Computer Networks 54 (2010) 1169–1182

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Criticality analysis of Internet infrastructure

Guanhua Yan [a,*,1], Stephan Eidenbenz [a], Sunil Thulasidasan [a], Pallab Datta [b,2], Venkatesh Ramaswamy [c,2]

[a] Los Alamos National Laboratory, Los Alamos, NM 87545, USA
[b] Neurosciences Institute, San Diego, CA 92121, USA
[c] Airvana, Inc., Chelmsford, MA 01824, USA

ARTICLE INFO

Article history:
Received 28 January 2009
Received in revised form 18 September 2009
Accepted 10 November 2009
Available online 17 November 2009
Responsible Editor: P. Dowd

Keywords:
Internet infrastructure
Network modeling
Infrastructure protection

ABSTRACT

The Internet has evolved into an indispensable component of our daily lives and protecting its critical infrastructure has thus become a crucial task. In this work, we present and compare different methods to assess the criticality of individual facilities of the Internet infrastructure at a national-level: graph-theoretical analysis, route-based analysis, traffic-based analysis, and consequence-based analysis. Our key observations are: (1) The geographical topology, which is derived from a national-level IP backbone network, has a power-law degree distribution and is a small-world network; (2) A few locations appear much more frequently among all paths in the IP backbone topology than others, and they also witness a high percentage of US Internet traffic. (3) Relative ranking of Internet facility locations from traffic-based analysis differs significantly from those derived from graph-theoretical analysis and route-based analysis, suggesting that a comprehensive, high-fidelity Internet model is necessary to assess critical Internet infrastructure facilities. (4) Consequence-based analysis, although computationally intense, cannot be replaced by other rankings, including traffic-based analysis. Conclusions drawn from this work extend our knowledge regarding the Internet and also shed lights on which critical Internet infrastructure facilities should be protected with limited resources.

Published by Elsevier B.V.

1. Introduction

As the Internet has permeated into almost every aspect of our lives, it is crucial to ensure that its infrastructure functions properly. The Internet infrastructure can suffer severe physical damages from natural disasters, such as hurricanes and earthquakes, or physical attacks. Meanwhile, a malicious cyber-attack (e.g., a distributed denial-of-service attack) can cause undesirable effects, if it disables a critical Internet infrastructure facility completely or even only makes it behave abnormally. As many other infrastructure sectors, such as power grids and transportation systems, become increasingly dependent on the Internet for their normal operations, it is vital to protect Internet infrastructure from severe physical damages and malicious cyber-attacks.

Given the vast scale of the Internet, it is a challenging task to decide where we should dedicate our resources to protect its infrastructure, especially when resources provided are only limited. This is the main theme of this paper: we perform criticality analysis of the Internet infrastructure at a national-level from four different methodological perspectives. *Graph-theoretical analysis* studies the structural properties of a geographical network derived from the Internet backbone topology, including its degree distribution, clustering structure, and also its centrality measures. Although theoretically appealing, graph-theoretical analysis ignores the hierarchical routing scheme of the real Internet and uses the shortest path routing scheme

* Corresponding author. Tel.: +1 505 667 0176.
E-mail address: ghyan@lanl.gov (G. Yan).
[1] Los Alamos National Laboratory Publication No. LA-UR-08-05874.
[2] This work was done when Pallab Datta and Venkatesh Ramaswamy were working at Los Alamos National Laboratory.

1389-1286/$ - see front matter Published by Elsevier B.V.
doi:10.1016/j.comnet.2009.11.002

due to its simplicity. *Route-based analysis*, instead, models realistic inter-domain and intra-domain routing schemes used in the Internet and then identifies those facilities that appear most frequently on paths in the Internet backbone topology. Realizing that route-based analysis still produces biased results because all paths are evenly weighted, we propose another approach, *traffic-based analysis*, which weighs each path by its traffic demands. To do this, we generate synthetic end devices and also session-level traffic among them. In contrast to previous efforts on Internet modeling, we attempt to achieve high-fidelity through a socio-technical modeling approach that uses realistic datasets such as US census data, computer usage statistics, and also market shares of Internet service providers. Based on a comprehensive Internet model, we perform *consequence-based analysis* to evaluate the importance of each Internet facility by measuring the amount of traffic lost after it is removed from the topology.

The key conclusions drawn from this paper are:

(1) The geographical topology, which is derived from a national-level IP backbone network, has a power-law degree distribution; it is a small-world network with a high clustering coefficient and a small characteristic path length. Moreover, the number of IP addresses at each location in the IP backbone network is also well characterized by a power-law distribution.
(2) A few locations appear much more frequently among all paths in the IP backbone topology than others, and these locations also witness a high percentage of US Internet traffic.
(3) Relative ranking of Internet facility locations from traffic-based analysis differs significantly from those derived from graph-theoretical analysis and route-based analysis, suggesting that a comprehensive, high-fidelity Internet model is necessary to assess critical Internet infrastructure facilities.
(4) We perform consequence-based analysis on Internet facilities by calculating the amount of unroutable traffic after each of them is removed. Consequence-based analysis, although computationally intense, cannot be replaced by other rankings, including traffic-based analysis.

The remainder of this paper is structured as follows. We first present related-work in Section 2. In Section 3, we discuss how we construct a geographical network from the Internet backbone topology and also analyze its structural property, including its degree distribution, clustering coefficient, and centrality measures. In Section 4, we present an algorithm that models realistic Internet routing from both inter-domain and intra-domain levels; after that, we analyze how frequently a backbone location is traversed among all paths in the backbone topology. In Section 5, we describe how to generate synthetic end devices and their traffic at a session-level; we also discuss how paths computed in the previous section, if weighted with their traffic demands, affect the relative rank of each location. In Section 6, we perform consequence-based analysis on Internet facilities using the Internet model built in the previous sections. Section 7 further gives the limitations of our work, such as imprecise datasets used in this study. We conclude this paper in Section 8.

2. Related work

There have been numerous efforts on analyzing the structural properties of the Internet topology, mostly at the AS-level and at the router level. Faloutsos et al. first observed that the Internet topology exhibits several power-law distributions [13]. This conclusion on the AS-level topology was later questioned in [9] and a recent study suggests that power-law distributions may result from sampling errors of traceroutes [22,3]. Using a method with solid statistical footing, we show that the skitter dataset at a national-level does not exhibit power-law degree distribution. Instead, we observe that the geographical topology condensed from it seems to fit well with the power-law distribution. The small-world property of the Internet topology has been investigated in [6,20], and spectral analysis of Internet topologies has been pursued in [7,40]. In this paper, we complement previous work along these lines by studying the geographical topology condensed from the IP backbone network.

A plethora of models have been proposed to characterize routing and traffic in the Internet, many of which were developed for simulation purposes [25,23,42,24,31]. Different from previous inter-domain routing models, we use an AS path inference algorithm to derive inter-domain paths that are used in the real Internet; our traffic model generates synthetic end-to-end sessions originating from end devices that statistically follow the observed distribution in the US. The level of authenticity carried in our model has rarely been pursued in the literature before.

3. Graph-theoretical analysis

In this section, we analyze the criticality of assets in the Internet infrastructure from a graph-theoretical perspective.

