

References

- [1] http://en.wikipedia.org/wiki/Naga_Jolokia. **xvi**
- [2] <http://www.qdrsram.com>. **6, 184, 189, 212, 238, 280, 293, 323**
- [3] <http://www.micron.com/products/dram>. **6, 184, 188, 212, 232, 271, 280, 323**
- [4] Cisco Systems Inc. Cisco Catalyst 6500 Series Router. http://www.cisco.com/en/US/products/hw/switches/ps708/\products_data_sheet0900aecd8017376e.html. **7, 90, 137, 187, 222, 259**
- [5] Cisco Systems Inc. Cisco HFR. <http://www.cisco.com/en/US/products/ps5763/>. **7, 90, 150, 175**
- [6] V. Cuppu, B. Jacob, B. Davis, and T. Mudge. A performance comparison of contemporary DRAM architectures. In *Proc. 26th International Symposium on Computer Architecture (ISCA '99)*, pages 222–233, Atlanta, Georgia, May 1999. **8**
- [7] <http://www.rambus.com/>. **8, 325**
- [8] <http://www.rldram.com>. **8, 189, 212**
- [9] <http://www.fujitsu.com/us/services/edevices/\microelectronics/memory/fcram>. **8, 189, 212**
- [10] R. R. Schaller. Moore's law: Past, present and future. *IEEE Spectrum*, 34(6):52–59, June 1997. **9**
- [11] N. McKeown, V. Anantharam, and J. Walrand. Achieving 100% throughput in an input queued switch. In *Proc. of IEEE INFOCOM '96*, volume 1, pages 296–302, March 1996. **10, 44, 113**
- [12] J. Dai and B. Prabhakar. The throughput of data switches with and without speedup. In *Proc. of IEEE INFOCOM '00*, pages 556–564, Tel Aviv, Israel, March 2000. **10, 43, 44, 102, 113, 114, 135**
- [13] S. Chuang, A. Goel, N. McKeown, and B. Prabhakar. Matching output queueing with a combined input/output queued switch. *IEEE J.Sel. Areas in Communications*, 17(6):1030–1039, June 1999. **10, 12, 19, 44, 46, 79, 104, 106, 110, 111, 112, 116, 117, 118, 126, 127**
- [14] N. McKeown. iSLIP: A scheduling algorithm for input queued switches. *IEEE Transactions on Networking*, 7(2), April 1999. **11**

- [15] Y. Tamir and H. C. Chi. Symmetric crossbar arbiters for VLSI communication switches. *IEEE Transactions on Parallel and Distributed Systems*, 4(1):13–27, January 1993. 11, 76, 332
- [16] J. N. Giacopelli, J. J. Hickey, W. S. Marcus, W. D. Sincoslie, and M. Littlewood. Sunshine: A high performance self-routing broadband packet switch architecture. *IEEE J. Select. Areas Commun.*, 9:1289–1298, 1991. 11
- [17] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queuing algorithm. In *ACM Computer Communication Review (SIGCOMM '89)*, pages 3–12, 1989. 14, 53, 229
- [18] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The single node case. *IEEE/ACM Transaction on Networking*, 1(3):344–357, June 1993. 14, 53, 230
- [19] L. Zhang. Virtual clock: A new traffic control algorithm for packet switching networks. *ACM Transactions on Computer Systems*, 9(2):101–124, 1990. 14, 230
- [20] M. Shreedhar and G. Varghese. Efficient fair queueing using deficit round robin. In *Proc. of ACM SIGCOMM '95*, pages 231–242, September 1995. 14, 230
- [21] J. Bennett and H. Zhang. WF2Q: Worst-case fair weighted fair queueing. In *Proc. of IEEE INFOCOM '96*, pages 120–128, San Francisco, CA, March 1996. 14, 230
- [22] B. Prabhakar and N. McKeown. On the speedup required for combined input and output queued switching. *Automatica*, 35:1909–1920, December 1999. 19, 96
- [23] P. Krishna, N. Patel, A. Charny, and R. Simcoe. On the speedup required for work-conserving crossbar switches. *IEEE J.Sel. Areas in Communications*, 17(6):1057–1066, June 1999. 19, 116
- [24] Cisco Systems Inc. Data obtained courtesy Global Commodity Management (GCM) Group. 21, 184, 223
- [25] Cisco Systems Inc. Personal communication, Power Engineering Group, DCBU. 21, 313
- [26] <http://en.wikipedia.org/wiki/EDRAM>. 25, 127, 173, 212, 280, 303, 304, 305
- [27] <http://en.wikipedia.org/wiki/SerDes>. 27, 212
- [28] Cisco Systems Inc. Cisco Telepresence. http://www.cisco.com/en/US/netsol/ns669/networking_solutions_solution_segment_home.html. 27, 228, 254, 371

