small scale weather phenomena can affect network availability, as one study by Schulman *et al.* [9] has shown that thunderstorms result in a three-fold increase in Internet outages.

Individual components of the RiskRoute methodology have been applied in prior studies. For example, we use kernel density estimations [27] to analyze disaster event probability. Kernel density approaches have been used in prior networking-related studies including IP geolocation [28], Point-of-Presence (PoP) identification [29], and network resiliency analysis [20]. With respect to prior work on network resiliency, our work differs in that we analyze historical disaster events, instead of arbitrary disaster likelihoods.

Active probe-based measurement techniques have been proposed for network fault and outage monitoring (*e.g.,* [30, 31, 32]). More recently, measurement-based methods have been proposed for the monitoring the Internet during disasters [33, 34]. Our study can inform the deployment and configuration of these kinds of monitoring efforts in order to make them more efficient and accurate. Data sets from these studies (in conjunction with external event observations) can also be used to help augment the RiskRoute framework to more accurately characterize outage risk for specific networks.

3. BACKUP ROUTES AND PROVISIONING

The focus of our work is to develop methods for mitigating the risk of outage threats to networks and thereby improve their operational availability. From a practical perspective, the issue of improving operational availability – especially in a large network – is complex, thus it is important to be clear on where RiskRoute fits in.

There are two general objectives for RiskRoute: to enable backup routes to be computed and to provide decision support for provisioning new network infrastructure. With respect to backup routes, prior work includes the study by Gao *et al.* [35], which describes an inherently safe model for

[2]While civil engineering literature has also focused on network robustness in the precence of natural disasters (*e.g.,* [26]), their specific focus on transportation networks introduces significant deviations with our study of Internet infrastructure under disasters.

backup routing at the inter-domain level that improves network reliability. RiskRoute complements that work by providing a higher level framework for selecting backup routes that mitigate risk. Furthermore, any routes specified by RiskRoute would require evaluation of *safety* using methods such as those described in [35]. For intra-domain routing, the notion of computing backup paths to address failures has been in practice for some time. For example, RFC 5714 specifies IP Fast Reroute - a mechanism for quickly repairing the effects of location failures without the need to invoke a routing update after the failure. RiskRoute fits very nicely into the IP Fast Reroute framework by offering an algorithm for backup/repair path calculation.

The problem of network provisioning the deployment of new resources or capabilities in a network can be considered in many ways depending on objectives. Standard methods and practices for provisioning vary widely, are typically proprietary and are based on both heuristics and optimization methods (*e.g.,* [36]). In this study, we consider network provisioning from the perspective of PoPs and the links that connect them. While this is certainly a coarse-grained perspective, we argue that it is the most relevant to mitigating risk of outage threats from natural disasters, which is why we adopt it in our study. RiskRoute does not preclude a more fine-grained approach that is more tuned to threats such as hardware failures within PoPs.

3.1 Putting RiskRoute into Practice

An objective of our work is to make the RiskRoute metric useful in practice for the purpose of establishing backup paths and/or alternative routes during outages in disaster scenarios. The most natural way to accomplish this objective is to identify ways in which the metric can be directly incorporated into the routing configuration process. In what follows, we describe general approaches to incorporating RiskRoute in routing configurations, but leave details to future work.

To address robustness to disasters within a single domain, the RiskRoute metric can be used directly in standard intra-domain routing protocols such as OSPF or ISIS. These protocols implement shortest path routing based on link weights. The problem of optimizing link weights for various operational objectives has been well studied in prior work (*e.g.,* [37]). The approach would simply be to create link weights that are a composite metric based on operational objectives and RiskRoute. Alternatively, backup configurations that use a composite link metric that includes RiskRoute can be computed off line following the method described in [38].

Another method for creating backup paths within a domain is based on the use of MPLS tunnels. Specifically, for domains that use MPLS, the *fast reroute* mechanism can be used to establish failover paths for single link or node failures (*e.g.,* [39]). While this is a somewhat limited failure model, the RiskRoute metric can be used directly to identify backup paths that can then be directly configured in a network.

Finally, to address robustness beyond a given domain or in PoPs where peering routers are located, the RiskRoute metric can be used to identify service providers that may be able to offer additional connectivity options. Over long time scales, RiskRoute-based analysis can lead to new provider or peering relationships. Over shorter time scales, RiskRoute could be used in conjunction with the proposed BGP "add paths" option as the basis for inter-domain fast path restoration [40].

4. DATA SETS

One of the key contributions of our work is the consideration of data assembled from primary sources (*i.e.,* ground truth). In the case of network topology, we use the latest maps from service providers. For historical disaster events and forecasts, we use several different US federal government archives. The use of primary source-based data enhances the quality and reliability of our results.

4.1 Networks

The ability to assess threats to real-world networks is dependent on detailed topological maps of Internet Service Providers that include accurate geolocation information. In this study, we rely on the publicly available Internet Topology Zoo [41] and Internet Atlas [42] projects. From these online repositories, we obtain detailed geographic information of 7 Tier-1 networks containing 354 total PoPs and 16 regional networks containing 455 total PoPs in the continental United States. Figure 1-(Left) and Figure 1-(Right) both show the geographic placement of this network infrastructre.

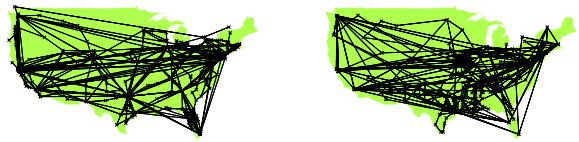

Figure 1: Network and population data sets - (Left) - Tier-1 infrastructure PoP locations and links, (Right) - Regional infrastructure PoP locations and links.

For the connectivity between PoPs, we use line-of-sight to place links. While the actual network fiber bundles do not follow line-of-sight (they are typically located along major thoroughfares such as highways, railways, etc. [43, 44]), their actual paths are often on reasonably direct routes between locations. With respect to between-AS relationships, we examine the CAIDA AS Relationship Dataset [45] which derives AS connectivity information from BGP announcements. Figure 2 shows the AS-level connectivity for the 23 networks that we consider in our study.

We do not argue for the completeness of the network coverage represented in this data corpus. We do, however, argue for the representativeness and accuracy of this data. The omission of certain nodes and links of either national or regional providers means that the density of network infrastructure is under-represented in our analysis, thus the overall risks are may be slightly overstated.