3.1. Internet backbone topology

3.1.1. Backbone topology

The Internet backbone consists of routers and links that are owned and operated by major Internet service providers such as AT&T, Sprint, and XO. Fig. 1 gives a geographical overview of 18,000 backbone equipment locations, which are shown in green circles. These locations house more than 291,000 unique backbone IP addresses in the US, which are extracted from the skitter dataset collected by the CAIDA project [8]. We note that the original skitter dataset does not produce a connected graph. Hence, two approaches are adopted to make it more connected. *First*, each IP address corresponds to a network interface at a backbone router and multiple IP addresses can belong to the same physical backbone router. This is the well-known *IP alias resolution* problem [34]. We use the alias clustering data provided by the

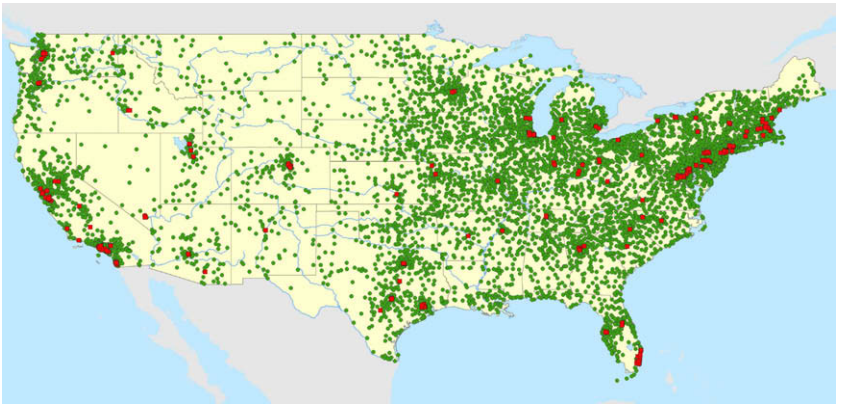

Fig. 1. Internet backbone in the US.

iPlane project [19] and for any two IP addresses in the skitter dataset that belong to the same physical router, we create a virtual link between them. We call such links *virtual alias links* to differentiate them from those realistic *observed links* in the skitter dataset. *Second*, we leverage the following heuristic: if two IP addresses belong to the same AS and are located at the same place, it is unlikely that traffic between them traverses through a different location. The geographical position of each backbone IP address, in the form of its longitude and latitude, is derived from the *ip2location* dataset [18]. We ensure that all co-located IP addresses that belong to the same AS are connected by creating virtual links among them. These links are termed *virtual co-located links* in our model. To avoid creating too many such links by, say, using a clique structure, we use a star structure instead to connect co-located IP addresses owned by the same AS. We call the center of such a star topology a *hub IP*. With these two techniques, we are able to produce a connected Internet backbone that covers more than 99.7% of the IP addresses in the skitter dataset.

3.1.2. Internet PoPs

Another important concept in the Internet infrastructure is *PoP* (Point of Presence). An Internet PoP is an access point to the Internet backbone, which is typically owned by an ISP, or located in an Internet exchange points or colocation centers. Large companies and institutions (e.g., IBM) and small Internet service providers are connected to the Internet backbone at these PoPs. We obtained a set of 543 PoPs from the *telegeography* colocation database [36], which lists operators present in each PoP. The locations of these 543 PoPs are also illustrated as red boxes in Fig. 1.

Next, we populate the 543 PoPs with backbone IP addresses in the skitter dataset: if the geo-location of a backbone IP address (i.e., its longitude and latitude) agrees with that of a PoP, we assign it to that PoP. We, however, find that such a simple assignment scheme leads to inconsistency: for an AS A_i present in PoP P_k according to the telegeography colocation database, it may not have any of its backbone IPs assigned to that PoP. To circumvent this problem, we create a virtual backbone IP address that belongs to AS A_i in PoP P_k; moreover, if there exists a geolocated backbone IP address that belongs to AS A_i within 15 miles, we connect the virtual IP address to it. In total, 1247 virtual backbone IP addresses have been created thereby, consisting of only 0.4% of all backbone IP addresses. For simplicity, we call a backbone IP address inside a PoP *a PoP IP*.

As the skitter dataset is generated from traceroute output, inter-AS links derived from it are incomplete and biased [16]. To mitigate this problem, we assume that all ASes present in the same PoP are internally connected. We create a virtual PoP IP inside each PoP, which is called *virtual inter-AS PoP IP*, and connect it to every hub IP in that PoP. Note that this process may introduce extra inter-AS links that do not exist in reality. For example, if two ASes do not have any business relationships, there may not be physical AS links between them. The routing scheme in our model, however, relies on inferred AS-level relationships from realistic BGP data to compute AS paths and these extra links thus will not be used in routing. We will explain it further in Section 4.

3.2. Analysis

3.2.1. Power-law?

We first investigate whether the degree distributions of the backbone topology in our model conform to power-law distributions. We consider two different IP topologies: graph G_{ip} generated directly from the skitter dataset and the full IP topology in our model G_{ip^+}. Their degree distributions, plotted on doubly logarithmic axes, are depicted in Fig. 2. An interesting question regarding the Internet topology is whether its degree distribution follows the power-law. Previous work along this line often uses a least-squares linear regression to determine whether a power-law exists [13]. This approach is however significantly biased and also often incorrect [10]. In the following, we perform a rigorous study on G_{ip} and G_{ip^+}, and pursue answers to what power-law models best fit these topologies or even whether the power-law hypothesis holds for these topologies at all.

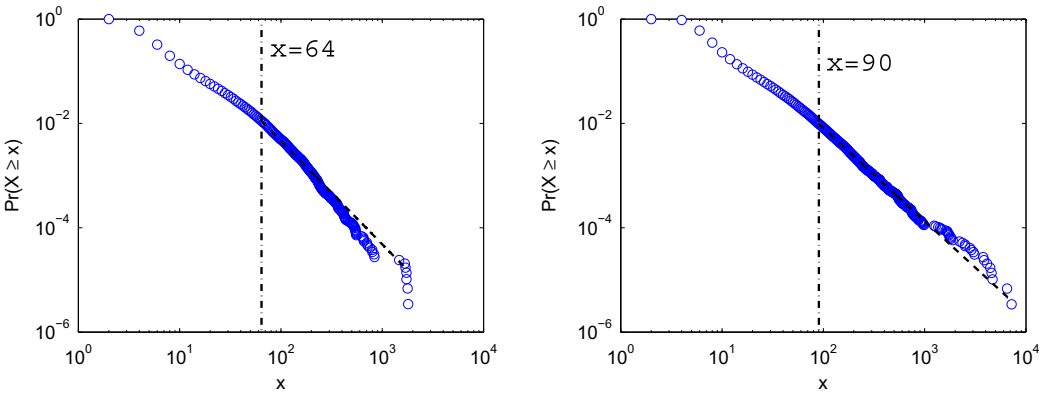

Fig. 2. Probability that a node's degree is no smaller than x (in log-scale).

We use the method proposed in [10], in which a power-law model has two parameters: scaling exponent α and cut-off β. In our context, the power-law hypothesis can be stated as follows: the frequency f_d of a degree d is proportional to d to the power of α, when d is no smaller than β. The reason for introducing β is simply to avoid the divergence of the density function at degrees close to 0. The method uses Maximum likelihood estimation and Kolmogorov–Smirnov statistics to estimate α and β; to test whether the power-law distribution models well the empirical data, the method computes the p-value. The results are summarized in Table 1. Based on the suggested p-value threshold in [10], which is 0.1, we can safely conclude that neither G_{ip} nor G_{ip^+} exhibits the power-law behavior, which is in contrast to previous conclusions made simply based on the least-squares linear regression approach.

As the main goal of this work is to identify key locations that house critical Internet infrastructure facilities, we now focus on analyzing the Internet topology from a geographical perspective. Based on the IP topology in our model, we further derive a location topology G_{loc} by aggregating all IP addresses that belong to the same location into a single location node; a link is added between two locations in the topology graph if and only if in the IP topologies there exist links that cross over between them. Note that the virtual alias links and co-located links introduced do *not* affect the location topology, because they are limited to connect backbone IPs at the same location. The degree distribution of G_{loc}, plotted on doubly logarithmic axes, is depicted in Fig. 3, and the goodness-of-fit is presented in Table 1. The p-value is higher than 0.1, suggesting that the location topology can be reasonably characterized by the power-law model, although we *cannot* rule out other alternatives with better fitting results. Here, it is worth noting that the IP topologies may miss some links in the real Internet due to sampling biases of traceroute, but such errors are decreased to a lesser degree in the location topology because links between IPs at two locations are aggregated into a single one in the latter.

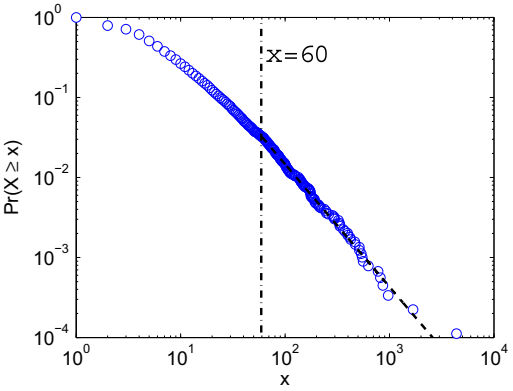

Fig. 3. Degree distribution of location topology G_{loc}.

