Bibliography

[1] H. Ahmadi and W.E. Denzel. “A survey of modern high-performance switching techniques,” *IEEE Journal of Selected Areas in Communications*, vol. 7, no. 7, pages 1091-103, September 1989.

[2] G. Apostolopoulos, D. Williams, S. Kamat, R. Guerin, A. Orda and T. Przygienda. “QoS routing mechanisms and OSPF extensions,” *RFC 2676*, http://www.ietf.org/rfc/rfc2676.txt, August 1999.

[3] F. Baker, editor. “Requirements for IP version 4 routers,” *RFC 1812*, http://www.ietf.org/rfc/rfc2676.txt, June 1995.

[4] M. de Berg, M. van Kreveld and M. Overmars. *Computational geometry: algorithms and applications*, Springer-Verlag, 2nd rev. ed. 2000.

[5] D.P. Bertsekas. *Nonlinear programming*, Athena Scientific, 1995.

[6] R. Braden, L. Zhang, S. Berson, S. Herzog and S. Jamin. “Resource reSerVation protocol (RSVP) — version 1 functional specification,” *RFC 2205*, http://www.ietf.org/rfc/rfc2205.txt, September 1997.

[7] M.M. Buddhikot, S. Suri and M. Waldvogel. “Space decomposition techniques for fast layer-4 switching,” *Proceedings of Conference on Protocols for High Speed Networks*, pages 25-41, August 1999.

[8] B. Chazelle. “Lower bounds for orthogonal range searching, i: the reporting case,” *Journal of the ACM*, vol. 37, no. 2, pages 200-12, April 1990.

[9] B. Chazelle. “Lower bounds for orthogonal range searching, ii: the arithmetic model,” *Journal of the ACM*, vol. 37, no. 3, pages 439-63, July 1990.

[10] G. Cheung and S. McCanne. “Optimal routing table design for IP address lookups under memory constraints,” *Proceedings of IEEE Infocom*, vol. 3, pages 1437-44, March 1999.

[11] T. Chiueh and P. Pradhan. “High performance IP routing table lookup using CPU caching,” *Proceedings of IEEE Infocom*, vol. 3, pages 1421-8, March 1999.

[12] KC Claffy. “Internet measurement: state of DeUnion,” presentation for Internet Measurement, Oct 99. Available at http://www.caida.org/outreach/presentations/Soa9911/mgp00027.html.

[13] T.H. Cormen, C.E. Leiserson and R.L. Rivest. *Introduction to algorithms*, MIT Press, 1990.

[14] T.M. Cover and J.A. Thomas. *Elements of information theory*, Wiley Series in Telecommunications, 1995.

[15] D. Decasper, Z. Dittia, G. Parulkar and B. Plattner. “Router plugins: a software architecture for next generation routers,” *Proceedings of ACM Sigcomm*, pages 229-40, October 1998.

[16] S. Deering and R. Hinden. “Internet Protocol, version 6 (IPv6) specification,” *RFC 1883*, http://www.ietf.org/rfc/rfc1883.txt, December 1995.

[17] M. Degermark, A. Brodnik, S. Carlsson and S. Pink. “Small forwarding tables for fast routing lookups,” *Proceedings of ACM Sigcomm*, pages 3-14, October 1997.

[18] A. Demers, S. Keshav and S. Shenker. “Analysis and simulation of a fair queueing algorithm,” *Internetworking: Research and Experience*, vol. 1, no. 1, pages 3-26, January 1990.

[19] W. Doeringer, G. Karjoth and M. Nassehi. “Routing on longest-matching prefixes,” *IEEE/ACM Transactions on Networking*, vol. 4, no. 1, pages 86-97, February 1996.

[20] F. DuPont. “Multihomed routing domain issues for IPv6 aggregatable scheme,” *Internet draft*, <draft-ietf-ngtrans-6bone-multi-01.txt>, June 1999.

[21] R.J. Edell, N. McKeown and P.P. Varaiya. “Billing users and pricing for TCP,” *IEEE Journal of Selected Areas in Communications, Special Issue on Advances in the Fundamentals of Networking*, vol. 13, no. 7, pages 1162-75, September 1995.