- [29] International Telecommunications Union. IPTV Focus Group. <http://www.itu.int/ITU-T/IPTV/>. 27, 228, 254, 371
- [30] <http://en.wikipedia.org/wiki/RAID>. 29, 313
- [31] Cisco Systems Inc. Personal communication, Central Engineering, Power Systems Group. 29, 223
- [32] Cisco Systems Inc. Skimmer, Serial Network Memory ASIC, DCBU. 29, 222, 304, 313
- [33] Cisco Systems Inc. Network Memory Group, DCBU. 30, 180, 184, 223, 280, 282, 300
- [34] T. Chaney, J. A. Fingerhut, M. Flucke, and J. Turner. Design of a gigabit ATM switching system. Technical Report WUCS-96-07, Computer Science Department, Washington University, February 1996. 42
- [35] S. Iyer, A. Awadallah, and N. McKeown. Analysis of a packet switch with memories running slower than the line rate. In *Proc. IEEE INFOCOM '00*, June 2000. 43, 148
- [36] S. Iyer and N. McKeown. Making parallel packet switches practical. In *Proc. IEEE INFOCOM '01*, volume 3, pages 1680–1687, 2001. 43, 114, 352
- [37] C. S. Chang, D. S. Lee, and Y. S. Jou. Load balanced Birkhoff-von Neumann switches, part I: one-stage buffering. In *IEEE HPSR Conference*, pages 556–564, Dallas, TX, May 2001. <http://www.ee.nthu.edu.tw/~cschang/PartI.ps>. 43, 44
- [38] I. Keslassy and N. McKeown. Maintaining packet order in two-stage switches. In *Proc. of the IEEE INFOCOM*, June 2002. 43, 44
- [39] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1949, 1992. 44, 113, 338
- [40] A. Prakash, S. Sharif, and A. Aziz. An $O(\log^2 n)$ algorithm for output queuing. In *Proc. IEEE INFOCOM '02*, pages 1623–1629, June 2002. 44, 59
- [41] B. Prabhakar and N. McKeown. On the speedup required for combined input and output queued switching. Technical Report STAN-CSL-TR-97-738, Stanford University, November 1997. 44, 112, 114, 116
- [42] R. B. Magill, C. Rohrs, and R. Stevenson. Output queued switch emulation by fabrics with limited memory. *IEEE Journal on Selected Areas in Communications*, 21(4):606–615, 2003. 44, 124, 130, 138

- [43] S. Chuang, S. Iyer, and N. McKeown. Practical algorithms for performance guarantees in buffered crossbars. In *Proc. IEEE INFOCOM '05*, pages 981–991, 2005. 43, 132, 135, 136
- [44] Image courtesy http://en.wikipedia.org/wiki/Pigeonhole_principle. 47
- [45] http://en.wikipedia.org/wiki/Hilbert's_paradox_of_the_Grand_Hotel. 47
- [46] N. Endo, T. Kozaki, T. Ohuchi, H. Kuwahara, and S. Gohara. Shared buffer memory switch for an ATM exchange. *IEEE Transactions on Communications*, 41(1):237–245, January 1993. 48
- [47] R. H. Hofmann and R. Muller. A multifunctional high-speed switch element for ATM applications. *IEEE Journal of Solid-State Circuits*, 27(7):1036–1040, July 1992. 48
- [48] H. Yamada, S. I. Yamada, H. Kai, and T. Takahashi. Multi-Purpose memory switch LSI for ATM-based systems. In *GLOBECOM*, pages 1602–1608, 1990. 48
- [49] M. Devault, J. Y. Cochenec, and M. Servel. The ‘PRELUDE’ ATD experiment: Assessments and future prospects. *IEEE Journal on Selected Areas in Communications*, 6(9):1528–1537, December 1988. 48
- [50] J. P. Coudreuse and M. Servel. PRELUDE: An asynchronous time-division switched network. In *ICC*, pages 769–772, June 1987. 48
- [51] M. A. Henrion, G. J. Eilenberger, G. H. Petit, and P. H. Parmentier. A multipath self-routing switch. *IEEE Communications Magazine*, pages 46–52, December 1993. 48
- [52] MMC Networks. ATMS2000: 5 Gb/s switch engine chipset, 1996. 48
- [53] S. Iyer and N. McKeown. Techniques for fast shared memory switches. Technical Report TR01-HPNG-081501, Computer Science Department, Stanford University, August 2001. 59
- [54] Y. Xu, B. Wu, W. Li, and B. Liu. A scalable scheduling algorithm to avoid conflicts in switch-memory-switch routers. In *Proc. ICCCN 2005*, pages 57–64, October 2005. 59, 83
- [55] A. Prakash. *Architectures and Algorithms for High Performance Switching*. Ph.D. Thesis Report, Univ. of Texas at Austin, August 2004. 59, 83
- [56] P. S. Sindhu, R. K. Anand, D. C. Ferguson, and B. O. Liencres. High speed switching device. United States Patent No. 5905725, May 1999. 66