The implication of a location graph with a power-law degree distribution is that a few locations are well connected to the other part of the network. To protect the Internet infrastructure, it is important to enhance the security and safety of these locations. As we shall discuss later, however, simply choosing those locations with the highest degrees may not function as effectively as other approaches.

The distribution of the number of IP addresses in the skitter dataset that are co-located at the same place is presented in Fig. 4. Similarly, we perform the hypothesis test on whether it can be modeled by a power-law distribution. The goodness-of-fit test has the following results: scaling exponent α is 2.05, cutoff β is 50, and the p-value is 0.9120. The high p-value suggests that the power-law distribution is a very good fit for the number of IPs within the same geographical location.

Table 1
Power-law model parameters and goodness-of-fit test.

Topology	α	β	p-Value
G_{ip}	2.98	64	0
G_{ip^+}	2.78	90	0.015
G_{loc}	2.53	60	0.302

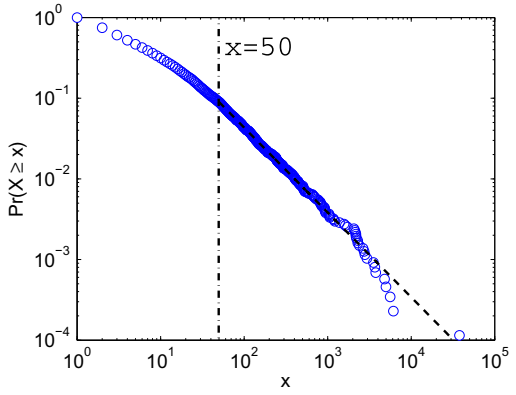

Fig. 4. Distribution of number of IPs at a location.

3.2.2. Small-world networks?

Small-world networks refer to a class of graphs in which nodes are highly clustered such that most nodes are reachable from each other by only a small number of hops. Internet topologies have been shown to exhibit small-world behavior at both AS and router levels [6,20]. In this study, we complement previous work by exploring whether the location graph is also a small network. Rigorously speaking, a small-world network can be characterized as a graph with a high *clustering coefficient* and a small *characteristic path length* [41]. The clustering coefficient ρ_v of a node v in an undirected graph $G(V,E)$ is defined by:

$$\rho_v = \frac{2H_v}{n_v(n_v-1)}, \tag{1}$$

where H_v is the number of edges between node v's neighbors and n_v is the number of neighbors node v has in graph G. The average clustering coefficient over all the nodes in graph G gives the probability that any two nodes sharing the same neighbor are directly connected.

To evaluate the degree to which nodes in the location graph are clustered, we compare clustering coefficients of the nodes in this graph (in non-increasing order) against those in an Erdos–Renyi random graph, which is not a small-world network, with a similar number of nodes and edges. The results, depicted in Fig. 5, show that nodes in the location graph are much more clustered than those in the Erdos–Renyi random graph: the average clustering coefficient is 0.3130 in the former, as opposed to 0.0015 in the latter. Actually, the average clustering coefficient of the location graph is very close to or even higher than those observed from social networks such as YouTube and Flickr [29]. In Fig. 6, we show how clustering coefficient of a node is correlated with its degree. It is obvious that nodes with low degrees tend to have high clustering coefficients, suggesting that there exists significant clustering among low-degree nodes.

Another feature of a small-world network is that it has a small characteristic path length, which is defined as the average length of the shortest paths between all pairs of vertices. The characteristic path length of the location topology is 2.9, which is even smaller than that of the Erdos–Renyi random graph, which is 3.9. In Fig. 7, we plot the fractions of the top 30 locations' appearances in the shortest paths among all vertex pairs. A close examination reveals that the site with the highest degree contributes to 55% of all shortest paths. To conclude, the small characteristic path length of the location topology, together with its high clustering coefficient, indeed makes it qualified as a small network [41].

A location graph with small-world behavior means that theoretically speaking, it is *possible* that packets are delivered to their destinations through only a small number of locations. Later, we will present how many locations are traversed on average in the context of realistic Internet routing.

3.2.3. Centrality

In graph theory, the importance of a vertex in a graph is measured by its *centrality*. The four centrality measures that are widely used in network analysis are degree centrality, betweenness, closeness, and eigenvector centrality [30]. We have already shown the degree distribution in Fig. 3, and the betweenness centrality of a vertex is essentially the fraction of all shortest paths on which it appears, which is provided in Fig. 7. Here, we further measure the eigenvector centrality of each location. The eigenvector

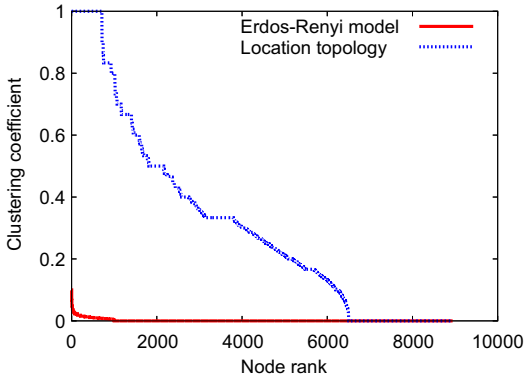

Fig. 5. Comparison of clustering coefficients.

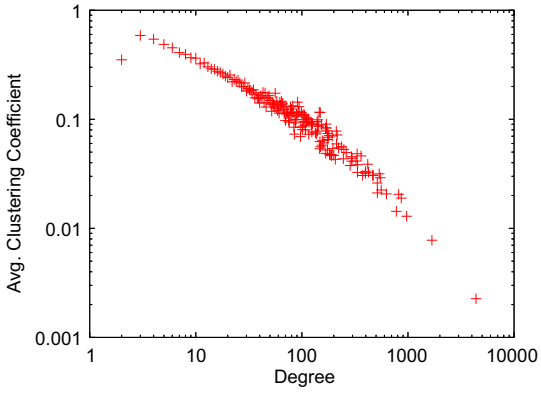

Fig. 6. Avarage clustering coefficient vs. degree in G_{loc}.

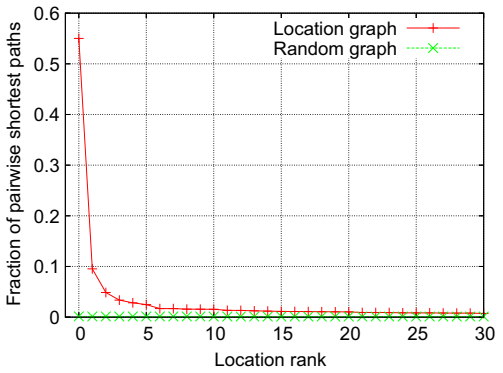

Fig. 7. Fraction of appearances in all shortest paths.

centrality of a graph is defined as the principal eigenvector of its adjacency matrix, i.e., the eigenvector that corresponds to the largest eigenvalue of its adjacency matrix. Instead of treating every connection in the network equally, the eigenvector centrality assigns a higher score to a node that has connections to neighbors which themselves have high scores.

In Fig. 8, we plot the eigenvector centrality of each node in the location topology against its rank. We observe that a few locations have much higher scores than the rest of locations. For instance, only 2.5% of the locations have an eigenvector centrality above 0.025. In the same graph, we also plot the correlation between eigenvector centrality and the other two rankings, degree and betweenness centrality rankings. We observe that the eigenvector centrality ranking is more consistent with the degree ranking than with the betweenness ranking. To further verify this, we define the *agreement ratio* among k top locations as the percentage of locations that appear among the top k in both rankings. For the degree ranking and the eigenvector centrality ranking, the agreement ratio among 100 top locations is 90%. For the betweenness centrality ranking and the eigenvector centrality ranking, however, the agreement ratio is only 66%. This suggests that if we decide to protect the top 100 locations, using the eigenvector centrality ranking or the betweenness centrality ranking will give us very different results. In the following sections, we will explore alternative ranking methods with better supporting evidence.

