designs, and limitations imposed due to available functionality in various devices.

8. RELATED WORK

Our study is motivated by, and complimentary to, prior work that highlights the complexity in configuring Class-of-Service [24] and BGP policies [16, 10] in ISP networks. The latter also highlights the relatively high incidence of errors in BGP configuration, and blames the poorly designed router configuration languages for many of the errors. Note that our study examines complexity *given* commonly-used configuration languages, but it does not examine how much of this complexity arises from the basic design of these languages themselves. We leave this for future work.

While we focus on complexity within an ISP's network configuration, prior works [7, 15, 18, 11] have focused on the causes for complexity in enterprises. These studies found core enterprise network designs to be complex due to the usage of VLANs, route-redistribution commands, and network security and filtering policies. For ISP network services, we find that a significant amount of complexity can be attributed to the configuration of customer virtual networks at the edge, which involves the usage of CoS, VRFs, and route control stanzas, while the core network itself is simple.

The models we employed in this study are similar to those used in [7] to study enterprise networks, but with key extensions to focus on network-based services and the unique functionality found in ISP networks (e.g., the extensive use of route-reflectors, MPLS and router virtualization).

Finally, our study adds to a growing body of work that has found mining configurations to be a useful way to obtain insights about networks. E.g., recent works [17, 25] have used configuration files to perform root cause analysis, and troubleshooting of anomalies and performance problems in ISPs.

9. CONCLUSIONS

The usage and deployment of network-based services is growing as ISPs aim to provide advanced features to residential and enterprise customers. In this paper, we present the first large-scale analysis of these services, focusing on their design and configuration. We study several years' worth of configuration data corresponding to five services in a tier-1 ISP. We decompose the configurations into various tasks performed by operators and systematically highlight the complexity underlying each. We show that complexity grows over time. We find that while the exact causes of complexity may differ across different services, two factors – customer provisioning differences and control-plane (BGP) configuration – consistently make the overall designs complex. We conclude the study by exploring ways to reduce complexity, specifically, by considering the alternatives in, and the trade-offs imposed by, factors such as the choice of vendors and resource scaling.

We hope that our findings can lead to a broader discussion on understanding service configurations, on the design of both service-specific and service-agnostic support mechanisms, and potentially on alternate network architectures and router mechanisms with much better intrinsic support for network-based services.

10. ACKNOWLEDGMENTS

We would like to thank the operations team at the tier-1 ISP for providing the data and documentation about the services. We would also like to thank Katerina Argyraki (our shepherd), Aaron Gember, Seungjoon Lee, Kobus van der Merwe, Dan Pei, Kevin D'Souza, Adrian Cepleanu, Stephen Hutnik, Maria Napierala and the anonymous reviewers for their insightful feedback. This work is supported in part by NSF grants CNS-0746531, CNS-1017545 and CNS-1050170. Theophilus Benson is supported by an IBM PhD Fellowship.

11. REFERENCES

[1] Cisco visual networking index: Forecast and methodology, 2009-2014. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html.

[2] Virtual routing and forwarding. http://www.cisco.com/en/US/docs/net_mgmt/active_network_abstraction/3.7/reference/guide/vrf.html.

[3] Virtual routing and forwarding. http://www.juniper.net/techpubs/software/junos/junos61/swconfig61-routing/html/instance-overview.html#1017937.

[4] T. Bates, R. Chandra, D. Katz, and Y. Rekhter. Multiprotocol Extensions for BGP-4. RFC 4760 (Draft Standard), Jan. 2007.

[5] T. Bates, E. Chen, and R. Chandra. BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP). RFC 4456 (Draft Standard), Apr. 2006.

[6] Z. ben Houidi and M. Meulle. A new VPN routing approach for large scale networks. In *Proc. IEEE ICNP*, 2010.

[7] T. Benson, A. Akella, and D. A. Maltz. Unraveling the complexity of network management. In *NSDI*, April 2009.

[8] D. Caldwell, A. Gilbert, J. Gottlieb, A. Greenberg, G. Hjalmtysson, and J. Rexford. The cutting edge of IP router configuration. In *In Proc. of Hotnets-II*, 2003.