- [57] D. König. Über Graphen und ihre Anwendung auf Determinantentheorie und Mengenlehre. *Math. Ann.*, 77:453–465, 1916. 71
- [58] D. König. Gráfok és mátrixok. *Matematikai és Fizikai Lapok*, 38:116–119, 1931. 71
- [59] S. Iyer, R. Zhang, and N. McKeown. Routers with a single stage of buffering. In *Proc. ACM SIGCOMM '02*, Pittsburg, PA, September 2002. 75
- [60] H. C. Chi and Y. Tamir. Decomposed arbiters for large crossbars with multiqueue input buffers. In *IEEE International Conference on Computer Design: VLSI in Computers and Processors*, pages 233–238, Cambridge, MA, 1991. 76
- [61] Technical Committee T11. Fiber Channel. <http://www.t11.org>. 90, 212, 254
- [62] Cisco Systems Inc. Cisco GSR 12000 Series Quad OC-12/STM-4 POS/SDH line card. http://www.cisco.com/en/US/products/hw/routers/ps167/products_data_sheet09186a00800920a7.html. 90, 187
- [63] D. Gale and L. S. Shapely. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962. 93, 105, 110, 121, 122, 126
- [64] M. Karol, M. Hluchyi, and S. Morgan. Input versus output queueing on a space-division switch. *IEEE Transactions on Communications*, 35(12):1347–1356, December 1987. 94, 111
- [65] A. Charny. *Providing QoS Guarantees in Input-buffered Crossbars with Speedup*. Ph.D. Thesis Report, MIT, September 1998. 94, 95, 96, 116, 117
- [66] S. Iyer and N. McKeown. Using constraint sets to achieve delay bounds in CIOQ switches. *IEEE Communication Letters*, 7(6):275–277, June 2003. 97
- [67] M. Akata, S. Karube, T. Sakamoto, T. Saito, S. Yoshida, and T. Maeda. A 250 Mb/s 32x32 CMOS crosspoint LSI for ATM switching systems. *IEEE J. Solid-State Circuits*, 25(6):1433–1439, December 1990. 97, 116
- [68] M. Karol, K. Eng, and H. Obara. Improving the performance of input queued ATM packet switches. In *Proc. of IEEE INFOCOM '92*, pages 110–115, 1992. 97, 116
- [69] H. Matsunaga and H. Uematsu. A 1.5 Gb/s 8x8 cross-connect switch using a time reservation algorithm. *IEEE J. Selected Area in Communications*, 9(8):1308–1317, October 1991. 97, 116
- [70] H. Obara, S. Okamoto, and Y. Hamazumi. Input and output queueing ATM switch architecture with spatial and temporal slot reservation control. *IEEE Electronics Letters*, pages 22–24, January 1992. 97, 116

- [71] E. Leonardi, M. Mellia, M. Marsan, and F. Neri. Stability of maximal size matching scheduling in input queued cell switches. In *Proc. ICC 2000*, pages 1758–1763, 2000. 102, 114
- [72] http://en.wikipedia.org/wiki/Stable_marriage_problem. 105
- [73] H. Mairson. The stable marriage problem. <http://www1.cs.columbia.edu/~evs/intro/stable/writeup.html>, 1992. 105
- [74] Personal communication with Da Chuang. 112
- [75] I. Iliadis and W. E. Denzel. Performance of packet switches with input and output queueing. In *Proceedings of ICC*, pages 747–753, 1990. 111
- [76] A. L. Gupta and N. D. Georganas. Analysis of a packet switch with input and output buffers and speed constraints. In *Proc. of INFOCOM '91*, pages 694–700, 1991. 111
- [77] Y. Oie, M. Murata, K. Kubota, and H. Miyahara. Effect of speedup in nonblocking packet switch. In *Proceedings of ICC*, pages 410–414, 1989. 111
- [78] J. S. C. Chen and T. E. Stern. Throughput analysis, optimal buffer allocation, and traffic imbalance study of a generic nonblocking packet switch. *IEEE Journal on Selected Areas Communication*, 9(3):439–449, April 2001. 111
- [79] T. Anderson, S. Owicki, J. Saxie, and C. Thacker. High speed switch scheduling for local area networks. *ACM Transactions of Computing Systems*, 11(4):319–352, 1993. 113
- [80] C. S. Chang, W. J. Chen, and H. Y. Huang. On service guarantees for input buffered crossbar switches: A capacity decomposition approach by Birkhoff and von Neumann. In *Proc. of IEEE INFOCOM '00*, Tel Aviv, Israel, 2000. 113
- [81] E. Altman, Z. Liu, and R. Righter. Scheduling of an input queued switch to achieve maximal throughput. *Probability in the Engineering and Informational Sciences*, 14:327–334, 2000. 113
- [82] G. Birkhoff. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucum an Rev. Ser. A*, 5:147–151, 1946. 113
- [83] J. von Neumann. *A certain zero-sum two-person game equivalent to the optimal assignment problem: Contributions to the Theory of Games*, volume 2. Princeton University Press, 1953. 113
- [84] T. Inukai. An efficient SS/TDMA time slot assignment algorithm. *IEEE Transactions on Communications*, 27:1449–1455, 1979. 113
- [85] J. E. Hopcroft and R. M. Karp. An $n^{2.5}$ algorithm for maximum matching in bipartite graphs. *Soc. Ind. Appl. Math. J.*, 2:225–231, 1973. 113

