

DomainImpute: Inferring Unseen Components in the Internet

Brian Eriksson
Boston University
Boston, MA
eriksson@cs.bu.edu

Paul Barford
Qualys and
University of Wisconsin
Madison, WI
pb@cs.wisc.edu

Joel Sommers
Colgate University
Hamilton, NY
jsommers@colgate.edu

Robert Nowak
University of Wisconsin
Madison, WI
nowak@ece.wisc.edu

Abstract—Despite many efforts over the past decade, the ability to generate topological maps of the Internet at the router-level accurately and in a timely fashion remains elusive. Mapping campaigns commonly involve `traceroute`-like probing that are usually non-adaptive and incomplete, thus revealing only a portion of the underlying topology. In this paper we demonstrate that standard probing methods yield datasets that implicitly contain information about much more than just the directly observed links and routers. Each probe, in addition to the underlying domain knowledge, returns information that places constraints on the underlying topology, and by integrating a large number of such constraints it is possible to accurately infer the existence of unseen components of the Internet. We describe `DomainImpute`, a novel data analysis methodology designed to accurately infer the unseen hop-count distances between observed routers. We use both synthetic and a large empirical dataset to validate the proposed methods. On our empirical real world dataset, we show that our methods can estimate over 55% of the unseen distances between observed routers to within a one-hop error.

I. INTRODUCTION

Topological connectivity is one of the most basic characteristics of network infrastructure. Connectivity can be specified in different ways, but one of the most important is *router-level connectivity*: the graph structure of routers and the links between them. While router-level connectivity within an individual ISP may be well known and relatively simple for the provider to monitor, understanding connectivity beyond a single domain is problematic.

Mapping the Internet’s router-level topology has been the subject of a large number of studies over the past decade (e.g., [1], [2], [3], [4], [5], [6]). The standard method for gathering data on Internet topology is to use `traceroute`-like tools to send TTL limited probes between a set of sources and destinations. One of the main challenges in recovering the topology from these measurements is *perspective* (i.e., limited visibility of the Internet structure captured by a set of measurements) [3]. The result of this challenge is that despite concerted efforts, we are still unable to generate accurate maps of the Internet’s router-level topology in a timely fashion.

This work was supported in part by NSF grants CNS-0716460, CNS-0831427 and CNS-0905186. Any opinions, findings, conclusions or other recommendations expressed in this material are those of the authors and do not necessarily reflect the view of the NSF.

In this paper, we address a subset of the general problem of router-level Internet topology mapping by introducing the *DomainImpute* methodology. Our methodology addresses the problem of inferring unseen route lengths in terms of the number of routers along the path (i.e., the *hop count* of the path) between observed routers. We exploit a modified matrix completion algorithm that is based only on hop counts between source-destination pairs extracted from `traceroute`-like probes, and readily available domain information of the routers (e.g., autonomous system information).

We evaluate the capabilities of our method by using both synthetic topology data sets generated from the BRITE topology generator [7], which allows for ground truth domain knowledge, and a real world dataset gathered from `traceroute` measurements between PlanetLab nodes [8] where domain knowledge is extracted from BGP. Our results on this real world dataset show that our `DomainImpute` inference methodology can estimate over 55% of the unseen distances between routers to within a one-hop error. We show that through the use of network domain knowledge, unseen hop estimation improvements of up to 40% can be found on synthetic datasets and 10% are found on real-world networks over off-the-shelf inference methods.

II. RELATED WORK

There have been many prior studies that have focused on developing methods for mapping Internet topology. While topology can be considered on different levels including application-specific connectivity or autonomous system connectivity, our focus in this work is on router-level connectivity. A great deal of prior work in this area has focused on using active `traceroute`-like probes as a basis (e.g., [1], [2], [3], [5]). In each case, these studies highlight several challenges associated with this kind of approach, including the need for widely distributed nodes from which probes can be sent (i.e., to address the need for a broad perspective) and the difficult problem of interface disambiguation.