[9] W. Enck, P. Mcdaniel, A. Greenberg, S. Sen, P. Sebos, S. Spoerel, and S. Rao. Configuration management at massive scale: System design and experience. In *In 2007 USENIX ATC*, pages 73–86, 2007.

[10] N. Feamster and H. Balakrishnan. Detecting BGP configuration faults with static analysis. In *Proceedings of USENIX NSDI*, pages 43–56, Berkeley, CA, USA, 2005.

[11] P. Garimella, Y.-W. E. Sung, N. Zhang, and S. Rao. Characterizing VLAN usage in an operational network. In *ACM INM '07*, pages 305–306, New York, NY, USA, 2007.

[12] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang. A clean slate 4D approach to network control and management. *SIGCOMM Comput. Commun. Rev.*, 35(5):41–54, 2005.

[13] T. Kamiya, S. Kusumoto, and K. Inoue. Ccfinder: a multilinguistic token-based code clone detection system for large scale source code. *IEEE Trans. Softw. Eng.*, 28(7), 2002.

[14] C. Kim, A. Gerber, C. Lund, D. Pei, and S. Sen. Scalable VPN routing via relaying. In *Proceedings of SIGMETRICS*, pages 61–72, New York, NY, USA, 2008. ACM.

[15] F. Le, G. G. Xie, D. Pei, J. Wang, and H. Zhang. Shedding light on the glue logic of the Internet routing architecture. In *Proceedings of ACM SIGCOMM*, pages 39–50, New York, NY, USA, 2008.

[16] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP misconfiguration. In *Proceedings of ACM SIGCOMM*, pages 3–16, New York, NY, USA, 2002.

[17] A. A. Mahimkar, H. H. Song, Z. Ge, A. Shaikh, J. Wang, J. Yates, Y. Zhang, and J. Emmons. Detecting the performance impact of upgrades in large operational networks. In *Proceedings of ACM SIGCOMM*, pages 303–314, New York, NY, USA, 2010.

[18] D. A. Maltz, G. Xie, J. Zhan, H. Zhang, G. Hjálmtýsson, and A. Greenberg. Routing design in operational networks: a look from the inside. In *Proceedings of ACM SIGCOMM*, pages 27–40, New York, NY, USA, 2004.

[19] M. Napierala. AT&T MPLS network and VPN services. PLNOG, 2008.

[20] T. Qiu, Z. Ge, S. Lee, J. Wang, J. Xu, and Q. Zhao. Modeling user activities in a large IPTV system. In *Proceedings of ACM IMC*, pages 430–441, New York, NY, USA, 2009.

[21] T. Qiu, Z. Ge, S. Lee, J. Wang, Q. Zhao, and J. Xu. Modeling channel popularity dynamics in a large IPTV system. In *Proceedings of ACM SIGMETRICS*, pages 275–286, New York, NY, USA, 2009.

[22] S. Raghunath and K. K. Ramakrishnan. Trade-offs in resource management for Virtual Private Networks. In *Proc. IEEE INFOCOM*, 2005.

[23] S. Raghunath, K. K. Ramakrishnan, and S. Kalyanaraman. Measurement-based characterization of IP VPNs. *IEEE/ACM Trans. Netw.*, 15:1428–1441, December 2007.

[24] Y.-W. E. Sung, C. Lund, M. Lyn, S. G. Rao, and S. Sen. Modeling and understanding end-to-end class of service policies in operational networks. In *Proceedings of SIGCOMM*, pages 219–230, New York, NY, USA, 2009. ACM.

[25] D. Turner, K. Levchenko, A. C. Snoeren, and S. Savage. California fault lines: understanding the causes and impact of network failures. In *Proceedings of ACM SIGCOMM*, pages 315–326, New York, NY, USA, 2010.

[26] Y. Zhao, Z. Zhu, Y. Chen, D. Pei, and J. Wang. Towards efficient large-scale VPN monitoring and diagnosis under operational constraints. In *Proc. IEEE INFOCOM*, pages 531–539, 2009.