

GeoTopo: A PoP-level Topology Generator for Evaluation of Future Internet Architectures

Yi Hu* Feixiong Zhang* K. K. Ramakrishnan[†] and Dipankar Raychaudhuri*

*WINLAB, Rutgers University
{yihu, feixiong, ray}@winlab.rutgers.edu

[†]University of California, Riverside
kk@cs.ucr.edu

Abstract—Network topology plays a critical role while designing and evaluating network protocols. Most existing topology generators are insufficient to reflect the real world network demands to a topology or to capture the Internet topology evolution such as the “flattening” Internet. They focused on the graph properties of a topology, thus, lacking of ability to model engineering features of the network. Some state-of-art topology generators that consider engineering factors fail to capture trends in both intra-AS and inter-AS connections, which are equally important for evaluating future network protocols.

We have developed a topology generator GeoTopo, which is to our best knowledge the first scalable topology generator modeling engineering factors for both intra-AS and inter-AS topology generation. The engineering factors that GeoTopo considers include demographic and geographic features as well as business interests of ASes. We use GeoTopo to create and study three classes of topologies: the topology characterized mainly by graph-properties (Status Quo topology), the topology driven by peering at Internet Exchange Points (IXP topology) and the topology characterized by country backbones (CB topology). The SQ topology follows the degree-based model and serves as a baseline for capturing topology features. The IXP and CB topologies model two major directions of the Internet “flattening”. The three classes of topologies enable us to analyze the impact of engineering factors on topology generation such as AS peering policies, IXP deployment and AS geo-settings.

GeoTopo’s ability to generate projected future Internet topologies make it a valuable tool for the design and evaluation of Future Internet Architectures that is currently under consideration in the research community. We use the evaluation of Global Name Resolution Service (GNRS), a key component shared by name-based network architectures, as an example application to demonstrate GeoTopo’s capability to capture the mobility of network entities, the locality of the traffic, and the impact of the evolving network.

I. INTRODUCTION

When designing and evaluating network architectures and protocols, it is desirable to map real world network demands to a topology for estimating performance, reliability and scalability. Such mapping requires the topology capture both global structures and local properties of the network, which rely on modeling engineering factors in the network design for both inter-AS connections and intra-AS topologies. For instance, to map the network traffic from web applications, we need to represent users’ page view traffic from each city/town to corresponding websites as source to destination pairs in a topology. This in turn requires modeling the Point of Presence (PoP) deployment of each AS, and intra-AS and inter-AS connections. To do this, topology generation needs to consider engineering factors including the geographic and demographic features of PoPs and AS economic interests since they drive the

formation of inter-AS connections and thus play a fundamental role in the eventual topology of networks.

However, simply mapping current world network demands to a topology is not enough for evaluating network architectures and protocols due to the rapid network evolution. It is important to model the topology evolution for projecting future network demands to an evolved topology. Figure 1 shows three major aspects of the Internet and their evolving directions. First, the geo-coverage of networks is quickly expanding. The studies on the Internet evolution [17], [20] conclude that major content distribution networks (CDNs) (e.g., Google, Akamai, Limelight) have expanded to almost every region of the developed world and have an ongoing trend to further extend their geo-coverage globally. Second, the formation of inter-AS peering links is becoming richer and denser, which is an important driving factor for the Internet “flattening” [8], [15], [16]. The third direction of the Internet evolution is the growing number of ASes, especially the emerging edge ASes, which further increases the number of links between edge ASes and existing provider ASes [28].

None of existing topology generators fully capture the real world network demands and the Internet topology evolution, especially the above three aspects. The majority of topology generators focus on graph properties of the network, such as hierarchical structures (e.g., [14], [18]) or degree distributions (e.g., [9], [10], [13], [22], [26]). While it is important to study the theoretic graph properties of network topologies, the lack of engineering considerations makes this line of generators unable to map real world network demands to a topology. Recently, a topology generator iGen [12] has been proposed to model the geo-locations of PoPs and intra-domain structures. Li et al. [24] also proposed to model router technologies and user demands for intra-AS topology generation. However, these methods are only effective on generating topology at a single AS level, incapable to capture inter-AS connections. Agent-based network formation models [17], [25] have been developed to generate inter-AS connections by modeling AS economic interests and traffic flows, but they are restricted to generate inter-AS connections between a limited number of ASes (e.g., 500) without considering PoP level intra-AS topologies. The scale is far insufficient to model the Internet consisting of more than 40000 ASes with a growing trend.

To fully capture the real world network demands and the Internet topology evolution, we developed a PoP-level topology generator GeoTopo (code is available at [3]). Specifically, GeoTopo considers engineering factors including geographic and demographic features of PoPs and economic interests of ASes. The three aspects of the Internet evolution are modeled by three groups of engineering factors. First, network geo-

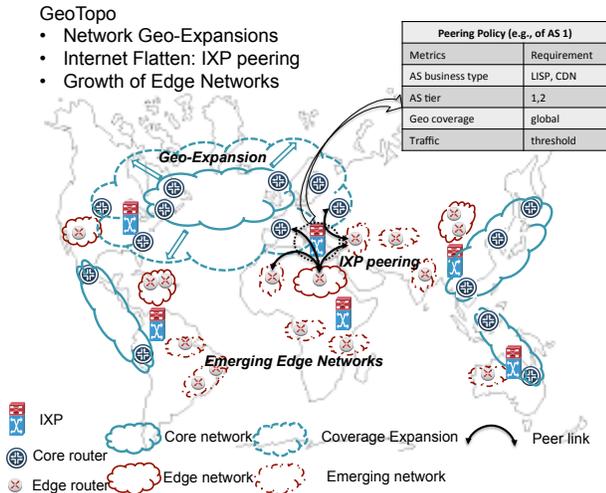


Fig. 1. GeoTopo Motivation

expansion is captured by modeling geographic, demographic and network infrastructure features of PoP deployment for each intra-AS topology. Second, the formation of inter-AS peering links is captured by modeling an AS’s hierarchical position, business type, peering policies, as well as the deployment of IXPs. Third, the emerging edge networks and their provider inter-AS connections are captured through our scalable topology generator with modeling on backbone ASes and customer-provider inter-AS connections.

To verify its applicability, we first use GeoTopo to generate synthetic topologies resembling a measured topology compared with a degree-based topology modeling method [26]. Then we use GeoTopo to study the evolving topologies and analyze the impact of AS peerings, IXP deployment and AS geo-scope in topology generation. As we focus on measuring the “flattening” of the Internet, we choose network expansion rates, link utilization ranks, and point-to-point latencies as metrics to measure the hierarchical and connectivity features of a topology. Finally, we explain how GeoTopo can be used to evaluate the upcoming protocols designed for Future Internet Architectures (e.g., [11], [21], [30]). Since FIAs aim to support high mobility in the Internet, GeoTopo is ideal to map out the mobility related components in a topology. We use GeoTopo to evaluate an implementation of global name resolution services (GNRS) [33] designed for Mobility First FIA. The major benefits of using GeoTopo include that PoP-level evaluations of GNRS remove the discrepancy in the measured inter-AS latency data from different AS-level topology measurement sources (e.g., DIMES and iPlane dataset) and enable examining the mobility and traffic locality issues in protocol designs.