4.2 Population Data

One component of the RiskRoute framework is incorporation of the impact, or effective size, of a particular network outage. For the focus of our study, we examine impact with respect to the estimated underlying population serviced by the network routes. Several studies, including [46], have correlated population density with Internet usage. To evaluate the size of the population serviced by each network resource, we consider information from the United

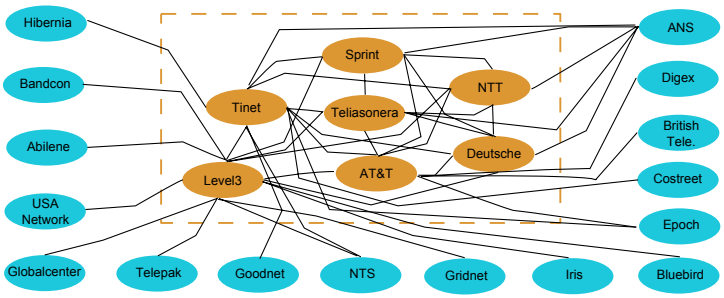

Figure 2: AS connectivity between all networks under consideration.

States Census survey [47]. The data set under evaluation is at the resolution of census-block level, returning population for 215,932 geographic partition regions in the continental US. Using kernel techniques, a heat map of the population (where lighter color indicates higher geo-spatial frequency of population and darker colors indicate lower geo-spatial frequency) from this data is shown in Figure 3-(Left). For each network containing multiple PoP locations, we consider a simple nearest-neighbor population assignment approach, where the population for a given census block is assigned to the nearest infrastructure location. An example of this assignment is shown in Figure 3-(Right) for the Teliasonera network.

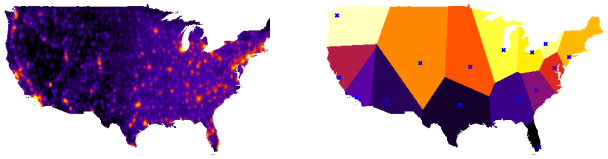

Figure 3: (Left) - Population density of the continental United States, (Right) - Nearest-neighbor assignment for Teliasonera network PoPs, with PoP locations denoted by "x" marker.

4.3 Natural Disaster Events

We also use weather-related disaster data provided by the US Federal Emergency Management Agency (FEMA) [48]. Due to our focus on events that affect Internet infrastructure, we only consider weather events that include severe storms, tornadoes, and hurricanes. For these particular events of interest, we observed 29,865 emergency declarations from FEMA between 1970 and 2010 specified at the United States county-level. Specifically, we find 20,623 severe storm declarations, 6,437 tornado declarations, and 2,805 hurricane declarations. While other events recorded by FEMA, such as volcanic eruptions, may also effect infrastructure, we ignore these events in this study due to their extreme rarity. For completeness, we also consider recorded events by the United States National Oceanic and Atmospheric Administration (NOAA) [49] between the years 1970 and 2010. Specifically, this data contains 143,847 wind damage events, and 2,267 earthquake events.

4.4 Forecasted Disaster Events

Finally, we perform natural language processing to parse NOAA weather forecasts [50]. We examine three weather events, Hurricanes Katrina, Irene, and Sandy, with 61, 70, and 60 public advisory weather forecasts, respectively. Each text forecast includes timestamps and the associated current center of the hurricane and the radius of tropical and hurricane force winds at the specified time. A selected portion of this forecast text for Hurricane Irene reads as follows:

> ...THE CENTER OF HURRICANE IRENE WAS LOCATED NEAR LATITUDE 35.2 NORTH...LONGITUDE 76.4 WEST. IRENE IS MOVING TOWARD THE NORTH-NORTHEAST NEAR 15 MPH...HURRICANE-FORCE WINDS EXTEND OUTWARD UP TO 90 MILES...150 KM...FROM THE CENTER...AND TROPICAL-STORM-FORCE WINDS EXTEND OUTWARD UP TO 260 MILES...415 KM...

By natural language parsing, we translate this text to a measure of forecasted outage risk in Section 5.3.

5. BIT-RISK MILES

For network traffic traveling between two PoPs (say, PoP i and PoP j), our goal is to define a measure of the total outage risk given the current routing. We define the bit-risk miles measure with respect to four properties:

1. *Geographic Distance* - The geographic distance traveled by the network traffic (defined as "bit-miles" in [10]).
2. *Outage Impact* - The estimated impact of an outage between the two PoPs.
3. *Historical Outage Risk* - The historical risk of encountering an outage.
4. *Immediate/Forecasted Outage Risk* - The current risk of encountering an outage.

We define the bit-risk miles as the geographic distance traveled by the packet ($d_{i,j}$) added to the impact scaled ($\gamma_{i,j}$, calculated here using population data) risk of network outage (historical risk, o_h and immediate risk, o_f, calculated here using natural disaster data). Therefore, we define the bit-risk miles, $r_{i,j}(\mathbf{p})$, as,

$$r_{i,j}(\mathbf{p}) = \sum_{x=2}^{K} \left(d_{p_x,p_{x-1}} + \gamma_{i,j} \left(\lambda_h o_h(p_x) + \lambda_f o_f(p_x) \right) \right). \quad (1)$$

Where the routing path between i and j is denoted as $\mathbf{p} = \{p_1, p_2, ..., p_K\}$, with $p_1 = i$ and $p_K = j$, and where p_k is the k-th PoP traversed by the path between PoPs i and j. The tuning parameters, $\lambda_h > 0$ and $\lambda_f > 0$, determine the contribution of historical and immediate outage risk to the bit-risk mile metric, respectively. This allows network operators to directly adjust the risk-averseness of

their network routing. The larger these tuning parameters, the longer the bit-miles of the routing paths and less likely the routing path traverses high outage risk infrastructure. In Section 7, we use the values 10^5 and 10^3 for λ_h and λ_f.

While it can certainly be the case that network physical infrastructure will survive light or moderate weather-related disasters, we do not consider the "hardness" of any infrastructure in our analysis. We simply argue that network physical infrastructure is vulnerable to severe forms of disasters (as previously seen in [4, 7, 8, 9]). In addition, while there exist outage threats that are unsuitable for accurate geospatial modeling (*e.g.,* fires, backhoe fiber cuts, etc.), we argue that focusing on natural disaster outages covers a wide range of potential network threats.

We note that network administrators could easily insert their own intuition about the risk and impact of outages of their network infrastructure. In practice, the outage risk value could also be derived from historical outage frequency information, known ability to recover from outage (*i.e.,* outage duration information), and/or peering information. Meanwhile, the impact of an outage could also be influenced by traffic flows between two PoPs, SLA information, or specific critical peering relationships. We leave this extension of risk characterization up to individual network administrators and offer the natural disaster and population analysis presented here as an instructional case study using the RiskRoute framework.

5.1 Outage Impact

From Census data, we estimate the amount of population serviced by each network resource. A simple technique for assigning population to network infrastructure is a *nearest-neighbor* approach, where the population for a geographic location is assigned to the closest network resource.

Using this nearest-neighbor model, we define c_i as the fraction of population serviced by PoP i, and the estimated impact of an outage between PoPs i and j as, $\gamma_{i,j} = c_i + c_j$. For geographically constrained regional networks, we only consider the population confined to the states where these networks have infrastructure.

