Arquitetura da Internet

Ítalo Cunha, Professor Adjunto, Departamento de Ciência da Computação

I. Objetivos

Este curso é sobre a arquitetura da Internet, uma das invenções mais significativas dos últimos 30 anos e de claro impacto na sociedade. O sucesso da Internet é devido em parte à sua arquitetura, que inclui soluções para problemas complexos de escalabilidade, desempenho, gerência, robustez e custo. Este curso irá cobrir teoria, heurísticas, algoritmos e técnicas utilizadas em várias soluções empregadas na arquitetura da Internet. O conteúdo do curso tem utilidade direta para auxiliar no desenvolvimento de aplicações, operação, gerência e resolução de problemas de rede.

II. Dinâmica e Avaliação

As aulas serão como um *workshop* de discussão. Cada semana terá uma lista de artigos para leitura. O professor irá fazer uma apresentação dos trabalhos em discussão.

A avaliação considerará aspectos como a qualidade de resenhas escritas pelos alunos sobre os artigos discutidos. Os alunos irão apresentar seminários sobre tópicos relacionados à disciplina; a qualidade dos seminários também será considerada na avaliação. Uma prova e participação nas discussões também estão entre as atividades avaliativas. Outras atividades avaliativas incluem trabalhos práticos e projetos orientados.

A. Prerequisitos

O aluno deve ter cursado Redes de Computadores (DCC023) ou disciplina equivalente; em particular o aluno deve conhecer os concentos fundamentais dos protocolos IP, UDP, UDP, DNS e protocolos de roteamento.

III. Programa

Textbooks. We will use Peterson and Davie as textbook for background information on routing and transport protocols [1], as well as other more specific references as shown below.

Design Principles. Internet architecture principles [2], [3], the end-to-end argument [4], [5], naming and binding [6], [7], provider competition and its impact on the Internet's architecture [8].

Internet Autonomous Systems. Inter-domain routing relationships [9]–[11], AS-level topology properties [12], AS-level topology evolution [13], AS-level topology modeling [14].

Internet Exchange Points. Infrastructure [15], [16], impact [17], [18], mapping [19], software-defined traffic engineering [20], [21].

Unicast Routing. Intradomain distance vector and link-state routing; BGP [22].[1] Mechanisms and practices: Layer-2 routes [23], [24], route redistribution [25], intradomain route dissemination [26], routing convergence [27], [28], hot-potato routing [29], flat-label routing [30]. Routing Control and Extensibility [31]–[36]. Properties and consequences: Path inflation [37], [38], synchronization problems [39], routing stability and oscillation [40], routing outages and causes [41], and routing granularity [42].

Tunneling and VPNs. MPLS [43], [44], VPN label distribution [45], [46].

Internet Hardware. High-speed switching and routing [47]–[49], optical switching [50], programmable routers [51], [52], software routing [53].

Content Distribution. Distributed content delivery [54], [55], request routing [56], traffic engineering [57]–[59], cost- and application-aware bandwidth allocation [60]–[62].

Transport and Congestion Control. Reliable end-to-end transmission and flow control [63], congestion avoidance and control [64]–[69], TCP throughput [70], and transport on low-latency high-throughput datacenter networks [71].

[1] BGP tutorials can be found at: http://www.academ.com/nanog/feb1997/BGPTutorial/ and http://www.ittc.ku.edu/EECS/EECS_800.ira/bgp_tutorial/.

Queue Management. Router buffer sizing [72], random early detection [73], fair queueing [74], [75], queueing delay [76], explicit congestion notification [77], and differentiated services [78], [79].

Multicast and Applications. Multicast routing [80]–[82], application-layer multicast [83], erasure codes and file distribution [84], and on-demand streaming [85].

Network Measurement and Characterization. Network characterization with network support [86], [87]. Measurement methodology [88], reverse engineering [89]. Topology mapping [19], [90]–[93], IP aliasing [94], performance measurements [95], [96], broadband measurements [97], [98], traffic characterization [99], [100], anomaly detection [101]–[103].

