Deriving Network Topologies from Real World Constraints

Mahmood A. Hameed, Abdul Jabbar, Egemen K. Çetinkaya, James P. G. Sterbenz†
Information and Telecommunication Technology Center
University of Kansas, Lawrence, Kansas 66045–7612
Email: {hameed, jabbar, ekc, jpgs}@ittc.ku.edu
†Corresponding author. Telephone: +1 785 864 7890

***Abstract*—Realistic network topologies are crucial for network research and are commonly used for the analysis, simulation, and evaluation of various mechanisms and protocols. In this paper, we discuss network topology models to generate physical topologies for backbone networks. In order to gain better understanding of current topologies and engineer networks for the future, it is necessary to generate realistic *physical* topologies that are governed by the infrastructure as opposed to only logical topologies that are governed by policy or higher-layer abstractions. The objective of this work is to present the principles that are key to node distributions of realistic topologies and the challenges involved. We argue that the dominant factors that influence the location of the PoPs are population density distribution and the technology penetration of a given region. Hence we implement a clustering algorithm to accurately predict the location of PoPs and later explore cost constrained models to generate realistic physical topologies.**

***Index Terms*—Network topology model, cost-constraint, geography, population, resilience, technology penetration.**

I. Introduction and Motivation

The motivation for network topology research is to understand the structure and evolution of the Internet as well as create realistic models for the purpose of experimentation and analysis. Therefore, it is necessary to model not only the graphical properties of a topology but also the processes that are fundamental to the growth of that topology. Moreover, the constraints that shape the physical topology heavily impact the properties of the resulting logical topologies. For example, the link connectivity of a network is dependent on the node locations. Realistic physical topology models enable us to accurately evaluate the performance of protocols and services and ultimately predict the topologies of the future Internet.

The ever-increasing importance and pervasiveness of communication networks also increases the expectations from these networks in terms of resilience and survivability. Practical networks are never fully resilient and hence it is necessary to do a systematic evaluation of topologies to know how a particular network reacts to challenges. Furthermore, such research is helpful for network architects to generate alternate topologies based on realistic constraints. To accurately model an existing ISP, we need to know where actual node locations are and how fiber infrastructure is laid out to connect these nodes. This information is key to accurately determining delay, capacity, and resilience characteristics of a network. On the other hand, simply placing nodes randomly in a given region and connecting them using an attachment model does not reflect realistic topologies.

The primary concern for a new ISP desiring to build a network in a region that already has network resources extensively deployed, is to know where infrastructure such as exchange points and fiber links are located in that region. Hence, we need a model to generate realistic topologies that accurately reflects the actual node and link locations. For certain regions that do not have a lot of fiber laid out already, this model could help predict the optimal node and link locations. Current examples of Sprint, AT&T, and Level 3 networks show that the blueprint of node locations is deeply embedded into their topologies. We claim that physical node locations combined with a realistic link generation model is critical for network evaluation, particularly for resilience [1], in which many challenges such as large-scale natural disasters and power failures are geographic in nature [2]. We present a model that generates realistic physical node locations for a given region based on real world constraints and later compare to the location of actual ISP PoPs (points of presence).

We begin our discussion with an overview of the proposed model in Section II. We compare random and population-based node distributions and consider the challenges faced by the network community in validating models due to the lack of real data. In Section III, we present our methodology discussing the real world constraints for node distribution and how we implement these in our simulations. In Section IV, we compare the location of backbone PoPs generated by our algorithm to actual ISP PoPs. Finally, in Section V, we discuss a simple topology example generated by linking cluster centers based on a cost constrained model.

II. Overview of Proposed Model

The main thrust of this work is to model and generate realistic topologies, with an emphasis on backbone networks in this paper. Therefore, the generation model should be representative of the actual network structure and evolution process. We seek the precise location of the backbone nodes for various countries and continents. Some of the well known, yet fundamental aspects that govern their placement are population density distribution and technology penetration. Our model has the ability to generate a specific number of cluster centers for a given geographical area while understanding that simply placing nodes in the order of most populated cities will not necessarily generate realistic topologies.