Our work is also informed by Eriksson *et al.* in [6], [9]. In those studies, the authors propose methods for establishing Internet maps based on passive observations of hop counts in packets. While the idea of using inference methods to estimate incomplete hop counts is similar, our work differs

in data (the use of active `traceroute`-based probes) and methods (Matrix Completion-based techniques). The work in [9] examines the problem of estimating pairwise hop counts using incomplete passive measurements to a set of landmarks. In contrast, our work will demonstrate a methodology for estimating pairwise hops using only massively incomplete pairwise measurements between the routers.

III. DATASETS

To validate our inference methodology we use both synthetic and real world datasets.

A. Synthetic Dataset

We rely on the BRITE topology generator [7] in order to create the synthetic topologies in this paper. While there may be synthetic topology generators that more accurately reflect real-world Internet topology at the router-level (e.g., Heuristically Optimized Topologies [4]), the BRITE methodology offers the unique ability to synthetically generate Internet router-level topologies from the top-down with autonomous system (AS) assignment. This domain knowledge will be exploited by our DomainImpute methodology in Section IV and is necessary to obtain the performance increases shown in our results. Three BRITE-generated synthetic topologies will be considered in this paper – (i) 10 ASes and 500 routers, (ii) 10 ASes and 1000 routers, and (iii) 25 ASes and 1250 routers. To synthesize hop counts on these synthetic topologies we will use shortest path routing between nodes.

B. Real World Empirical Dataset

The goal of our empirical data collection effort was to establish a data set from which our inference methods could realistically be evaluated. Specifically, we required a representative corpus of Internet routers with disambiguated interfaces (to the extent possible using available techniques) that could act as ground truth for our work.

To generate our empirical dataset, we collected `traceroute` data between a full mesh of Planetlab hosting sites [10]. At the time of our measurement collection there were over 900 *hosts* that are part of Planetlab, but there were only about 375 distinct *sites*. Of these sites, only a subset are available at any given time due to host maintenance or other issues. To perform each `traceroute`, we used the Paris `traceroute` tool [11]. We collected the full mesh of Planetlab `traceroutes` three times (roughly evenly distributed). Due to Planetlab site and host transience, we were able to use approximately 216 Planetlab sites for each of the three rounds of full mesh probing. We gathered three complete measurement data sets beginning on December 11, 2008, December 22, 2008 and January 6, 2009.

A standard problem in `traceroute`-based topological studies is the issue of IP interface disambiguation, which is also referred to in the literature as *alias resolution*. That is, Internet routers are typically assigned multiple IP addresses (e.g., each interface on a router may have a unique IP address assigned to it). Identifying which addresses correspond to the

same physical router is the role of alias resolution. To identify the routers (i.e., de-alias) our data set, we used a router alias database published by the iPlane project, which builds on previous published alias resolution methodologies, including those used by the Rocketfuel project [2].

IV. ESTIMATING UNSEEN CONNECTIVITY USING DOMAINIMPUTE

A natural way to represent the router-level topology of the Internet is by using a *hop count matrix* $\mathbf{H}^{(I)}$. For the entire IPv4 Internet, $\mathbf{H}^{(I)}$ is a $2^{32} \times 2^{32}$ matrix with each element $h_{i,j}^{(I)}$ representing the number of links between IP address i and IP address j . If the matrix $\mathbf{H}^{(I)}$ is known, then the router-level topology of the Internet in all places is completely resolved. In this paper, we will focus on reconstructing a portion of this full hop count matrix (\mathbf{H}) given a subset of observed routers from our measurement datasets. To fill in even this portion of the hop count matrix for N routers completely would require an infeasible N^2 probing of the Internet. Instead, we examine how the `traceroute` measurements (i.e., paths containing labels of intermediate nodes) can be used to improve our connectivity estimates.

To construct a hop count matrix from `traceroute` measurements, consider a single `traceroute` probe sent between two of the Planetlab nodes (p_1, p_2) that returns the path,

$$p_1 \rightarrow r_1 \rightarrow r_2 \rightarrow r_3 \rightarrow r_4 \rightarrow p_2 \quad (1)$$

From this single path many hop elements of \mathbf{H} can be observed, with r_1 being one hop away from r_2 , two hops away from r_3 , etc. Intuitively, this has a multiplicative effect on our ability to populate the hop count matrix \mathbf{H} versus a single ping-style hop count measurement. In this fashion, we use the large set of `traceroute` measurements to extract the hop count matrix partition in Table I and fill in the hop count matrix \mathbf{H} .