Overlay Networks. Overlay networks [104], distributed hash tables [105]–[107], P2P content distribution and incentives [108], [109], content distribution networks [110].

Security and Privacy. Denial-of-service attacks [111], botnets [112], [113], onion routing [114], [115], ethics [116], differentially-private analysis [117], routing security [118], [119], prefix hijacks [120]–[122], network neutrality [123], [124], malicious activity [125].

Alternate Architectures and Solutions. Delay-tolerant networking [126], [127], content-based networking [128], network virtualization [129], extensible architectures [130].

REFERENCES

[1] L. Peterson and B. Davie, *Computer Networks: a Systems Approach*. Morgan Kaufmann, 2012.

[2] D. Clark, "The Design Philosophy of the DARPA Internet Protocols," in *Proc. ACM SIGCOMM*, 1988.

[3] S. Shenker, "Fundamental Design Issues for the Future Internet," *IEEE J. Selected Areas in Communications*, vol. 13, no. 7, pp. 1176–1188, 2006.

[4] J. Saltzer, D. Reed, and D. Clark, "End-to-end Arguments in System Design," *ACM Trans. Comput. Syst.*, vol. 2, no. 4, pp. 277–288, 1984.

[5] V. Cerf and R. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Trans. on Comm.*, vol. 22, no. 5, pp. 637–648, 1974.

[6] J. Saltzer, "On the Naming and Binding of Network Destinations," in *IFIP Symposium on Local Computer Networks*, 1982.

[7] P. Mockapetris and K. J. Dunlap, "Development of the Domain Name System," in *Proc. ACM SIGCOMM*, 1988.

[8] D. Clark, J. Wroclawski, K. Sollins, and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," *IEEE/ACM Trans. Netw.*, vol. 13, no. 3, pp. 462–475, 2005.

[9] M. Luckie, B. Huffaker, K. Claffy, A. Dhamdhere, and V. Giotsas, "AS Relationships, Customer Cones, and Validation," in *Proc. ACM IMC*, 2013.

[10] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, 2001.

[11] R. Anwar, H. Niaz, D. R. Choffnes, I. Cunha, P. Gill, and E. Katz-Bassett, "Investigating Interdomain Routing Policies in the Wild," in *Proc. ACM IMC*, 2015.

[12] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (in)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM Trans. Netw.*, vol. 18, no. 1, pp. 109–122, 2010.

[13] A. Dhamdhere and C. Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem," *IEEE/ACM Trans. Netw.*, vol. 19, no. 5, pp. 1420–1433, 2011.

[14] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, "Building an AS-topology Model That Captures Route Diversity," in *Proc. ACM SIGCOMM*, 2006.

[15] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger, "Peering at Peerings: On the Role of IXP Route Servers," in *Proc. ACM IMC*, 2014.

[16] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a Large European IXP," in *Proc. ACM SIGCOMM*, 2012.

[17] N. Chatzis, G. Smaragdakis, J. Boettger, T. Krenc, and A. Feldmann, "On the Benefits of Using a Large IXP as an Internet Vantage Point," in *Proc. ACM IMC*, 2013.

[18] N. Feamster, "Revealing Utilization at Internet Interconnection Points," *CoRR*, vol. abs/1603.03656, 2016.

[19] B. Augustin, B. Krishnamurthy, and W. Willinger, "IXPs: Mapped?" in *Proc. ACM IMC*, 2009.

[20] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett, "SDX: A Software Defined Internet Exchange," in *Proc. ACM SIGCOMM*, 2014.

[21] A. Gupta, R. MacDavid, R. Birkner, M. Canini, N. Feamster, J. Rexford, and L. Vanbever, "An Industrial-scale Software Defined Internet Exchange Point," in *Proc. USENIX NSDI*, 2016.

[22] J. W. Stewart, *BGP4. Inter-Domain Routing in the Internet*. Addison-Wesley, 1998.

[23] R. Perlman, "An Algorithm for Distributed Computation of a Spanningtree in an Extended LAN," in *Proc. ACM SIGCOMM*, 1985.