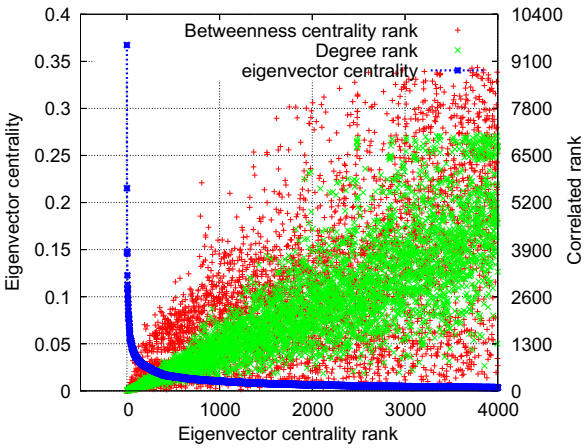

Fig. 8. Eigenvector centrality and rank comparison.

4. Route-based analysis

Ranking locations according to their eigenvector centrality, although mathematically appealing, lacks a strong physical foundation: what does an eigenvector centrality really mean? Ranking locations based on their betweenness centrality scores, however, ignores the fact that Internet routing is not always using shortest paths. Internet routing, instead, is strictly hierarchical: inter-domain routing protocols (e.g., BGP) regulate Internet traffic among different ASes, and intra-domain routing protocols (e.g., OSPF and RIP) specifies how traffic is routed within the same AS. In this section, we discuss how to assess critical Internet assets in the context of more realistic Internet routing schemes.

4.1. Internet routing

BGP is the de-facto inter-domain routing protocol used in the current Internet. Due to its complexity and the fact that commercial relationships between ASes are generally unavailable to the public, we use AS-level paths inferred from existing BGP routing tables for inter-domain routing. We use the AS path inference algorithm in [32], which is able to infer AS-level paths with 95% accuracy. With regard to intra-domain routing, we simply use the shortest path algorithm.

Algorithm 1 is used to compute the route between any two PoP IPs in our model. In the algorithm, we let $\Gamma(AS_i, AS_j)$ denote the set of PoPs where both ASes AS_i and AS_j have presence. We introduce $\phi(X_k)$ to denote the virtual inter-AS PoP IP of PoP X_k (see Section 3). We also use $G(AS_i)$ to denote the network formed by all backbone IPs inside AS AS_i, those virtual inter-AS PoP IPs that they are connected to, and all links between these IPs. Given the source and destination PoP IPs, the algorithm first derives an AS-level path $AS_1^{\circ}AS_2^{\circ}\cdots^{\circ}AS_n$ between them. As the algorithm in [32] provides multiple alternative AS paths, we choose the shortest one in our study. We then iteratively work on each AS_i from $i = 1$ to n. Once the first backbone IP inside AS_i (or a source virtual inter-AS PoP IP) has been decided, we calculate the shortest path in graph $G(AS_i)$, starting from it to any virtual inter-AS PoP IP in set $\Gamma(AS_i, AS_{i+1})$.[3] Next, we start from that (virtual) destination inter-AS PoP IP and find the shortest path in graph $G(AS_{i+1})$ to any virtual inter-AS PoP IP in set $\Gamma(AS_{i+1}, AS_{i+2})$. This process repeats until the last AS along the path. Inside the last AS, we simply find the shortest path to the destination IP on the route.

[3] If the source and the destination are both the same virtual inter-AS PoP IP, we make sure at least one backbone IP inside AS AS_i is traversed on the path.

Algorithm 1. Compute the route from PoP IP E_a to E_b

$\phi \leftarrow$ AS number of E_a
Use the algorithm in [32] to obtain the AS-level path P with origin AS number ϕ and destination E_b, where $P = AS_1^{\circ}AS_2^{\circ}\cdots^{\circ}AS_n$
$src \leftarrow E_a$
$Z \leftarrow \emptyset$
for $k = 1$ to $n-1$ **do**
  $\alpha_{min}(k) \leftarrow \infty, \beta_{min}(k) \leftarrow null$
  for each virtual inter-AS PoP IP $dst \in \Gamma(AS_k, AS_{k+1})$ **do**
    $Q \leftarrow$ the shortest path from src to dst in $G(AS_k)$
    if $\alpha_{min}(k) > |Q|$ **then**
      $\alpha_{min}(k) \leftarrow |Q|, \beta_{min}(k) \leftarrow dst$
      $Q_{min}(k) \leftarrow Q$
    end if
  end for
  $src \leftarrow \beta_{min}(k)$
  $Z \leftarrow Z^{\circ}Q_{min}(k)$
end for
$Q \leftarrow$ the shortest path from src to E_b in $G(AS_n)$
$Z \leftarrow Z^{\circ}Q$
Output route Z

How realistic is the path as calculated by Algorithm 1? Actually, this algorithm is very similar to routing policy (*Min AS path, early exit*) described in [27]. It has been shown that such a policy can achieve more than 78% accuracy due to BGP's default object function as minimizing AS path length and the default early exit intra-domain routing policy. In the future, we plan to perform a more rigorous study to validate routes derived from Algorithm 1 against realistic paths observed from the Internet.

We compute IP-level paths between every pair of PoP IPs in the Internet backbone topology. There are 60,506 PoP IPs in our model and calculating routes between each pair of them still imposes high computation cost. We thus simplify the routing model by computing routes only between a source hub IP and a destination PoP IP. As there are only 4916 hub IPs in the model, the overall computation cost is reduced by more than one order of magnitude. We assume that packets from an arbitrary PoP IP to a destination PoP IP traverse the hub IP of that source PoP IP first and use the precomputed path from the hub IP to the destination PoP IP.

For some destination PoP IPs, the algorithm in [32] fails to infer AS-level paths to them. In such circumstances, we derive their AS numbers and use www.fixedorbit.com to obtain a list of prefixes for each of them. We then use the algorithm in [32] again to infer AS-level paths to these prefixes. These derived AS-level paths are further used to compute the IP-level paths to these destination PoP IPs.

4.2. Analysis

We now analyze paths derived between all PoP IP pairs. Fig. 9 depicts the frequency histogram of the number of locations that appear on a path. We observe that on average a path between any two PoP IPs traverses 4.7 different locations (including both source and destination), which is 60% longer than the characteristic path length of the location graph. Also, it is very rare that a path traverses more than 10 different locations. Here, it is worth mentioning that the number of locations discussed here differs from what we see from a traceroute output. First, multiple IPs reported from a traceroute command can actually locate at the same place. Second, the output from a traceroute command includes all IPs seen on the path, which includes not only backbone IP addresses but also IPs on Internet access routers.

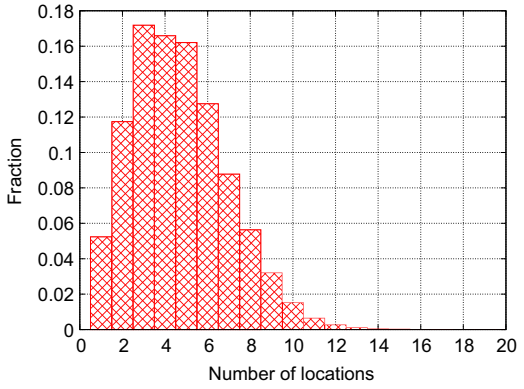

Fig. 9. Frequency histogram of the number of locations in a path.

We rank the 543 PoPs based on the frequency at which they appear on a path between any two PoP IPs and show the top 100 PoPs in Fig. 10. We observe that one PoP appears on almost 40% of paths between all PoP IP pairs and 10 PoPs appear on more than 5% of paths between all PoP IP pairs. Such high frequency suggests that these PoPs are important for network connectivity in the context of Internet hierarchical routing.