[22] D. Estrin and D. J. Mitzel. “An assessment of state and lookup overhead in routers,” *Proceedings of IEEE Infocom*, vol. 3, pages 2332-42, May 1992.

[23] A. Feldmann and S. Muthukrishnan. “Tradeoffs for packet classification,” *Proceedings of IEEE Infocom*, vol. 3, pages 1193-202, March 2000.

[24] D.C. Feldmeier. “Improving gateway performance with a routing-table cache,” *Proceedings of IEEE Infocom*, pages 298-307, March 1988.

[25] S. Floyd and V. Jacobson. “Random early detection gateways for congestion avoidance,” *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pages 397-413, August 1993.

[26] V. Fuller, T. Li, J. Yu and K. Varadhan. “Classless inter-domain routing (CIDR): an address assignment and aggregation strategy,” *RFC 1519*, http://www.ietf.org/rfc/rfc1519.txt, September 1993.

[27] M.R. Garey. “Optimal binary search trees with restricted maximal depth,” *SIAM Journal on Computing*, vol. 3, no. 2, pages 101-10, March 1974.

[28] A.M. Garsia and M.L. Wachs, “A new algorithm for minimum cost binary trees,” *SIAM Journal on Computing*, vol. 6, no. 4, pages 622-42, December 1977.

[29] E. Gerich. “Guidelines for management of IP address space,” *RFC 1466*, http://www.ietf.org/rfc/rfc1466.txt, May 1993.

[30] E.N. Gilbert. “Codes based on inaccurate source probabilities,” *IEEE Transactions on Information Theory*, vol. 17, no. 3, pages 304-14, May 1971.

[31] P. Gupta, S. Lin and N. McKeown. “Routing lookups in hardware at memory access speeds,” *Proceedings of IEEE Infocom*, vol. 3, pages 1240-7, April 1998.

[32] P. Gupta and N. McKeown, “Packet classification on multiple fields,” *Proceedings ACM Sigcomm*, pages 147-60, September 1999.

[33] C. Hedrick. “Routing information protocol,” *RFC 1058*, http://www.ietf.org/rfc/rfc1058.txt, June 1988.

[34] J. Hennessey and D. Patterson. *Computer architecture: a quantitative approach*, Morgan Kaufmann Publishers, 2nd edition, 1996.

[35] Y. Horibe. “An improved bound for weight-balanced tree,” *Information and Control*, vol. 34, no. 2, pages 148-51, June 1977.

[36] T.C. Hu. *Combinatorial Algorithms*, Addison-Wesley, 1982.

[37] T.C. Hu and A.C. Tucker. “Optimal computer search trees and variable length alphabetic codes,” *SIAM Journal on Applied Mathematics*, vol. 21, no. 4, pages 514-32, December 1971.

[38] N.F. Huang, S.M. Zhao, Jen-Yi Pan and Chi-An Su. “A Fast IP routing lookup scheme for gigabit switching routers,” *Proceedings of IEEE Infocom*, vol. 3, pages 1429-36, March 1999.

[39] D.A. Huffman. “A method for the construction of minimum redundancy codes,” *Proceedings Inst. Radio Engineers*, vol. 40, no. 10, pages 1098-101, September 1952.

[40] L. Hyafil and R.L. Rivest. “Constructing optimal binary decision trees is NP-complete,” *Information Processing Letters*, vol. 5, no. 1, pages 15-7, May 1976.

[41] A. Itai. “Optimal alphabetic trees,” *SIAM Journal on Computing*, vol. 5, no. 1, pages 9-18, March 1976.

[42] S. Jamin, P.B. Sanzig, S.J. Shenker and L. Zhang, “A measurement-based admission control algorithm for integrated service packet networks,” *IEEE/ACM Transactions on Networking*, vol. 5, no. 1, pages 56-70, February 1997.

[43] T. Kijkanjanarat and H.J. Chao. “Fast IP routing lookups for high performance routers,” *Computer Communications*, vol. 22, no. 15-16, pages 1415-22, September 1999.