5.2 Historical Outage Risk

The historical outage risk can be considered a prior on the likelihood that physical infrastructure (*i.e.,* a PoP) at a specific location encounters an outage. For our study, we construct geo-spatial outage probability estimates from historical geospatial disaster data using nonparametric kernel density estimates. For a set of observed disaster events $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_N\}$, such that $\mathbf{x}_i$ is the latitude/longitude coordinates of event i, we obtain the kernel likelihood probability, $\widehat{p}(\mathbf{y})$, the estimated probability of disaster at latitude/longitude location $\mathbf{y}$, where,

$$\widehat{p}(\mathbf{y}) = \frac{1}{\sigma N} \sum_{i=1}^{N} K\left(\frac{(\mathbf{x}_i - \mathbf{y})}{\sigma}\right) \quad (2)$$

We use a Gaussian kernel, where, $K(\mathbf{z}) = \frac{1}{2\pi} \exp\left(\frac{-1}{2}\mathbf{z}^T\mathbf{z}\right)$. While we acknowledge that many of the disaster events have strong seasonal correlations (*e.g.,* tornados, hurricanes), for simplicity, here we only consider a single outage probability distribution for each disaster event type.

One of the benefits of this approach is that for each disaster likelihood, the only tuning parameter is the bandwidth (or variance) of the kernel function, σ, the geospatial contribution of each event in our dataset. By properly setting this bandwidth value, we avoid both overfitting or underfitting to the observed historical events. To determine the optimal bandwidth value, we use 5-way cross validation [27] (where the best bandwidth is found from 80% of the observed events to fit the remaining 20%). The distance metric we consider is the KL divergence [51]. The optimized kernel bandwidth values are shown in Table 1 for all seven disaster data sets. Intuitively, these bandwidth values return the level of confidence in prediction for each event type, ranging from events with the highest predictive resolution (*i.e.,* NOAA wind data) to events with the lowest resolution (*i.e.,* NOAA earthquake data). This value is, of course, dependent on the number of historical events in our data set.

Table 1: Trained kernel density bandwidths for FEMA and NOAA data.

Event Type	Number of Entries	Optimal Kernel Bandwidth
FEMA Hurricane	2,805	71.56
FEMA Tornado	6,437	59.48
FEMA Storm	20,623	24.38
NOAA Earthquake	2,267	298.82
NOAA Wind	143,847	3.59

Using the optimized kernel bandwidths, we construct the kernel density estimates from our database of natural disaster events. Figure 4 shows the resulting geo-spatial likelihoods for the three FEMA data sets and two NOAA data sets. As expected, hurricanes are more prevalent along the Gulf Coast region, while severe storms are prevalent in the central plain states, and earthquakes dominate the west coast.

From Equation 1, we distill the historical outage risk down to a single risk value for every router PoP location. We consider the aggregated historical risk to be the sum of all five outage probabilities, where for PoP location i the aggregate risk, $o_h(i)$, is defined as the sum of all outage probabilities. We note that individual events that network operators find to be particularly disruptive for network service (*e.g.,* flooding events for network infrastructure that lies on the first floor of a building, etc.) could be emphasized using this risk metric calculation via user-defined weights. We leave these network specific extensions to future work.

5.3 Forecasted Outage Risk

While the use of NOAA and FEMA data allows us to assess the historical risk of a disaster event at a specified physical infrastructure location (*i.e.,* a PoP), there remains the issue of determining the *current* risk to the network. Events like tornados, hurricanes, and even earthquakes have a time delay between knowledge of likely event occurrence and impact. We examine tick-by-tick disaster forecasts, in terms of both the disaster *location* and *scope*, to simulate observing outage risks as disaster events happen.

In this study, we focus on historical records of hurricane events to assess forecast outage risk. Using the NOAA forecast text data described in Section 4.4, we parse this text corpus to extract both the current center and intensity of

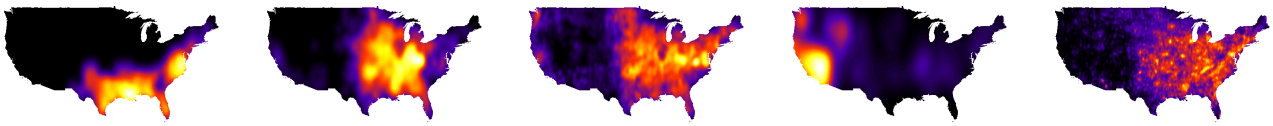

(A) (B) (C) (D) (E)

Figure 4: Bandwidth-optimized kernel density estimates of NOAA and FEMA data from 1970 - 2010. (A) - Hurricane likelihood, (B) - Tornado likelihood, (C) - Storm likelihood, (D) - Earthquake likelihood, (E) - Damaging wind likelihood.

hurricane disaster events. Disaster-dependent measures of risk are then defined given these two properties.

Specifically for hurricane events, we construct forecasted risk likelihood using the two extracted disaster data points, the radius of tropical-force wind and the radius of hurricane-force wind. We declare the forecasted risk of an area under tropical-force wind as γ_t, and the risk of an area under hurricane-force winds as γ_h, with the probability of outage given hurricane-force winds is higher than the outage probability for tropical-storm force winds, $\gamma_h > \gamma_t$ (in Section 7 we use $\gamma_t = 50$ and $\gamma_h = 100$). An example weather forecast risk is shown in Figure 5 for Hurricane Irene given our dataset of advisory-dervied knowledge of the geographic area of the disaster. As seen in the figure, at different time-steps the current geographic impact-region of the storm is tracked.

Figure 5: Geo-spatial Disaster Forecast for Hurricane Irene where the lighter blue area represents tropical storm force winds and the darker red area represents hurricane force winds. - (Left) - Forecast for 11:00 AM 8/25/2011, (Center) - 5:00 PM 8/26/2011, and (Right) - Forecast for 8:00 AM 8/28/2011.

The final geo-spatial scope of all three hurricane events considered in this paper can be seen in Figure 6. We remark that the varying geographic location of these disasters allows us to more fully evaluate the performance of network routing during disasters.

6. RISKROUTE METHODOLOGY

A consequence of the bit-risk miles concept is the ability to estimate the risk for specified routes through the network's physical infrastructure. This leads to two questions. *(1)* - For two PoPs, can we determine the routing through the infrastructure that has the minimum bit-risk miles? and *(2)* - What additional infrastructure (with respect to network connectivity) would best decrease the outage risk in the network? The RiskRoute methodology answers these two questions for both the intradomain case (*i.e.,* routing through a single ISP) and the interdomain case (*i.e.,* routing through multiple ISPs).