Similarly, we rank all the backbone IP locations based on the frequency at which they are visited by a path between any two PoP IPs. Fig. 11 depicts the frequency histogram for the top 100 locations. We note that the top location appears on almost 70% of the paths between all pairs of PoP IPs. Although surprising, this is actually consistent with earlier reports about the crucial role of some locations in the US [38,39]. We will further

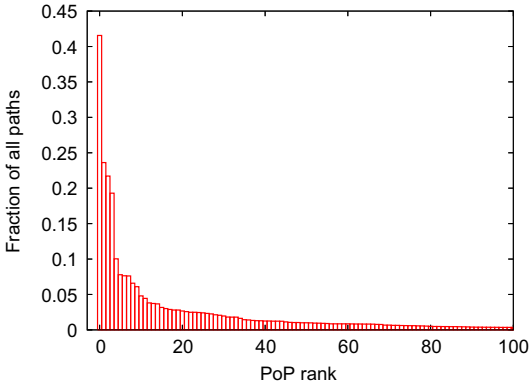

Fig. 10. Rank of PoPs based on their appearances in derived paths.

discuss this when we present results from traffic-based analysis.

Fig. 12 presents the scatterplot of the correlation between the route-based rank of a location and its rank if its degree or eigenvector centrality is used. Clearly, these rankings are not consistent: a location with a high route-based rank can have a low-degree rank or eigenvector centrality rank. The agreement ratio of the top 100 locations is only 49% between the route-based rank and both the degree rank and the eigenvector centrality rank. One may wonder whether route-based ranking and betweenness centrality ranking produces similar results because both of them focus on routing. Fig. 13 depicts the scatter plot of such correlation, which clearly shows that ranking results from them are not strongly correlated. Actually, the agreement ratio for the top 100 is only 40%.

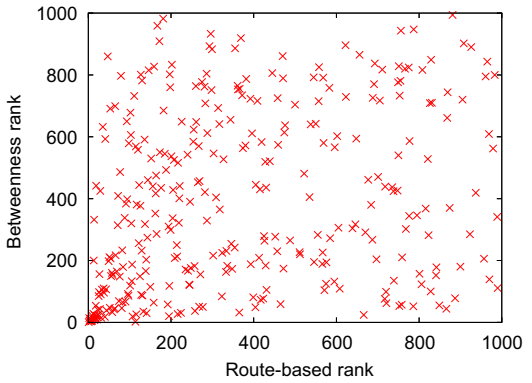

Fig. 13. Correlation between route-based and betweenness centrality rankings.

5. Traffic-based analysis

Route-based criticality analysis, although taking realistic Internet routing schemes into account, is still biased because it does not consider traffic demands between different PoP IP pairs. For instance, if a PoP location houses a large number of IPs, paths among them may significantly improve its rank, although there is little traffic flowing among these IPs in reality. In this section, we present a more complicated ranking scheme that improves route-based ranking by considering traffic demands between PoP IP pairs. In this approach, we first generate synthetic end devices, including both residential and business computers, and then connect them to the Internet backbone topology. We further generate synthetic sessions, including HTTP, Email, P2P, and streaming traffic, for every computer within 24 h.

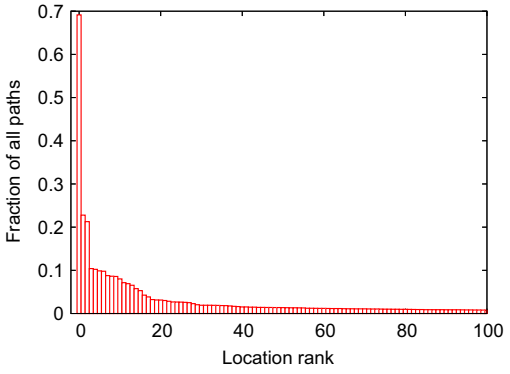

Fig. 11. Rank of locations based on their appearances in derived paths.

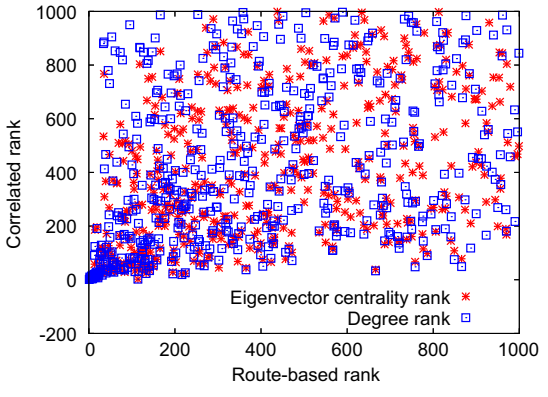

Fig. 12. Correlation between route-based ranking and other methods.

5.1. End devices and access routers

5.1.1. End devices

We distinguish residential and business computers in our model. To generate residential computers, we use a US Census data source that gives the census-block level population data in each 250×250 m^2 grid in the entire US for both day and night time [28]. From this dataset, we synthesize the average number of households in each grid. We then derive the income distribution per household in each grid from the US census-block-group dataset, which provides the fraction of households whose annual income falls into each range in every block group. Also, the US census data provides the percentage of residential computer usages for each annual family income category. Based on these data, we synthesize the number of computers that are located in each grid. In total, we have generated 73,884,296 residential computers in the US. This number is close to the 73 million computers reported in a Yankee group survey.

To generate business computers, we use the Dun & Bradstreet (D&B) dataset, which provides information about all companies in the US, including their headquarter locations, numbers of employees, and SIC (Standard Industrial Classification) codes. A SIC code has four digits and indicates the business type of a company. The US census data presented in [11] gives us computer penetration ratios in different business categories. Based on this, we synthesize the number of business computers in each company. In total, we have generated 58,923,964 business computers in the US, which is close to the 65 million business computers reported by the US Department of Commerce.

5.1.2. Access routers

Internet access routers are used to connect end devices to the Internet backbone. To generate these routers, we need to know which companies provide Internet access services in each region. Currently we only consider three types of Internet access services, dial-up, DSL, and Cable, as they are mostly widely used in US. According to the Home Broadband Adoption 2006 report by Pew Internet & American Life Project, there are in total 48 million dial-up users and 84 million broadband users in the US, and among all broadband users, the market shares for DSL and Cable are 50% and 41%, respectively. These numbers are used to assign the Internet access type of each end device.

For the dial-up service, we collect a list of aggregators for each zip code from the Internet Service Provider Directory [14] and for each of these aggregators we create an Internet access router. A dial-up access router is located at the central office of the corresponding area.[4] For the other two types of services, we model the entire market by using the subscriber numbers of the top nine companies for each service that collectively cover more than 50% of the respective market [26]. For each of these companies that provide broadband Internet access service, we collect a list of zip codes or area codes that it provides broadband Internet access services (i.e., DSL or Cable) and also create an Internet access router for each zip code within its service coverage. Each DSL access router is located at the central office of the corresponding area and each Cable access router is located at the closest business office that the cable company has in the corresponding area.

5.1.3. Connections

For each end device, we connect it to an Internet access router. We first randomly choose the type of Internet access service based on the market shares of all Internet access services. If the chosen access service is dial-up, we randomly assign the end device to an aggregator for the zip code where the device is located. If the access type is DSL or Cable, we randomly choose an Internet broadband access router based on the market shares of the top broadband companies. After an Internet access router is chosen, we create a link between it and the end device.

Recall that there are 543 PoPs in the backbone topology and each of them has a list of backbone IPs. Also, each PoP IP is associated with an AS number. Given an Internet access router, we decide which PoP IP it connects to based on the following algorithm. *First*, we sort all PoPs according to their distances from the Internet access router. *Second*, starting from the closest PoP, we check whether it has a PoP IP that peers with the ISP company owning that Internet access router. This can be done by checking whether the AS (Autonomous System) number of the PoP IP connects to any one of the AS numbers owned by the ISP company in the AS-level graph. If we cannot find it, we try the second closest PoP. This process repeats until one such PoP IP is found. Thereafter, we create a link between it and the Internet access router.

[4] A central office is a building that houses telephone switches in telecommunication networks.

5.2. Sessions

Due to the complexity and dynamics of Internet traffic, any attempt to characterize it accurately will be arduous, if not impossible. For simplicity, we model Internet traffic at the granularity of sessions, as they closely reflect the behaviors of Internet users, such as web browsing, file downloading, and Email communications. We ignore effects of transport layers (such as TCP or UDP) or layers below on the traffic characteristics. Currently, HTTP, P2P, Email, and streaming traffic constitutes the majority of the Internet traffic [21,1,2], so it suffices to generate these four types of sessions in our model. From a high level, our session generation algorithm works as follows: for each Internet session, based on the relative occurrence of different types of sessions, we assign it a session type, and then we choose its origin and destination. In total, we have generated around 1.14 billion sessions for the entire US population within a 24-h period.