[44] S. Kirpatrick, M. Stahl and M. Recker. “Internet Numbers,” *RFC 1166*, http://www.ietf.org/rfc/rfc1166.txt, July 1990.

[45] D. Knox and S. Panchanathan. “Parallel searching techniques for routing table lookup,” *Proceedings of IEEE Infocom*, vol. 3, pages 1400-5, March 1993.

[46] D.E. Knuth. *The art of computer programming, vol. 3: sorting and searching*, Addison-Wesley, 3rd edition, 1998.

[47] C. Labovitz, G. Malan and F. Jahanian. “Internet routing instability,” *Proceedings of ACM Sigcomm*, pages 115-26, September 1997.

[48] T.V. Lakshman and D. Stiliadis. “High-speed policy-based packet forwarding using efficient multi-dimensional range matching”, *Proceedings of ACM Sigcomm*, pages 191-202, September 1998.

[49] B. Lampson, V. Srinivasan and G. Varghese. “IP lookups using multiway and multicolumn search,” *Proceedings of IEEE Infocom*, vol. 3, pages 1248-56, April 1998.

[50] L.L. Larmore and T.M. Przytycka. “A fast algorithm for optimum height-limited alphabetic binary trees,” *SIAM Journal on Computing*, vol. 23, no. 6, pages 1283-312, December 1994.

[51] G. Malkin. “RIP version 2. Carrying additional information,” *RFC 1723*, http://www.ietf.org/rfc/rfc1723.txt, November 1994.

[52] A.J. McAuley and P. Francis. “Fast routing table lookup using CAMs,” *Proceedings of IEEE Infocom*, vol. 3, pages 1382-91, April 1993.

[53] S. McCanne and V. Jacobson. “A BSD packet filter: a new architecture for user-level packet capture,” *Proceedings of Usenix Winter Conference*, pages 259-69, January 1993.

[54] E.J. McCluskey. *Logic design principles: with emphasis on testable semicustom circuits*, Prentice-Hall, Englewood Cliffs, New Jersey, 1986.

[55] N. McKeown. “Scheduling algorithms for input-queued cell switches,” Ph.D. thesis, University of California at Berkeley, 1995.

[56] N. McKeown, M. Izzard, A. Mekkittikul, B. Ellersick and M. Horowitz. “The Tiny Tera: a packet switch core,” *Proceedings* of *Hot Interconnects V,* Stanford University, pages 161-73, August 1996.

[57] A. Mekkittikul. “Scheduling non-uniform traffic in high speed packet switches and routers,” PhD thesis, Stanford University, 1998.

[58] A. Mekkittikul, N. McKeown and M. Izzard. “A small high-bandwidth ATM switch.” *Proceedings of the SPIE*, vol. 2917, pages 387-97, November 1996.

[59] R.L. Mildiu and E.S. Laber. “Warm-up algorithm: A lagrangean construction of length restricted huffman codes,” *Monografias em Ciencia da Computacao*, no. 15, January 1996.

[60] R.L. Mildiu and E.S. Laber. “Improved bounds on the inefficiency of length-restricted prefix codes,” *unpublished manuscript*.

[61] R.L. Mildiu, A.A. Pessoa and E.S. Laber. “Efficient implementation of the Warm-up algorithm for the construction of length-restricted prefix codes,” *Proceedings of the ALENEX*, Baltimore, Maryland and in vol. 1619, *Lecture Notes in Computer Science,* Springer-Verlag, January 1999.

[62] A. Moestedt and P. Sjodin. “IP address lookup in hardware for high-speed routing,” *Proceedings Hot Interconnects VI,* August 1998.

[63] S. Morgan and M. Delaney. “The Internet and the local telephone network: conflicts and opportunities,” *XVI International Switching Symposium*, pages 561-9, September 1997.

[64] D.R. Morrison. “PATRICIA — practical algorithm to retrieve information coded in alphanumeric,” *Journal of the ACM*, vol. 15, no. 14, pages 514-34, October 1968.