- [86] R. E. Tarjan. *Data Structures and Network Algorithms*. Bell laboratories, 1983. 113
- [87] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand. Achieving 100% throughput in an input queued switch. *IEEE Transactions on Communications*, 47(8):1260–1267, August 1999. 113
- [88] A. Mekkittikul and N. McKeown. A practical scheduling algorithm to achieve 100% throughput in input queued switches. In *Proc. of IEEE INFOCOM '98*, volume 2, pages 792–799, San Francisco, CA, April 1998. 113, 114
- [89] T. Weller and B. Hajek. Scheduling nonuniform traffic in a packet switching system with small propagation delay. *IEEE/ACM Transactions on Networking*, 5(6):813–823, 1997. 113
- [90] S. Iyer and N. McKeown. Maximum size matchings and input queued switches. In *Proceedings of the 40th Annual Allerton Conference on Communications*, October 2002. 113
- [91] C. S. Chang, D. S. Lee, and Y. Jou. Load balanced Birkhoff-von Neumann switches, part I: one-stage buffering. *Computer Communications - special issue on Current Issues in Terabit Switching*, 2001. 114, 171
- [92] C. S. Chang, D. S. Lee, and C. Lien. Load balanced Birkhoff-von Neumann switches, part II: multi-stage buffering. *Computer Communications - special issue on Current Issues in Terabit Switching*, 2001. 114, 171
- [93] I. Keslassy and N. McKeown. Analysis of scheduling algorithms that provide 100% throughput in input queued switches. In *Proceedings of the 39th Annual Allerton Conference on Communications*, 2001. 114
- [94] L. Tassiulas. Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In *Proc. of IEEE INFOCOM '98*, pages 533–539, New York, NY, 1998. 114
- [95] P. Giaccone, B. Prabhakar, and D. Shah. Towards simple, high performance schedulers for high-aggregate bandwidth switches. In *Proc. of IEEE INFOCOM '02*, New York, NY, 2002. 114
- [96] I. Stoica and H. Zhang. Exact emulation of an output queueing switch by a combined input output queueing switch. In *Proc. of IEEE IWQoS '98*, 1998. 116
- [97] <http://www.cs.berkeley.edu/~istoica/IWQoS98-fix.html>. 116
- [98] A. Firoozshahian, A. Manshadi, A. Goel, and B. Prabhakar. Efficient, fully local algorithms for CIOQ switches. In *Proc. of IEEE INFOCOM '07*, pages 2491–2495, Anchorage, AK, May 2007. 116

- [99] Rojas-Cessa, E. Oki, Z. Jing, and H. J. Chao. CIXB-1: Combined input-one-cell-crosspoint buffered switch. In *IEEE Workshop on High Performance Switching and Routing*, 2001. 122
- [100] Rojas-Cessa, E. Oki, Z. Jing, and H. J. Chao. CIXOB-k: Combined input-crosspoint-output buffered packet switch. In *IEEE Globecom*, pages 2654–2660, 2001. 122
- [101] L. Mhamdi and M. Hamdi. MCBF: A high-performance scheduling algorithm for buffered crossbar switches. *IEEE Communications Letters*, 7(9):451–453, September 2003. 122
- [102] K. Yoshigoe and K. J. Christensen. Design and evaluation of a parallel-pollled virtual output queued switch. In *Proceedings of the IEEE International Conference on Communications*, pages 112–116, 2001. 122
- [103] D. Stephens and H. Zhang. Implementing distributed packet fair queueing in a scalable switch architecture. In *Proc. of INFOCOM '98*, pages 282–290, 1998. 122
- [104] N. Chrysos and M. Katevenis. Weighted fairness in buffered crossbar scheduling. In *IEEE Workshop on High Performance Switching and Routing*, 2003. 122
- [105] T. Javidi, R. B. Magill, and T. Hrabik. A high throughput scheduling algorithm for a buffered crossbar switch fabric. In *Proceedings of IEEE International Conference on Communications*, pages 1586–1591, 2001. 124
- [106] R. B. Magill, C. Rohrs, and R. Stevenson. Output queued switch emulation by fabrics with limited memory. *IEEE Journal on Selected Areas in Communications*, 21(4):606–615, May 2003. 124, 130, 138
- [107] N. McKeown, C. Calamvokis, and S. Chuang. A 2.5Tb/s LCS switch core. In *Hot Chips '01*, August 2001. 126
- [108] Shang-Tse Chuang. *Providing Performance Guarantees in Crossbar-based routers*. Ph.D. Thesis Report, Stanford University, January 2005. 127
- [109] Cisco Systems Inc. Personal communication, Catalyst 4K Group, GSBU. 137, 257, 259
- [110] C. Clos. A study of non-blocking switching networks. *The Bell System Technical Journal*, 32:406–424, 1953. 148, 174, 312
- [111] Nevis Networks Inc. <http://www.nevisnetworks.com/>. 150, 175
- [112] Juniper E Series Router. http://www.juniper.net/products_and_services/m_series_routing_portfolio/. 150, 175

