half-title page, prepared by publisher

blank page

Cluster Computing for Robotics and Computer Vision

Damian M. Lyons

full-title page, prepared by publisher

Dedication

To Jeanie Shippey Lyons
and Jeani Shannon Lyons

Preface

Robotic technology is expanding into many new fields including the military, professional service, home/consumer and educational fields. Keeping up with the computational demands of these new applications requires harnessing the latest computational technologies. Cluster computing is an approach to providing supercomputer performance from a collection of off-the-shelf computer systems. The popularization of multi-core processors means that even single computer systems can furnish several processor cores for a cluster.

The objective of this text is to give professionals and students working in the cluster computing field, or in the robotics and computer vision field, a concrete view of the strong synergy between these two areas. A second objective is to spur further fruitful exploitation of this connection.

The book is written at a level appropriate for an advanced undergraduate (or robot enthusiast) or graduate student. To make the text useful to a wide audience in these fields, the key concepts in robotics, computer vision, and cluster computing are introduced before being used.

I chose the algorithms and applications covered in the text because they were either easily accessible to a robotics person looking at cluster computing for the first time, or a cluster computing person looking at robotics for the first time. The first sections are necessarily therefore fairly simple robotics and cluster topics. However, the chapters build on each other and more complex robotics and cluster computing topics are covered in later chapters.

The reader will notice that each chapter ends with a bibliography section containing the references cited in that chapter. The reference numbers in the text refer to this section. I hope this system will offer the reader a convenience in locating references, though it is at the cost of some duplication of references. A merged bibliography is presented at the end of the text for completeness. The Appendices contain a summary of MPI related material from the text and from the MPI V2.2 specification.

It would be very difficult to provide a comprehensive list of robotics and computer vision algorithms and their cluster algorithm designs. My hope is that if a reader does not find the algorithm or analysis they need here, at least they find an entry point into the topic from which they can prototype their own solutions.

I have many to thank for their inspiration, insight and help in producing this text. I start by thanking Michael Arbib and Ken Overton for taking an engineering student with a distributed computing background into the Laboratory for Perceptual Robotics at UMass, kick-starting my interest in concurrency and robotics. I have been lucky to know and work with many brilliant and insightful professionals including Ron Arkin, Paul Benjamin, Tomas Brodsky, Eric Cohen-Solal, Teun Hendricks, MiSuen Lee, Yun-ting Lin, Tom Murphy, and many others. I owe a special debt of gratitude to Frank Hsu, who has been a mentor and friend in my years at Fordham University. Thanks are due to Stephen Fox, Gary Weiss and Arthur G. Werschulz for many comments on an early draft of the text. I also want to thank the many students who took my Parallel Computation course and my Robotics course at Fordham over the years and on which much of this material was dry-run including Sirhan Chaudhry, Sothearith Chanty, Jeremy Drysdale, Jose DeLeon, Giselle Isner, Andy Palumbo, Kiran Pamnany, and many others.

Finally, my family deserves all my thanks and more for graciously putting up with my many absences over the past few years while holed up in my office or lab assembling this text.

Damian M. Lyons
December 2010

Contents

List of Tables

List of Figures

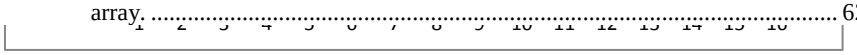

Bibliography

1. Agre, P. and Chapman, D., PENGI: An implementation of a theory of activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)***,** pages 268-272, Seattle WA, 1987.
2. Andersson, L.A.A. and Nygards, J., On multi-robot map fusion by inter-robot observations, *12th International Conference on Information Fusion (FUSION '09)*, Seattle WA, 2009.
3. Arbib, M.A., Iberall, T. and Lyons, D., Coordinated Control Programs for Movements of the Hand, in *Hand Function and the Neocortex*, A.W. Goodwin and I. Darian-Smith (Eds), Berlin, Springer-Verlag (1985).
4. Arbib, M.A., Iberall, T. and Lyons, D.M., Coordinated Control Programs for Movements of the Hand, *Exp. Brain Res. #10*, Springer-Verlag, 1985
5. Arbib, M.A., *The Metaphorical Brain 2*. Wiley & Sons 1989.
6. Arkin, R.C., *Behavior-Based Robotics*. MIT Press 1998.
7. Asbeck, A.T., Kim, S., Cutkosky, M.R., Provancher, W.R. and Lanzetta, M., Scaling hard vertical surfaces with compliant microspine arrays, *Proceedings, Robotics Science and Systems*, Cambridge MA, 2005.
8. B. Szatmary, J. Fleischer, D. Hutson, D. Moore, J. Snook, G. M. Edelman and J. Krichmar, A Segway-based human-robot soccer team, *IEEE International Conference on Robotics and Automation*, Orlando FL, 2006,
9. Baase, S. and van Gelder, A., *Computer Algorithms*. Addison-Wesley 1999.
10. Bar-Shalom, Y. and Fortmann, T., *Tracking and Data Association*. Academic Press 1988.
11. Bekey, G.A., *Autonomous Robots*. MIT Press 2005.