[65] J. Moy. “OSPF version 2,” *RFC 2328*, http://www.ietf.org/rfc/rfc2328.txt, April 1998.

[66] M. Naldi. “Size estimation and growth forecast of the Internet,” *Centro Vito Volterra preprints*, University of Rome, October 1997.

[67] P. Newman, G. Minshall, T. Lyon and L. Huston. “IP switching and gigabit routers,” *IEEE Communications Magazine*, vol. 35, no. 1, pages 64-9, January 1997.

[68] P. Newman, T. Lyon and G. Minshall. “Flow labelled IP: a connectionless approach to ATM”, *Proceedings of IEEE Infocom*, vol. 3, pages 1251-60, April 1996.

[69] S. Nilsson and G. Karlsson. “IP-address lookup using LC-tries,” *IEEE Journal of Selected Areas in Communications*, vol. 17, no. 6, pages 1083-92, June 1999.

[70] H. Obara. “Optimum architecture for input queueing ATM switches,” *IEEE Electronics Letters*, pages 555-7, March 1991.

[71] H. Obara and Y. Hamazumi. “Parallel contention resolution control for input queueing ATM switches,” *IEEE Electronics Letters*, vol. 28, no. 9, pages 838-9, April 1992.

[72] H. Obara, S. Okamoto and Y. Hamazumi. “Input and output queueing ATM switch architecture with spatial and temporal slot reservation control,” *IEEE Electronics Letters*, vol. 28, no. 1, pages 22-4, January 1992.

[73] M.H. Overmars and A.F. van der Stappen. “Range searching and point location among fat objects,” *Journal of Algorithms*, vol. 21, no. 3, pages 629-56, November 1996.

[74] 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 Transactions on Networking*, vol. 1, no. 3, pages 344-57, June 1993.

[75] A.K. Parekh and R.G. Gallager. “A generalized processor sharing approach to flow control in integrated services networks: the multiple node case,” *IEEE/ACM Transactions on Networking*, vol. 2, no. 2, pages 137-50, April 1994.

[76] 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, G.D. Troxel, D. Waitzman and S. Winterble. “A 50-Gb/s IP router,” *IEEE/ACM Transactions on Networking*, vol. 6, no. 3, pages 237-48, June 1998.

[77] C. Partridge. “Locality and route caches,” *NSF workshop on Internet Statistics Measurement and Analysis*, San Diego, February 1996. Also see http://moat.nlanr.net/ISMA/Positions/partridge.html

[78] M. Patyra and W. Maly. “Circuit design for a large area high-performance crossbar switch,” *Proceedings International Workshop on Defect and Fault Tolerance on VLSI Systems*, pages 32-45, November 1991.

[79] F. Preparata and M. I. Shamos. *Computational geometry: an introduction*, Springer-Verlag, 1985.

[80] Y. Rekhter and T. Li. “A border gateway protocol 4 (BGP-4),” *RFC 1771*, http://www.ietf.org/rfc/rfc1771.txt, March 1995.

[81] Y. Rekhter and T. Li. “An architecture for IP address allocation with CIDR.” *RFC 1518*, http://www.ietf.org/rfc/rfc1518.txt, September 1993.

[82] Y. Rekhter, B. Davie, D. Katz, E. Rosen and G. Swallow. “Cisco systems’ tag switching architecture overview,” *RFC 2105*, http://www.ietf.org/rfc/rfc2105.txt, February 1997.

[83] C. Rose. “Rapid optimal scheduling for time-multiplex switches using a cellular automaton,” *IEEE Transactions on Communications*, vol. 37, no. 5, pages 500-9, May 1989.

[84] H. Samet. *The design and analysis of spatial data structures*, Addison-Wesley, 1990.

[85] B. Schieber. “Computing a minimum weight k-link path in graphs with the concave Monge property,” *Journal of Algorithms*, vol. 29, no. 2, pages 204-22, November 1998.

[86] R. Sedgewick and P. Flajolet. *An introduction to the analysis of algorithms*, Addison-Wesley, 1996.