TABLE I
AN EXAMPLE HOP COUNT MATRIX USING OBSERVED HOP ELEMENTS FROM THE SINGLE `traceroute` PATH $p_1 \rightarrow r_1 \rightarrow r_2 \rightarrow r_3 \rightarrow r_4 \rightarrow p_2$ (WHERE “-” REPRESENTS AN UNKNOWN ELEMENT).

| | r_1 | r_2 | r_3 | r_4 |
|-------|-------|-------|-------|-------|
| r_1 | 0 | 1 | 2 | 3 |
| r_2 | - | 0 | 1 | 2 |
| r_3 | - | - | 0 | 1 |
| r_4 | - | - | - | 0 |

It is obvious from Table I that many of the matrix elements have no observed information (indicated here as a “-”). Using only `traceroute` probes to fill in the router-to-router hop elements will result in a hop count matrix that could be highly incomplete depending on the perspective afforded by the `traceroute` campaign. Given an assumed incomplete hop count observation matrix, our objective is to impute (or “fill in”) the missing observations to resolve the router-level topology.

From the observed elements of the incomplete hop count matrix, \mathbf{H} , we will now estimate the unseen hop elements. We

appeal to recent work on *Matrix Completion* methodologies in [12], [13] that frames the estimation of incomplete matrices as the problem of inferring the matrix with the lowest rank (the number of nonzero eigenvalues) that agrees with the observed elements.

$$\begin{aligned} & \text{minimize} && \text{rank}(\mathbf{X}) \\ & \text{subject to} && x_{i,j} = h_{i,j} \quad \text{s.t. } h_{i,j} \text{ was observed} \end{aligned}$$

This NP-hard problem can be efficiently solved using the Matrix Completion algorithm from [13].

Motivation for using Matrix Completion to impute hop counts is two-fold. First, given k observations from the $N \times N$ matrix \mathbf{H} (for N number of end hosts), it was found in [12] that using the Matrix Completion algorithm, the matrix can be *exactly* reconstructed given the number of observations, k , satisfies $k \sim O(rN \log N)$ (where r is the rank of the matrix). Therefore, even for massively incomplete low rank matrices, we may be able to accurately reconstruct the unseen elements. Second, prior work in [9] shows that a hop count matrix can be accurately represented by a low-rank approximation, indicating Matrix Completion should perform well on these matrices.

While prior work has shown that low-rank *approximations* can accurately represent hop count matrices, further examination shows that these matrices are not actually low rank. While approximation of the rank of our `traceroute`-derived real world hop count matrix is problematic due to missing elements, hop matrices derived from synthetic topologies allow us to examine all hop elements. From the Eigen-Decomposition in Table II, we see the number of eigenvectors needed to bound the approximation error of the full hop count matrix for a synthetic BRITE-derived topology with 10 autonomous systems and 1,000 total routers. As seen in the table, over 156 eigenvectors are needed to reconstruct the hop matrix with approximation error less than 10%. This indicates that hop matrices are not effectively low rank and significant detail of the hop matrix would be excluded by estimating this matrix with a low rank approximation.

TABLE II
SYNTHETIC HOP COUNT MATRIX EIGEN-DECOMPOSITION - 10 AS, 1000 ROUTER TOPOLOGY

| Approximation Error < | Number of Eigenvectors Needed | |
|-----------------------|-------------------------------|------------------------|
| | Full Hop Matrix | Inter-domain Submatrix |
| 40% | 4 | 1 |
| 30% | 7 | 2 |
| 20% | 40 | 3 |
| 10% | 156 | 6 |

While full hop count matrices are of high rank, we identify an approach for reducing the rank of this problem by separately considering network paths from *inter*-domain and from *intra*-domain routes. Inter-domain routes consist of routes from one domain to another domain in the Internet. Consider paths from one domain (denoted as AS_A) to other domains (AS_B, AS_C, AS_D). As depicted in Figure 1, these inter-domain routes are commonly trafficked through a common

border router between the domains [3], resulting in a shared path for all routes past the border router.