12. Besl P., and N. McKay. A Method for Registration of 3–D Shapes, *IEEE Trans. on PAMI*, 14(2):239 – 256, February 1992.
13. Bilmes, J., A Gentle Tutorial on the EM Algorithm Including Gaussian Mixtures and Baum-Welch, *ICSI Technical Report TR-97-021*, May 1997.
14. Birchfield S.T., and Sriram Rangarajan, Spatial Histograms for Region-Based Tracking, *ETRI Journal*, V29, N5, Oct. 2007.
15. Bookman, C., *Linux Clustering*. New Riders 2003.
16. Borenstein, J., Experimental Results from Internal Odometry Error Correction with the OmniMate Mobile Robot, *IEEE Trans. on Robotics and Automation, Vol. 14, No. 6*, December 1998, pp. 963-969.
17. Bradski, G. and Kaehler, A., *Learning OpenCv*. O'Reilly 2008.
18. Braunl, T., Feyrer, S., Rapf, W. and Reinhardt, M., *Parallel Image Processing*. Springer-Verlag, 2001.
19. Brooks, R. A., A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14-23, 1986.
20. Brummit, B. and Stentz, A., Dynamic Mission Planning for Multiple Mobile Robots, *Proc. IEEE Int. Conf. on Robotics and Automation*, Minneapolis MN, 1996.
21. Chamberlain, B.L., Sung-Eun Choi, E Christopher Lewis, Calvin Lin, Lawrence Snyder and W. Derrick Weathersby. ZPL: A machine independent programming language for parallel computers. *IEEE Transactions on Software Engineering*, 6(3):197-211, March 2000.
22. Chapman, B., Jost, G. and van der Pas, R., *Using OpenMP: Portable Shared Memory Parallel Programming.* MIT Press 2008.
23. Cheng, D.Y., A Study of Parallel Programming Languages and Tools. *NASA Ames Research Center Report RND-93-005*, March 1993.
24. Choset, H., Lynch, K., Hutchinson, S., Kantor, G., Burgard, W., Kavraki, L. and Thrun, S., *Principles of Robot Motion*. MIT Press 2005.
25. Cox, I.J. and Hingorani, S.L. An Efficient Implementation and Evaluation of Reid's Multiple Hypothesis Tracking Algorithm for Visual Tracking. *Int. Conf. on Pattern Recog.* (1994) pp.437-442.