[24] J. McCauley, M. Zhao, E. J. Jackson, B. Raghavan, S. Ratnasamy, and S. Shenker, "The Deforestation of L2," in *Proc. ACM SIGCOMM*, 2016.

[25] F. Le, G. Xie, D. Pei, J. Wang, and H. Zhang, "Shedding Light on the Glue Logic of the Internet Routing Architecture," in *Proc. ACM SIGCOMM*, 2008.

[26] N. Gvozdiev, B. Karp, and M. Handley, "LOUP: The Principles and Practice of Intra-Domain Route Dissemination," in *Proc. USENIX NSDI*, 2013.

[27] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet Routing Convergence," *IEEE/ACM Trans. Netw.*, vol. 9, no. 3, pp. 293–306, 2001.

[28] T. G. Griffin and G. Wilfong, "An Analysis of BGP Convergence Properties," in *Proc. ACM SIGCOMM*, 1999.

[29] R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford, "Dynamics of Hot-potato Routing in IP Networks," in *Proc. ACM SIGMETRICS*, 2004.

[30] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, I. Stoica, and S. Shenker, "ROFL: Routing on Flat Labels," in *Proc. ACM SIGCOMM*, 2006.

[31] R. R. Sambasivan, D. Tran-Lam, A. Akella, and P. Steenkiste, "Bootstrapping evolvability for inter-domain routing with D-BGP," in *Proc. ACM SIGCOMM*, 2017.

[32] S. Vissicchio, O. Tilmans, L. Vanbever, and J. Rexford, "Central Control Over Distributed Routing," in *Proc. ACM SIGCOMM*, 2015.

[33] R. Hartert, S. Vissicchio, P. Schaus, O. Bonaventure, C. Filsfils, T. Telkamp, and P. Francois, "A Declarative and Expressive Approach to Control Forwarding Paths in Carrier-Grade Networks," in *Proc. ACM SIGCOMM*, 2015.

[34] J. Reich, C. Monsanto, N. Foster, J. Rexford, and D. Walker, "Modular SDN Programming with Pyretic," *;login Magazine*, vol. 38, pp. 128–134, 2013.

[35] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat, "B4: Experience with a Globally-deployed Software Defined Wan," in *Proc. ACM SIGCOMM*, 2013.

[36] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving High Utilization with Software-driven WAN," in *Proc. ACM SIGCOMM*, 2013.

[37] N. Spring, R. Mahajan, and T. Anderson, "The Causes of Path Inflation," in *Proc. ACM SIGCOMM*, 2003.

[38] K. Zarifis, T. Flach, S. Nori, D. Choffnes, R. Govindan, E. Katz-Bassett, Z. M. Mao, and M. Welsh, "Diagnosing Path Inflation of Mobile Client Traffic," in *Proc. PAM*, 2014.

[39] S. Floyd and V. Jacobson, "The Synchronization of Periodic Routing Messages," *IEEE/ACM Trans. Netw.*, vol. 2, no. 2, pp. 122–136, 1994.

[40] A. Shaikh, A. Varma, L. Kalampoukas, and R. Dube, "Routing Stability in Congested Networks: Experimentation and Analysis," in *Proc. ACM SIGCOMM*, 2000.

[41] F. Wang, Z. M. Mao, J. Wang, L. Gao, and R. Bush, "A Measurement Study on the Impact of Routing Events on End-to-end Internet Path Performance," in *Proc. ACM SIGCOMM*, 2006.

[42] W. Mühlbauer, S. Uhlig, B. Fu, M. Meulle, and O. Maennel, "In Search for an Appropriate Granularity to Model Routing Policies," in *Proc. ACM SIGCOMM*, 2007.

[43] J. Sommers, P. Barford, and B. Eriksson, "On the Prevalence and Characteristics of MPLS Deployments in the Open Internet," in *Proc. ACM IMC*, 2011.

[44] A. Pathak, M. Zhang, Y. C. Hu, R. Mahajan, and D. Maltz, "Latency Inflation with MPLS-based Traffic Engineering," in *Proc. ACM IMC*, 2011.