To connect these nodes we propose a cost-constrained connectivity model. High resilience can be achieved at unacceptably high costs, however unrealistic. If costs were not constrained, networks would be full meshes. Hence realistic generators must have the ability to produce feasible topologies at a finite cost.

A. Node Distribution

The physical topologies of networks are highly constrained by the geographic location of its components. It is obviously inaccurate to assume that the distribution of nodes is uniform. In fact, the distribution is very irregular for any given geographic area due to the constraints imposed by population, terrain such as mountains and bodies of water, and policy issues. Several works in the past [3], [4] have shown that the router-level topology shows a very high correlation to the population density. However, we are not aware of any existing models that apply such realistic constraints to *deduce* node placement. While it is obvious that distribution of network infrastructure is driven by population, it is certainly not the only factor governing the spatial distribution of nodes. For certain regions where the demand for Internet access is non-uniform, we need to consider *effective population*, that is the number of Internet users in a sub-region as opposed to the absolute population. Typically, for a well developed nation, there is very little difference between effective and absolute population. However, the difference is pronounced in an economically diverse region such as India.

B. Challenges

One of the fundamental challenges in developing a physical topology model is the lack of real data for validation. The physical topology of commercial networks including the Internet is not readily available because of competition among ISPs and security concerns. Previous research has considered several inference mechanisms to determine geographic node locations and physical link distances [5], [6], [3], but despite these efforts, the inference of physical topologies remains an open problem. There are, however, a few educational and research networks such as GÉANT2 and Internet2 for which the physical topology is available for validation[1].

III. Methodology

In this section, we discuss the principles which are key to the topologies generated by our model. We discuss their implementation for geographical areas including the USA, Europe, India, and Africa.

A. Clustering Algorithm

We use the *k-means* clustering algorithm [7] to find optimal locations to place a backbone PoP; the location at which an ISP terminates services and hands off to the local service provider. This algorithm uses gridded population density data sets from the *Center for International Earth Science Information Network* (CIESIN) [8] . The data is available for all areas across the globe for the year 2000 and is in raster format with a 1 km^2 grid. K-means is an iterative clustering method that works in two phases. The goal is to minimize the sum of the distances between all data points to cluster centers for all clusters. The initial selection of the cluster centers is random. The first phase, generally referred to as the *batch* phase, recomputes the cluster centers by re-associating each data point to its nearest cluster center. This phase provides an approximate but fast computation of cluster centers.

The second phase is generally called the *on-line* phase that uses the output of the first phase as the initial cluster centers and re-associates points to a different cluster only if doing so reduces the sum of distances. Cluster centers are recomputed after each re-association. Each iteration requires one pass through all data points. This is computationally complex and time consuming, especially for such large data sets.

The two inputs to our algorithm are the population data and the number of cluster centers. From the inferred topologies obtained from Rocketfuel [9], we note the number of PoPs for various ISPs for the various geographical areas considered. For example, the Sprint backbone has 27 nodes spatially distributed across the USA. We use this number as the input to our algorithm and generate an equal number of population centers. We consider multiple ISPs so that we can aggregate across major tier-1 providers, so as to not neglect certain parts of a country which may not be serviced by a specific ISP.

B. Technology Penetration

The other fundamental aspect governing the location of the PoPs is *technology penetration*. We argue that the location of the backbone PoPs is highly dependent on the number of Internet users in a given area. We denote the technology penetration factor as γ, defined as the fraction of *Internet* users to the total population in a particular area. Intuitively, this factor is uniform for a developed country like USA, for which we consider γ=1. This factor particularly has significant influence on a developing country such as India, where there are many densely populated areas along the river Ganges with very few Internet users. Hence, placing network resources solely based on the population density data set would not lead to a realistic network deployment. We use the quarterly report released by *Telecom Regulatory Authority of India* [10] to get the state-wise list of broadband subscribers in India. We incorporate technology penetration into our model by weighting the population of each grid in an area by corresponding γ and then clustering the resulting data set.