Fig. 1. Example of inter-domain paths through the Internet traversing a common border router.

As demonstrated in [14], this common shared path structure of routes will result in low rank hop counts between the two domains. Therefore, a partition of the hop matrix consisting of only inter-domain routes should effectively be low rank. To test this hypothesis, the Eigen-Decomposition for both a full hop count matrix and a partitioned inter-domain hop count submatrix can be seen in Figure 2 for a synthetic BRITE generated topology with 10 ASes and 1,000 routers. The 500×500 -sized inter-domain submatrix was generated by partitioning the 10 AS domains into two groups of 5 ASes each (the subset of routers $\mathcal{I}_A, \mathcal{I}_B$) and considering only the hop elements between the two domain partitions (*i.e.*, the submatrix $\mathbf{H}_{\mathcal{I}_A, \mathcal{I}_B}$). As seen in Figure 2 and Table II, significantly fewer eigenvectors are required to reduce the approximation error of the inter-domain hop matrix compared with the full hop count matrix. In the inter-domain matrix, only 5 eigenvectors are required to obtain an approximation error of less than 10%, while the full hop matrix (with both intra-domain and inter-domain paths) requires over 156 eigenvectors to obtain this same error rate.

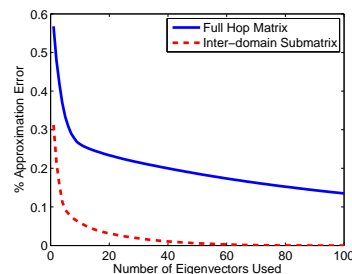


Fig. 2. Approximation error of both the full hop matrix and partitioned inter-domain submatrix with respect to the number of eigenvectors used. BRITE generated synthetic topology with 10 ASes, 1000 routers.

This low rank property of the inter-domain submatrix motivates a modified Matrix Completion approach we call *Domain-Impute*. Consider the set of domain indices partitioning the routers in the hop count matrix (denoted as $\{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_M\}$, where \mathcal{I}_i is the set of routers in domain i). By bisecting this set of domain indices into $\mathcal{I}_A = \{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n\}$ and $\mathcal{I}_B = \{\mathcal{I}_{n+1}, \dots, \mathcal{I}_M\}$, we can construct a potentially low rank inter-domain submatrix ($\mathbf{H}_{\mathcal{I}_A, \mathcal{I}_B}$) and perform Matrix Completion on this submatrix. This single bisection will only resolve at most $\frac{1}{4}$ of the full hop count matrix, which can be

considered the upper right corner of the matrix in Figure 3. In order to infer all unseen inter-domain elements of the matrix, consider performing multiple random bisections of our set of domains (where at random we choose which domains are in \mathcal{I}_A and \mathcal{I}_B). The final inter-domain hop count estimates are found by averaging across our multiple imputations of each inter-domain element. Finally, to resolve the small number of intra-domain components, we perform Matrix Completion on the entire hop count matrix of both intra-domain and inter-domain hop elements. The full DomainImpute methodology is described in Algorithm 1.

$$\begin{array}{c}
 \mathcal{I}_1 \quad \mathcal{I}_2 \quad \mathcal{I}_3 \quad \mathcal{I}_4 \\
 \begin{array}{c}
 \mathcal{I}_1 \\
 \mathcal{I}_2 \\
 \mathcal{I}_3 \\
 \mathcal{I}_4
 \end{array}
 \begin{bmatrix}
 \mathbf{H}_{\mathcal{I}_1, \mathcal{I}_1} & \mathbf{H}_{\mathcal{I}_1, \mathcal{I}_2} & \mathbf{H}_{\mathcal{I}_1, \mathcal{I}_3} & \mathbf{H}_{\mathcal{I}_1, \mathcal{I}_4} \\
 \mathbf{H}_{\mathcal{I}_2, \mathcal{I}_1} & \mathbf{H}_{\mathcal{I}_2, \mathcal{I}_2} & \mathbf{H}_{\mathcal{I}_2, \mathcal{I}_3} & \mathbf{H}_{\mathcal{I}_2, \mathcal{I}_4} \\
 \mathbf{H}_{\mathcal{I}_3, \mathcal{I}_1} & \mathbf{H}_{\mathcal{I}_3, \mathcal{I}_2} & \mathbf{H}_{\mathcal{I}_3, \mathcal{I}_3} & \mathbf{H}_{\mathcal{I}_3, \mathcal{I}_4} \\
 \mathbf{H}_{\mathcal{I}_4, \mathcal{I}_1} & \mathbf{H}_{\mathcal{I}_4, \mathcal{I}_2} & \mathbf{H}_{\mathcal{I}_4, \mathcal{I}_3} & \mathbf{H}_{\mathcal{I}_4, \mathcal{I}_4}
 \end{bmatrix}
 \end{array}$$