6.1 Intradomain RiskRoute

Consider a single ISP network with physical infrastructure consisting of N PoPs. In this regime, for a given pair of PoPs, RiskRoute examines the physical infrastructure (the set of PoP nodes and links) to find the path through the network such that these two specified PoPs are connected and the bit-risk miles associated with this path is minimized. We define this RiskRoute optimized path between PoPs i and j, $\mathbf{p}^{rr}_{i,j}$, as,

$$\mathbf{p}^{rr}_{i,j} = \arg \min_{\mathbf{p} \subset \mathbb{P}_{i,j}} r_{i,j}(\mathbf{p}). \quad (3)$$

Where $r_{i,j}(\mathbf{p})$ is defined in Equation 1 and $\mathbb{P}_{i,j}$ is the set of all possible paths through the network between PoPs i and j. We discuss and give examples of how the tuning parameter values affect RiskRoute routing paths in Section 7.

6.2 Interdomain RiskRoute

We also focus on the more challenging problem of assessing bit-risk miles when a packet is routed through multiple networks. In contrast to the single-domain case, when considering multiple networks we encounter the issue that we do not have control over the routing of traffic in other networks. Because of this, we characterize multi-network bit-risk miles with respect to two factors: shortest-path routing throughout all peering networks between the two PoPs and the RiskRoute path given all possible PoPs.

First, consider finding the minimum geographic distance routing (or shortest path) between all considered networks. While other routing schemes, such as hot potato routing [52], may result in worse bit-risk miles routing paths, we consider the shortest path approach to reveal an upper bound on the bit-risk miles of a reasonable routing path between the two PoPs.

Second, if we have the ability to control every routing decision in every network, then we can consider using the RiskRoute methodology to find the best routing path that minimizes the bit-risk of packets sent between the two PoPs. Given this idealized case, we consider this a lower bound on the bit-risk miles between the two PoPs. The ratio between these upper and lower bit-risk miles bounds is considered in Section 7.

6.3 Robustness Analysis

In addition to the optimal risk-averse routing, the RiskRoute framework can also reveal how to augment existing networks to best reduce the bit-risk miles of traffic.

For the intradomain case, consider the ability to add a single additional link to a specified network. We consider the subset $\mathcal{E}^C$, the collection of all links that currently do not appear in the network and do not overlap greatly with

Figure 6: Final geo-spatial scope of historical disaster events considered where the lighter blue area represents tropical storm force winds and the darker red area represents hurricane force winds. (Left) - Hurricane Irene, (Center) - Hurricane Katrina, (Right) - Hurricane Sandy.

existing links[3]. Given our RiskRoute framework of bit-risk analysis, this robustness analysis consists of finding the edge, $e \in \mathcal{E}^C$, that results in the largest reduction of bit-risk miles throughout all possible paths in the network,

$$\widehat{e} = \arg \min_{e \in \mathcal{E}^C} \sum_{i=1}^{N} \sum_{j=i+1}^{N} \min_{\mathbf{p} \subset \mathbb{P}_{i,j}} r_{i,j}(\mathbf{p}). \quad (4)$$

Of course, instead of a single link, we may be more interested in finding the impact of adding some number of additional links. When this occurs, we can consider a greedy methodology, where to add the k-th best additional link we examine the network with $k-1$ added links and minimize the total aggregated bit-risk miles in the current network.

While the single-domain case allows to consider adding additional links to the infrastructure, in the multi-domain case we will not have this ability. Therefore, we examine the best possible new peering or additional multihoming egress point to minimize the RiskRoute measure of the lower bit-risk mile bound. For each specified network, we define "candidate peers" as the collection of PoPs in other networks which are co-located with infrastructure from the specified network, but for which there is no previously known peering relationship. Then, the best candidate peer is found such that the RiskRoute paths have the smallest lower-bound bit-risk miles.

6.4 Optimization and Computational Complexity

For each route, solving for the best RiskRoute path in Equation 3 requires constructing a graph structure where the node are PoPs and the link weights consist of the bit-risk miles between infrastructure locations. Using this constructed risk graph, to find the minimum bit-risk miles route consists of solving a shortest-path problem between the two specified PoPs.

The RiskRoute framework does not consider other objectives in routing such as SLAs which are a central consideration in ISP network configuration and management. However, the RiskRoute framework could easily be expanded to include multiple objective functions that would balance risk and SLA-related issues such as latency in route calculations. The impact would be additional computational complexity in route calculation.

[3]Here we consider only links that would results in >50% reduction in bit-miles between the two PoPs. This eliminates impractical cross-country links from consideration.

7. EXPERIMENTS

We now present an evaluation of the RiskRoute framework on real-world networks and historical disaster outage risk. The first set of experiments examines the application of RiskRoute to both tier-1 and regional networks in the continental United States. The second experiment evaluates network robustness with respect to network link structure. We conclude with case studies of RiskRoute performance during real-world hurricane disaster events. Our goal is to highlight networks that are particularly vulnerable to outages and supply network operators with provisioning recommendations.

The raw bit-risk mile numbers may be difficult to interpret and highly dependent on tuning parameters. As a result, we present the results as a series of "ratio values" with comparison to the shortest path routing over the same network. The *risk reduction ratio* is defined as the fractional decrease of the average bit-risk miles for RiskRoute compared with the average bit-risk miles of shortest path routing. Without complete knowledge of each ISP's routing table, we consider shortest-path routing a reasonable approximation of true network routes.

Evaluating with respect to intradomain routing, for a network with N PoPs, the risk reduction ratio is defined as,

$$r_r = 1 - \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \frac{r\left(\mathbf{p}_{i,j}^{rr}\right)}{r\left(\mathbf{p}_{i,j}^{shortest}\right)}. \quad (5)$$

Where $p_{i,j}^{rr}$ is the RiskRoute path evaluated using Equation 3, and $\mathbf{p}_{i,j}^{shortest}$ is the shortest path between PoPs i and j through the topology. For example, a risk reduction ratio of 0.2 implies that using RiskRoute reduces the bit-risk miles of a routing path by 20% compared with shortest path routing.

The *distance increase ratio*, d_r, is defined as fractional increase in average bit-miles for RiskRoute paths compared with the average bit-miles of shortest path routing paths. Given $d(\mathbf{p})$ is the bit-miles length of routing path $\mathbf{p}$, we define the distance increase ratio as,

$$d_r = \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \frac{d\left(\mathbf{p}_{i,j}^{rr}\right)}{d\left(\mathbf{p}_{i,j}^{shortest}\right)} - 1. \quad (6)$$

For example, a distance increase ratio of 0.2 implies that using RiskRoute increases the bit-miles of a routing path by 20% compared with shortest path routing.

In the interdomain regime, the distance and risk reduction ratios for a specific regional network is evaluated such that each PoP in the regional network is considered a path source, and path destinations is the set of all PoPs in the

collection of 16 regional networks. This allows us to analyze the performance of routing between regional networks and through tier-1 peers.