C. Cost-Constrained Connectivity Model

Cost constraints significantly impact network design and evolution. The resilience and survivability of networks [1] is almost always limited by the cost, therefore, realistic models must incorporate cost constraints. For simplicity, we assume that all nodes in the backbone network have equal cost, denoted by C_b. The link cost $C_{i,j}$ of a link i, j is calculated

[1]Note that these networks are much smaller than many commercial ISPs.

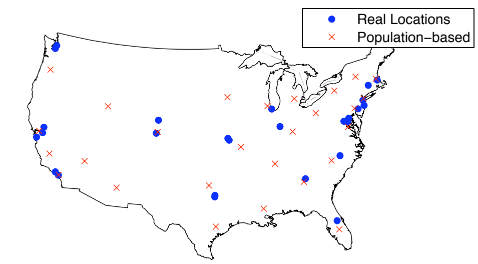

(a) Sprint

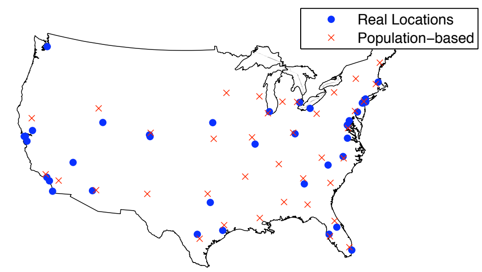

(b) Level 3

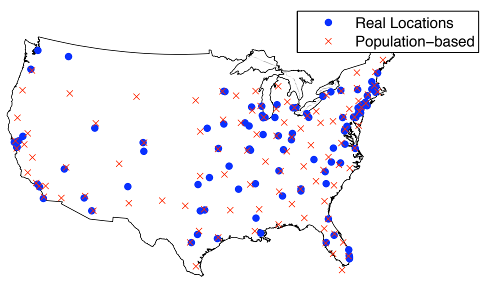

(c) AT&T

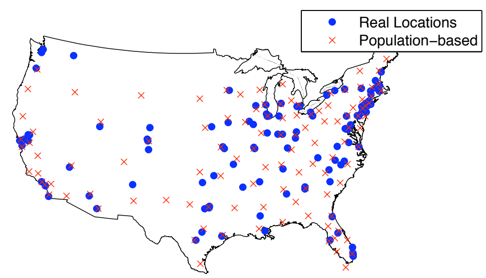

(d) Combined ISPs

Fig. 1. Relative node locations for ISPs in USA

(a) Predicted PoPs (20)

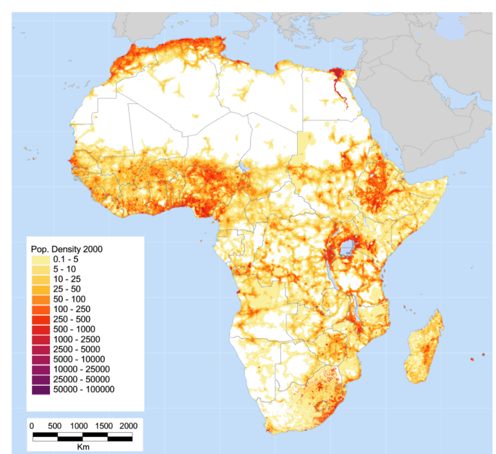

(b) Actual population density chart [8] (reprinted with permission)

Fig. 2. Cluster centers for Africa

as:$C_{i,j} = f + v \times d_{i,j}$, where f is the fixed cost associated with link, v is the variable cost per unit distance for the link and $d_{i,j}$ is the length of the link. The Internet is commonly modeled as pure preferential attachment but here we consider just the *backbone network*. The nodes in our model are connected using a cost-constrained Waxman model, which accurately represents link connectivity in backbone networks [3]. While it is generally agreed that backbone networks are mesh like [11], there is some contention as to exact relationship between link probability and distance. While some works claim that this is exponential [3], others claim that this is linear [4].

