References

Abts, D. and Felderman, B., (2012), "A guided tour through data-center networking", Queue, Vol.10, No.5, pp.10, 2012

Al-Fares, M., Loukissas, A., and Vahdat, A., (2008), "A scalable, commodity data center network architecture", ACM SIGCOMM Computer Communication Review, Vol.38, No.4, pp.63–74, 2008

Al-fares, M., Radhakrishnan, S., Raghavan, B., Huang, N., and Vahdat, A., (2010), "Hedera: Dynamic flow scheduling for data center networks", In In Proc. of Networked Systems Design and Implementation (NSDI) Symposium

Alizadeh, M., Edsall, T., Dharmapurikar, S., Vaidyanathan, R., Chu, K., Fingerhut, A., Matus, F., Pan, R., Yadav, N., Varghese, G., et al., (2014), "CONGA: Distributed congestion-aware load balancing for datacenters", In Proceedings of the 2014 ACM conference on SIGCOMM, pp. 503–514, ACM

Alizadeh, M., Greenberg, A., Maltz, D. A., Padhye, J., Patel, P., Prabhakar, B., Sengupta, S., and Sridharan, M., (2010), "Data Center TCP (DCTCP)", In Proceedings of the ACM SIGCOMM 2010 Conference, SIGCOMM '10, pp. 63–74, New York, NY, USA, ACM

Alizadeh, M., Javanmard, A., and Prabhakar, B., (2011), "Analysis of DCTCP: stability, convergence, and fairness", In Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems, pp. 73–84, ACM

Alizadeh, M., Kabbani, A., Edsall, T., Prabhakar, B., Vahdat, A., and Yasuda, M., (2012), "Less is More: Trading a Little Bandwidth for Ultra-low Latency in the Data Center", In Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation, NSDI'12, Berkeley, CA, USA, USENIX Association

Alizadeh, M., Yang, S., Sharif, M., Katti, S., McKeown, N., Prabhakar, B., and Shenker, S., (2013), "pFabric: Minimal Near-optimal Datacenter Transport", In Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13, pp. 435–446, New York, NY, USA, ACM

Barré, S., Paasch, C., and Bonaventure, O., (2011), "Multipath TCP: from theory to practice", NETWORKING 2011,, pp.444–457, 2011

Chen, K., Hu, C., Zhang, X., Zheng, K., Chen, Y., and Vasilakos, A. V., (2011), "Survey on routing in data centers: insights and future directions", IEEE network, Vol.25, No.4, 2011

Chen, L., Li, B., and Li, B., (2013), "On meeting deadlines in datacenter networks", Tsinghua Science and Technology, Vol.18, No.3, pp.273–285, 2013

Chowdhury, M., Zaharia, M., Ma, J., Jordan, M. I., and Stoica, I., (2011), "Managing Data Transfers in Computer Clusters with Orchestra", In Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11, pp. 98–109, New York, NY, USA, ACM

Cisco, (2015), "Cisco Global Cloud Index: Forecast and Methodology, 2015–2020", 2015

Cui, W. and Qian, C., (2014), "Difs: Distributed flow scheduling for adaptive routing in hierarchical data center networks", In Proceedings of the tenth ACM/IEEE symposium on Architectures for networking and communications systems, pp. 53–64, ACM

Divakaran, D. M., Altman, E., Post, G., Noirie, L., and Primet, P. V.-B., (2009), "From packets to XLFrames: sand and rocks for transfer of mice and elephants", In INFOCOM Workshops 2009, IEEE, pp. 1–6, IEEE

Dixit, A., Prakash, P., Hu, Y. C., and Kompella, R. R., (2013), "On the impact of packet spraying in data center networks", In INFOCOM, 2013 Proceedings IEEE, pp. 2130–2138, IEEE

Dukkipati, N. and McKeown, N., (2006), "Why flow-completion time is the right metric for con-

gestion control", ACM SIGCOMM Computer Communication Review, Vol.36, No.1, pp.59–62, 2006