We believe that these aggregated ratio values demonstrate the benefits of RiskRoute and clearly state the trade-offs of its performance against shortest path routing.

7.1 Routing Analysis

We begin by analyzing the bit-risk miles of Tier-1 ISPs in the continental United States. Using no defined forecasted risk and the historical outage risk described in Section 5, we plot two routes (RiskRoute and shortest path) between the Houston, TX and Boston, MA PoPs in the Level3 tier-1 ISP in Figure 7. The figure shows that as the tuning parameter λ_h grows, the routing become more risk averse and therefore deviates significantly from the shortest path routing.

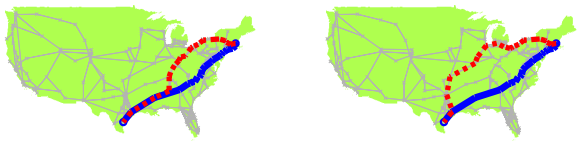

Figure 7: RiskRoute applied to the Level3 Network topology between Houston, TX and Boston, MA PoPs. Routing with the solid line represents the constant shortest path route and the dotted line represents RiskRoute inferred routing. - (Left) - $\lambda_h = 10^4$, (Right) - $\lambda_h = 10^5$

To assess the trade-offs in using RiskRoute we show in Table 2 the average deviation in both bit-risk miles and bit-miles for intradomain routing in the seven tier-1 networks using RiskRoute. As expected, increasing the tuning parameter (*i.e.,* requesting more risk-averse routes) results in both smaller average bit-risk miles and larger bit-mile routing necessary to avoid outage risk. We find that the much larger Level3 network results in the smallest risk reduction ratio (*i.e.,* smallest bit-risk miles improvement using RiskRoute over shortest path routing), while the smaller Deutsche Telekom (*i.e.,* DT) and Teliasonera networks have the largest risk reduction ratios across the various tuning parameter values.

Next, we focus on the avoidance of outage risk in interdomain routing for 16 smaller, regional networks. We expect outage risk to distribute itself differently when network resources are confined to a smaller geographic area (compared with nationwide tier-1 networks). We present the results using RiskRoute for a single historical outage tuning parameter value ($\lambda_h = 10^5$) in Figure 8. This plot shows that most of the regional networks considered have distance increase ratios similar to the resulting risk reduction ratio (*i.e.,* the reduction in bit-risk miles over shortest path routing is equivalent to the inflation of bit-miles miles for the RiskRoute defined paths), such as Iris, USA Network, and Epoch. For a subset of networks, we find that the bit-risk reduction ratio is significantly decreased with a relatively minor increase in the distance increase ratio. For example, the Digex, Gridnet, Hibernia and Bandcon networks all find a roughly 20% measured decrease in bit-risk using RiskRoute (compared with shortest path), while the resulting routing paths only increase the bit-miles by about 10% on average. These results point to the specific regional networks which would have the largest benefit for using RiskRoute.

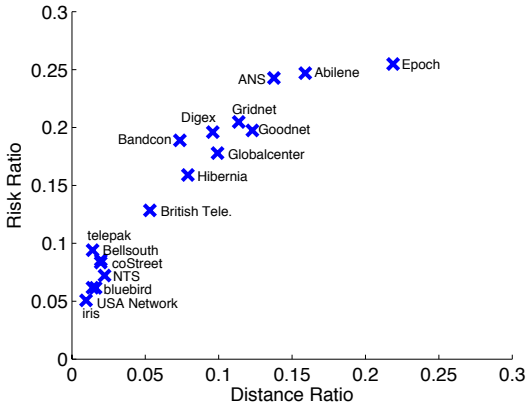

Figure 8: Interdomain RiskRoute experiments. Comparison of distance increase ratio and risk reduction ratio for regional networks.

7.1.1 *Network Characteristic Study*

The characteristics of these regional networks, such as the number of PoPs and geographic scope of the network, affect the performance of RiskRoute. Six network characteristics are presented in Table 3 and presented with respect to the R^2 coefficient of determination [27], the quality of the data to a linear trend (where $R^2 = 0$ indicates that no trend exists and $R^2 = 1$ indicates that the data is completely described by a linear trend). As the table shows, some characteristics, like number of PoPs and geographic footprint size of the network (taken as the largest geographic distance between two PoPs in each network), are very well correlated with the observed risk reduction ratio measurements. Meanwhile, other characteristics, like number of peering relationships and average outdegree of network PoPs, have relatively little correlation with our results. The lack of correlation of risk reduction ratios and the average PoP risk of network infrastructure may seem at first counter-intuitive, but we note that we are comparing performance of RiskRoute against shortest path routing, meaning that any unavoidable outage risk at network PoPs will be canceled out.

Table 3: Regional Network routing performance coefficient of determination (R^2) with respect to network characteristics.

Network Characteristic	Risk Reduction Ratio R^2	Distance Increase Ratio R^2
Geographic Footprint	0.618	0.243
Average PoP Risk	0.104	0.064
Average Outdegree	0.116	0.106
Number of PoPs	0.552	0.405
Number of Links	0.531	0.361
Number of Peers	0.155	0.002

Table 2: Tier-1 Networks Analysis of Bit-Risk to Bit-Miles using RiskRoute.

		$\lambda_h = 10^5$		$\lambda_h = 10^6$	
Network Name	# PoPs	Risk Reduct. Ratio	Distance Incr. Ratio	Risk Reduct. Ratio	Distance Incr. Ratio
Level3	233	0.075	0.015	0.258	0.136
AT&T	25	0.207	0.045	0.340	0.168
DT	10	0.245	0.130	0.384	0.446
NTT	12	0.187	0.040	0.295	0.127
Sprint	24	0.222	0.079	0.352	0.191
Tinet	35	0.177	0.045	0.347	0.195
Teliasonera	15	0.223	0.068	0.336	0.226

7.2 Robustness Analysis

Using the methodology described in Section 6.3, we use the RiskRoute methodology to resolve additional links for a target network that best reduce the total aggregated intradomain bit-risk miles. In Figure 9 we show the ten best additional links for three of the tier-1 networks presented as the fraction with respect to the original network's RiskRoute bit-risk miles. As expected, these suggested links best add connectivity to avoid areas of high outage risk.

Not all networks are equal in terms of their ability to decrease risk via new link infrastructure. In Figure 10, we show how the aggregated bit-risk miles decay given additional links added to the tier-1 networks. We find that the Level3 network, with its high level of existing connectivity, has the least improvement given additional links. Meanwhile, the networks with less existing connectivity, such as Sprint and TeliaSonera, show a marked improvement by adding only a few additional links.