According to the Waxman model [12] the probability that two nodes *u* and *v* have a link between them is given by

$$P(u, v) = \beta e^{\frac{-d(u,v)}{L\alpha}}$$

where $0 < \alpha, \beta \leq 1$ and L is the maximum distance between any two nodes. The Waxman parameters α and β are controlled by the cost. A high value of α corresponds to a high fraction of short links to long links and β is directly proportional to the link density; *d* is the Euclidian distance between the two nodes. We use the node locations based on realistic constraints and connect them using Waxman model for a realistic backbone topology.

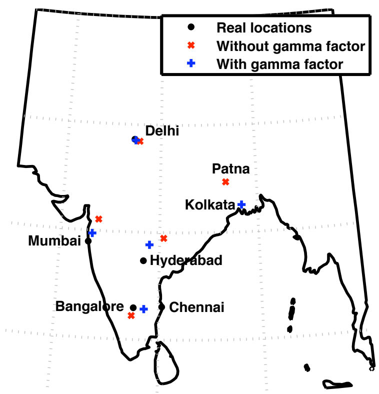

Fig. 3. Illustration of γ factor for India

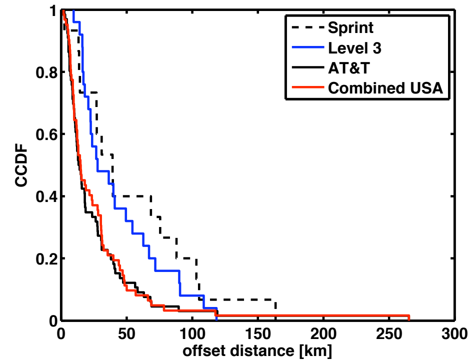

Fig. 4. CCDF of offset distance

IV. Comparison of Node Locations

In this section we compare the physical location of real ISP backbone PoPs to the location of PoPs generated using our population based clustering algorithm. We present results pertaining to uniform and non-uniform technology penetration factors.

A. Uniform Technology Penetration

As discussed earlier, a well developed country has relatively uniform technology penetration. For illustration purposes, we discuss results for various ISPs in the USA. Since the focus is on backbone topologies, we consider major tier-1 providers. Figure 1 shows the comparison of Sprint, Level 3, and AT&T backbone PoP locations to 27, 38, and 106 nodes respectively, generated by performing clustering on the US Population density data set. A visual inspection (Figure 1a, 1b, 1c) shows that for very densely populated regions like the east coast, the locations are very well matched. However, there are a few outliers. This is because not all regions across the nation are serviced by a particular ISP. It is therefore necessary to aggregate multiple tier-1 ISPs to make an appropriate comparison. We combine the PoPs for all ISPs while limiting to one PoP per city. This results in 112 unique points as shown in Figure 1d. We then generate 112 cluster centers using our population based clustering algorithm for comparison.

TABLE I
Comparing offset distance with existing PoPs in kms

Network (POPs)	Mean	σ	Min.	Max.
Sprint (27)	54.2	45.3	2.6	163.6
AT&T (106)	26.5	37.3	1.1	265.2
Level 3 (38)	43.4	31.7	9.6	118.6
GÉANT2 (34)	101.5	54.1	20.2	252.3
Ebone (27)	56.3	27.9	17.7	131.1
Tiscali (47)	34.8	22.3	2.47	80.6
VSNL (5)	26.7	34.9	2.6	265.1

We quantify the distance between inferred PoP locations and population based cluster centers as the *offset distance* for a pair of nodes. To provide a rigorous analysis of such a comparison, we plot the complementary cumulative distribution function (CCDF) of the offset distance for individual and combined ISPs. We note from Figure 4 that when we combine all ISPs, almost 90% of the nodes generated by our algorithm are within 50 km offset distance. A very small percentage of the nodes are outliers.