[45] Z. Ben-Houidi and M. Meulle, "A New VPN Routing Approach for Large Scale Networks," in *Proc. IEEE ICNP*, 2010.

[46] E. Rosen and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)," 2006.

[47] N. McKeown, "The iSLIP Scheduling Algorithm for Input-queued Switches," *IEEE/ACM Trans. Netw.*, vol. 7, no. 2, pp. 188–201, 1999.

[48] C. Partridge, P. P. Carvey, E. Burgess, I. Castineyra, T. Clarke, L. Graham, M. Hathaway, P. Herman, A. King, S. Kohalmi, T. Ma, J. Mcallen, T. Mendez, W. C. Milliken, R. Pettyjohn, J. Rokosz, J. Seeger, M. Sollins, S. Storch, B. Tober, and G. D. Troxel, "A 50-Gb/s IP Router," *IEEE/ACM Trans. Netw.*, vol. 6, no. 3, pp. 237–248, 1998.

[49] A. Singh, J. Ong, A. Agarwal, G. Anderson, A. Armistead, R. Bannon, S. Boving, G. Desai, B. Felderman, P. Germano, A. Kanagala, J. Provost, J. Simmons, E. Tanda, J. Wanderer, U. Hölzle, S. Stuart, and A. Vahdat, "Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network," in *Proc. ACM SIGCOMM*, 2015.

[50] B. Mukherjee, "WDM Optical Communication Networks: Progress and Challenges," *IEEE J. Selected Areas in Communications*, vol. 18, no. 10, pp. 1810–1824, 2000.

[51] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker, "P4: Programming Protocol-independent Packet Processors," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 87–95, 2014.

[52] A. Sivaraman, A. Cheung, M. Budiu, C. Kim, M. Alizadeh, H. Balakrishnan, G. Varghese, N. McKeown, and S. Licking, "Packet Transactions: High-Level Programming for Line-Rate Switches," in *Proc. ACM SIGCOMM*, 2016.

[53] H. Asai and Y. Ohara, "Poptrie: A Compressed Trie with Population Count for Fast and Scalable Software IP Routing Table Lookup," 2015.

[54] J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, and B. Weihl, "Globally Distributed Content Delivery," *IEEE Internet Computing*, vol. 6, no. 5, pp. 50–58, 2002.

[55] T. Böttger, F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig, "Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN," *CoRR*, vol. abs/1606.05519, 2016. [Online]. Available: http://arxiv.org/abs/1606.05519

[56] M. T. Fangfei Chen, Ramesh K. Sitaraman, "End-User Mapping: Next Generation Request Routing for Content Delivery," in *Proc. ACM SIGCOMM*, 2015.

[57] K.-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain, V. Lin, C. Rice, B. Rogan, A. Singh, B. Tanaka, M. Verma, P. Sood, M. Tariq, M. Tierney, D. Trumic, V. Valancius, C. Ying, M. Kallahalla, B. Koley, and A. Vahdat, "Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering," in *Proc. ACM SIGCOMM*, 2017.

[58] B. Schlinker, H. Kim, T. Cui, E. Katz-Bassett, H. V. Madhyastha, I. Cunha, J. Quinn, S. Hasan, P. Lapukhov, and H. Zeng, "Engineering Egress with Edge Fabric: Steering Oceans of Content to the World," in *Proc. ACM SIGCOMM*, 2017.

[59] P. Sun, L. Vanbever, and J. Rexford, "Scalable Programmable Inbound Traffic Engineering," in *Proc. ACM SOSR*, 2015.

[60] V. Jalaparti, I. Bliznets, S. Kandula, B. Lucier, and I. Menache, "Dynamic Pricing and Traffic Engineering for Timely Inter-Datacenter Transfers," in *Proc. ACM SIGCOMM*, 2016.

[61] A. Kumar, S. Jain, U. Naik, A. Raghuraman, N. Kasinadhuni, E. C. Zermeno, C. S. Gunn, J. Ai, B. Carlin, M. Amarandei-Stavila, M. Robin, A. Siganporia, S. Stuart, and A. Vahdat, "BwE: Flexible, Hierarchical Bandwidth Allocation for WAN Distributed Computing," in *Proc. ACM SIGCOMM*, 2015.