Table 2 summarizes the relative occurrence of each type of sessions originating from both residential and business computers. The relative occurrence of each session is obtained from the traffic mix observed in the current Internet traffic as presented in [17,1,2,37]. The size of each session is selected based on the following studies. The literature suggests that approximately 80% of Web document transfers are less than 100 kB in size [5], though there is a significant heavy tail to the distribution [5,12]. The average size of Email sessions is taken as 100 K, which is chosen based on the average size of all Emails in the inbox of various employees in a large institution. Similarly, the size of HTTP session is computed by downloading a number of web pages and finding the average of these downloaded web pages. The average streaming rate of streaming sessions is 200 kB/s (kBps) [15] and the average duration of streaming session is approximately 125 s [35]. This gives the average size of streaming session as approximately 30 M. The average size of P2P session is computed by observing the history of already completed transfers in a P2P client [33]. The size of each session is drawn from exponential distribution with the average size given in Table 2.

We have a distribution that gives us the probability of each type of session for any particular hour of a day for sessions originating at home and at work. We iterate through every second of the day, and compute the number of sessions that will be generated for a particular second following a Normal Distribution. Once we assign the session a particular type, we decide the origin of the session. The origin of a session can be either a business location or a home

Table 2
Session types and size parameters.

Session type	Percentage of home sessions (%)	Percentage of work sessions (%)	Avarage size (in bytes)
HTTP	25.14	14.86	25 K
Email	11.87	18.13	100 K
P2P	18.71	1.29	10 M
Streaming	6.29	3.71	30 M

location. This assignment is again based on the proportion of sessions that originate from different locations. For instance, the probability of Email sessions originating from work locations are higher during office hours (8 a.m. to 5 p.m.) and those originating from other locations are higher after office hours. For both P2P sessions and streaming sessions, they originate mainly after office hours, as people tend to watch news or download music either from home or after office hours if they are engaged in those activities from the office. For HTTP sessions originating from work, the activity is mainly during lunch hours and during the end of the day, and for those originating from other locations, they are mainly after office hours.

We characterize the end-points of the sessions based on the type of the session. We need to assign an originating device and a destination device or server to the session. For an Email and P2P session, we assume that the end-points of the session are end devices residing either at home or business locations. For HTTP and streaming traffic we assume that the source of the session is an end-device whereas the destination of the session is a server. To pick an end device as either the source or destination, we pick a device from a state based on the percentage of devices in that state. When the end-device is a server (for HTTP and streaming sessions), we pick a server from one of the top 100 servers that are most visited, based on the proportion of web access hits they receive [4].

Many web servers are located in the technological centers of Silicon Valley and Washington, D.C., as well as a few smaller centers mostly in metropolitan areas. Most major websites use the services of content distribution networks, such as the one offered by Akamai, Inc. These overlay networks cause HTTP traffic to be distributed evenly across different content sites. Unfortunately, very little data is available on the infrastructure and geographical distribution of these overlay networks or content distribution networks (CDNs). We however geo-locate the websites that are most widely visited (according to [4]) and split the sessions uniformly across those locations for each incoming session request to a particular server.

5.3. Analysis

In Fig. 14, we plot the average number of *transit ASes* that each byte traverses on a path. Here, we only consider ASes associated with backbone IPs and hence the number *does not* include the ASes associated with access routers. From the graph, we observe that the majority of the traffic traverses only two ASes in the Internet backbone topology. The average number of ASes traversed by each byte is 1.9. We further present in Fig. 15 the frequency histogram of the number of locations that each byte visits in the backbone topology. Although the shape of the curve is similar to what we have observed in Fig. 9, in which the frequency histogram is calculated over all pairs of PoP IP pairs, we note that paths with a single location appearance are used less frequently than their fraction among paths between all pairs of PoP IPs. This suggests that traffic-based analysis helps eliminate biases introduced by the route-based counterpart when a PoP location houses a large number of IPs. The average number of locations that a byte visits in the backbone topology is 5.2, which is slightly larger than the number derived from the route-based analysis and 80% longer than the characteristic path length of the location graph.

Figs. 16 and 17 depict the portion of traffic that each PoP sees against its rank and the portion of traffic that traverses each location against its rank, respectively. We note that the top PoP sees about 60% of the total traffic, and the top location is traversed by about 70% of the total traffic. Although very surprising, this result is actually in concert with earlier observations that some locations witness more

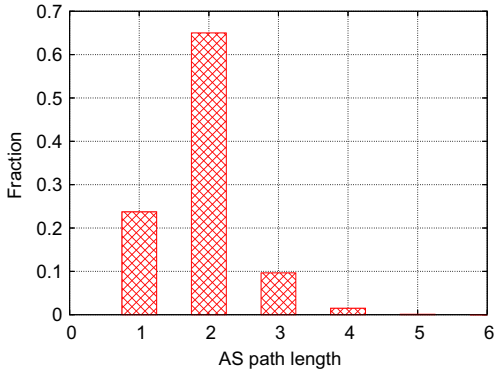

Fig. 14. Average number of ASes that a byte visits on its path.

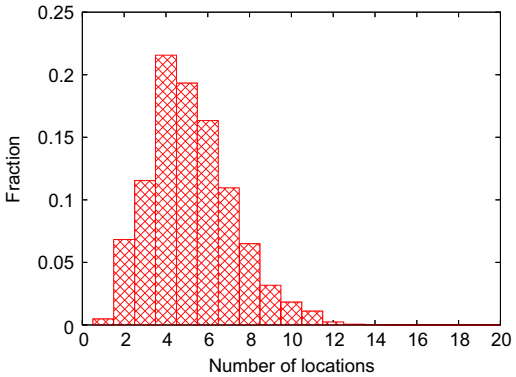

Fig. 15. Number of locations that a byte visits on its path.

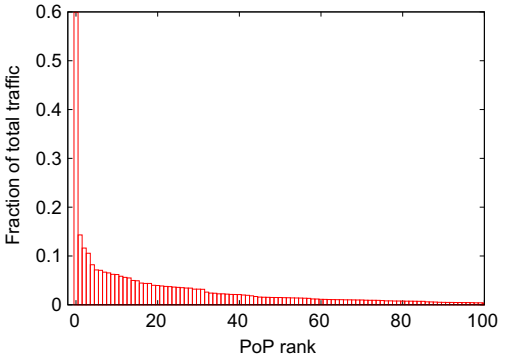

Fig. 16. Rank of PoPs based on their appearances on a byte's path.

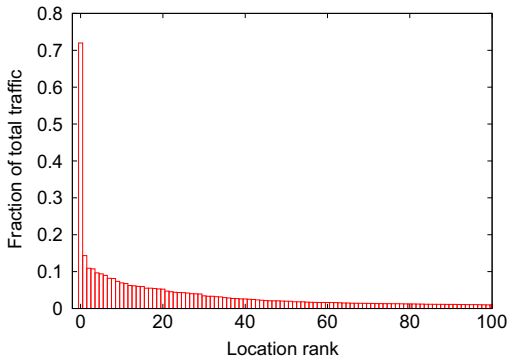

Fig. 17. Rank of locations based on their appearances on a byte's path.

than 50% of the Internet traffic in the US [38,39]. Admittedly, our model still may overestimate the amount of traffic traversing a single location. This is because our model only allows traffic from different ASes to be exchanged at those 543 PoPs, which may leave out many local IXPs (Internet Exchange Points), where regional ISPs also exchange their traffic in reality. Nevertheless, our model puts more emphasis on *all* 543 PoPs, many of which are indeed critical Internet assets. Hence, the fidelity of our model suffices to evaluate the relative importance of Internet infrastructure facilities.

In Fig. 18, we present a scatterplot of the route-based ranking and the eigenvector centrality ranking against the traffic-based ranking. The graph shows that no strong correlation exists for both of them. In Table 3, we present the agreement ratios for the top 100 between the traffic-based ranking and other methods. The results reveal that the relative importance of each location or PoP among the top 100, if traffic-based analysis is applied, differs significantly from what we have derived from the other methods. This suggests that a comprehensive, high-fidelity Internet model is indeed necessary to evaluate the criticality of Internet infrastructure facilities.