Farrington, N., Porter, G., Radhakrishnan, S., Bazzaz, H. H., Subramanya, V., Fainman, Y., Papen, G., and Vahdat, A., (2010), "Helios: a hybrid electrical/optical switch architecture for modular data centers", ACM SIGCOMM Computer Communication Review, Vol.40, No.4, pp.339–350, 2010

Greenberg, A., Hamilton, J., Maltz, D. A., and Patel, P., (2008), "The Cost of a Cloud: Research Problems in Data Center Networks", SIGCOMM Comput. Commun. Rev., Vol.39, No.1, pp.68–73, December 2008

Greenberg, A., Hamilton, J. R., Jain, N., Kandula, S., Kim, C., Lahiri, P., Maltz, D. A., Patel, P., and Sengupta, S., (2009), "VL2: a scalable and flexible data center network", In ACM SIGCOMM computer communication review, volume 39, pp. 51–62, ACM

Grosvenor, M. P., Schwarzkopf, M., Gog, I., Watson, R. N., Moore, A. W., Hand, S., and Crowcroft, J., (2015), "Queues don't matter when you can jump them!", In NSDI, pp. 1–14

Guo, C., Lu, G., Li, D., Wu, H., Zhang, X., Shi, Y., Tian, C., Zhang, Y., and Lu, S., (2009), "BCube: a high performance, server-centric network architecture for modular data centers", ACM SIGCOMM Computer Communication Review, Vol.39, No.4, pp.63–74, 2009

Guo, C., Wu, H., Tan, K., Shi, L., Zhang, Y., and Lu, S., (2008), "Dcell: a scalable and fault-tolerant network structure for data centers", In ACM SIGCOMM Computer Communication Review, volume 38, pp. 75–86, ACM

Handley, M., Raiciu, C., Agache, A., Voinescu, A., Moore, A. W., Antichi, G., and Wójcik, M., (2017), "Re-architecting datacenter networks and stacks for low latency and high performance", In Proceedings of the Conference of the ACM Special Interest Group on Data Communication, pp. 29–42, ACM

He, K., Rozner, E., Agarwal, K., Felter, W., Carter, J., and Akella, A., (2015), "Presto: Edge-based load balancing for fast datacenter networks", In Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, pp. 465–478, ACM

T. Hoff,Latency is everywhere and it costs you sales - How to crush it,

Scott Hogg,Does your network support Jumbo Frames and should you enable it?,

Hong, C.-Y., Caesar, M., and Godfrey, P., (2012), "Finishing flows quickly with preemptive scheduling", ACM SIGCOMM Computer Communication Review, Vol.42, No.4, pp.127–138, 2012

Hopps, C. E., (2000), "Analysis of an equal-cost multi-path algorithm", 2000

Jalaparti, V., Bodik, P., Kandula, S., Menache, I., Rybalkin, M., and Yan, C., (2013), "Speeding Up Distributed Request-response Workflows", In Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13, pp. 219–230, New York, NY, USA, ACM

Joy, S. and Nayak, A., (2015), "Improving flow completion time for short flows in datacenter networks", In Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on, pp. 700–705, IEEE

Kabbani, A., Vamanan, B., Hasan, J., and Duchene, F., (2014), "FlowBender: Flow-level Adaptive Routing for Improved Latency and Throughput in Datacenter Networks", In Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies, pp. 149–160, ACM

Kant, K., (2009), "Data center evolution: A tutorial on state of the art, issues, and challenges", Computer Networks, Vol.53, No.17, pp.2939–2965, 2009

Krevat, E., Vasudevan, V., Phanishayee, A., Andersen, D. G., Ganger, G. R., Gibson, G. A., and Seshan, S., (2007), "On application-level approaches to avoiding TCP throughput collapse in cluster-based storage systems", In Proceedings of the 2nd international workshop on Petascale data storage: held in conjunction with Supercomputing'07, pp. 1–4, ACM