The rest of the paper is organized as follows. Section II describes the design of GeoTopo. Section III presents two use cases. The first case verifies how well topologies generated by GeoTopo can capture a measured topology. And the second case studies the impacts of important settings in topology generation. Section IV demonstrates the applicability of GeoTopo through the evaluation of a global name resolution service for FIA. Section V reviews the related works. We conclude the paper in Section VI.

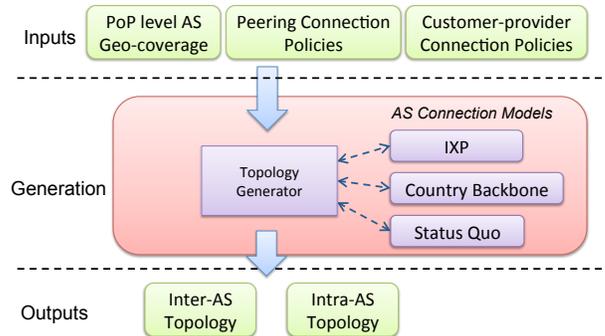


Fig. 2. GeoTopo Overview

II. DESIGN OF GEOTOPO

In this section, we present the architecture of GeoTopo and its functional elements in detail.

A. Architecture

An overview of GeoTopo is shown in Figure 2. Corresponding to the three motivations, GeoTopo takes three groups of inputs as building blocks allowing users to specify design features of their intended topology which may be a typical realistic topology or a projected future topology. We carefully condensed the inputs to the most essential factors to avoid burdening users of feeding the detailed topological facts that are difficult to obtain. The topology generation module generates PoP-level intra-AS and inter-AS topologies from the input settings. Three classes of inter-AS connection generation models are provided in GeoTopo to capture the current Internet topology and its evolution directions. The resulting topologies are termed as status-quo (SQ) topologies, IXP topologies, and country backbone (CB) topologies. In the rest of this section, we explain the input settings and the topology generation.

B. Input to Topology Generation

We choose the most essential factors that affect the network topology as inputs to GeoTopo. Realistic intra-AS PoP connections are resulted from design optimizations of PoP locations to minimize operation costs while maintaining high availability and low latencies (i.e., small world properties) [23], [29]. An important feature of PoP deployment is that there is a significant geographic locality in AS connections. That is the inter-AS connections and cross-country intra-AS connections are commonly located in a small number of cities which are covered by most of ASes. For example, there are 10 large metropolitan cities are covered by more than 85% of ASes in US. The PoPs deployed in such cities are called gateway PoPs, where inter-AS links and international intra-AS links are usually formed. Further, domestic intra-AS links often cluster PoPs to physically nearby gateway PoPs for achieving small world properties. Therefore, in GeoTopo we model the demographic and geographic features of PoPs, select gateway PoPs and generate intra-AS topologies to achieve small world properties. IXP distributions and AS peering policies determine where and how AS peering links are formed. Thus, in GeoTopo, we model IXP distributions and a set of business, hierarchical and geographic features of ASes to manifest AS peering policies and to model distinctive growing features of

each type of ASes. A class of IXP topologies is generated in GeoTopo to model IXP centered AS peering connections in the Internet topology. Another topology evolution trend is found by studies [5], [29] showing that individual edge ASes are connecting to multiple service provider ASes for improved availability and performance (i.e., “site multi-homing”) and prefer to connect to country backbone networks. The reason is that highly centralized country backbone structures facilitate governments to control their local areas of the Internet, like filtering offensive sites, surveilling user activities, and limiting users’ access to the Internet. The measurement [5] shows that the number of ASes that have access to more than 90% traffic within a given country is often only a few. Such ASes are called country backbone ASes. Specifically, in more than 75% countries, the Internet traffic is covered by no more than 5 country backbones. In GeoTopo, a class of CB topologies is generated to model customer-provider inter-AS connections featured by country backbones in the Internet topology. Finally, we provide a class of SQ topologies modeling the degree based graph properties of the topology for comparison.

PoP level AS geo-coverage includes the following three inputs for generating intra-AS topologies. *Geography and demography database* defines the geographic and demographic statistics of the countries and cities where PoPs may be located. The densities of AS coverage at a country and a city are also specified in the database. *AS geo-scope distributions* define the geo-scope levels of ASes from each tier. The geo-scope of an AS is the minimal geographic scope containing the majority (e.g., 95%) of its PoPs. We categorize geo-scope of an AS to four levels: metro, regional, country and global-level. We follow the conventional three-tier hierarchy for its prevalent usage. *Geographic connection constraints* define the enforced and excluded network connections between countries. These rules are used to enforce geographic, political, economic preferences and constraints. For example, users may specify no direct link between Europe and Australia as measured in [31].

Peering connection policies consist of inputs for defining where and how to generate peering links. *IXP deployment* defines the geo-locations of IXPs, which takes inputs in the form of geography or demography criteria or direct geo-locations. The AS peering policies are expressed by requirements on a AS’s hierarchical position, business type and geo-coverage. *AS hierarchical distribution* defines the number of ASes in each tier of the three-tier Internet hierarchy. *AS business type distributions* define the business type distribution of ASes in each tier. In GeoTopo, four business types are provided: large ISPs (LISP) (e.g., ComCast, AT&T), small ISPs (SISP) (e.g., Cable One, Juno), hosting and content distribution networks (HCDN) (e.g., Google, Akamai) and academic and enterprise networks (AEN) as in [8]. We group ASes of the same tier, business type and geo-coverage level to adopt a set of peering policies defined as the threshold of their desired peer ASes’ hierarchical positions, business types and geo-coverages.

Customer-provider connection policies consist of two sets of inputs for generating customer-provider AS links. *Backbone model* specifies whether globally shared backbone ASes or country backbone ASes are adopted. The number of backbone ASes for each country can be specified in the geography and demography database. *Minimal provider AS number* specifies the minimal number of customer-provider

links for a non-backbone AS.

C. Topology Generation

GeoTopo generates a topology instance in three steps. First, PoPs of each AS are deployed in the topology map, then individual intra-AS topologies are generated, finally the inter-AS connections are formed using three topology models – SQ, IXP, and CB. The detailed procedure is described as follows.

(1) Deploy intra-AS PoPs in a topology map. We begin with probabilistically selecting a base country for each AS, where the probability is weighted by the corresponding coverage density rank of each country defined in *geography and demography database*. Higher the rank, larger the probability of being selected. Let $\Pi(r_i)$ denote the probability for choosing country i with corresponding coverage density rank r_i . $\Pi(r_i)$ is defined as

$$\Pi(r_i) = r_i^\alpha / \sum_{j \in C} r_j^\alpha \quad (1)$$

, where C is the set of countries in the database, $\alpha > 0$ is a parameter indicating the preference for a high rank country to be covered. Then, we probabilistically select a base PoP from gateway cities in the base country, where the probability is weighted by coverage density ranks of candidate gateway cities. Once the base PoP is determined, we instantiate other PoPs according to the AS’s geo-scope, its intra-PoP number, and coverage density ranks of candidate cities. For example, in a metro-scope AS, PoPs are selected within the diameter of a metro area from the base PoP. Similarly, in a regional-scope or country-scope AS, PoPs are selected within the regional-diameter or the country boundary from the base PoP, respectively. In a global-scope AS, we first determine the number of continents it covers based on *continent coverage distribution* (e.g., N_1), then determine the number of countries it covers based on the *country coverage distribution* (e.g., N_2). We evenly selected N_2 countries from N_1 continents. Finally, we instantiate PoPs in each selected country the same way as in a country-scope AS. The number of intra-PoPs is evenly distributed among each selected country.