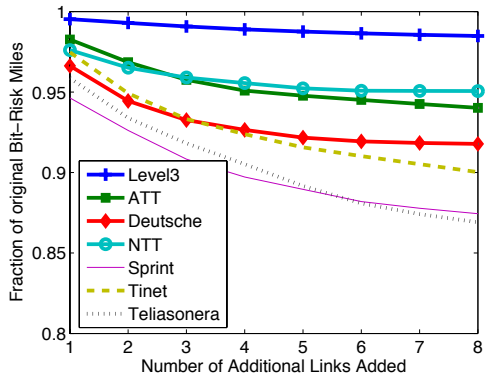

Figure 10: Estimated risk reduction with added links.

With respect to interdomain studies, we can analyze which additional peering links will best improve connectivity. For the entire corpus of networks considered, we find the best additional peering relationships is presented graphically in Figure 11. Our analysis suggests that a majority of the regional networks shown choose to peer with either the AT&T or the Tinet tier-1 networks in order to most effectively reduce the risk of outages. These results mirror the ratio performance of these tier-1 ISPs in Table 2.

7.3 Disaster Case Studies

Finally, we examine historical disasters and evaluate the performance of RiskRoute during three hurricane disasters

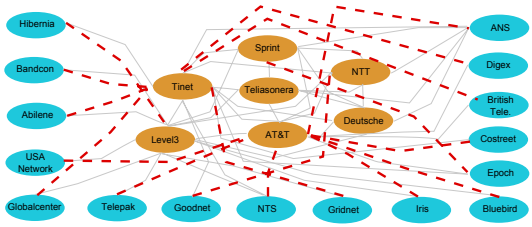

Figure 11: Robust Experiments - The best additional peering relationship (indicated by dotted red links) for each regional network.

(Irene, Katrina, and Sandy). The complete geographic scope of these three disasters can be seen in Figure 6. For each of these three events, we evaluate both shortest-path and RiskRoute routing in our corpus of networks for the duration of National Weather Service advisory forecasts in our dataset[4].

First we examine the tier-1 physical infrastructure in the path of these extreme weather events. With respect to hurricane force winds, we find 86 PoPs for Irene, 8 PoPs for Katrina, and 115 PoPs for Sandy. Analyzing the bit-risk miles reduction ratio with respect to the individual weather forecasts observed, the time-series results on intradomain routing of tier-1 networks is shown in Figure 12. An initial high-level observation is that the risk reduction ratio is relatively small for the Katrina event, compared with much larger risk reduction ratios for the Hurricane Sandy and Hurricane Katrina events. This can be explained by the relatively little infrastructure affected by Hurricane Katrina, with a larger percentage of PoPs encountering damaging weather from Hurricane Sandy and Hurricane Irene, therefore the risk reduction ratio (the reduction in bit-risk miles using RiskRoute) is intuitively smaller.

We now examine each hurricane event individually, starting with the tier-1 networks shown under the Hurricane Irene event in Figure 12-(Left). We found that the Level3 network has the largest number of PoPs in the scope of this event, while the Sprint network has the largest percentage of their network PoPs in the Irene's scope. As the figure demonstrates, the raw number of network PoPs in the path of the

[4]Specifically, we use all NOAA advisory reports [50] from 11:00 AM EDT Monday October 22th 2012 to 1100 PM EDT Monday October 29th 2012 for Hurricane Sandy, 5 PM EDT Tuesday August 23rd 2005 to 10 AM CDT Tuesday August 30th 2005 for Hurricane Katrina, and 700 PM EDT Saturday August 20th 2011 to 1100 PM EDT Sunday August 28th 2011 for Hurricane Irene.

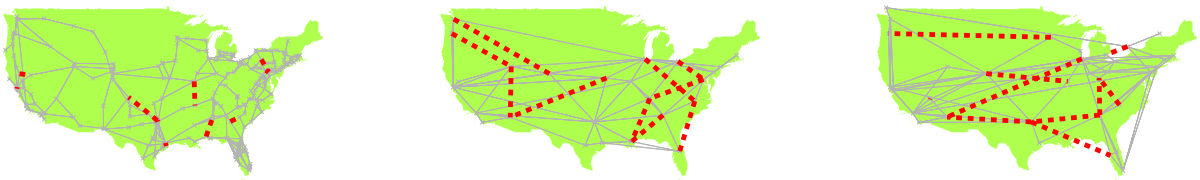

(A) (B) (C)

Figure 9: Tier-1 Network RiskRoute Robustness Suggestions. The 10 best additional links found using the RiskRoute methodology. Solid blue lines represent existing links, dotted red lines represent suggested robust additional links. (A) Level3 Network, (B) AT&T Network, (C) Tinet Network.

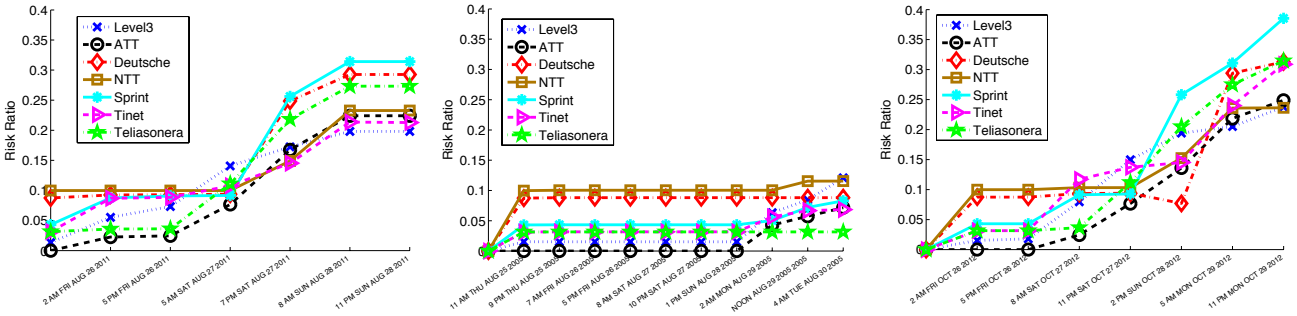

Figure 12: Tier-1 network case study (Left) - Hurricane Irene, (Center) - Hurricane Katrina, (Right) - Hurricane Sandy.

storm has relatively little correlation with the risk reduction ratio performance, as Level3 actually has the smallest decrease in bit-risk miles when considering the entire scope of the hurricane event (*i.e.,* the last observation on Saturday, August 27th). Meanwhile, the Sprint network, with the highest percentage of PoPs affected, results in the largest improvement using RiskRoute.

Different observations can be made for the Hurricane Katrina event in Figure 12-(Center). Again, we find that the network with the highest percentage of infrastructure in the scope of the hurricane, NTT, shows the largest improvement using RiskRoute. Most other networks (*e.g.,* AT&T, Tinet, Teliasonera) show relatively little effect or need for RiskRoute during Hurricane Katrina. We believe that this confirms real-world observations of network availability during Katrina in [7], which showed the affects of this weather event were relatively localized to the storm's geospatial scope.