Fig. 3. Hop Matrix with domains $\{\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3, \mathcal{I}_4\}$ with a highlighted inter-domain submatrix using domain sets $\{\mathcal{I}_1, \mathcal{I}_2\}$ and $\{\mathcal{I}_3, \mathcal{I}_4\}$. Where $\mathbf{H}_{\mathcal{I}_1, \mathcal{I}_2}$ is the submatrix of hop counts between routers in domain \mathcal{I}_1 and routers in domain \mathcal{I}_2 .

Algorithm 1 - DomainImpute - Domain-based Matrix Completion

Given:

- \mathbf{H} - $[N \times N]$ -sized Incomplete hop count matrix.
- $\mathcal{I} = \{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_M\}$ - Domain indices (with $\mathcal{I}_1 \cup \mathcal{I}_2 \cup \dots \mathcal{I}_M = \mathcal{I}$ and $\mathcal{I}_1 \cap \mathcal{I}_2 \cap \dots \mathcal{I}_M = \emptyset$)
- K - Number of iterations.

Methodology:

- 1) For $k = \{1, 2, \dots, K\}$ iterations
 - a) Randomly order the set of domains \mathcal{I} and partition into two sets of indices, $\mathcal{I}_A, \mathcal{I}_B$ (where $\mathcal{I} = \mathcal{I}_A \cup \mathcal{I}_B$), revealing an incomplete inter-domain hop matrix $\mathbf{H}_{\{\mathcal{I}_A, \mathcal{I}_B\}}^{(k)}$
 - b) Using Matrix Completion [13], obtain the inter-domain hop estimates $\hat{\mathbf{H}}_{\{\mathcal{I}_A, \mathcal{I}_B\}}^{(k)}$
 - 2) Resolve each inter-domain hop count element by averaging the hop estimates for each element.
 - 3) For intra-domain elements, perform Matrix Completion on the entire incomplete matrix to resolve elements.
-

V. EXPERIMENTS

In addition to our DomainImpute methodology and the off-the-shelf Matrix Completion algorithm [13], we consider the following simple approach to impute the missing hop counts. Given the set of observed hop indices, \mathcal{S}_{known} , and the set of held out (*i.e.* unobserved) hop indices, $\mathcal{S}_{unknown}$, one can estimate each missing hop count using the mean of the hop counts that have been observed. Where, $\hat{h}_{i,j}^{mean} = \frac{1}{|\mathcal{S}_{known}|} \sum_{\{m,n\} \in \mathcal{S}_{known}} h_{m,n}$. For all imputation methodologies, the error metric used to assess the estimation accuracy is the Root Mean Squared Error

(RMSE) defined as:

$$RMSE(\hat{\mathbf{H}}) = \sqrt{\frac{1}{|\mathcal{S}_{unknown}|} \sum_{\{i,j\} \in \mathcal{S}_{unknown}} (h_{i,j} - \hat{h}_{i,j})^2}$$

If our estimator has an RMSE of 1, then we can estimate the hop distance (on average) within a single hop of the true hop distance.