[87] F. Shafai, K.J. Schultz, G.F.R. Gibson, A.G. Bluschke and D.E. Somppi. “Fully parallel 30-Mhz, 2.5-Mb CAM,” *IEEE Journal of Solid-State Circuits*, vol. 33, no. 11, pages 1690-6, November 1998.

[88] D. Shah and P. Gupta. “Fast incremental updates on ternary-CAMs for routing lookups and packet classification,” *Proceedings of Hot Interconnects VIII*, August 2000. To also appear in *IEEE Micro January/February 2001*.

[89] M. Shreedhar and G. Varghese. “Efficient fair queuing using deficit round-robin,” *IEEE/ACM Transactions on Networking*, vol. 4, no. 3, pages 375-85, June 1996.

[90] K. Sklower. “A tree-based packet routing table for berkeley unix,” *Proceedings of the 1991 Winter Usenix Conference*, Dallas, pages 93-9, 1991.

[91] F. Solensky and F. Kastenholz. “A revision to IP address classifications,” *work in progress*, Internet draft, March 1992

[92] F. Solensky and F. Kastenholz. “Definition of class E IP addresses,” *Internet draft*, available at http://gnietf.vlsm.org/17.txt, August 1991.

[93] V. Srinivasan and G. Varghese. “Faster IP lookups using controlled prefix expansion,” *Sigmetrics*, 1998. Also in *ACM Transactions on Computer Systems*, vol. 17, no. 1, pages 1-40, February 1999.

[94] V. Srinivasan and G. Varghese. “A survey of recent IP lookup schemes,” *Proceedings of Conference on Protocols for High Speed Networks*, pages 9-23, August 1999.

[95] V. Srinivasan, S. Suri, G. Varghese and M. Waldvogel. “Fast and scalable layer four switching,” *Proceedings of ACM Sigcomm*, pages 203-14, September 1998.

[96] V. Srinivasan, S. Suri and G. Varghese. “Packet classification using tuple space search”, *Proceedings of ACM Sigcomm*, pages 135-46, September 1999.

[97] V. Srinivasan. “Fast and efficient Internet lookups,” PhD thesis, Washington University, 1999.

[98] W.R. Stevens and G.R. Wright. *TCP/IP Illustrated, vol. 2, the implementation*, Addison-Wesley, 1995.

[99] H. Suzuki, H. Nagano, T. Suzuki, T. Takeuchi, S. Iwasaki. “Output-buffer switch architecture for asynchronous transfer mode,” *IEEE International Conference on Communications*, vol. 1, pages 99-103, June 1989.

[100] Y. Tamir and G. Frazier. “High performance multi-queue buffers for VLSI communication switches,” *Proceedings of the 15th Annual Symposium on Computer Architecture*, pages 343-54, June 1988.

[101] Y. Tamir and H.C. Chi. “Symmetric crossbar arbiters for VLSI communication switches,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 4, no. 1, pages 13-27, January 1993.

[102] R.E. Tarjan. “Data structures and network algorithms,” *Society for Industrial and Applied Mathematics*, Pennsylvania, November 1983.

[103] L. Tassiulas and A. Ephremides. “Stability properties of constrained queueing systems and scheduling policies for maximum throughput in mul-

tihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pages 1936-48, December 1992.

[104] F.A. Tobagi. "Fast packet switch architectures for broadband integrated services digital networks," *Proceedings of the IEEE*, vol. 78, no. 1, pages 133-67, January 1990.

[105] T.P. Troudet and S.M. Walters. "Neural network architecture for crossbar switch control," *IEEE Transactions on Circuits and Systems,* vol. 38, no. 1, pages 42-57, January 1991.

[106] P. Tsuchiya. "A search algorithm for table entries with non-contiguous wildcarding," *unpublished report, Bellcore*.

[107] H.H.Y. Tzeng and T. Przygienda. "On fast address-lookup algorithms," *IEEE Journal of Selected Areas in Communications*, vol. 17, no. 6, pages 1067-82, June 1999.

[108] M. Waldvogel, G. Varghese, J. Turner and B. Plattner. "Scalable high-speed IP routing lookups," *Proceedings of ACM Sigcomm*, pages 25-36, October 1997.