Finally, we consider the recent Hurricane Sandy event on the United States east coast in Figure 12-(Right). We find that all networks have a significant improvement in bit-risk miles using the RiskRoute framework. These results are reinforced by the wide-spread outages encountered after Hurricane Sandy [53].

7.3.1 *Interdomain Regional Network Results*

The interdomain routing performance of regional networks is shown in Figure 13. Because many of the regional networks in our corpus contain no locations in the scope of the three hurricane events, for each event we only consider the regional networks that include more than 20% of their PoPs in locations contained in the scope of each event. Similar to the Tier-1 network experiments, we again find that Katrina has a relatively minor affect on these networks in comparison to Sandy and Irene. We also notice a much larger deviation in RiskRoute performance as the event persists, with some networks showing very large improvements upwards of 40% (with respect to reduction in bit-risk miles), while other show minor improvements of around 10%.

This is highly dependent on the fraction of the infrastructure in the path of each event. For the Hurricane Irene event, we find that the network with the largest improvement, Digex, has a relatively smaller percentage, 22.2% of their PoPs in the path of the hurricane, while the network with the smallest improvement, Globalcenter, has a vast majority, 87.5% of their PoPs in the path of the hurricane. This suggests that RiskRoute returns the largest improvement when a majority of the infrastructure is not under outage risk and therefore can be reliably used to reroute network traffic through areas of low outage risk.

8. CONCLUSIONS

High availability is a central focus in network design, provisioning and operations. This paper introduces the concept of bit-risk miles, the outage risk weighted distance of network routes. We use bit-risk miles to develop RiskRoute, a generalized framework that minimizes the bit-risk miles of routes in a network infrastructure. RiskRoute can be used to inform network operators with respect to backup paths, route changes, provisioning recommendations, and new peering connections which minimize the risk of outage threats. To assess and demonstrate the capabilities of RiskRoute, we use maps of US ISPs, population and weather related events. The results of our analyses highlight current risks of network infrastructures and how those risks can, in some cases, be significantly mitigated using RiskRoute recommendations. Our future work includes ex-

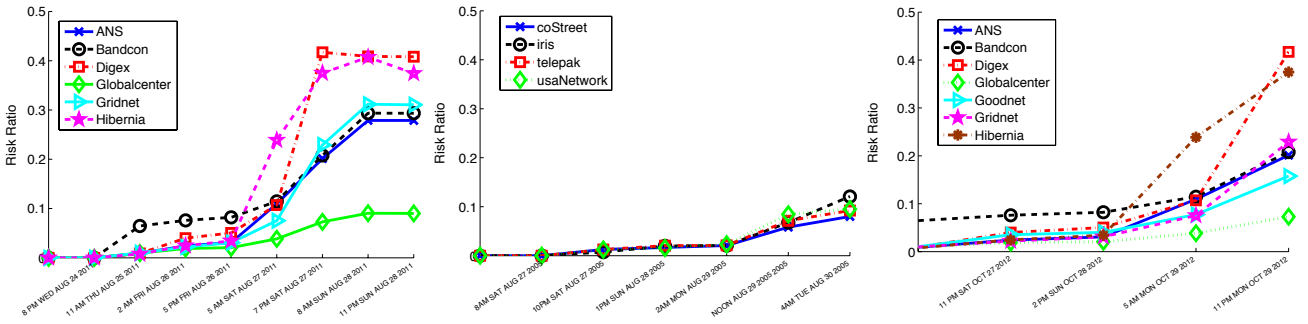

Figure 13: Regional network case study (Left) - Hurricane Irene, (Center) - Hurricane Katrina, (Right) - Hurricane Sandy.

tending the RiskRoute framework to consider operational objectives such as SLAs, assessing shared risk between multiple ISPs using RiskRoute, extending our case studies to specific networks with known outage events, and making this RiskRoute analysis framework openly available to the community.

9. ACKNOWLEDGMENTS

The authors would like to thank our shepherd Ken Calvert for his helpful reviews and comments. This work was supported in part by NSF grants CNS-0831427, CNS-0905186, ARL/ARO grant W911NF1110227 and the DHS PREDICT Project. Any opinions, findings, conclusions or other recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF, ARO or DHS.

10. REFERENCES

[1] K. McGrattan and A. Hamins, "Numerical Simulation of the Howard Street Tunnel Fire," *Fire Technology*, vol. 42, pp. 273–281, October 2006.

[2] E. Zmijewski, "Mediterranean cable break," *Renesys Blog*, January 2008.

[3] P. Hunter, "Pakistan YouTube Block Exposes Fundamental Internet Security Weakness," *Computer Fraud and Security*, no. 4, pp. 10 – 11, April 2008.

[4] N. R. C. Committee on the Internet Under Crisis Conditions: Learning from the Impact of September 11, *The Internet Under Crisis Conditions:Learning from September 11*. The National Academies Press, 2003.

[5] J. Foster, E. Gjelde, W. Graham, R. Hermann, H. Kluepfel, R. Lawson, G. Soper, L. Wood, and J. Woodard, "Report of the Commission to Assess the Threat to the United States from Electromagnetic Pulse (EMP) Attack," *Critical National Infrastructures Report*, vol. 1, April 2004.

[6] A. Dainotti, C. Squarecella, E. Aben, K. Claffy, M. Chiesa, M. Russo, and A. Pescapè, "Analysis of Country-wide Internet Outages Caused by Censorship," in *Proceedings of ACM IMC*, November 2011.

[7] J. Cowie, A. Popescu, and T. Underwood, *Impact of Hurricane Katrina on Internet Infrastructure*. Renesys Corporation Report, 2005.

[8] 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 ACM Special Workshop on Internet and Disasters (SWID)*, December 2011.

[9] A. Schulman and N. Spring, "Pingin' in the Rain," in *Proceedings of the ACM IMC*, November 2011.

[10] L. 3, *IP Traffic Exchange Policy*, 2012. [Online]. Available: http://www.level3.com/en/legal/ip-traffic-exchange-policy/

[11] E. Katz-Bassett, C. Scott, D. R. Choffnes, I. Cunha, V. Valancius, N. Feamster, H. V. Madhyastha, T. E. Anderson, and A. Krishnamurthy, "LIFEGUARD: Practical Repair of Persistent Route Failures," in *Proceedings of ACM SIGCOMM Conference*, August 2012.

[12] H. Wang, Y. R. Yang, P. H. Liu, J. Wang, A. Gerber, and A. Greenberg, "Reliability as an Interdomain Service," in *Proceedings of ACM SIGCOMM Conference*, August 2007.