(2) Generate intra-AS connections. The intra-AS links are created to achieve the small world properties that keep local traffic staying local while wide-area traffic encountering fewer AS and PoP hops [13], [31]. Specifically, such properties are obtained by clustering non-gateway PoPs around nearby gateway PoPs, and densely connecting gateway PoPs with each other [24]. Therefore, in GeoTopo, if an AS spans multiple countries, the gateway PoPs of each country are connected to form a mesh. Within a country, non-gateway PoPs are clustered to the nearest gateway PoP. When an average intra-PoP connection degree requirement is set, additional intra-AS PoP connections may be created using the generalized linear preference method [13].

(3) Generate inter-AS connections following one of the three models. A **Status-quo topology** is initialized by connecting all tier-1 ASes as a mesh and each tier-2 and tier-3 AS to a minimal number of provider ASes from upper tiers. Additional inter-AS connections are created using the Generalized Linear Preference (GLP) method [13] according to the inter-AS connection degree distribution. The GLP method

is chosen because of its capability to achieve both the power-law distribution and small-world properties. Most degree based generators do not consider small-world properties, which represent local robustness of the graph and path diversity with practical impacts such as efficiency of searching and routing. In the GLP method, a new inter-AS connection is created by probabilistically selecting two endpoint ASes based on the current AS connection degree. Let $\Pi(d_i)$ denote the probability for choosing AS i when adding a new inter-AS link, where d_i is the connection degree of AS i . $\Pi(d_i)$ is defined as

$$\Pi(d_i) = (d_i - \beta) / \sum_j (d_j - \beta) \quad (2)$$

, where $\beta \in (-\infty, 1)$ is a parameter indicating the preference for a new link connecting to more popular ASes. Figure 3 shows an example SQ topology. At the initialization, AS1, AS2 and AS3 are fully connected as they are tier-1 ASes, and other ASes are connected to the minimal number (1 in this example) provider-AS. Additional inter-AS connections are generated by the GLP method (we do not show all the additional inter-AS connections in the figure.)

An **IXP topology** is initialized in the same way as a SQ topology, while, additional inter-AS links are created from ASes peering with each other at IXPs instead of using the GLP method. Figure 4 shows an example IXP topology. We first deploy IXPs in cities based on *Geography and demography database* input. Each country is ensured to have at least one IXP. If there is no IXP deployed in a country from the database, we choose the city with the largest population in the country to deploy an IXP. Each AS participates in every IXP which is co-located in the same city as one of its intra-AS PoPs. At each IXP, all ASes participating in this IXP peer with each other based on peering criteria specified in *Peering connection policies*. An inter-AS connection is created between a pair of peering ASes.

A **Country backbone topology** is initialized as each country has its backbone AS(es) and the backbone ASes from tier-1 countries are fully connected. Each backbone AS of a non tier-1 country is connected to a number of tier-1 country backbones, specified as the minimal provider-connection number in *Customer-provider connection policies*. Each tier-2 AS is connected to a number of backbone ASes of the country/countries it covers and each tier-3 AS is connected to a number of tier-2 ASes specified as the minimal provider-connection number. Figure 5 is an example of CB topology. For country backbones, we initialize the inter-AS connections by connecting tier-1 country backbones (e.g., CB1, CB2 and CB3) as a mesh and each tier-2 country backbone (e.g., CB4) to a minimal number (1 in this example) of tier-1 country backbones. Additional inter-AS connections are created through AS peering at IXPs the same way as in an IXP topology.

D. Output of Topology Generation

GeoTopo chooses to generate topologies at the PoP level for two major reasons. First, when studying the Internet, PoP-level topologies give a better level of aggregation than router level or IP level topologies while having the ability to examine important features of ASes such as intra-AS PoP deployments, geographic coverages, routing path diversities. Second, PoP level topologies enable us to annotate PoPs with

geographic, demographic and economic information to develop a realistic topology generator and a reliable forecast framework predicting the growth of the Internet.

III. CASE STUDY

In this section, we present two experiments of GeoTopo. In the first experiment (Section III-A), we use GeoTopo to generate synthetic topologies to model a measured topology with comparison to a degree based topology generation method [26]. We aim to verify that by properly applying the IXP model in GeoTopo to generating inter-AS topologies, GeoTopo outperforms the degree based method because of the capability to model geo-locations and policies that affect inter-AS connections. In the second experiment (Section III-B), we analyze the impact of evolving topologies on network structures and performance. We show that by adjusting the inputs, GeoTopo can be a useful tool for users to test their network designs on the outcome topologies.

A. Synthetically Model Measured Topology

The goal of this experiment is to show that GeoTopo is able to capture a measured topology by extracting the geographic settings and modeling inter-AS connection policies. We compare with a degree based topology generation method [26], which is the most advanced degree based method and can accurately reproduce the expected inter-AS degree of each AS through modeling the joint degree distribution. The reason why we choose to compare with a degree based method is because the only state-of-art geo-location enabled topology generator iGen [12] requires manual inputs for generating inter-AS topologies. As iGen requires inputs of an AS connection graph and link connections for each pair of connected ASes, making it unsuitable for comparison.

The measured topology we use in this experiment is from the RocketFuel data [31], which contains intra-AS and inter-AS topologies over 1900 PoPs of 70 ASes (i.e., 21 tier-1 ASes, 36 tier-2 ASes, and 13 tier-3 ASes). We perform a sanity check on the RF topology data: (1) For each inter-AS link, we ensure the two endpoint PoPs exist in the corresponding AS's intra topology. (2) We ensure no partition in any intra AS topology.

1) *Experimental Methodology*: We focus on modeling the inter-AS topology and choose the IXP model in GeoTopo to generate synthetic topologies. We find that in RF topology over 50% of the inter-AS links are between PoPs co-located in a set of cities. Such inter-AS connections can be well modeled by the IXP topology. We extract the IXP deployed cities from RF topology and for each AS we extract their peering policies based on their tier information. Following the study in [8], we assume that ASes are open to peer with other ASes from the same or upper tiers but selective to peer with ASes from lower tiers. For example, a tier-2 AS is open to peer with tier-1 and tier-2 ASes but selective on peering with tier-3 ASes. Define x_{ij} for each AS in tier i to be the percentage of tier j ASes with whom this AS is willing to peer, where $i, j \in 1, 2, 3$. For $j \leq i$, $x_{ij} = 1$. We set x_{ij} for $j > i$ from RF topology. The set of x_{ij} of each AS is used as its peering matrix at each IXP where it participates in. The intra-AS topologies are directly set based on RocketFuel data. For degree based method we profile the inter-AS degree distribution and the joint

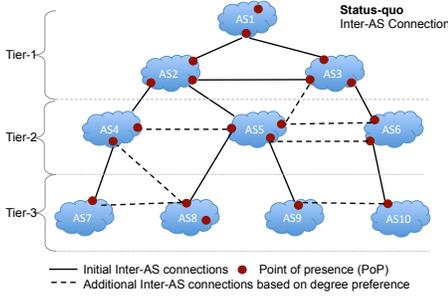


Fig. 3. An example of SQ topology.