We consider a geographic area like the continent of Africa, currently with few network resources limited to countries like South Africa, to predict the location of network infrastructure. We generate the optimal location of backbone PoPs that can be used by an ISP desiring to have a continent-wide topology. Figure 2a shows the predicted location of 20 PoPs. We provide a population density chart (Figure 2b) [13] to make a visual comparison. It should be noted that since there is no continent-wide ISP in Africa, we cannot compare node locations with real data, but use this as a tool to predict future PoP locations.

B. Non-Uniform Technology Penetration

As discussed earlier, γ is a significant factor, generally, for less developed nations. We take a simple and obvious case of India to illustrate this factor. India is highly populated in the northern belt of the river Ganges. However, the number of Internet users for this region is small compared to the absolute population. We consider the inferred topology of the VSNL network in India [9]. VSNL has only 5 PoPs in Delhi, Mumbai, Hyderabad, Bangalore, and Chennai. We run our clustering algorithm both on the absolute population data set as well as the effective population data set (γ weighted per grid) to make visual comparison. Figure 3 shows that 4 of the PoPs match closely. Instead of a PoP near Chennai, we end up with one in Patna (γ=1) for two reasons: a) Patna is much denser in population than Chennai and, b) the PoP placed near Bangalore is close enough to Chennai for our algorithm to place another PoP. After correcting for γ [10], the four PoPs which matched earlier get closer to the real locations, while the one in Patna moves to Kolkata as it is one of the metropolitan areas with a high number of Internet users.

We provide Table I as summary of our results pertaining to the locations of PoPs for various ISPs for various geographic areas. It is to be noted that all of them are inferred topologies except for GÉANT2 [14], which is a research network in Europe. Our predictions match very well with ISPs with large infrastructure. For example, in the case of AT&T, the mean separation between real and clustered nodes is 26.5 km and the closest match is with an offset of 1.1 km. Initial experimentation has produced promising results and we discuss our road map for future work in Section VI.

V. Sample Synthetic Graph

In this section, we demonstrate the ability to generate a realistic 27 node topology based on US population density data set. We use a cost-constrained Waxman model to connect the backbone nodes. The objective is to go from realistic node locations to understanding realistic topologies and evaluate resilience of synthetic graphs. Figure 5 shows the topology generated by our model. We are unable to do a comparison between synthetically generated physical topologies to real ISP topologies due to lack of validation data as discussed in Section II-B. Tools such as Rocketfuel provide us with inferred topologies that are logical and are not sufficient to validate synthetic physical topologies.

We computed betweenness, average node degree and clustering coefficient along with other metrics for graph shown in Figure 5. Betweenness is the number of shortest paths through a particular node or link [15]. The node betweenness values were calculated as (max:124, min:1, avg:32) and link betweenness of (max:35.1, min:2.9, avg:11.3). A higher average node degree value (mean number of links connected to a node) generally indicates that a graph is better connected and is more robust [15]. We observe that average node degree is 5.23. Clustering coefficient is the measure of how well neighbors of a node are connected and is calculated as 0.28.

Fig. 5. Synthetic topology for 27 nodes

VI. Conclusion and Future Work

We have provided a model that precisely generates the optimal locations to place backbone PoPs in a given region. However, a new ISP wanting to layout a network will consider the location of existing fiber infrastructure. Hence we plan to constrain the location of PoPs to existing infrastructure by doing a *snap-to-grid*. For the link model, we plan to constrain links to the existing fiber routes which were in turn constrained by deployable routes such as railway lines and highways. On the flip side, for countries that have not yet extensively laid out their fiber lines, such a model could provide the answer to *Where should the fiber infrastructure go?* This gives us the ability to plan and engineer future networks such that resources are efficiently deployed. We have laid out key aspects of realistic physical topologies, however they are by no means exhaustive and require further research. Working with large data sets is computationally time consuming. To increase efficiency of our algorithm, we are in the process of re-implementing the weighted clustering algorithms without compromising accuracy.

