inflation needs more study; we only note here that a small fraction of prefixes are responsible for most route changes; and those prefixes receive comparatively little traffic [11].

Additional Approaches. Finally, we ask whether there are other ways ISPs could attack this problem.

One direction, as mentioned in Section 2, is for ISPs to make note of the routes that they learn through BGP. How many entries in $\widehat{T}$ could the observer fill in using BGP-learned routes? From BGP, the observer learns a path to each destination from each of its neighbors. Each path as well contains sub-paths. This implies that the number of $\widehat{T}$ elements that can be learned in this way for any given destination is equal to the number of unique ASes found in all neighbor paths to the destination. Experience with BGP suggests that this is not usually a large number, except for the small set of ASes with very high degree. For example, for each of the Routeviews monitors, the average number of elements per destination that can be learned by such a method is $\sim$40. In comparison, for ASes in the `Core-1000` set, the average number of elements per destination potentially visible through traffic observation is $\sim$500. Thus, we believe that BGP-learned paths are likely to add only incremental improvement to estimation of $\widehat{T}$.

Another approach would be for ISPs to use information about the *volume* of traffic observed. The general idea is to assume a distributional model for traffic – one which includes a nonzero probability that a traffic element is zero. One then models the observed data as a mixture of the values taken from two sources: the assumed traffic distribution (which generates false zeros), and an additional source of (true) zeros. Such 'zero-inflated' models are used, for example in ecology, in settings loosely analogous to the VISIBILITY-INFERENCE problem [9]. Taking this approach requires imposition of modeling assumptions, with all the difficulties that accompany it: addressing the model selection problem, estimating parameters, and assessing confidence in the results. However the potential exists to estimate the *number* of false zeros via this sort of method. While this approach presents many hurdles, a considerable amount of theory has been developed around how to do this in various settings, and some issues relevant to traffic models are being addressed [6].

9. CONCLUSIONS

We started from the simple question 'what routes pass through a network?' We showed that using the network's traffic data to answer this question is equivalent to identifying source-destination pairs that are *not* communicating at a given time.

Answering these questions prompted us to look for ways to identify sets of source-destination paths that are *routed similarly in general.* This can be thought of as an considerable generalization of the notion of BGP atoms: rather than groups of prefixes that are routed *identically*, we look for groups of paths that are routed *similarly*. Surprisingly, we show that such groups of paths can be identified in a large set of representative locations in the Internet.

A key enabling idea has been the definition of a new distance metric for network prefixes: routing-state distance. Using it we were able to extrapolate from the relatively small amount of information available in publicly accessible BGP tables to estimate routing similarity between any two prefix pairs in the Internet.

We realized these ideas in the form of a family of classifiers; applying these classifiers to traffic measurements we showed that one can generally answer our motivating question with a high degree of accuracy.

While a number of challenges remain and considerable additional work is warranted, we are hopeful that the methods we develop here can improve knowledge of route visibility and further inform both operators and researchers.

Acknowledgements. This work was supported by NSF grants CNS-1017529, CNS-0905565, CNS-1018266, CNS-1012910, CNS-1117039, by a GAANN Fellowship, and by grants from Microsoft, Yahoo!, and Google. The authors thank the SIGCOMM referees and shepherd for their help in improving the paper.

10. REFERENCES

[1] J. I. Alvarez-Hamelin, L. Dall'Asta, A. Barrat, and A. Vespignani. K-core decomposition: a tool for the visualization of large scale networks. Technical report, Arxiv, 2005.

[2] V. Bharti, P. Kankar, L. Setia, G. Gürsun, A. Lakhina, and M. Crovella. Inferring invisible traffic. In *CoNEXT*, Philadelphia, PA, 2010.

[3] A. Broido and k. claffy. Analysis of RouteViews BGP data: policy atoms. In *NRDM*, 2001.

[4] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717–772, 2009.

[5] S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir. A model of internet topology using k-shell decomposition. In *PNAS*, 2007.

[6] D.-L. Couturier and M.-P. Victoria-Feser. Zero-inflated truncated generalized Pareto distribution for the analysis of radio audience data. *Annals of Applied Statistics*, 2010.

[7] T. Fawcett. An introduction to ROC analysis. *Pattern Recogn. Lett.*, 2006.

[8] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz. Towards an accurate AS-level traceroute tool. In *Proceedings of ACM SIGCOMM*, August 2003.

[9] T. G. Martin, B. A. Wintle, J. R. Rhodes, P. M. Kuhnert, S. A. Field, S. J. Low-Choy, A. J. Tyre, and H. P. Possingham. Zero tolerance ecology: improving ecological inference by modelling the source of zero observations. *Ecology Letters*, 2005.

[10] W. Mühlbauer, S. Uhlig, B. Fu, M. Meulle, and O. Maennel. In search for an appropriate granularity to model routing policies. In *SIGCOMM*, 2007.

[11] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. BGP routing stability of popular destinations. In *SIGCOMM Workshop on Internet measurment*, 2002.

[12] RIPE routing information service raw data project. `http://www.ripe.net/data-tools/stats/ris/ris-raw-data`.

[13] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush. 10 lessons from 10 years of measuring and modeling the internet's autonomous systems. *IEEE Journal on Selected Areas in Communications*, 2011.

[14] University of Oregon route views project. `http://www.routeviews.org/`.

[15] Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. Technical report, Rice University, 2010.

[16] Y. Zhang, M. Roughan, W. Willinger, and L. Qiu. Spatio-temporal compressive sensing and internet traffic matrices. *SIGCOMM Comput. Commun. Rev.*, 2009.