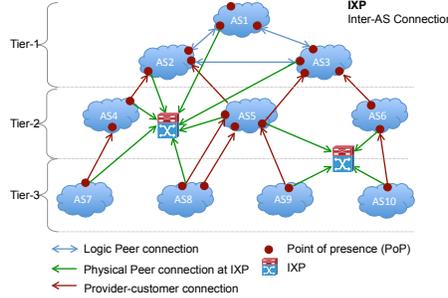


Fig. 4. An example of IXP topology

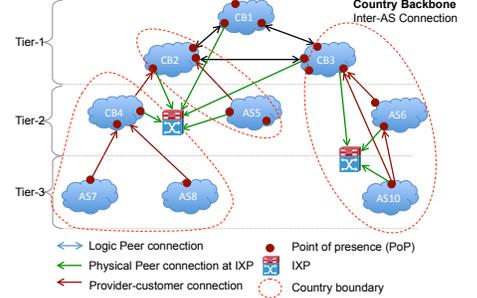


Fig. 5. An example of CB topology

degree distribution and follow the method in [26] to generate synthetic topologies. Due to their lack of modeling inter-AS links between a pair of connected ASes, we choose the pair of nearest PoPs from each AS to form the inter-AS link. We repeated the synthetic topology generation experiments for 50 runs to plot the comparison results below.

2) *Experimental Metrics*: We select three representative network metrics [32] to analyze network topological features: network connectivity expansion rates, inter-AS link ranks, and point-to-point inter AS latencies.

To understand the extent of hierarchical structure a network has, we measure the network connectivity expansion rate of an AS/PoP and the rank of each inter-AS link. As the notion of hierarchy is that a set of backbone ASes connecting with each other and serving as a core network to provide network access for edge ASes. Thus, the network expansion rate for each AS is an indicator of the network hierarchy as backbone ASes have faster expansion rates than edge ASes. Specifically, *network connectivity expansion rate* $E(h)$ of an AS is the fraction of ASes that can be reached within h hops from an AS. Similarly, *network connectivity expansion rate* $E(h)$ of a PoP is the fraction of PoPs that can be reached within h hops from a PoP. We categorize ASes/PoPs based on their tier information and plot the min, 25 percentile, median, 75 percentile and max values for each tier. In addition to measure AS connectivity, the links connecting backbone ASes may carry traffic from a larger number of source-destination pairs than the links connecting edge ASes. Thus, for each link, we measure its utilization by a *link rank*. The rank of a link i is defined as the percentage of PoP source-destination pairs that traverse through link i using the shortest path routing. For example, given a network with n PoPs, there is r_i number of source-destination pairs' shortest paths traversing through link i , then link i 's rank is $\frac{r_i}{n(n-1)}$. A backbone link has a higher rank than a peripheral link.

To have a basic understanding about the network delay between source-destination points in the network, we measure point-to-point inter-AS latencies. For each source-destination pair of ASes the *point-to-point inter AS latency* is defined as the latency between each PoP in the source AS and that in the destination AS. We plot the cumulative density function (CDF) results. The shortest-path routing is used to calculate point-to-point paths between pairs of PoPs in both RocketFuel topology and synthetic topologies. The latency of each path is calculated as the geo-latency between the two endpoint PoPs, whose geo-locations are provided in RocketFuel topology.

3) *Experimental Results*: The comparison results are shown in Figures 6, 7, 8. GeoTopo outperforms degree based method in modeling the network hierarchy and network latencies. As shown by the inter-AS link ranks in Figure 7, GeoTopo captures the network hierarchy phenomenon that a small number of core links carrying massive amounts of network traffic aggregated from a large number of edge links. Specifically, both GeoTopo and RF topologies have less than 0.06% of links with ranks higher than 0.1 and over one third of links of ranks less than or equal to 10^{-7} . However, degree based method fails to capture such distinctive features between core and edge links. Also, the network delay between PoPs in GeoTopo matches those in RF topology as shown by the overlapping plots in Figure 8 (a)-(c). We select tier-1 to tier-1, tier-1 to tier-2, and tier-1 to tier-3 AS latencies (denoted as T1-T1, T1-T2, and T1-T3) to represent the inter-AS latencies among core networks, between core networks and edge networks. Whereas, the delay results of the degree based method show a gap behind those of RF topology. While both GeoTopo and degree based method are able to capture the network expansion feature of RF topology as shown by PoP expansion rates in Figure 6. GeoTopo's better performance comes from the inter-AS link modeling. Degree based method accurately models the inter-AS degree, but fails to model the links between a pair of connected ASes as no consideration of geo-locations or formations of inter-AS links. However, GeoTopo is able to capture the geo-location of inter AS links by modeling the IXP locations and capture the formation of inter-AS links by modeling the AS peering policies.

For inter-AS link ranks, GeoTopo's results have around 7% more links with extreme low ranks than RocketFuel's. The reason is that we enforce each AS to participate in every IXP within the AS coverage. However, in RocketFuel topology an AS may not participate in every IXP within its reach. Thus, synthetic topologies have more inter-AS links than the input topology. Given the same total number of PoPs and ASes, the sum of link ranks over all links is the same. Both topologies have similar number of links with high ranks. This causes the percentage of low rank links in synthetic topologies is slightly higher than that in RocketFuel topology.

Overall, this experiment demonstrates that by embedding the geo-information and modeling IXP featured inter-AS connections, GeoTopo synthetic topologies can well capture the network hierarchical and delay features of an input topology, while degree based method only can capture the inter-AS expansion by modeling the inter-AS degree features but unable to model the inter-AS link formations due to lack of considering

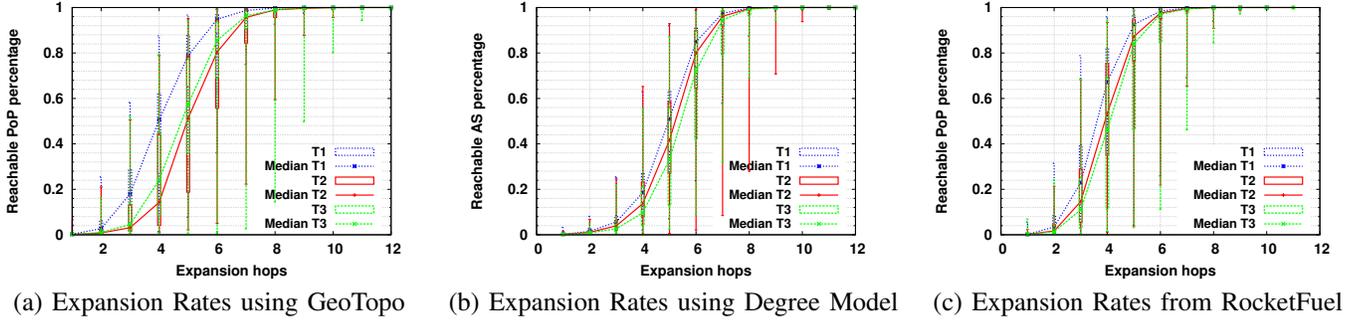


Fig. 6. Synthetic and Measured Topology Comparison – AS Expansion Rates

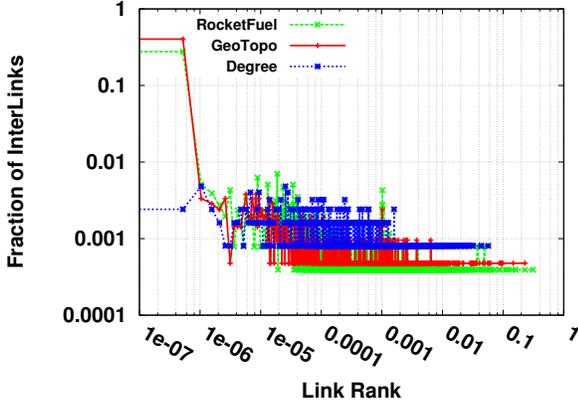


Fig. 7. Synthetic and Measured Topology Comparison – Inter-AS Link Ranks

other essential topology features.