[62] H. H. Liu, R. Viswanathan, M. Calder, A. Akella, R. Mahajan, J. Padhye, and M. Zhang, "Efficiently Delivering Online Services over Integrated Infrastructure," in *Proc. USENIX NSDI*, 2016.

[63] V. Cerf and R. Khan, "A Protocol for Packet Network Intercommunication," *IEEE Trans. on Comm.*, vol. 22, no. 5, pp. 637–648, 1974.

[64] V. Jacobson, "Congestion Avoidance and Control," in *Proc. ACM SIGCOMM*, 1988.

[65] K. K. Ramakrishnan and R. Jain, "A binary feedback scheme for congestion avoidance in computer networks," *ACM Trans. Comput. Syst.*, vol. 8, no. 2, pp. 158–181, 1990.

[66] L. S. Brakmo, S. W. O'Malley, and L. L. Peterson, "TCP Vegas: New Techniques for Congestion Detection and Avoidance," in *Proc. ACM SIGCOMM*, 1994.

[67] K. Winstein and H. Balakrishnan, "TCP Ex Machina: Computer-generated Congestion Control," in *Proc. ACM SIGCOMM*, 2013.

[68] A. Langley, A. Riddoch, A. Wilk, A. Vicente, C. Krasic, D. Zhang, F. Yang, F. Kouranov, I. Swett, J. Iyengar, J. Bailey, J. Dorfman, J. Roskind, J. Kulik, P. Westin, R. Tenneti, R. Shade, R. Hamilton, V. Vasiliev, W.-T. Chang, and Z. Shi, "The QUIC Transport Protocol: Design and Internet-Scale Deployment," in *Proc. ACM SIGCOMM*, 2017.

[69] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, "BBR: Congestion-Based Congestion Control," *Queue*, vol. 14, no. 5, pp. 50:20–50:53, 2016.

[70] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Throughput: A Simple Model and Its Empirical Validation," in *Proc. ACM SIGCOMM*, 1998.

[71] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan, "Data Center TCP (DCTCP)," in *Proc. ACM SIGCOMM*, 2010.

[72] N. Beheshti, Y. Ganjali, M. Ghobadi, N. McKeown, and G. Salmon, "Experimental Study of Router Buffer Sizing," in *Proc. ACM IMC*, 2008.

[73] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Trans. Netw.*, vol. 1, no. 4, pp. 397–413, 1993.

[74] A. Demers, S. Keshav, and S. Shenker, "Analysis and Simulation of a Fair Queueing Algorithm," in *Proc. ACM SIGCOMM*, 1989.

[75] I. Stoica, S. Shenker, and H. Zhang, "Core-stateless Fair Queueing: A Scalable Architecture to Approximate Fair Bandwidth Allocations in High-speed Networks," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 33–46, 2003.

[76] K. Nichols and V. Jacobson, "Controlling Queue Delay," *Queue*, vol. 10, no. 5, pp. 20–34, 2012.

[77] D. Katabi, M. Handley, and C. Rohrs, "Congestion Control for High Bandwidth-delay Product Networks," in *Proc. ACM SIGCOMM*, 2002.

[78] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," 1994.

[79] B. Davie, "Deployment Experience with Differentiated Services," in *SIGCOMM Workshop on Revisiting IP QoS*, 2003.

[80] S. E. Deering, "Multicast Routing in Internetworks and Extended LANs," in *Proc. ACM SIGCOMM*, 1988.

[81] S. Floyd, V. Jacobson, C.-G. Liu, S. McCanne, and L. Zhang, "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing," *IEEE/ACM Trans. Netw.*, vol. 5, no. 6, pp. 784–803, 1997.

[82] S. Ratnasamy, A. Ermolinskiy, and S. Shenker, "Revisiting IP Multicast," in *Proc. ACM SIGCOMM*, 2006.

[83] Y.-H. Chu, S. G. Rao, S. Seshan, and H. Z., "A Case for End System Multicast," *IEEE J. Selected Areas in Communications*, vol. 20, no. 8, pp. 1456–1471, 2002.