Kumar, R., (2014), "The Mathematical way to Decide", 2014

Lebiednik, B., Mangal, A., and Tiwari, N., (2016), "A Survey and Evaluation of Data Center Network Topologies", CoRR, Vol.abs/1605.01701„ 2016

Lee, C., Park, C., Jang, K., Moon, S. B., and Han, D., (2015), "Accurate Latency-based Congestion Feedback for Datacenters.", In USENIX Annual Technical Conference, pp. 403–415

Liu, S., Xu, H., and Cai, Z., (2013), "Low Latency Datacenter Networking: A Short Survey", CoRR, Vol.abs/1312.3455,, 2013

Mai, L., Rupprecht, L., Alim, A., Costa, P., Migliavacca, M., Pietzuch, P., and Wolf, A. L., (2014), "NetAgg: Using Middleboxes for Application-specific On-path Aggregation in Data Centres", In Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies, CoNEXT '14, pp. 249–262, New York, NY, USA, ACM

Mudigonda, J., Yalagandula, P., Al-Fares, M., and Mogul, J. C., (2010), "Spain: Cots data-center ethernet for multipathing over arbitrary topologies.", In NSDI, volume 10, pp. 18–33

Munir, A., Baig, G., Irteza, S. M., Qazi, I. A., Liu, A. X., and Dogar, F. R., (2014), "Friends, Not Foes: Synthesizing Existing Transport Strategies for Data Center Networks", In Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14, pp. 491–502, New York, NY, USA, ACM

Munir, A., Qazi, I. A., Uzmi, Z. A., Mushtaq, A., Ismail, S. N., Iqbal, M. S., and Khan, B., (2013), "Minimizing flow completion times in data centers", In INFOCOM, 2013 Proceedings IEEE, pp. 2157–2165, IEEE

Murray, D., Koziniec, T., Lee, K., and Dixon, M.,, "Large MTU's and internet performance.", In 13th IEEE Conference on High Performance Switching and Routing (HPSR 2012), pp. 82–87

Niranjan Mysore, R., Pamboris, A., Farrington, N., Huang, N., Miri, P., Radhakrishnan, S., Subramanya, V., and Vahdat, A., (2009), "Portland: a scalable fault-tolerant layer 2 data center network fabric", In ACM SIGCOMM Computer Communication Review, volume 39, pp. 39–50, ACM

Perry, J., Balakrishnan, H., and Shah, D., (2017), "Flowtune: Flowlet Control for Datacenter Networks", In 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17), pp. 421–435, USENIX Association

Perry, J., Ousterhout, A., Balakrishnan, H., Shah, D., and Fugal, H., (2014), "Fastpass: A Centralized "Zero-queue" Datacenter Network", In Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14, pp. 307–318, New York, NY, USA, ACM

Prakash, P., Dixit, A., Hu, Y. C., and Kompella, R., (2012), "The TCP Outcast Problem: Exposing Unfairness in Data Center Networks", In 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12), pp. 413–426, San Jose, CA, USENIX

Prakash, P., Lee, M., Hu, Y. C., Kompella, R. R., et al., (2013), "Jumbo frames or not: That is the question!", 2013

Raiciu, C., Barre, S., Pluntke, C., Greenhalgh, A., Wischik, D., and Handley, M., (2011), "Improving datacenter performance and robustness with multipath tcp", ACM SIGCOMM Computer Communication Review, Vol.41, No.4, pp.266–277, 2011

Ramaboli, A. L., Falowo, O. E., and Chan, A. H., (2012), "Bandwidth aggregation in heterogeneous wireless networks: A survey of current approaches and issues", Journal of Network and Computer Applications, Vol.35, No.6, pp.1674–1690, 2012

Rojas-Cessa, R., Kaymak, Y., and Dong, Z., (2015), "Schemes for fast transmission of flows in data center networks", IEEE Communications Surveys & Tutorials, Vol.17, No.3, pp.1391–1422, 2015