[13] D. Madory, *Hurricane Sandy: Global Impacts*. Renesys Blog, 2012. [Online]. Available: http://www.renesys.com/blog/2012/11/sandys-global-impacts.shtml

[14] R. Bush, O. Maennel, M. Roughan, and S. Uhlig, "Internet Optometry: Assessing the Broken Glasses in Internet Reachability," in *Proceedings of the ACM IMC Conference*, November 2009.

[15] E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. Anderson, "Studying Black Holes in the Internet with Hubble," in *Proceedings of USENIX NSDI Conference*, April 2008.

[16] L. Quan, J. Heidemann, and Y. Pradkin, "Detecting Internet Outages with Precise Active Probing (extended)," in *USC Technical Report*, February 2012.

[17] E. Glatz and X. Dimitropoulos, "Classifying Internet One-way Traffic," in *Proceedings of ACM ICM*, New York, NY, USA, 2012.

[18] P. E. Heegaard and K. S. Trivedi, "Network Survivability Modeling," vol. 53, no. 8, June 2009, pp. 1215–1234.

[19] J. Wu, Y. Zhang, Z. M. Mao, and K. G. Shin, "Internet Routing Resilience to Failures: Analysis and Implications," in *Proceedings of ACM CoNEXT Conference*, December 2007.

[20] P. Agarwal, A. Efrat, S. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "The Resilience of WDM Networks to Probabilistic Geographical Failures," in *Proceedings of IEEE INFOCOM Conference*, April 2011.

[21] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient Overlay Networks," in *Proceedings of the ACM SOSP Conference*, October 2001.

[22] Y. Zhu, A. Bavier, N. Feamster, S. Rangarajan, and J. Rexford, "UFO: A Resilient Layered Routing Architecture," *SIGCOMM CCR*, vol. 38, no. 5, pp. 59–62, 2008.

[23] A. F. Hansen, A. Kvalbein, T. Cicic, and S. Gjessing, "Resilient Routing Layers for Network Disaster Planning," in *Proceedings of the International Conference on Networking*, 2005.

[24] K. P. Gummadi, H. V. Madhyastha, S. D. Gribble, H. M. Levy, and D. Wetherall, "Improving the Reliability of Internet Paths with One-hop Source Routing," in *Proceedings of USENIX OSDI Conference*, December 2004.

[25] S. Gorman, *Networks, Security And Complexity: The Role of Public Policy in Critical Infrastructure Protection*. Edward Elgar, 2005.

[26] H. Sakakibara, Y. Kajitani, and N. Okada, "Road network robustness for avoiding functional isolation in disasters," *Journal of Transportation Engineering*, vol. 130, no. 5, pp. 560–567, 2004.

[27] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer, 2001.

[28] B. Eriksson, P. Barford, , J. Sommers, and R. Nowak, "A Learning-based Approach for IP Geolocation," in *Proceedings of Passive and Active Measurements Conference*, April 2010.

[29] A. H. Rasti, N. Magharei, R. Rejaie, and W. Willinger, "Eyeball ASes: From Geography to Connectivity," in *Proceedings of ACM IMC*, November 2010.

[30] J. Sommers, P. Barford, N. Duffield, and A. Ron, "Network Performance Anomaly Detection and Localization," in *Proceedings of IEEE INFOCOM Conference*, August 2009.

[31] Y. Bejerano and R. Rastogi, "Robust Monitoring of Link Delays and Faults in IP Networks," in *Proceedings of IEEE INFOCOM Conference*, April 2003.

[32] C. D. A. Dhamdhere, R. Teixeira and C. Diot, "NetDiagnoser: Troubleshooting Network Unreachabilities Using End-to-end Probes and Routing data," in *Proceedings of ACM CoNEXT Conference*, December 2007.

[33] Z. S. Bischof, J. S. Otto, and F. Bustamante, "Distributed Systems and Natural Disasters: BitTorrent as a Global Witness," in *Proceedings of the ACM Special Workshop on Internet and Disasters (SWID)*, December 2011.

[34] J. Li and S. Brooks, "I-seismograph: Observing and Measuring Internet Earthquakes," in *Proceedings of IEEE INFOCOM Conference*, April 2011.

[35] L. Gao, T. Griffin, and J. Rexford, "Inherently Safe Backup Routing with BGP," in *Proceedings of IEEE INFOCOM '01*, Anchorage, AK, April 2001.

[36] I. Houidi, W. Louati, W. Ameur, and D. Zeghlache, "Virtual Network Provisioning Across Multiple Substrate Networks," *Computer Networks*, vol. 55(4), March 2011.

[37] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS weights in a changing world," *IEEE Journal of Selected Areas of Communications*, vol. 20, no. 4, May 2002.

[38] A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, and O. Lysne, "Fast IP Network Recovery using Multiple Routing Configurations," in *Proceedings of IEEE INFOCOM*, Barcelona, Spain, April 2006.

[39] P. Pan, G. Swallow, and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels," http://www.ietf.org/rfc/rfc4090.txt, 2005.

[40] D. Walton, A. Retana, and E. Chen, "Advertisement of Multiple Paths in BGP," http://tools.ietf.org/html/draft-ietf-idr-add-paths-08, 2005.

[41] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The Internet Topology Zoo," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 29, pp. 1765–1775, October 2011. [Online]. Available: http://www.topology-zoo.org

[42] R. Durairajan, S. Ghosh, X. Tang, P. Barford, and B. Eriksson, "Internet atlas: a geographic database of the internet," in *Proceedings of ACM HotPlanet*, 2013, pp. 15–20. [Online]. Available: http://atlas.wail.wisc.edu/

[43] "Sprint Nextel Network History," http://en.wikipedia.org/wiki/Sprint_Nextel, 2013.

[44] "The National Broadband Plan," http://www.broadband.gov/plan/, 2013.

[45] "The CAIDA AS Relationships Dataset," June 2012. [Online]. Available: http://www.caida.org/data/active/as-relationships/

[46] John B. Horrigan, *Broadband Adoption and Use in America*. Federal Communications Commission, 2010.

[47] "U.S. Census Bureau." [Online]. Available: http://www.census.gov/

[48] "Federal Emergency Management Agency, Disasters and Maps." [Online]. Available: http://www.fema.gov/hazard/

[49] "National Oceanic and Atmospheric Administration, GIS Data." [Online]. Available: http://www.srh.noaa.gov/gis/kml/

[50] "National Weather Service, National Hurricane Center." [Online]. Available: http://www.nhc.noaa.gov/

[51] L. Wasserman, "All of Nonparametric Statistics (Springer Texts in Statistics)." Springer, May 2007.

[52] R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford, "Dynamics of Hot-Potato Routing in IP Networks," in *Proceedings of ACM SIGMETRICS Conference*, New York, NY, USA, June 2004.

[53] J. Cowie, *Hurricane Sandy: Outage Animation*. Renesys Blog, 2012. [Online]. Available: http://www.renesys.com/blog/2012/10/hurricane-sandy-outage-animati.shtml