-
- [113] S. Iyer. Analysis of a packet switch with memories running slower than the line rate. M.S. Thesis Report, Stanford University, May 2000. 157
- [114] V. E. Benes. *Mathematical theory of connecting network and telephone traffic*. Academic Press, New York, 1965. 158
- [115] J. Hui. *Switching and traffic theory for integrated broadband network*. Kluwer Academic Publications, Boston, 1990. 158
- [116] A. Jajszczyk. Nonblocking, repackable, and rearrangeable Clos networks: fifty years of the theory evolution. *IEEE Communications Magazine*, 41:28–33, 2003. 159
- [117] H. Adishesu, G. Parulkar, and George Varghese. A reliable and scalable striping protocol. In *Proc. ACM SIGCOMM '96*, 1996. 173
- [118] P. Fredette. The past, present, and future of inverse multiplexing. *IEEE Communications*, pages 42–46, April 1994. 174
- [119] J. Duncanson. Inverse multiplexing. *IEEE Communications*, pages 34–41, April 1994. 174
- [120] J. Frimmel. Inverse multiplexing: Tailor made for ATM. *Telephony*, pages 28–34, July 1996. 174
- [121] J. Turner. Design of a broadcast packet switching network. *IEEE Trans. on Communications*, pages 734–743, June 1988. 174
- [122] H. Kim and A. Leon-Garcia. A self-routing multistage switching network for broadband ISDN. *IEEE J. Sel. Areas in Communications*, pages 459–466, April 1990. 174
- [123] I. Widjaja and A. Leon-Garcia. The helical switch: A multipath ATM switch which preserves cell sequence. *IEEE Trans. on Communications*, 42(8):2618–2629, August 1994. 174
- [124] F. Chiussi, D. Khotimsky, and S. Krishnan. Generalized inverse multiplexing of switched ATM connections. In *Proc. IEEE Globecom '98 Conference. The Bridge to Global Integration*, Sydney, Australia, 1998. 174
- [125] F. Chiussi, D. Khotimsky, and S. Krishnan. Advanced frame recovery in switched connection inverse multiplexing for ATM. In *Proc. ICATM '99 Conference*, Colmar, France, 1999. 174
- [126] Bandwidth ON Demand INTERoperability Group. Interoperability requirements for nx56/64 kbit/s calls, 1995. ISO/IEC 13871. 174

- [127] D. Hay and H. Attiya. The inherent queuing delay of parallel packet switches. *IEEE Transactions on Parallel and Distributed Systems*, 17(8), August 2006. 174
- [128] H. Attiya and D. Hay. Randomization does not reduce the average delay in parallel packet switches. *SIAM Journal on Computing*, 37(5):1613–1636, January 2008. 174, 352
- [129] S. Mneimneh, V. Sharma, and Kai-Yeung Siu. Switching using parallel input–output queued switches with no speedup. *IEEE/ACM Transactions on Networking*, 10(5):653–665, 2002. 174
- [130] S. Iyer. Personal consultation with Nevis Networks Inc., 2003. 175
- [131] Cisco Systems Inc. Personal communication, Low Latency Ethernet Product Group. 175, 212
- [132] FocalPoint in Large-Scale Clos Switches. www.fulcrummicro.com/product_library/applications/clos.pdf. 175
- [133] <http://www.cisco.com/>. 180
- [134] M. Sanchez, E. Biersack, and W. Dabbous. Survey and taxonomy of IP address lookup algorithms. *IEEE Network*, 15(2):8–23, 2001. 186, 229, 267, 289
- [135] Juniper E Series Router. <http://juniper.net/products/eseries>. 187
- [136] Force 10 E-Series Switch. <http://www.forcel10networks.com/products/pdf/prodoverview.pdf>. 187
- [137] Foundry BigIron RX-series Ethernet switches. http://www.foundrynet.com/about/newsevents/releases/\pr5_03_05b.html. 187
- [138] P. Chen and David A. Patterson. Maximizing performance in a striped disk array. In *ISCA*, pages 322–331, 1990. 194, 219
- [139] B. R. Rau, M. S. Schlansker, and D. W. L. Yen. The Cydra 5 stride-insensitive memory system. In *In Proc. Int Conf. on Parallel Processing*, pages 242–246, 1989. 194, 219
- [140] S. Kumar, P. Crowley, and J. Turner. Design of randomized multichannel packet storage for high performance routers. In *Proceedings of Hot Interconnects*, August 2005. 194, 219
- [141] Distributed Denial-of-Service (DDoS) Attack. http://en.wikipedia.org/wiki/Denial-of-service_attack#Incidents. 206
- [142] K. Houle and G. Weaver. Trends in Denial of Service Attack Technology. http://www.cert.org/archive/pdf/DoS_trends.pdf, 2001. 206

