

Exploiting the Intrinsic Structures of Simultaneous Localization and Mapping

by

Kasra Khosoussi

Submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

at the

Centre for Autonomous Systems
Faculty of Engineering and Information Technology
University of Technology Sydney

March 2017

Declaration

I certify that the work in this thesis has not previously been submitted for a degree nor has it been submitted as part of requirements for a degree except as fully acknowledged within the text.

I also certify that the thesis has been written by me. Any help that I have received in my research work and the preparation of the thesis itself has been acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

Signature:

Date:

Abstract

Imagine a *robot* which is assigned a *complex task* that requires it to navigate in an *unknown environment*. This situation frequently arises in a wide spectrum of applications; e.g., surgical robots used for diagnosis and treatment operate inside the body, domestic robots have to operate in people's houses, and Mars rovers explore Mars. To navigate safely in such environments, the robot needs to constantly estimate its location and, simultaneously, create a consistent representation of the environment by, e.g., building a map. This problem is known as simultaneous localization and mapping (SLAM). Over the past three decades, SLAM has always been a central topic in mobile robotics. Tremendous progress has been made during these years in efficiently solving SLAM using a variety of sensors in large-scale environments. However, some of the intrinsic structures of SLAM are yet to be fully understood and utilized. This thesis is devoted to the study of some of these structures.

This thesis is comprised of two main parts:

In the first part, we examine the specific structure of the standard measurement models in pose-graph and feature-based SLAM. These standard models are affine with respect to the position of robot and landmarks. Consequently, given the orientation of the robot, the conditionally-optimal estimate for the robot's and landmarks' positions can be recovered instantly by solving a sparse linear least squares problem. We propose an algorithm to exploit this intrinsic property of SLAM, by stripping the problem down to its nonlinear core, while maintaining its natural sparsity. By taking advantage of the separable structure of SLAM, we gain a global perspective that guides the conventional local search algorithms towards a potential solution, and hence achieve a fast and stable convergence.

In the second part of this thesis, we investigate the impact of the graphical structure of SLAM and several other estimation-over-graph problems, on the estimation error covariance. We establish several connections between various graph-connectivity measures and estimation-theoretic concepts. For example, we prove that, under certain conditions, the

determinant of the estimation error covariance matrix of the maximum likelihood estimator can be estimated by counting the weighted number of spanning trees in the corresponding graph, without solving the optimization problem or using any information about the “geometry” of the scene (e.g., numerical values of the measurements). This surprising result shows that the graphical structure of SLAM provides a compact, but rich representation of the underlying estimation problem, that can be used—e.g., in decision making and active SLAM—to predict and reason about the resulting estimation error covariance. Consequently, we tackle the combinatorial optimization problem of designing sparse graphs with the maximum weighted number of spanning trees. Characterising graphs with the maximum number of spanning trees is an open problem in general. To tackle this problem, we establish several new theoretical results, including the monotone log-submodularity of the weighted number of spanning trees in connected graphs. By exploiting these structures, we design a complementary pair of near-optimal efficient approximation algorithms with provable performance guarantees and near-optimality certificates. We discuss several applications of our results and apply our algorithms on the measurement selection problem in SLAM to design sparse near-D-optimal (determinant-optimal) SLAM problems.

Acknowledgements

The past five years have been a mysterious journey, filled with countless moments of joy and despair, triumph and failure, optimism and pessimism, and so many other contradictory emotions. Five years ago, I couldn't have imagined the adventures, people and places that this roller coaster would take me to.

First and foremost, I would like to thank my Ph.D. supervisors, Shoudong Huang and Gamini Dissanayake, for sharing this amazing journey alongside me. Their vision, wisdom, and guidance along the way have helped me to stay on the right track. Thank you for asking the right questions at the right time, for your enduring trust, patience, and unconditional support. No words can describe how grateful I am for the invaluable lessons I have learnt from them as great mentors, and wonderful colleagues.