B. Topology Evolution Model Comparison

The goal of this experiment is to show the impact of the Internet evolution on network topologies. Specifically, through the three inter-AS connection models in GeoTopo, we show two directions of the Internet “flattening”: denser inter-AS connections due to AS peering at IXPs and the expansion of core networks through country backbone networks. The comparison of outcome topologies reflect the impact of the corresponding input settings, which serve as examples to show how users can use GeoTopo to test their network designs by adjusting input settings and analyzing outcome topologies.

1) *Experiment Setup:* We use the status-quo topology (SQ topology) as a baseline for studying the evolving Internet. The Internet Exchange Point topology (IXP topology) and the country backbone topology (CB topology) present the two directions of the Internet “flattening”. In a SQ topology, inter-AS connections are generated to follow the power-law distributions of AS connection degrees and the small world properties of the inter-AS topology using the GLP method [13]. We set the parameters in the GLP method from the measurement data in [28]. Different from a SQ topology, in an IXP topology inter-AS peering connections are generated following ASes’ peering policies at IXPs. The IXP topology emphasizes the Internet evolution driven by flexible and dense peering connections between ASes, which relieve edge ASes’ reliance on customer-provider connections with core ASes. We use the world-wide IXP locations and AS peering policies measured in [8], [16] for generating the IXP topology. We compare the

IXP topology with the SQ topology to study how the inter-AS peering methods affect the topologies.

To examine the effects of IXP deployment, we deploy additional IXPs in each country, denoted as the IXP+ topology and compare the IXP topology with the IXP+ topology to examine benefits from more prevalent IXP deployment.

The CB topology features the Internet “flattening” resulted from the expansion of core networks. The difference between an IXP topology and a CB topology is the model of backbone ASes. In an IXP topology, similar to a SQ topology, we use the existing 21 global backbone ASes as the core networks which connect with each other as a mesh. In contrast, in a CB topology we model each country has its own backbone ASes, replacing the 21 global ones. It is infeasible for all countries’ backbones peering with each other, thus, we classify countries into tier-1 countries and tier-2 countries, where tier-1 countries’ backbones are tier-1 ASes. The tier-1 countries are countries that currently have core network coverage. There is in total 43 tier-1 countries [31]. Each tier-2 country backbone connects to at least one tier-1 country backbone to ensure international network access. Non-backbone ASes connect to its country’s backbone ASes directly or indirectly. We compare the IXP topology with the CB topology to study the impact of expanded core networks.

We use 21 tier-1 ASes, 2000 tier-2 ASes, and 20000 tier-3 ASes for synthetic topologies. And each result is from 50 runs of the experiment.

2) *Impact of AS Peering:* We first compare AS connectivity expansion rates in the SQ topology and the IXP topology as shown in Figure 9 (a) and (b). We notice two changes: first, the difference between three tiers is reduced in the IXP topology from that in the SQ topology. In both topologies, the three tier hierarchical network structure is reflected in AS connectivity expansion rates as tier-1 ASes have the fastest expansion, tier-2 ASes in the middle and tier-3 ASes have the slowest. However, the hierarchical difference become less obvious in the IXP topology than in the SQ topology. For example, within 2 hops, the three rectangles representing the major range of ASes from three tiers (i.e., 25 percentile to 75 percentile values) are clearly distant from each other in the SQ topology, while they are much closer in the IXP topology. The reduced AS hierarchical difference in the IXP topology is due to edge ASes participating at IXPs with open peering policy. Second, the connectivity expansion rate of each tier well reflects the AS peering policy. Most benefits of IXP peering come to edge ASes with open peering policy, that is, most tier-3 ASes and some of tier-2 ASes in the IXP topology

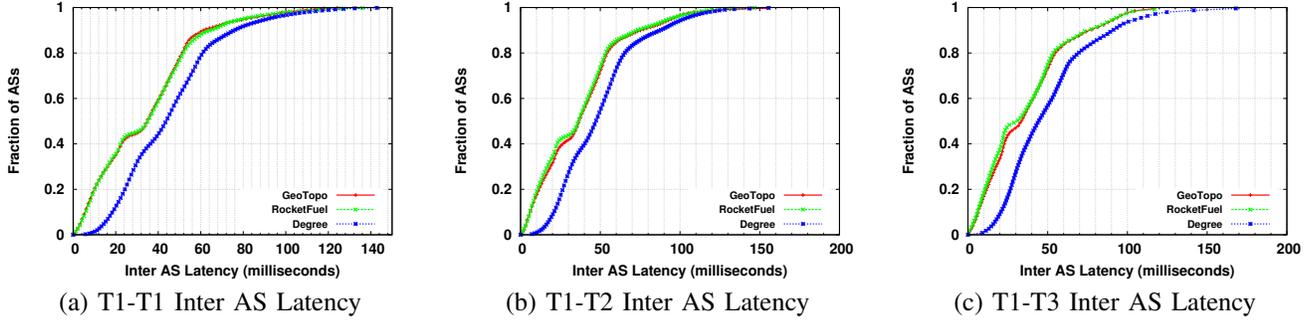


Fig. 8. Synthetic and Measured Topology Comparison – Inter-AS Latencies

have improved connectivity expansion rates as shown in Figure 9 (b) and reduced inter-AS latencies as shown in Figure 10 (a) (b) (c). For instance, within 2 hops, the maximum reachability of tier-3 ASes increases from 44% in the SQ topology to 95% in the IXP topology. This shows that with open peering policies, edge ASes have the chance to expand network connectivity faster relying on peering connections. On the other hand, the peering policy of tier-1 ASes is much more selective, where only 10% of tier-1 ASes peer with certain tier-2 or tier-3 ASes and the rest tier-1 ASes only peer with each other in the same way as in SQ topology, which makes tier-1 ASes' connectivity expansion rates remain the same in both topologies.

We then compare inter-AS latencies as shown in Figure 10. In both topologies, the core ASes (tier-1 ASes) are connected with each other, their inter-AS latencies are almost the same. Consistent with the observation from expansion rate results, ASes with more open peering policies gain larger improvements in inter-AS latencies. Specifically, tier-3 ASes have reduced latencies to other edge networks and to core networks as shown in Figure 10 (b) and (c). Tier-2 ASes also have improvement but less than that of tier-3 ASes comparing Figure 10 (a) and (c).

3) *Impact of IXP Deployment:* We deploy one more IXP in each country by selecting the most populated city in the country where IXP has not been deployed yet. We denoted this topology as IXP+1 topology. Keeping other settings unchanged, adding more IXPs improves AS connectivity for all tiers, while the improvement margins differ. We show the inter-AS latency results to represent the improvement margins. The tier 1-to-tier 1 (T1-T1) inter AS latency as shown in Figure 12 (a) has the largest improvement, the improvement of T2-T2 inter AS latency is in the middle as shown in Figure 12 (b) while that of T3-T3 as shown in Figure 12 (c) is the smallest. The reason is that geo-scope affects the extent of benefits attained from extra deployed IXPs. Tier-1 ASes have the largest geo-scope. Thereby they have the highest chance that newly deployed IXPs are within their geo-coverage so that they can participate to get more peer connections. Most tier-2 ASes are of country-scope. Thus they get fewer benefits from newly-deployed IXPs. Tier-3 ASes have the smallest geo-scope, that is, the majority is of regional or metro level. Therefore they have the least chance of benefiting from newly deployed IXPs.