1) *Synthetic Topology - Missing Hop Count Estimation Results:* The results of estimating missing hop counts on our synthetic topologies can be seen in Figure 4. Three synthetic topologies were generated using the BRITE topology generator [7] with 10 ASes and 500 routers, 10 ASes and 1,000 routers, and 25 ASes and 1,250 routers. Hop counts were then generated for each topology using shortest path routing on the synthetic graphs. Hop imputation performance was assessed for varying levels of incompleteness where a specified number of elements in the hop count matrix were observed-at-random, and all unobserved hop elements were imputed. As seen in the figures, the new DomainImpute methodology (with iterations, $K = 40$) outperforms standard off-the-shelf Matrix Completion in terms of hop reconstruction accuracy on all three synthetic topologies. When only 5% of the hop counts are observed, for the synthetic topology with 10 autonomous systems and 500 routers we observe an RMSE of 1.50 hops for the Matrix Completion algorithm and only RMSE of 1.05 hops for our DomainImpute methodology, an error rate reduction of 30%. For the larger synthetic topology of 25 autonomous systems and 1,250 routers, we observe an RMSE of 0.94 for Matrix Completion, but only 0.57 for our DomainImpute methodology, an error rate reduction of almost 40%.

2) *Real World Topology - Missing Hop Count Estimation Results:* For our Planetlab-based dataset, 2,211 routers were found such that each router was located in an autonomous system with at least 10 other observed routers from RouteViews-based BGP data [15]. This dataset is massively incomplete, with only 1.42% of the hop elements observed. To assess our estimation error rate, 20% of the observed router-to-router hop elements (chosen completely at random) were held out of the dataset and used to validate the performance of our imputation.

Table III shows the error rate for estimating the held-out hop elements in the $2,211 \times 2,211$ router hop count matrix. Using 40 iterations, the new DomainImpute algorithm estimates the missing route length elements with accuracy over 2 hops better than the naive mean estimation procedure on average and almost 10% more accurate than the current state-of-the-art Matrix Completion methodology. In Table IV, the deviation of the estimated held-out hop element values from the true hop element values are shown with respect to the three imputation methods. As evident in the table, the improvements of the new DomainImpute methodology are shown as 55.7% of the held-out route lengths can be imputed within a single hop, while the Matrix Completion methodology will only impute 51.7% of the held out route lengths within a single hop and the simple mean procedure only estimates 20.0% of the lengths within a single hop.

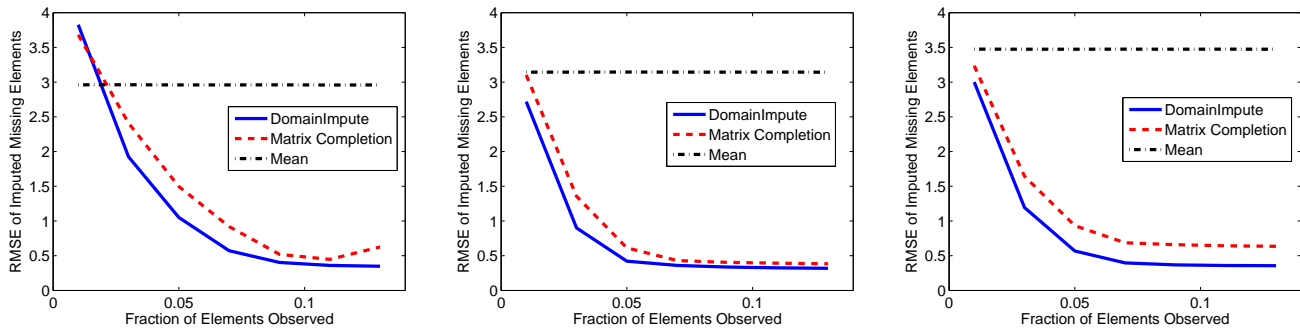


Fig. 4. Root Mean Squared Error (RMSE) of hop imputation for synthetic BRITE generated topologies. (Left) Number of ASes = 10, Number of routers = 500, (Center) Number of ASes = 10, Number of routers = 1000, (Right) Number of ASes = 25, Number of routers = 1250

TABLE IV
DIVISION OF MATRIX COMPLETION ERRORS FOR HOLDOUT DATA

| Error Range | DomainImpute Percent | Matrix Completion Percent | Mean Percent Imputation |
|--------------------------|----------------------|---------------------------|-------------------------|
| Less than one hop off | 55.7% | 51.7% | 20.0% |
| Less than two hops off | 80.8% | 79.0% | 38.2% |
| Less than three hops off | 90.7% | 90.1% | 54.1% |
| More than three hops off | 9.3% | 9.9% | 45.9% |