[84] J. W. Byers, M. Luby, M. Mitzenmacher, and A. Rege, "A Digital Fountain Approach to Reliable Distribution of Bulk Data," in *Proc. ACM SIGCOMM*, 1998.

[85] D. L. Eager, M. K. Vernon, and J. Zahorjan, "Bandwidth Skimming: a Technique for Cost-effective Video on Demand," 1999.

[86] D. Turner, K. Levchenko, A. Snoeren, and S. Savage, "California Fault Lines: Understanding the Causes and Impact of Network Failures," in *Proc. ACM SIGCOMM*, 2010.

[87] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. N. Chuah, Y. Ganjali, and C. Diot, "Characterization of Failures in an Operational IP Backbone Network," *IEEE/ACM Trans. Netw.*, vol. 16, no. 4, pp. 749–762, 2008.

[88] V. Paxson, "Strategies for Sound Internet Measurement," in *Proc. ACM IMC*, 2004.

[89] N. Spring, D. Wetherall, and T. Anderson, "Reverse Engineering the Internet," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 1, pp. 3–8, 2004.

[90] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP Topologies with Rocketfuel," *IEEE/ACM Trans. Netw.*, vol. 12, no. 1, pp. 2–16, 2004.

[91] I. Cunha, R. Teixeira, D. Veitch, and C. Diot, "Predicting and Tracking Internet Path Changes," in *Proc. ACM SIGCOMM*, 2011.

[92] E. Katz-Bassett, H. Madhyastha, V. Adhikari, C. Scott, J. Sherry, P. van Wesep, A. Krishnamurthy, and T. Anderson, "Reverse Traceroute," in *Proc. USENIX NSDI*, 2010.

[93] H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane: an Information Plane for Distributed Services," in *Proc. USENIX OSDI*, 2006.

[94] K. Keys, Y. Hyun, M. Luckie, and K. Claffy, "Internet-scale IPv4 Alias Resolution with MIDAR," *IEEE/ACM Trans. Netw.*, vol. 21, no. 2, pp. 383–399, 2013.

[95] R. Kapoor, L.-J. Chen, L. Lao, M. Gerla, and M. Y. Sanadidi, "CapProbe: A Simple and Accurate Capacity Estimation Technique," in *Proc. ACM SIGCOMM*, 2004.

[96] J. Sommers, P. Barford, N. Duffield, and A. Ron, "Accurate and Efficient SLA Compliance Monitoring," in *Proc. ACM SIGCOMM*, 2007.

[97] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson, "Netalyzr: Illuminating The Edge Network," in *Proc. ACM IMC*, 2010.

[98] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè, "Broadband Internet Performance: a View from the Gateway," in *Proc. ACM SIGCOMM*, 2011.

[99] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet Inter-domain Traffic," in *Proc. ACM SIGCOMM*, 2010.

[100] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot, "Packet-level Traffic Measurements from the Sprint IP Backbone," *IEEE Network Magazine*, vol. 17, no. 6, pp. 6–16, 2003.

[101] M. Dischinger, M. Marcon, S. Guha, K. P. Gummadi, R. Mahajan, and S. Saroiu, "Glasnost: Enabling End Users to Detect Traffic Differentiation," in *Proc. USENIX NSDI*, 2010.

[102] A. Lakhina, M. Crovella, and C. Diot, "Diagnosing Network-wide Traffic Anomalies," in *Proc. ACM SIGCOMM*, 2004.

[103] F. Silveira, C. Diot, N. Taft, and R. Govindan, "ASTUTE: Detecting a Different Class of Traffic Anomalies," in *Proc. ACM SIGCOMM*, 2010.

[104] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient Overlay Networks," in *Proc. ACM SOSP*, 2001.

[105] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan, "Chord: a Scalable Peer-to-peer Lookup Protocol for Internet Applications," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 17–32, 2003.

[106] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph, "Tapestry: A Fault-tolerant Wide-area Application Infrastructure," *SIGCOMM Comput. Commun. Rev.*, vol. 32, no. 1, pp. 81–81, 2002.