4) *Impact of Expanded Core Networks:* Comparing the IXP topology and the CB topology, both use the IXP peering model for creating inter-AS connections. A major difference comes from expanded core networks, which is set using AS

geo-scope inputs. In the IXP topology, the AS geo-scope distribution is based on the measurement data of the Internet topology [29]. According to the data, the majority of tier-1 ASes are of global-level, most of tier-2 ASes are either of global-level or country-level, and most of tier-3 ASes are either of regional or metro-level while there are a few of global- and country-level tier-3 ASes. However, in the CB topology we set the AS geo-scope from our projections of country backbone featured Internet structure. In other words, all tier-1 ASes are of country scope, the most majority (e.g., 90%) of tier-2 ASes are of country scope, and the rest (e.g., 10%) of tier-2 ASes are of regional scope. And tier-3 ASes are equally split between regional- and metro-level scope. Overall, tier-1 ASes in the CB topology have more restricted geo-scopes than those in the IXP topology. Tier-3 ASes also have reduced geo-scopes but the reduction is less than that of tier-1 ASes. Tier-2 ASes in the CB topology have more country level and less global level than those in the IXP topology. The geo-scope changes are clearly reflected in the network expansion rates as shown in Figure 9 (b) and (c) as tier-1 and tier-3 ASes in the CB topology have slightly lower connectivity expansion rates than those in the IXP topology due to their reduced geo-scopes. The results of tier-2 ASes in the CB topology have less variety than those in the IXP topology because their geo-scopes become more uniform. Each country having its own backbone networks improves inter-AS latencies from edge ASes to backbone ASes. The T1-T2 and T1-T3 inter-AS latencies are shortened in the CB topology as shown in Figure 13 (a) and (b) because backbone networks are more widely spread. As tier-2 countries' backbone networks are tier-2 ASes, the T2-T2 inter-AS latency is also improved, as shown in Figure 13 (c).

IV. APPLICATIONS

This section shows the applicability of GeoTopo for FIA research. We give a detailed example of using GeoTopo to evaluate a name resolution service in Mobility First FIA.

A. Evaluation of GNRS

As an application example, we use GeoTopo to evaluate a Global Name Resolution Service (GNRS) [33] in Mobility First FIA. The GNRS protocol works as follows: each network entity has a globally unique identifier (GUID) for identification and network address(es) (NA) as its locators for routing. A GUID is published to all users that allow to connect with the entity. NA of an entity may be changed due to mobility, hence, a user needs to query GNRS using the GUID to get

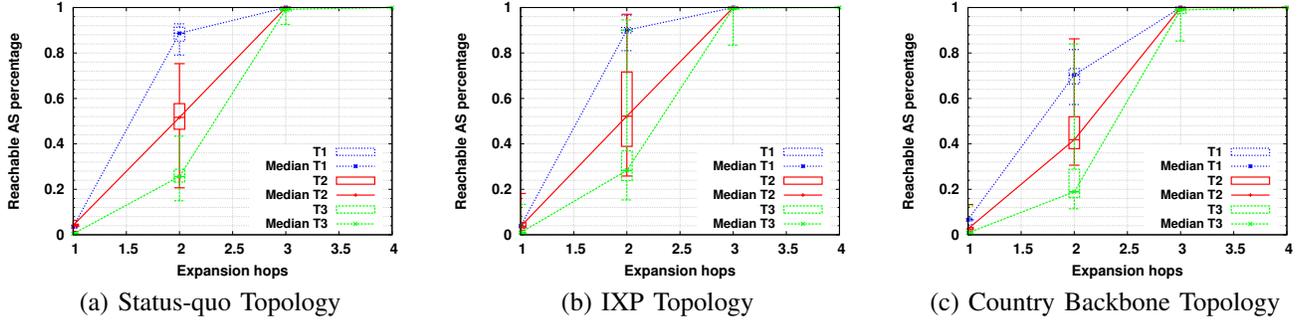


Fig. 9. AS Expansion Rate Comparison

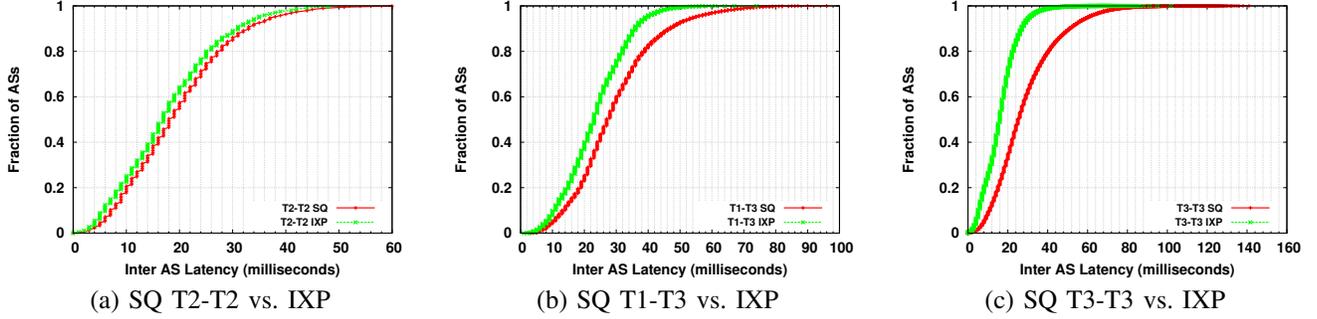


Fig. 10. Inter AS Latency Comparison SQ vs. IXP Topology

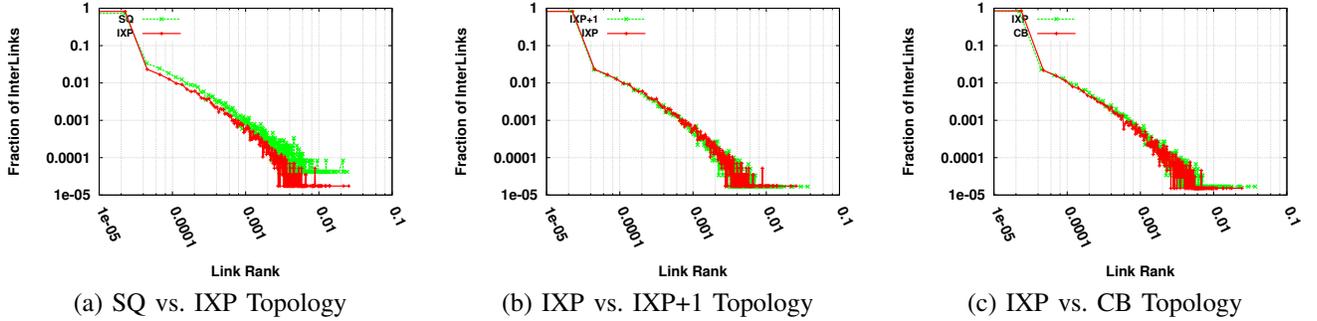


Fig. 11. Link Rank Comparison

the correct NA before connecting to the entity. One way to implement GNRS is to replicate a mapping of GUID to NA to K ASes (called host ASes), where every update of the mapping is applied to K replicas and every query is answered by one of the K replicas. Using consistent hash functions to determine the K host ASes, users can correctly calculate the host ASes using GUID and the hash functions. Therefore, users can query the nearest host AS to obtain the mapping information.