- [143] M. Handley and E. Rescorla. RFC 4732: Internet denial-of-service considerations. <http://tools.ietf.org/html/rfc4732>, 2006. 206
- [144] <http://www.spirentcom.com>. 206
- [145] <http://www.agilent.com>. 206
- [146] R. Bhagwan and B. Lin. Fast and scalable priority queue architecture for high-speed network switches. In *Proc. of IEEE INFOCOM '00*, 2000. 217
- [147] Y. Joo and N. McKeown. Doubling memory bandwidth for network buffers. In *Proc. IEEE INFOCOM '98*, pages 808–815, San Francisco, CA, 1998. 217, 219
- [148] J. Corbal, R. Espasa, and M. Valero. Command vector memory systems: High performance at low cost. In *Proceedings of the 1998 International Conference on Parallel Architectures and Compilation Techniques*, pages 68–77, October 1998. 219
- [149] B. K. Mathew, S. A. McKee, J. B. Carter, and A. Davis. Design of a parallel vector access unit for SDRAM memory systems. In *Proceedings of the Sixth International Symposium on High-Performance Computer Architecture*, January 2000. 219
- [150] S. A. McKee and W. A. Wulf. Access ordering and memory-conscious cache utilization. In *Proceedings of the First International Symposium on High-Performance Computer Architecture*, pages 253–262, January 1995. 219
- [151] S. Rixner, W. J. Dally, U. J. Kapasi, P. Mattson, and J. D. Owens. Memory access scheduling. In *Proceedings of the 27th Annual International Symposium on Computer Architecture*, pages 128–138, June 2000. 219
- [152] T. Alexander and G. Kedem. Distributed prefetch-buffer/cache design for high performance memory systems. In *Proceedings of the 2nd International Symposium on High-Performance Computer Architecture*, pages 254–263, February 1996. 219
- [153] W. Lin, S. Reinhardt, and D. Burger. Reducing DRAM latencies with an integrated memory hierarchy design. In *Proc. 7th Int symposium on High-Performance Computer Architecture*, January 2001. 219
- [154] S. I. Hong, S. A. McKee, M. H. Salinas, R. H. Klenke, J. H. Aylor, and W. A. Wulf. Access order and effective bandwidth for streams on a direct rambus memory. In *Proceedings of the Fifth International Symposium on High-Performance Computer Architecture*, pages 80–89, January 1999. 219
- [155] D. Patterson and J. Hennessy. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann Publishers, San Francisco, CA, 2nd edition, 1996. 219, 313

- [156] L. Carter and W. Wegman. Universal hash functions. *J. of Computer and System Sciences*, 18:143–154, 1979. 219
- [157] R. Impagliazzo and D. Zuckerman. How to recycle random bits. In *Proc. of the Thirtieth Annual Symposium on the Foundations of IEEE*, 1989. 219
- [158] G. Shrimali and N. McKeown. Building packet buffers with interleaved memories. *Proceedings of IEEE Workshop on High Performance Switching and Routing*, pages 1–5, May 2005. 219
- [159] A. Birman, H. R. Gail, S. L. Hantler, and Z. Rosberg. An optimal service policy for buffer systems. *Journal of the Association for Computing Machinery*, 42(3):641–657, May 1995. 220
- [160] H. Gail, G. Grover, R. Guerin, S. Hantler, Z. Rosberg, and M. Sidi. Buffer size requirements under longest queue first. In *Proceedings IFIP '92*, volume C-5, pages 413–424, 1992. 220
- [161] G. Sasaki. Input buffer requirements for round robin polling systems. In *Proceedings of 27th Annual Conference on Communication Control and Computing*, pages 397–406, 1989. 220
- [162] I. Cidon, I. Gopal, G. Grover, and M. Sidi. Real-time packet switching: A performance analysis. *IEEE Journal on Selected Areas in Communications*, SAC-6:1576–1586, December 1988. 220
- [163] A. Birman, P. C. Chang, J. Chen, and R. Guerin. Buffer sizing in an ISDN frame relay switch. Technical Report RC14286, IBM Research Report, Aug 1989. 220
- [164] S. Iyer, R. R. Kompella, and N. McKeown. Analysis of a memory architecture for fast packet buffers. In *Proc. IEEE HPSR*, Dallas, TX, 2001. 220
- [165] S. Iyer, R. R. Kompella, and N. McKeown. Techniques for fast packet buffers. In *Proceedings of GBN 2001*, Anchorage, AK, April 2001. 220
- [166] S. Iyer, R. R. Kompella, and N. McKeown. Designing packet buffers for router line cards. *IEEE Transactions on Networking*, 16(3):705–717, June 2008. 220, 361
- [167] S. Iyer, R. R. Kompella, and N. McKeown. Designing packet buffers for router line cards. Technical Report TR02-HPNG-031001, Computer Science Department, Stanford University, March 2002. 220
- [168] A. Bar-Noy, A. Freund, S. Landa, and J. Naor. Competitive on-line switching policies. *Algorithmica*, 36:225–247, 2003. 220, 221