[109] R.L. Wessner. "Optimal alphabetic search trees with restricted maximal height," *Information Processing Letters*, vol. 4, no. 4, pages 90-4, January 1976.

[110] N. Weste and K. Eshraghian. *Principles of CMOS VLSI design: a systems perspective*, Addison-Wesley, Massachusetts, 1993.

[111] D.E. Willard. "Log-logarithmic worst-case range queries are possible in space $\Theta(N)$," *Information Processing Letters*, vol. 17, no. 2, pages 81-4, 1983.

[112] T. Y. C. Woo, "A modular approach to packet classification: algorithms and results," *Proceedings of IEEE Infocom*, vol. 3, pages 1203-22, March 2000.

[113] T. Yamagata, M. Mihara, T. Hamamoto, Y. Murai, T. Kobayashi, M. Yamada and H. Ozaki. "A 288-kb fully parallel content addressable mem-

ory using a stacked-capacitor cell structure," *IEEE Journal of Solid-state Circuits*, vol. 27, no. 12, December 1992.

[114] R. W. Yeung. "Alphabetic codes revisited," *IEEE Transactions on Information Theory*, vol. 37, no. 3, pages 564-72, May 1991.

[115] H.S. Yoon. "A large-scale ATM switch: analysis, simulation and implementation," *Proceedings ICATM-98*, pages 459-64, 1998

[116] H. Zhang. "Service disciplines for guaranteed performance service in packet-switching networks," *Proceedings of the IEEE*, vol. 83, no. 10, pages 1374-96, October 1995.

[117] L. Zhang. "VirtualClock: a new traffic control algorithm for packet-switched networks," *ACM Transactions on Computer Systems*, vol. 9, no. 2, pages 101-24, May 1991.

[118] All vBNS routes snapshot, at http://www.vbns.net/route/Allsnap.rt.html.

[119] Arrow Electronics Inc., at http://www.arrow.com.

[120] Cisco 12000 series GSR, at http://www.cisco.com/warp/public/cc/pd/rt/12000/index.shtml.

[121] Histogram of packet sizes, at http://www.nlanr.net/NA/Learn/plen.970625.hist.

[122] IBM ASIC Standard Cell/Gate Array Products, at http://www.chips.ibm.com/products/asics/products/edram/index.html.

[123] Information Processing Systems — Open Systems Interconnection, "Transport Service Definition," International Organization for Standardization, International Standard 8072, June 1986.

[124] "Internet routing table statistics," at http://www.merit.edu/ipma/routing_table.

[125] IPng working group in the IETF, at http://www.ietf.org/html.charters/ipngwg-charter.html.

[126] Juniper Networks Inc., Internet Processor II: Performance without compromise, at http://www.juniper.net/products/brochures/150006.html.

[127] Lara Networks, at http://www.laratech.com.

[128] Lucent NX64000 multi-terabit switch/router, at http://www.lucent.com/ins/products/nx64000.

[129] Memory-memory, at http://www.memorymemory.com.

[130] Mosaid, at http://www.mosaid.com/semiconductor/networking.htm.

[131] Netlogic microsystems, at http://www.netlogicmicro.com. CIDR products at http://www.netlogicmicro.com/products/cidr/cidr.html.

[132] Network Address Translators IETF working group, at http://www.ietf.org/html.charters/nat-charter.html.

[133] "Nortel Networks breaks own 'land speed record' using light — redefines speed of Internet and networking," press release, Nortel Networks, http://www.nortelnetworks.com/corporate/news/newsreleases/1999d/10_12_9999633_80gigabit.html, October 12, 1999.

[134] NLANR Network Analysis Infrastructure, at http://moat.nlanr.net.

[135] Sibercore Technologies, at http://www.sibercore.com.

[136] Telstra Internet — AS 1221 BGP table data, at http://www.telstra.net/ops/bgptable.html.

[137] Toshiba America Electronic Components, at http://www.toshiba.com/taec.

[138] VTune(TM) Performance Analyzer Home Page, at http://developer.intel.com/vtune/analyzer/index.htm.