GeoTopo provides two major benefits in GNRS evaluations. First, we can evaluate the point-to-point latency of GNRS operations at PoP level in addition to AS level. Here we choose the query latency as the example metric. The measured inter-AS latency from different resources show large discrepancies which affect the validity of GNRS evaluation. Figure 14 (a) shows the measurement data of the inter-AS latency provided by DIMES [2] and iPlane [4] datasets. As the median value in iPlane dataset is 4 times larger than that in DIMES, using which dataset to evaluate GNRS will greatly alter the latency results. Given that DIMES and iPlane use different methods in collecting data, it is difficult for us to decide which one is more reliable to use. Therefore, we need

GeoTopo to provide a PoP-level topology embedding the geolocations of PoPs so that the GNRS query latency can be evaluated more reliably through the routing between source-destination PoPs rather than directly using inter-AS latencies from either of the dataset. For example, we use the web traffic of top 10K popular websites from Alexa Data [1] as the GNRS query workload and use GeoTopo to generate SQ, IXP and CB topologies with PoPs located in 18000 cities over more than 80 countries. The inter-AS latency calculated through shortest path routing in the GeoTopo topologies is shown in Figure 14 (a), whose values are in between DIMES and iPlane data.

Second, the PoP-level evaluation enabled by GeoTopo allow us to examine locality-related GNRS design and performance issues. Specifically, we need PoP-level topologies to map the real world workload from cities to PoPs to test how efficiently GNRS can utilize the workload locality to improve performance. We compare the GNRS query latency with the total number of replica $K = 5$ in three classes of GeoTopo topologies as shown in Figure 14 (b). The results in Figure 14 (b) are consistent with the inter-AS latency comparison between the three classes of GeoTopo topologies as shown in

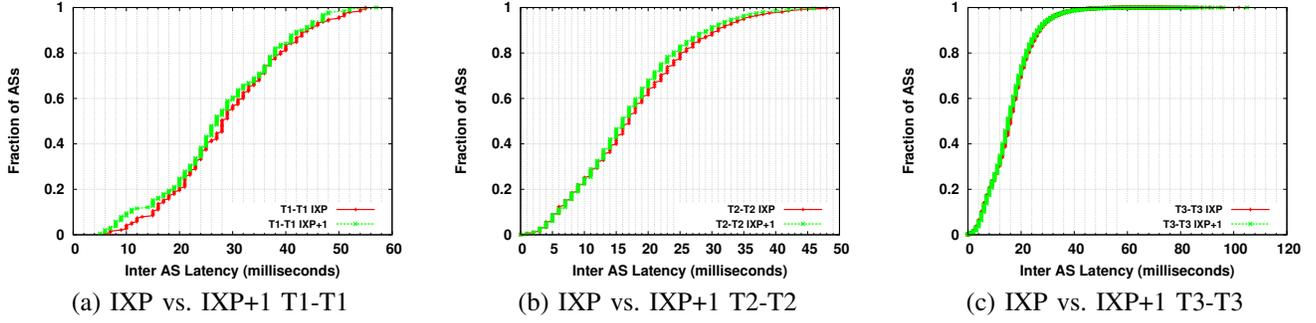


Fig. 12. Inter AS Latency Comparison IXP vs. IXP+1 Topology

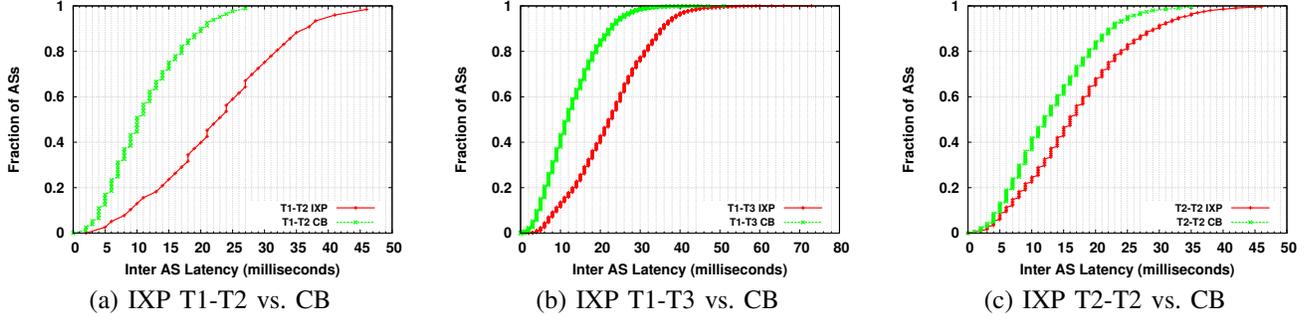


Fig. 13. Inter AS Latency Comparison IXP Vs. CB Topology

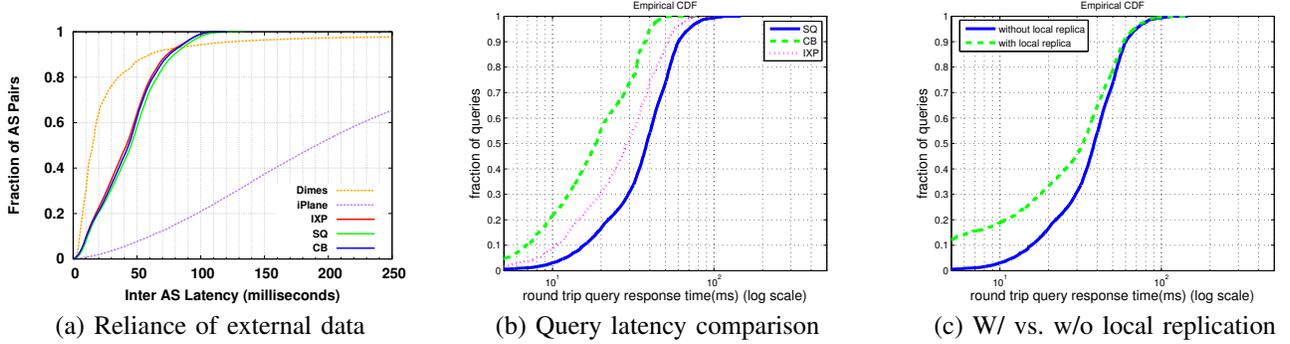


Fig. 14. GNRS Query Latency Comparison

Figure 10 and Figure 13. To examine the benefits of taking advantage of spatial locality exhibited in query workload, we deploy an extra replica at the local AS where the GUID connects to the network. Figure 14 (c) shows the query latency improvement in SQ topology. The extra local replica improves the percentage of queries that can be served extremely quickly. For example, without local replica only 5% of the queries can be served within 10ms and local replica improves that to 20%.

In sum, GeoTopo provides a tool for mapping the real workload from cities to topologies and enables the examination of locality issues.