TABLE III
HOP MATRIX RECONSTRUCTION ERROR RATES. THE RMSE OF ROUTER TO ROUTER HOP DISTANCES HELD OUT.

| Method | RMSE |
|-------------------|-----------|
| Mean | 4.13 hops |
| Matrix Completion | 2.08 hops |
| DomainImpute | 1.89 hops |

VI. CONCLUSIONS

Generating timely and accurate maps of the Internet has been a compelling objective for some time, but remains beyond our grasp. The scope, diversity and dynamics of the infrastructure along with the fact that service providers often actively thwart measurement by standard methods all complicate the issue. In this paper we addressed a subset of the Internet mapping problem, by restricting our focus to discovering unseen route lengths in the Internet by developing DomainImpute, a novel estimation methodology using a modified matrix completion technique that exploits network structure. We demonstrated the capabilities of our methodology using both synthetic topologies and a traceroute dataset collected between PlanetLab nodes. On a real world dataset, we show that our novel DomainImpute methodology is able to estimate over 55% of the links within one hop of actual and over 80% of the links within two hops of their actual value. This material fits into the larger framework of inferring additional unseen components of the Internet (*e.g.*, unseen links and unseen routers), which we look to explore in future work.

REFERENCES

[1] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet Map Discovery," in *Proceedings of IEEE INFOCOM '00*.

[2] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP topologies with Rocketfuel," *IEEE/ACM Transactions on Networking*, vol. 12, no. 1, 2004.

[3] P. Barford, A. Bestavros, J. Byers, and M. Crovella, "On the Marginal Utility of Network Topology Measurements," in *Proceedings of ACM Internet Measurement Workshop (IMW '01)*, October 2001.

[4] D. Alderson, L. Li, W. Willinger, and J. Doyle, "Understanding Internet Topology: Principles, Models and Validation," *IEEE/ACM Transactions on Networking*, vol. 13, no. 6, December 2005.

[5] R. Sherwood, A. Bender, and N. Spring, "DisCarte: A Disjunctive Internet Cartographer," in *Proceedings of ACM SIGCOMM*, Seattle, WA, August 2008.

[6] B. Eriksson, P. Barford, and R. Nowak, "Network Discovery from Passive Measurements," in *Proceedings of ACM SIGCOMM*, Seattle, WA, August 2008.

[7] A. Medina, A. Lakhina, I. Matta, and J. Byers, "Brite: An approach to universal topology generation," 2001.

[8] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "PlanetLab: An Overlay Testbed for Broad-Coverage Services," *SIGCOMM CCR*, vol. 33, no. 3, pp. 3–12, 2003.

[9] B. Eriksson, P. Barford, and R. Nowak, "Estimating Hop Distance between Arbitrary Host Pairs," in *Proceedings of IEEE Infocom Conference 2009, Rio de Janeiro, Brazil*, April 2009.

[10] A. Bavier, M. Bowman, B. Chun, D. Culler, S. Karlin, S. Muir, L. Peterson, T. Roscoe, T. Spalink, and M. Wawrzoniak, "Operating System Support for Planetary-Scale Network Services," in *USENIX NSDI '04*, March 2004.

[11] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding Traceroute Anomalies with Paris Traceroute," in *Proceedings of ACM Internet Measurement Conference '06*, October 2006.

[12] E. J. Candes and B. Recht, "Exact Matrix Completion via Convex Optimization," To appear in *Found. of Comput. Math.*, 2008.

[13] B. Recht, M. Fazel, and P. Parrilo, "Guaranteed Minimum Rank Solutions to Linear Matrix Equations via Nuclear Norm Minimization," Submitted to *SIAM Review*, 2007.

[14] B. Eriksson, P. Barford, and R. Nowak, "Network Discovery from Passive Measurements," in *Proceedings of ACM SIGCOMM '08*, Seattle, Washington, August 2008.

[15] "Route Views Project," <http://www.routeviews.org/>.