26. Dempster, P., N. M. Laird and D. B. Rubin, Maximum Likelihood from Incomplete Data via the EM Algorithm, *Journal of Royal Statistics Society*, vol. B-39, 1977.
27. de Souza, P.S.L., Sabourin, R., de Souza, S.R.S. and Borges, D.L., A low-cost parallel K-means VQ algorithm using cluster computing, *Seventh International Conference on Document Analysis and Recognition (ICDAR'03)*, Edinburgh Scotland, 2003.
28. Dellaert_F., Fox, D., Burgard W. and Thrun, S., Monte Carlo Localization for Mobile Robots, *IEEE International Conference on Robotics and Automation (ICRA99)*, Detroit MI, May 1999.
29. Diosi, A. and Kleeman, L., Laser Scan Matching in Polar Coordinates. *Proc. Int. Robots and Systems (IROS)*, Edmonton Canada, 2005.
30. Dudek, G. and Jenkin, M., *Computational Principles of Mobile Robotics*. Cambridge Press 2000.
31. Elfes, A., Using Occupancy Grids for Mobile Robot Perception and Navigation, *IEEE Computer*, June 1989.
32. Ellore, B. K., Dynamically expanding occupancy grids. *Master's thesis, Texas Tech University*, 2002.
33. Engelberger, J.F., *Robotics in Practice*. MIT Press 1989.
34. Fikes, R. E. and Nilsson, N., STRIPS: A new approach to the application of theorem proving to problem solving. Artificial Intelligence, 5(2):189-208, 1971.
35. Firby, J. A., An investigation into reactive planning in complex domains. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, pages 202-206, Milan Italy, 1987.
36. Flynn, M., Some Computer Organizations and Their Effectiveness, *IEEE Transactions on Computing*, Vol. C-21, pp. 948, 1972
37. Fox, D., J. Ko, K. Konolige, B. Limketkai, D. Schulz and B. Stewart, Distributed Multi-Robot Exploration and Mapping**,** *Proc. of the IEEE: Special Issue, Multi-Robot Systems V94 N7, 2006.*
38. Goto Y. and Stenz A., The CMU System for Mobile Robot Navigation, *IEEE Int. Conf. on Robotics and Automation*, Raleigh NC, pp. 99-105, 1987.
39. Graham, J. H., Special Computer Architectures for Robotics: Tutorial and Survey, *IEEE Trans. on Robotics and Automation, vol.5, no. 5*, 1989.

40. Gropp, E., Lusk, E. and Sterling, T., Beowulf Cluster Computing with Linux. *MIT Press* 2003.
41. Henrich, D., and Honiger, T., Parallel Processing Approaches in Robotics, *IEEE International Symposium on Industrial Electronics (ISIE'97)*, Guimarães Portugal, July 7-11, 1997.
42. Hillis, W.D., *The Connection Machine.* MIT Press 1989.
43. Hoare, C.A.R., *Communicating Sequential Processes.* Prentice-Hall 1985.
44. Hough, P., Methods and means for recognizing complex patterns, *US Patent #3,069,654*. 1962.
45. Isard M. and A. Blake: Condensation -- Conditional density propagation for visual tracking. *International Journal of Computer Vision* 29 (1998) 5-28.
46. Jeannerod, M., M. A. Arbib, G. Rizzolatti and H. Sakata. Grasping objects: the cortical mechanisms of visuomotor transformation, *Trends Neurosci*. 18: 314–320, 1995.
47. Jones, J., *Robot Programming: A Practical guide to Behavior-Based Robotics*. McGraw-Hill 2004.
48. Joshi, M.N., Parallel K-Means Algorithm on Distributed Memory Multiprocessors, *Computer*, 2003.
49. Kerridge, J., *Occam Programming: A Practical Approach.* Wiley-Blackwell, 1987.
50. Klechenov , A., Gupta, A.k., Wong, W.F., Ng, T.K. and Leow, W.K., Real-time Mosaic for Multi-Camera Videoconferencing, *MIT-Singapore Alliance Symposium*, Singapore 2003.
51. Kuhn, M., Parallel Image Registration in Distributed Memory Environments, *M.S. Thesis*, Swiss Federal Institute of Technology, 2004.
52. Lakaemper, R., Adluru, N., Latecki, L.J. and Madhavan, R., Multi Robot Mapping using Force Field Simulation. *Journal of Field Robotics,* 24(8/9), 2007, pp.747-762.
53. Lars Ailo Bongo, Otto J. Anshus and John Markus Bjørndalen**,** Evaluating the Performance of the Allreduce Collective Operation on Clusters: Approach and Results, *Technical Report 2004-48*, University of Tromso, 2004.
54. Lewis, M. A., Fagg, A. H. and Bekey, G. A., The USC Autonomous Flying Vehicle: an Experiment in Real-Time Behavior-Based Control, *Proceedings of the IEEE Conference on Robotics and Automation,* May 1993, Atlanta Georgia.
55. López de Teruel P. E., José M. García and Manuel E.Acacio. The Parallel EM Algorithm and its Applications in Computer

Vision, *Proceedings Parallel and Distributed Processing Techniques and Applications*, Las Vegas Nevada, 1999.