Acknowledgment

We would like to acknowledge Jing Han, Yufei Cheng, Shi Qian, Justin Rohrer, and other members of the ResiliNets group for discussions on this work. This research was supported in part by the National Science Foundation FIND (Future Internet Design) Program under grant CNS-0626918 (Postmodern Internet Architecture) and the EU FP7 FIRE programme ResumeNet project (grant agreement no. 224619). We would also like to acknowledge CIESIN at Columbia University for permitting us to reprint Africa map.

References

[1] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, "Resilience and Survivability in Communication Networks: Strategies, Principles, and Survey of Disciplines," *Computer Networks: Special Issue on Resilient and Survivable Networks (COMNET)*, vol. 54, no. 8, pp. 1245–1265, June 2010.

[2] E. K. Çetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, and J. P. Sterbenz, "A Comprehensive Framework to Simulate Network Attacks and Challenges," in *RNDM'10 - Second International Workshop on Reliable Networks Design and Modeling*, Moscow, Russia, October 2010.

[3] A. Lakhina, J. Byers, M. Crovella, and I. Matta, "On the geographic location of Internet resources," *IEEE JSAC*, vol. 21, no. 6, pp. 934–948, Aug. 2003.

[4] S. Yook, H. Jeong, and A. Barabasi, "Modeling the internet's large-scale topology," *PNAS*, vol. 99, no. 21, pp. 13 382–13 386, 2002.

[5] P. Francis and et al., "Idmaps: a global internet host distance estimation service," *IEEE/ACM TON*, vol. 9, no. 5, pp. 525–540, 2001.

[6] T. S. E. Ng and H. Zhang, "Predicting internet network distance with coordinates-based approaches," in *INFOCOM*, 2001, pp. 170–179.

[7] J. MacQueen, "Some Methods for classification and Analysis of Multivariate Observations," *Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, pp. 281–297, 1967.

[8] Center for International Earth Science Information Network (CIESIN), Columbia University; and Centro Internacional de Agricultura Tropical (CIAT). 2005. Gridded Population of the World Version 3 (GPWv3): Population Density Grids. Palisades, NY: Socioeconomic Data and Applications Center (SEDAC), Columbia University. [Online]. Available: http://sedac.ciesin.columbia.edu/gpw

[9] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP Topologies with Rocketfuel," *IEEE/ACM TON*, vol. 12, no. 1, pp. 2–16, Feb 2004.

[10] Telecom Regulatory Authority of India, "The Indian Telecom Services Performance Indiactors April–June 2008," *Quarterly press release*, vol. 109, October 7th 2008.

[11] J. Doyle, D. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger, "The "robust yet fragile" nature of the Internet," *PNAS*, vol. 102, no. 41, pp. 14 497–14 502, 2005.

[12] B. Waxman, "Routing of multipoint connections," *IEEE JSAC*, vol. 6, pp. 1617–1622, Dec 1988.

[13] Africa: Population Density, 2000. Center for International Earth Science Information Network (CIESIN), Columbia University; and Centro Internacional de Agricultura Tropical (CIAT). 2005. Gridded Population of the World Version 3 (GPWv3). Palisades, NY: Center for International Earth Science Information Network (CIESIN), Columbia University. Available at http://sedac.ciesin.columbia.edu/gpw.

[14] "The GÉANT2 Website," http://www.geant2.net/.

[15] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, K. Claffy, and A. Vahdat, "The Internet AS-level topology: three data sources and one definitive metric," *ACM SIGCOMM CCR*, vol. 36, no. 1, pp. 17–26, 2006.