Salyers, D., Jiang, Y., Striegel, A., and Poellabauer, C., (2007), "JumboGen: dynamic jumbo frame generation for network performance scalability", ACM SIGCOMM Computer Communication Review, Vol.37, No.5, pp.53–64, 2007

Singla, A., Hong, C.-Y., Popa, L., and Godfrey, P. B., (2012), "Jellyfish: Networking Data Centers, Randomly.", In NSDI, volume 12, pp. 17–17

Vamanan, B., Hasan, J., and Vijaykumar, T., (2012), "Deadline-aware datacenter tcp (d2tcp)", ACM SIGCOMM Computer Communication Review, Vol.42, No.4, pp.115–126, 2012

Vanini, E., Pan, R., Alizadeh, M., Taheri, P., and Edsall, T., (2017), "Let It Flow: Resilient Asymmetric Load Balancing with Flowlet Switching.", In NSDI, pp. 407–420

Vasudevan, V., Phanishayee, A., Shah, H., Krevat, E., Andersen, D. G., Ganger, G. R., Gibson, G. A., and Mueller, B., (2009), "Safe and effective fine-grained TCP retransmissions for datacenter communication", In ACM SIGCOMM computer communication review, volume 39, pp. 303–314, ACM

Wang, G., Andersen, D. G., Kaminsky, M., Papagiannaki, K., Ng, T., Kozuch, M., and Ryan, M., (2010), "c-Through: Part-time optics in data centers", In ACM SIGCOMM Computer Communication Review, volume 40, pp. 327–338, ACM

Wang, T., Su, Z., Xia, Y., and Hamdi, M., (2014), "Rethinking the data center networking: Architecture, network protocols, and resource sharing", IEEE access, Vol.2„ pp.1481–1496, 2014

Wilson, C., Ballani, H., Karagiannis, T., and Rowtron, A., (2011), "Better never than late: Meeting deadlines in datacenter networks", In ACM SIGCOMM Computer Communication Review, volume 41, pp. 50–61, ACM

Wischik, D., Handley, M., and Braun, M. B., (2008), "The resource pooling principle", ACM SIGCOMM Computer Communication Review, Vol.38, No.5, pp.47–52, 2008

Wu, H., Feng, Z., Guo, C., and Zhang, Y., (2013), "ICTCP: Incast congestion control for TCP in data-center networks", IEEE/ACM transactions on networking, Vol.21, No.2, pp.345–358, 2013

Wu, X. and Yang, X., (2012), "Dard: Distributed adaptive routing for datacenter networks", In Distributed Computing Systems (ICDCS), 2012 IEEE 32nd International Conference on, pp. 32–41, IEEE

Xia, W., Zhao, P., Wen, Y., and Xie, H., (2017), "A survey on data center networking (DCN): infrastructure and operations", IEEE Communications Surveys & Tutorials, Vol.19, No.1, pp.640–656, 2017

Xu, H. and Li, B., (2013), "RepFlow: Minimizing flow completion times with replicated flows in data centers", arXiv preprint arXiv:1307.7451,, 2013

Zats, D., Das, T., Mohan, P., Borthakur, D., and Katz, R., (2012), "DeTail: reducing the flow completion time tail in datacenter networks", ACM SIGCOMM Computer Communication Review, Vol.42, No.4, pp.139–150, 2012

Zhang, H., Zhang, J., Bai, W., Chen, K., and Chowdhury, M., (2017), "Resilient datacenter load balancing in the wild", In Proceedings of the Conference of the ACM Special Interest Group on Data Communication, pp. 253–266, ACM

Zhang, T., Wang, J., Huang, J., Huang, Y., Chen, J., and Pan, Y., (2016), "Adaptive marking threshold method for delay-sensitive TCP in data center network", Journal of Network and Computer Applications, Vol.61„ pp.222–234, 2016