56. Lowe, D., Distinctive image features from scale-invariant keypoints, *International Journal of Computer Vision*, 60, 2 (2004), pp. 91-110.
57. Lucke, R.W., *Building Clustered Linux Systems.* Prentice Hall 2005.
58. Lyons, D., Selection and Recognition of Landmarks Using Terrain Spatiograms, *IEEE/RSJ International Conference on Intelligent RObots and Systems (IROS)*, Tapei, Taiwan, October 2010.
59. Lyons, D. and Hsu, D.F., Rank-based Multisensory Fusion in Multitarget Video Tracking, *IEEE Intr. Conf. on Advanced Video & Signal-Based Surveillance*. Como Italy, 2005.
60. Lyons, D.M. and Arbib, M.A., A Formal Model of Computation for Sensory-based Robotics, *IEEE Transactions on Robotics and Automation 5(3)* Jun. 1989.
61. Lyons, D.M., Sharing Landmark Information using Mixture of Gaussian Terrain Spatiograms, *IEEE/RSJ International Conference on Intelligent RObots and Systems (IROS)*, St. Louis MO, October 2009.
62. Lyons, D.M. and Arkin, R., Towards Performance Guarantees for Emergent Behavior, *IEEE International Conference on Robotics and Automation*, New Orleans LA, 2004.
63. Lyons, D.M., Detection and Filtering of Landmark Occlusions using Terrain Spatiograms. *IEEE Int. Conference on Robotics and Automation*, Anchorage Alaska, May 2010.
64. Lyons, D.M. and Isner, G.R., Evaluation of a Parallel Algorithm and Architecture for Mapping and Localization, *7th International Symposium on Computational Intelligence In Robotics and Automation, CIRA 2007*, Jacksonville FL, June 20-23, 2007.
65. Lyons, D.M., Representing and Analyzing Action Plans as Networks of Concurrent Processes, *IEEE Transactions on Robotics and Automation* 9(3), June 1993.
66. McGraw, J., S. Skedzielewski, S. Allan, D. Grit, R. Oldehoeft, J. Glauert, I. Dobes, and P. Hohensee. SISAL-Streams and Iterations in a Single Assignment Language, Language Reference Manual, version 1. 2. *Technical Report TR M-146, University of California - Lawrence Livermore Laboratory*, March 1985.

67. Medeiros, H., Gao, X.T., Kleihorst, R.P., Park, J. and Kak, A.C., A parallel color-based particle filter for object tracking **,** *The 6th ACM Conference on Embedded Networked Sensor Systems*, Raleigh NC, Nov. 5 - Nov. 7, 2008.
68. Message Passing Interface Forum, MPI: A Message-Passing Interface Standard, Version 2.2, 2009.
69. Mobilerobots Inc. Pioneer 3 & Pioneer 2 H8-Series Operations Manual, 2003.
70. Montemerlo, M. et al., Junior: The Stanford Entry in the Urban Challenge**,** *Journal of Field Robotics* Volume 25, Issue 9 (September 2008).
71. Moravec, H. and Elfes, A, High Resolution Maps from Wide Angle Sonar, *IEEE International Conference on Robotics and Automation*, Washington DC, 1985.
72. Mundhenk, N., et al., Low cost, high performance robot design utilizing off-the-shelf parts and the Beowulf concept, The Beobot project, *Proceedings, SPIE Intelligent robots and computer vision XXI : algorithms, techniques, and active vision*, Providence RI, 28-29 October 2003.
73. Nof, S.Y., *Handbook of Industrial Robotics*. Wiley 1999.
74. Nuchter, A., Parallelization of Scan Matching for Robotic 3D Mapping, *Proc. 3rd European Conf. on Mobile Robots*, 2007.
75. Oliveira, P. and du Buf, H., SPMD Image Processing on Beowulf Clusters: Directives and Libraries, *Proceedings of the 17th International Symposium on Parallel and Distributed Processing*, 2003.
76. Ouerhani, N., von Wartburg, R., Hugli, H. and Muri, R., Empirical Validation of the Saliency-based Model of Visual Attention, *Electronic Letters on Computer Vision and Image Analysis* 3(1):13-24, 2004.
77. Parker, J., *Algorithms for Image Processing and Computer Vision*. Wiley 1997.
78. Pfister, G., *In Search of Clusters*. *2nd Ed.* Prentice Hall, 1998.
79. Patarasuk P., and Xin Yuan, Bandwidth optimal all-reduce algorithms for clusters of workstations, *Journal of Parallel and Distributed Computing* Volume 69, Issue 2 (February 2009).
80. Pjesivac-Grbovic, J., Angskun, T., Bosilca, G., Fagg, G. E., Gabriel, E. and Dongarra, J. J., Performance analysis of MPI collective operations, *Cluster Computing*, V10, N2, 2007, pp. 127-143.