- [169] R. Fleischer and H. Koga. Balanced scheduling toward loss-free packet queuing and delay fairness. *Algorithmica*, 38:363–376, 2004. 220, 221
- [170] P. Damaschek and Z. Zhou. On queuing lengths in on-line switching. *Theoretical Computer Science*, 339:333–343, 2005. 220, 221
- [171] J. Garcia, J. Corbal, L. Cerda, and M. Valero. Design and implementation of high-performance memory systems for future packet buffers. *Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 372–384, 2003. 221
- [172] G. Shrimali, I. Keslassy, and N. McKeown. Designing packet buffers with statistical guarantees. In *HOTI '04: Proceedings of the High Performance Interconnects*, pages 54–60. IEEE Computer Society, 2004. 221
- [173] M. Arpaci and J. Copeland. Buffer management for shared-memory ATM switches. *IEEE Comm. Surveys and Tutorials*, 3(1):2–10, 2000. 221, 267
- [174] M. L. Irland. Buffer management in a packet switch. *IEEE Trans. Communication*, COM-26(3):328–337, March 1978. 221, 267
- [175] Cisco Systems Inc. http://www.cisco.com/en/US/products/hw/switches/ps708/\products_data_sheet0900aecd801459a7.htm. 222
- [176] Cisco Systems Inc. Cisco nexus 5000 series switch. <http://www.cisco.com/en/US/products/ps9670/index.html>. 222, 259
- [177] Cisco Systems Inc. Cisco nexus 7000 series switch. <http://www.cisco.com/en/US/products/ps4902/index.html>. 222, 259
- [178] Cisco Systems Inc. Data Center Ethernet. http://www.cisco.com/en/US/netsol/ns783/networking_solutions_package.html. 222, 259
- [179] Cisco Systems Inc. Personal communication, Mid Range Routers Business Unit. 222
- [180] Cisco Systems Inc. Personal communication, Core Router Project. 222
- [181] Cisco Systems Inc. Personal communication, DC3 VOQ Buffering ASIC Group. 222, 313
- [182] Cisco Systems Inc. Personal communication, DC3 Storage ASIC Group. 222, 313
- [183] Pablo Molinero. *Circuit Switching in the Internet*. Ph.D. Thesis Report, Stanford University, June 2003. 228
- [184] P. Gupta and N. McKeown. Algorithms for packet classification. *IEEE Network*, 15(2):24–32, 2001. 229, 267, 289

- [185] S. Fide and S. Jenks. A survey of string matching approaches in hardware. http://spds.ece.uci.edu/~sfide/String_Matching.pdf. 229
- [186] R. Geurin and V. Peris. Quality-of-Service in packet networks: Basic mechanisms and directions. *Computer Networks*, 31(3):169–189, February 1999. 230, 267
- [187] S. Iyer and Da Chuang. Designing packet schedulers for router line cards. In preparation for IEEE INFOCOM '09. 234
- [188] <http://en.wikipedia.org/wiki/Freenet>. 254
- [189] J. Oikarinen and D. Reed. RFC 1459: Internet relay chat protocol. <http://tools.ietf.org/html/rfc1459>, 1993. 254
- [190] J. Rosenberg and H. Schulzrinne. RFC 2871: A framework for telephony routing over IP. <http://tools.ietf.org/html/rfc2871>, 2000. 254
- [191] <http://en.wikipedia.org/wiki/Multiplayer>. 254
- [192] [http://en.wikipedia.org/wiki/Doom_\(video_game\)](http://en.wikipedia.org/wiki/Doom_(video_game)). 254
- [193] Cisco Systems Inc. Personal communication, DC3 and Catalyst 6K Group, ISBU, DCBU. 257
- [194] D. Newman. RFC 2647: Benchmarking terminology for firewall performance. <http://tools.ietf.org/html/rfc2647>, 1999. 265, 287
- [195] S. Waldbusser. RFC 2819: Remote network monitoring management information base. <http://tools.ietf.org/html/rfc2819>, 2000. 265
- [196] Cisco Systems Inc. <http://www.cisco.com/warp/public/732/Tech/netflow>. 265, 268
- [197] Juniper Networks. www.juniper.net/techcenter/appnote/350003.html. 265, 268
- [198] Huawei Inc. Technical Whitepaper for Netstream. <http://www.huawei.com/products/datacomm/pdf/view.do?f=65>. 265
- [199] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja, and X. Xiao. RFC 3272: Overview and principles of Internet traffic engineering. <http://tools.ietf.org/html/rfc3272>, 2002. 265
- [200] D. Shah, S. Iyer, B. Prabhakar, and N. McKeown. Analysis of a statistics counter architecture. *IEEE Hot Interconnects 9*, August 2001. 267
- [201] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proc. ACM SIGCOMM '01*, pages 75–80, 2001. 267