[107] A. I. T. Rowstron and P. Druschel, "Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems," in *Proceedings of the IFIP/ACM International Conference on Distributed Systems*, 2001.

[108] M. Piatek, T. Isdal, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "Do Incentives Build Robustness in Bit Torrent," in *Proc. USENIX NSDI*, 2007.

[109] M. Piatek, A. Krishnamurthy, A. Venkataramani, R. Yang, D. Zhang, and A. Jaffe, "Contracts: Practical Contribution

Incentives for P2P Live Streaming," in *Proc. USENIX NSDI*, 2010.

[110] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy, "An Analysis of Internet Content Delivery Systems," in *Proc. USENIX OSDI*, 2002.

[111] A. Hussain, J. Heidemann, and C. Papadopoulos, "A Framework for Classifying Denial of Service Attacks," in *Proc. ACM SIGCOMM*, 2003.

[112] S. Staniford, V. Paxson, and N. Weaver, "How to Own the Internet in Your Spare Time," in *USENIX Security Symposium*, 2002.

[113] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, "Understanding the mirai botnet," in *USENIX Security Symposium*, 2017.

[114] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The Second-generation Onion Router," in *USENIX Security Symposium*, 2004.

[115] I. Clarke, S. G. Miller, T. W. Hong, O. Sandberg, and B. Wiley, "Protecting Free Expression Online with Freenet," *IEEE Internet Computing*, vol. 6, no. 1, pp. 40–49, 2002.

[116] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, "The Menlo Report," *IEEE Security and Privacy*, vol. 10, no. 2, pp. 71–75, 2012.

[117] F. McSherry and R. Mahajan, "Differentially-private Network Trace Analysis," in *Proc. ACM SIGCOMM*, 2010.

[118] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira, "Jumpstarting BGP Security with Path-End Validation," in *Proc. ACM SIGCOMM*, 2016.

[119] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal, "RAPTOR: Routing Attacks on Privacy in Tor," in *Proc. USENIX CSS*, 2015.

[120] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush, "iSPY: Detecting IP Prefix Hijacking On My Own," in *Proc. ACM SIGCOMM*, 2008.

[121] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," in *Proc. ACM SIGCOMM*, 2007.

[122] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking Bitcoin: Routing Attacks on Cryptocurrencies," in *IEEE Symposium on Security and Privacy (to appear)*, 2017.

[123] T. Flach, P. Papageorge, A. Terzis, L. Pedrosa, Y. Cheng, T. Karim, E. Katz-Bassett, and R. Govindan, "An Internet-Wide Analysis of Traffic Policing," in *Proc. ACM SIGCOMM*, 2016.

[124] Y. Yiakoumis, S. Katti, and N. McKeown, "Neutral Net Neutrality," in *Proc. ACM SIGCOMM*, 2016.

[125] M. Konte, R. Perdisci, and N. Feamster, "ASwatch: An AS Reputation System to Expose Bulletproof Hosting ASes," in *Proc. ACM SIGCOMM*, 2015.

[126] K. Fall and S. Farrell, "DTN: An Architectural Retrospective," *IEEE J. Selected Areas in Communications*, vol. 26, no. 5, pp. 828–836, 2008.

[127] K. Fall, "A Delay-tolerant Network Architecture for Challenged Internets," in *Proc. ACM SIGCOMM*, 2003.

[128] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking Named Content," in *Proc. ACM CoNEXT*, 2009.

[129] T. Anderson, L. Peterson, S. Shenker, and J. Turner, "Overcoming the Internet Impasse Through Virtualization," *IEEE Computer*, vol. 38, no. 4, pp. 34–41, 2005.

[130] D. Han, A. Anand, F. Dogar, B. Li, H. Lim, M. Machado, A. Mukundan, W. Wu, A. Akella, D. G. Andersen, J. W. Byers, S. Seshan, and P. Steenkiste, "XIA: Efficient Support for Evolvable Internetworking," in *Proc. USENIX NSDI*, 2012.