81. Pomerleau, D., Neural Networks for Intelligent Vehicles, *Proceedings of the Intelligent Vehicles Conference*, Tokyo, Japan 1993.
82. Quinn, M.J., *Parallel Programming in C with MPI and OpenMP*. McGraw-Hill 2004.
83. Rao, B., *Data Association Methods for Tracking Systems*, in *Active Vision*, A. Blake, Yuille, A., Editor. 1992, MIT Press. pp. 91-106.
84. Rasmussen, C. and Hager, G., Joint Probabilistic Techniques for Tracking Multi-Part Objects. *Proc. Computer Vision & Pattern Recognition*, Santa Barbara CA, (1998) pp.16-21.
85. Rosselot, D. and Hall, E., Processing real-time stereo video for an autonomous robot using disparity maps and sensor fusion, *SPIE Intelligent Robots and Computer Vision XXII: Algorithms, Techniques, and Active Vision*. pp. 70-78 (2004).
86. Sack, D. and Burgard, W., A comparison of methods for line extraction from range data, *5th IFAC Symposium on Intelligent Autonomous Vehicles*, Lisbon, Portugal, 2004.
87. Se, S., David Lowe, Jim Little, Local and Global Localization for Mobile Robots using Visual Landmarks, *Proceedings of the International Conference on Intelligent Robots and Systems* (IROS), pp.414—420, 2001.
88. Sermanet, P., Raia Hadsell, Marco Scoffier, Matt Grimes, Jan Ben, Ayse Erkan, Chris Crudele, Urs Muller and Yann LeCun: A Multi-Range Architecture for Collision-Free Off-Road Robot Navigation, *Journal of Field Robotics*, 26(1):58-87, January 2009.
89. Sloan, J., *High performance Linux Clusters*. O'Reilly 2005.
90. Squyres, J., Lumsdaine, A. and Stevenson, R., A toolkit for parallel image processing, *Proceedings of the SPIE Conference on Parallel and Distributed Methods for Image Processing*, 1998.
91. Steenstrup, M., M.A. Arbib and E.G. Manes, Port automata and the algebra of concurrent process, *J. Computer System Sciences* 27 (1), 29-50 (1983).
92. Thrun, S., Burgard, W. and Fox, D., A Probabilistic Approach to Concurrent Mapping and Localization for Mobile Robots. *Machine Learning and Aut. Robots* 31/5 1998, pp. 1-25.
93. Thrun, S., D. Fox, W. Burgard and F. Dellaert, Robust monte carlo localization for mobile robots. *Artificial Intelligence 128*(1-2), 99–141, 2001.

94. Thrun. S., *Robotic mapping: A survey*. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millennium*. Morgan Kaufmann, 2002.
95. Usher, K., Obstacle avoidance for a non-holonomic vehicle using occupancy grids, *Proceedings of the 2006 Australasian Conference on Robotics & Automation*, Auckland, New Zealand, Dec. 2006.
96. Venkataraman S.T. and. Iberall, T., (Eds.) Dexterous Robot Hands. *New York: Springer Verlag*, 1990.
97. Wilkes, D., Dudek, G., Jenkin, M. and Milios, E., Modeling Sonar Range Sensors, in: (Archibald and Petriu, Eds) *Advances in Machine Vision: Strategies and Applications*, World Scientific Press, Singapore 1992.
98. Wilkinson, B., Alen, M., *Parallel Programming.* 2nd Ed. Prentice-Hall 2005.
99. Yuh, J., Design and Control of Autonomous Underwater Robots: A Survey, *Autonomous Robots*, Volume 8, Number 1, January 2000.

Index