6. Consequence-based analysis

In the previous section, we rank each location based on how much traffic traverses each location under *normal* operational circumstances. Another way of measuring the

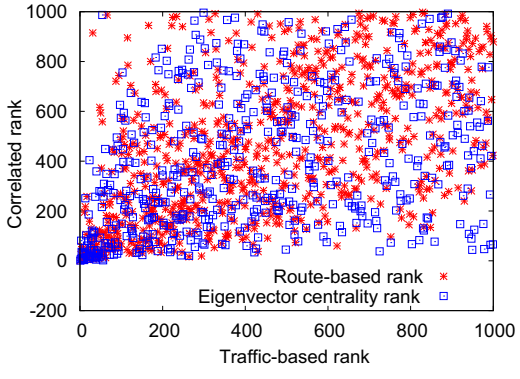

Fig. 18. Correlation of traffic-based ranking and other methods.

Table 3
Agreement ratios of consequence-based analysis with other rankings for the top 100 locations.

Rank 1	Rank 2	Agreement ratio (%)
PoPs (traffic-based)	PoPs (route-based)	67
Locations (traffic-based)	Locations (route-based)	56
Locations (traffic-based)	Locations (degree)	48
Locations (traffic-based)	Locations (eigenvector centrality)	49
Locations (traffic-based)	Locations (betweenness centrality)	38

importance of a location is how much impact it would cause if we remove this location from the network. This kind of consequence-based analysis is crucial for us to prepare for unexpected incidents such as natural disasters (e.g., hurricanes and earthquakes) and physical attacks against Internet facilities. In this section, we consider ranking Internet assets based on how much traffic would be lost if a location is removed.

The routing and traffic models performed in consequence-based analysis are the same as those developed in Sections 4 and 5. Basically, when we evaluate the importance of a location, we remove all IP addresses that belong to this location and also edges associated with these IP addresses. We then use the routing algorithm illustrated in Algorithm on the remaining topology to route the traffic generated with the models in Section 5. But due to the removal of a location, for some traffic that is routable before the removal, it is possible that we cannot find a path for it in the remaining topology any more. By counting the number of bytes that become unroutable after each location is removed, we use it to measure its importance in the topology.

It is worth noting that performing consequence-based analysis on a large topology like the Internet infrastructure is computationally expensive. This is because we have to redo path computation after each of the thousands of locations is removed from the graph. To accelerate consequence-based analysis, we, when evaluating the importance of a location, consider only those sessions that traverse this location under normal operational circumstances. We then attempt to find a path for each of these sessions on the topology without the location under analysis. If it fails, we add the number of bytes in the session to the lost traffic amount without this location.

Fig. 19 depicts the fraction of total traffic that is lost due to removal of the top 100 locations. We note that removing the top location that carries the most traffic as shown in Fig. 17 leads to loss of about 22% of the total traffic. This further confirms the importance of this location. We show the scatterplot of the consequence-based ranking and the traffic-based ranking in Fig. 20, and it is clear that there is strong correlation between these two rankings.

We calculate the agreement ratios of the ranking based on consequence-based analysis and other rankings and the results are demonstrated in Table 4. We observe that con-

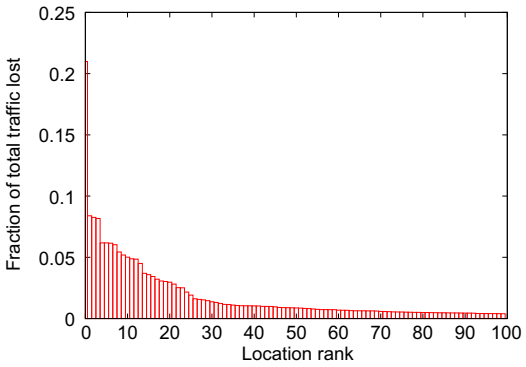

Fig. 19. Rank of locations based on the amount of traffic lost if a location is removed.

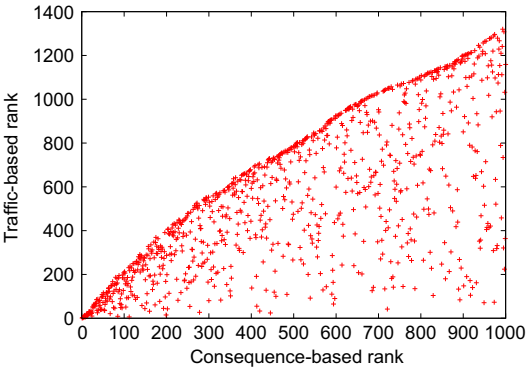

Fig. 20. Scatterplot of the traffic-based ranking against the consequence-based ranking.

sequence-based ranking and traffic-based ranking have about two thirds in common for the top 100 locations, much higher than those between consequence-based ranking and other rankings. The results still suggest that consequence-based ranking, although computationally expensive, cannot be replaced by other rankings, including traffic-based ranking.

Table 4
Agreement ratios of consequence-based analysis with other rankings for top 100 locations.

Rank 1	Rank 2	Agreement ratio (%)
Locations (consequence-based)	Locations (traffic-based)	67
Locations (consequence-based)	Locations (route-based)	39
Locations (consequence-based)	Locations (degree)	36
Locations (consequence-based)	Locations (eigenvector centrality)	39
Locations (consequence-based)	Locations (betweenness centrality)	31

7. Scope of our work

In this work, we attempt to build a comprehensive Internet model for the purpose of evaluating the relative importance of Internet assets. As these models are abstracted from datasets collected from the real-world Internet, the correctness of our analysis is contingent on how accurately they characterize the state and behavior of the Internet. It is, however, known that modeling Internet at high-fidelity is a notoriously daunting, if not impossible, task. Our work presented previously relies on a backbone topology derived from traceroute outputs. Such sampling may introduce a distorted IP topology due to missed links. Moreover, our model only incorporates IP addresses that reside in the US and is thus not representative of the full Internet structure. More importantly, the IP geo-location software used in our study sometimes fails to provide high-resolution results. Lack of street-level IP geo-location precision, for instance, results in some IPs belonging to different but close PoPs that are mapped onto the same location. This inevitably affects the structural analysis performed in this study.

Another challenge of our work is the dynamic nature of the Internet. It is possible that some datasets used in our study are outdated and thus do not reflect the current state of the Internet. For instance, we notice that some IPs in the skitter dataset we used cannot be reached any more, possibly because they have been decommissioned. Also, the AS-level graph may vary over the years due to changes of business relationships among ISPs, which poses another challenge for us to obtain a consistent view of the Internet backbone topology.

This work is also limited to discover those facilities that are crucial for data transmissions. It does not, however, consider other components of the Internet that are also indispensable for its normal operation, such as DNS servers and BGP routers. Moreover, our traffic generation model in this work is still preliminary, considering the diverse types of Internet traffic nowadays. We will continue to improve it in our future work.

8. Conclusions

The main focus of this paper is to evaluate the criticality of assets in the Internet infrastructure. Towards this end, we first analyze the structural property of the geographical network derived from the Internet backbone topology, using standard graph-theoretical tools. We then model realistic Internet routing and compare the frequencies at which Internet locations appear on the paths in the backbone topology. We further improve our ranking results by weighing these paths with session-level traffic demands that are generated from synthetic end devices. Finally, we perform consequence-based analysis on Internet facilities by computing the amount of unroutable traffic after each of them is removed. The contributions made in this paper extend our knowledge regarding the Internet and also shed lights on which critical Internet infrastructure facilities should be protected with limited resources.

Acknowledgments

We thank the CAIDA project team for their Skitter dataset, Jian Qiu and Lixin Gao at the University of Massachusetts for their AS path inference software.