- [202] Cisco Systems Inc. Personal communication, Ethernet Address and Route Lookup Group, DC3. 280, 281, 282, 313
- [203] S. Ramabhadran and G. Varghese. Efficient implementation of a statistics counter architecture. In *Proc. ACM SIGMETRICS*, pages 261–271, 2003. 281, 313
- [204] Q. G. Zhao, J. J. Xu, and Z. Liu. Design of a novel statistics counter architecture with optimal space and time efficiency. In *Proc. SIGMetrics/Performance*, pages 261–271, 2006. 281, 300, 313
- [205] Y. Lu, A. Montanari, B. Prabkakar, S. Dharmapurikar, and A. Kabbani. Counter Braids: A novel counter architecture for per flow measurement. In *Proc. ACM Sigmetrics*, 2008. 281, 313
- [206] P. Gupta and D. Shah. Personal communication. 281
- [207] S. Iyer and N. McKeown. High speed packet-buffering system. Patent Application No. 20060031565, 2006. 281, 411
- [208] K. Egevang and P. Francis. RFC 1631: The IP network address translator (NAT). <http://tools.ietf.org/html/rfc1631>, 1994. 287, 289
- [209] http://www.cisco.com/en/US/docs/ios/12_4t/qos/configuration/guide/qsnbar1.html. 287
- [210] S. Iyer. Maintaining state for router line cards. In preparation for IEEE Communication Letters. 291
- [211] Y. Joo and N. McKeown. Doubling memory bandwidths for network buffers. In *Proc. IEEE INFOCOM '98*, pages 808–815, 1998. 295
- [212] http://en.wikipedia.org/wiki/Double_buffering. 295
- [213] Cisco Systems Inc. Personal communication, ISBU and DCBU ASIC Engineering Groups. 300
- [214] Cisco Systems Inc., Network Memory Group. Skimmer Serial Network Memory Protocol. 313
- [215] <http://en.wikipedia.org/wiki/ACID>. 313
- [216] H.J. Kushner. *Stochastic Stability and Control*. Academic Press, 1967. 338
- [217] G. Fayolle. On random walks arising in queuing systems: ergodicity and transience via quadratic forms as Lyapunov functions - Part I. *Queueing Systems*, 5:167–184, 1989. 338

- [218] E. Leonardi, M. Mellia, F. Neri, and M.A. Marsan. On the stability of input queued switches with speed-up. *IEEE Transactions on Networking*, 9(1):104–118, 2001. 338
- [219] R. L. Cruz. A Calculus for Network Delay: Part I. *IEEE Transactions on Information Theory*, 37:114–131, January 1991. 350
- [220] Personal communication with David Hay. 352
- [221] B. Cain, S. Deering, I. Kouvelas, B. Fenner, and A. Thyagarajan. RFC 3376: Internet group management protocol, v3. <http://tools.ietf.org/html/rfc3376>, 2002. 371
- [222] Z. Albanna, K. Almeroth, D. Meyer, and M. Schipper. RFC 3171: IANA guidelines for IPv4 multicast address assignments. <http://tools.ietf.org/html/rfc3171>, 2001. 371
- [223] R. Hinden and S. Deering. RFC 2375: IPv6 multicast address assignments. <http://tools.ietf.org/html/rfc2375>, 1998. 371
- [224] D. Thaler. RFC 2715: Interoperability rules for multicast routing protocols. <http://tools.ietf.org/html/rfc2715>, 1999. 371
- [225] S. Deering. RFC 1112: Host extensions for IP multicasting. <http://tools.ietf.org/html/rfc1112>, 1998. 371
- [226] Cisco Systems Inc. Personal communication, DC3 100G MAC ASIC Group. 371
- [227] Gua, Ming-Huang, and R.S. Chang. Multicast ATM switches: survey and performance evaluation. *Computer Communication Review*, 28(2), 1998. 373
- [228] J. Turner and N. Yamanaka. Architectural choices in large scale ATM switches. *IEICE Trans. Communications*, E81-B(2):120–137, February 1998. 373
- [229] N. F. Mir. A survey of data multicast techniques, architectures, and algorithms. *IEEE Communications Magazine*, 39:164–170, 2001. 373
- [230] S. Iyer and N. McKeown. On the speedup required for a multicast parallel packet switch. *IEEE Comm. Letters*, 2001. 376, 377
- [231] Cisco Systems Inc. Personal communication, DCBU. 380, 381
- [232] S. Iyer and N. McKeown. High speed memory control and I/O processor system. Patent Application No. 20050240745, 2005. 411