V. RELATED WORK

Topology generator. A long line of topology generators have been designed for modeling the Internet topology. The first widely used topology generator [34] for protocol simulations was a variant of the classical Erdos-Renyi random graph. After that, structural network topology generators, such as Transit-Stub [14] and Tiers [18], were designed to generate topologies based on hierarchical structural of networks. Falout-

os et al. [19] found that the node degree distributions of both router-level and AS-level internet graphs follow a power-law distribution. A line of degree-based topology generators [9], [10], [13], [22] were then designed to produce network topologies matching the Internet’s power law degree distribution. Mahadevan et al. [26] defines dK-series to capture the joint node-degree distribution of a topology and construct random graphs that reproduce a given topology with graphic property metrics. Extensive studies [27], [32] have been devoted to compare the two categories of generators. Validated through data of the measured AS-level [6] and PoP-level Internet topology [7], it is found that degree-based generators can better capture the measured topology than structure generators.

A missing piece in degree-based topology generators is the consideration of engineering factors because the real networks are the outcome of network optimization design processes and constrained by technical and economical settings. Li et al. [24] argued the importance of incorporating router technology constraints and economic considerations in studying network topologies. However, their method focused on the principles of designing network topology of a single ISP. The only

existing topology generator iGen [12] that is able to model certain engineering aspects such as geographic features, intra-AS structures, requires inter-AS connections of a topology as input, which negates the purpose of topology generation. In contrast, GeoTopo extends the engineering modeling to generate both intra-AS topologies and inter-AS connections. An agent-based model to study the inter-AS relationships was presented in [25], however, their model lacks of PoP-level information and only scales to around 500 ASes, which cannot fully capture the Internet consisting of more than 40K ASes or the Internet growth featured by emerging edge networks.

Internet studies. A series of works [8], [16], [28], [29] have been devoted to studying the Internet evolution trends. The growth of the Internet size is examined in [28], where the authors developed an empirical model to analyze how fast the Internet topology changes. Their model shows that the growth is largely due to increasing number of customer networks while the number of provider networks remains relatively stable. In regards to the evolution of inter-AS connection, the analysis on world-wide IXPs [8], [16] show that the number of AS peering connections is orders of magnitude larger than what we assumed before, and IXP peering connections are actively used for carrying daily traffic instead of being backup. The diversity of participating networks at an IXP and their peering matrices are also investigated. Another aspect of the Internet evolution is network geography features, the authors in [29] analyze the geo-footprint of eyeball networks (i.e., ASes that directly provide services to end-users) and their connections to the rest of the Internet. The study uncovers the important geography features of ASes and the impacts of geography on intra-AS and inter-AS connections. Based on these studies, GeoTopo provides a tool for applying the major topology evolution trends in synthetic topology generation.

VI. CONCLUSIONS

In this paper we developed a network topology generator “GeoTopo” for modeling the Internet topology and its evolution. GeoTopo provides three models of generating inter-AS connections, which can be used to capture the Internet “flattening” with the real world geographic, demography and business settings. Through validating the GeoTopo synthetic topologies with a measured Internet topology, we show that IXP peering captures the major features of Inter-AS connections. Our experiments also show that GeoTopo enables us to evaluate mobility and locality issues in protocols designed for medium to long term usage in the Internet.

REFERENCES

- [1] Alexa database. <http://www.alexa.com/>.
- [2] The dimes project. <http://www.netdimes.org/>.
- [3] Geotopo: An evaluation tool for future internet study. <http://www.winlab.rutgers.edu/~yihu/geotopo/>.
- [4] iplane: Datasets. <http://iplane.cs.washington.edu/data/data.html>.
- [5] Mapping local internet control. <http://cyber.law.harvard.edu/netmaps/>.
- [6] Route views project. <http://www.routeviews.org/>.
- [7] Traceroute. <http://www.traceroute.org/>.
- [8] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a large european ixp. In *ACM SIGCOMM*, 2012.
- [9] W. Aiello, F. Chung, and L. Lu. A random graph model for massive graphs. In *ACM STOC*, 2000.
- [10] R. Albert and A.-L. Barabási. Topology of evolving networks: Local events and universality. *Physical Review Letters*, 85, 2000.
- [11] A. Anand, F. Dogar, D. Han, B. Li, H. Lim, M. Machado, W. Wu, A. Akella, D. G. Andersen, J. W. Byers, S. Seshan, and P. Steenkiste. Xia: An architecture for an evolvable and trustworthy internet. In *ACM HotNets-X*, 2001.
- [12] Q. Bruno, V. den Schrieck, F. Pierre, and B. Olivier. Igen: Generation of router-level internet topologies through network design heuristics. In *21st International Teletraffic Congress*, 2009.
- [13] T. Bu and D. Towsley. On distinguishing between internet power-law generators. In *IEEE INFOCOM*, 2002.
- [14] K. Calvert, M. Doar, and E. Zegura. Modeling internet topology. *IEEE Communications Magazine*, 35(6), 1997.
- [15] H. Chang, S. Jamin, and W. Willinger. To peer or not to peer: Modeling the evolution of the internet’s as-level topology. In *IEEE INFOCOM*, 2006.
- [16] N. Chatzis, G. Smaragdakis, J. Bottger, T. Krenc, and A. Feldmann. On the benefits of using a large ixp as an internet vantage point. In *ACM IMC*, 2013.
- [17] A. Dhamdhere and C. Dovrolis. The internet is flat: Modeling the transition from a transit hierarchy to a peering mesh. In *ACM CoNEXT*, 2010.
- [18] M. Doar. A better model for generating test networks. In *IEEE GLOBECOM*, 1996.
- [19] C. Faloutsos, M. Faloutsos, and P. Faloutsos. On power-law relationships of the internet topology. In *ACM SIGCOMM*, 1999.
- [20] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *PAM*, 2008.
- [21] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard. Networking named content. In *ACM CoNEXT*, 2009.
- [22] C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Technical Report CSE-TR-433-00, EECS Department, University of Michigan, 2000.
- [23] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan. The internet topology zoo. *IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS*, 29, 2011.
- [24] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the internet’s router-level topology. In *ACM SIGCOMM*, 2004.
- [25] A. Lodhi, A. Dhamdhere, and C. Dovrolis. Genesis: An agent-based model of interdomain network formation, traffic flow and economics. In *IEEE INFOCOM*, 2012.
- [26] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat. Systematic topology analysis and generation using degree correlations. In *ACM SIGCOMM*, 2006.
- [27] A. Medina, I. Matta, and J. Byers. On the origin of power laws in internet topologies. *ACM SIGCOMM Computer Communication Review*, 30(2), 2000.
- [28] R. Oliveira, B. Zhang, and L. Zhang. Observing the evolution of internet as topology. In *ACM SIGCOMM*, 2007.
- [29] A. H. Rasti, N. Magharei, R. Rejaie, and W. Willinger. Eyeball ases: From geography to connectivity. In *ACM IMC*, 2010.
- [30] D. Raychaudhuri, K. Nagaraja, and A. Venkataramani. Mobilityfirst: A robust and trustworthy mobilitycentric architecture for the future internet. *MC²R*, 16, 2012.
- [31] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of path inflation. In *ACM SIGCOMM*, 2003.
- [32] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network topology generators: Degree-based vs. structural. In *ACM SIGCOMM*, 2002.
- [33] T. Vu, A. Baid, Y. Zhang, T. D. Nguyen, J. Fukuyama, R. P. Martin, and D. Raychaudhuri. Dmap: A shared hosting scheme for dynamic identifier to locator mappings in the global internet. In *IEEE ICDCS*, 2012.
- [34] B. Waxman. Routing of multipoint connections. *IEEE Journal of Selected Areas in Communications*, 6(9), 1988.