References

[1] <http://www.readwriteweb.com/archives/p2p_growth_trend_watch.php>.
[2] <http://www.dslreports.com/shownews/85022>, 2007.
[3] D. Achlioptas, A. Clauset, D. Kempe, C. Moore, on the bias of traceroute sampling: or, power-law degree distributions in regular graphs, in: Proceedings of the ACM STOC'05, 2005.
[4] <http://www.alexa.com/>, 2008.
[5] M. Arlitt, C. Williamson, Internet web servers: workload characterization and performance implications, IEEE/ACM Transactions on Networking 5 (5) (1997).
[6] T. Bu, D. Towsley, On distinguishing between internet power law topology generators, in: Proceedings of the IEEE INFOCOM'02, 2002.
[7] C., Gkantsidis, M. Mihail, E. Zegura, Spectral analysis of internet topologies, in: Proceedings of the IEEE INFOCOM, 2003.
[8] <http://www.caida.org/tools/measurement/skitter/>.
[9] Q. Chen, H. Chang, R. Govindan, J. Jamin adn S. Shenker, W. Willinger, The origin of power laws in internet topologies revisited, in: Proceedings of the IEEE INFOCOM, 2002.
[10] A. Clauset, C.R. Shalizi, M.E.J. Newman, Power-law distributions in empirical data, June 2007. Available from: <arXiv:0706.1062>.
[11] Computer and internet use in the united states: 2003. <http://www.census.gov/prod/2005pubs/p23-208.pdf>.
[12] M. Crovella, A. Bestavros, Self-similarity in world wide web traffic: evidence and possible causes, IEEE/ACM Transactions on Networking (5) (1997).
[13] M. Faloutsos, P. Faloutsos, C. Faloutsos, On power-law relationships of the internet topology, in: Proceedings of the ACM SIGCOMM, 1999.
[14] <http://www.findanisp.com>.
[15] L. Guo, E. Tan, S. Chen, Z. Xiao, O. Spatscheck, X. Zhang, Delving into internet streaming media delivery: a quality and resource utilization perspective, in: Proceedings of the 6th ACM SIGCOMM on Internet measurement, 2006.
[16] Y. He, G. Siganos, M. Faloutsos, S.V. Krishnamurthy, A systematic framework for unearthing the missing links: measurements and impact, in: Proceedings of the 4th USENIX Symposium on Networked Systems Design and Implementation (NSDI'07), 2007.
[17] <http://www.internettrafficreport.com/>, 2008.
[18] <http://www.ip2location.com>.
[19] <http://iplane.cs.washington.edu/data/alias_lists.txt>.
[20] S. Jin, A. Bestavros, Small-world characteristics of internet topologies and implications on multicast scaling, Computer Networks: The International Journal of Computer and Telecommunications Networking 50 (5) (2006) 648–666.
[21] M.J. Karam, F.A. Tobagi, On traffic types and service classes in the internet, in: Proceedings of the Globecom'00, 2000.
[22] A. Lakhina, J.W. Byers, M. Crovella, P. Xie, Sampling biases in ip topology measurements, in: Proceedings of the INFOCOM'03, 2003.
[23] J. Leguay, M. Latapy, T. Friedman, K. Salamatian, Describing and simulating internet routes, in: Fourth International IFIP-TC6 Networking Conference, pp. 2–6.
[24] W.E. Leland, M.S. Taqqu, W. Willinger, D.V. Wilson, On the self-similar nature of ethernet traffic extended version, IEEE/ACM Transactions on Networking 2 (1994) 1–15.
[25] M. Liljenstam, D.M. Nicol, On-demand computation of policy based routes for large-scale network simulation, in: Proceedings of the 2004 Winter Simulation Conference.
[26] <http://www.leichtmanresearch.com/press/081108release.html>.
[27] H.V. Madhyastha, T. Anderson, A. Krishnamurthy, N. Spring, A. Venkataramani, A structural approach to latency prediction, in: Proceedings of the 6th ACM SIGCOMM conference on Internet measurement, Rio de Janeriro, Brazil, 2006.
[28] T.N. McPherson, M.J. Brown, Estimating daytime and nighttime population distributions in US cities for emergency response activities, Preprint, 84th AMS Annual Meeting, Seattle, WA.
[29] A. Mislove, M. Marcon, K.P. Gummadi, P. Druschel, B. Bhattacharjee, Measurement and analysis of online social networks, in: Proceedings of the ACM IMC'07, 2007.
[30] M.E.J. Newman, Mathematics of networks, The New Palgrave Encyclopedia of Economics, 2008.
[31] V. Paxson, S. Floyd, Wide-area traffic: the failure of poisson modeling, IEEE/ACM Transactions on Networking 3 (1995) 226–244.
[32] J. Qiu, L. Gao, As path inference by exploiting known as paths, in: Proceedings of the Globecom'06, 2006.
[33] S. Saroiu, P.K. Gummadi, S.D. Gribble, A measurement study of peer-to-peer file sharing system, in: Proceedings of the Multimedia Computing and Networking, 2002.
[34] N. Spring, M. Dontcheva, M. Rodrig, D. Wetherall, How to resolve ip aliases, Technical Report UW-CSE-TR 04-05-04, Department of Computer Science and Engineering, University of Washington, 2004.
[35] W. Tan, W. Cui, J.G. Apostolopoulos, Playback-buffer equalization for streaming media using stateless transport prioritization, Packet Video, 2003.
[36] Telegeography, a research division of primetrica, inc. <http://www.telegeography.com/>.
[37] <http://www.jcho.de/jc/Pubs/issls2000-col.pd>.
[38] <http://www.usipv6.com/6sense/2005/sep/03.htm>.
[39] <http://www.foxnews.com/story/0,2933,347035,00.html>.
[40] D. Vukadinovic, P. Huang, T. Erlebach, A spectral analysis of the internet topology, in: Proceedings of the DIMACS Workshop on Internet and WWW Measurement, Mapping, and Modeling, 2001.
[41] D. Watts, S. Strogatz, Collective dynamics of 'small-world' networks, Nature 393 (1998) 440–442.
[42] M. Zukerman, T.D. Neame, R.G. Addie, Internet traffic modeling and future technology implications, in: Proceedings of the IEEE Infocom'03.

Guanhua Yan obtained his Ph.D. degree in Computer Science from Dartmouth College, USA, in 2005. From 2003 to 2005, he was a visiting graduate student at the Coordinated Science Laboratory in the University of Illinois at Urbana-Champaign. He is now working as a Technical Staff Member in the Information Sciences Group (CCS-3) at the Los Alamos National Laboratory. His research interests are cyber-security, networking, and large-scale modeling and simulation techniques. He has contributed about 20 articles in these fields.

Stephan Eidenbenz received his Ph.D. degree in Computer Science from the Swiss Federal Institute of Technology (ETH) in Zurich in 2000. He is now a team leader in the Information Sciences Group (CCS-3) at the Los Alamos National Laboratory, where he leads the Multi-scale Integrated Information and Telecommunications System (MIITS) project that models and simulates large-scale communication networks. His research interests are in wire-line and wireless networking, sensor networks, selfish networking, infrastructure modeling, discrete event simulation, computational geometry, and algorithms. He has published about 50 articles in these fields.

Sunil Thulasidasan is a Technical Staff Member in the Computational Sciences Division of Los Alamos National Laboratory, where he develops distributed memory simulation software for High-Performance Computing platforms. He holds a Masters degree in Computer Science from the University of Southern California.

Pallab Datta is presently working as a Research Engineer in The NeuroSciences Institute in San Diego, CA. He was a Technical Staff Member in Computer, Computational, and Statistical Sciences Division at the Los Alamos National Laboratory of Department of Energy. He received his Ph.D. degree from Iowa State University in 2005 in Computer Engineering, and B.E. degree from the University of Allahabad, India in 1999 in Electronics Engineering. He has been conducting research in the areas of Optical Networks, Fault tolerance, High-Performance Computing, Optimization techniques, Graph algorithms and Wireless and Mobile communications since 1999. Prior to joining Los Alamos, he was working as a visiting researcher in INRIA, Sophia-Antipolis, France.

Venkatesh Ramaswamy is a senior Research Engineer at Airvana, Inc. in Chelmsford, MA, where he is currently involved in the design of next generation cellular networks. In 2005–2007 he spent two years at the Los Alamos National Laboratory, Los Alamos, NM, where he was actively involved in congestion control for high speed networks. Ramaswamy holds a Ph.D. degree in Electrical Engineering from the University of Mississippi.