I am grateful to my external examiners, Professors Simon Julier and Stefan Williams, for carefully reviewing this thesis and providing me with helpful comments. I am also thankful to the Australian Research Council and UTS for funding my Ph.D. studies.

I would like to express my sincere gratitude to Gaurav S. Sukhatme for giving me the incredible opportunity to collaborate with him during my 8 memorable months of stay at USC. My interaction with Gaurav has been extremely rewarding and enjoyable for me. I have been inspired by his passion for tackling intellectually challenging problems, and I admire him for his clarity of thought. I am forever indebted to him for providing me with guidance, his generosity with his time, and long lasting support.

During the past five years I have been lucky to collaborate with many talented researchers. Special thanks go to Gibson Hu, Heng Wang, Udo Frese, and Teresa Vidal-Calleja. I would like to thank Luca Carlone for bringing many exciting aspects of SLAM to my attention through his influential contributions to this field during the past few years. Since the RSS 2015 in Rome, I have been lucky to have him as a wonderful friend.

I am grateful to John J. Leonard and Gabriel Agamennoni for hosting my seminars at MIT and ETH Zürich. I truly enjoyed the experience and my fruitful discussions with the talented members of the Marine Robotics Group and the Autonomous Systems Lab.

I was lucky to be surrounded by so many amazing friends during my stay in Los Angeles. Special thanks go to my friend Mohammad Reza Rajati whom I will never forget his generosity and kindness. I would also like to thank Alireza Divsalar for our "cross-departmental" collaborations. I am thankful

to the members of the Robotic Embedded Systems Lab for so many enjoyable moments. Particularly thanks to Lantao Liu, Karol Hausman, Stephanie Kemna, Artem Molchanov, James Preiss, Christian Potthast, Oliver Kroemer, and David Kim.

I have always been blessed with the best flatmates in the world. In Sydney, thanks go to Ainsley and Lachlan, and Kaori and Hideki. In Los Angeles, thanks go to Elisabetta Missinato and to my extraordinarily talented musician friend, Mario Sevayega, for creating countless unforgettable moments, our jam sessions, authentic Italian cuisines, and weekly movie nights.

After five magnificent years, I will be departing Sydney in the coming days. Thanks to my good friend, Syamak Castle, for a warm welcome. Special thanks to my amazing friend, Maani Ghaffari Jadidi, with whom I have shared five years of unspeakable memories. I would like to express my gratitude to my wonderful friends and labmates in CAS. In particular, special thanks to Gibson Hu, Lakshitha Dantanarayana (a.k.a. Lakipedia¹), Raphael Falque, Kanzhi Wu, Liye Sun, Buddhi Wijerathna, Teng Zhang, Andrew To, Daobilige Su, Mahdi Hassan, Philip Quin, Linh Van Nguyen, Leo Shi, Kim Doang Nguyen, Behnam Maleki, Hasti Hayati, Teresa Vidal-Calleja, Brad Skinner, Liang Zhao, Alen Alempijevic, Jaime Valls Miro, and many others.

My journey really began when I joined the ARAS Lab in Iran as an undergraduate student. For that, I am always grateful to Hamid D. Taghirad and Ali Agha for their patience, commitment to excellence, and trust in me. I am also grateful to the SLAM community for sharing their codes and datasets through community-driven efforts such as Radish and OpenSLAM. In particular, special thanks to Rainer Kümmerle and his colleagues for g2o which played a significant role in my work.

Finally, I would like to express my deepest gratitude to my family for their selflessness, unconditional love, and never-ending support. Thanks to: my magnificent mom, Kamila, for being a ceaseless source of inspiration; my extraordinary dad, Hamid, for being my best friend; my precious sister, Sadaf, to whom I owe everything; and finally, Kourosh who is the most wonderful brother I could ever ask for. Staying apart from you has been the most difficult part of this journey. I wouldn't have been here if it weren't for your patience, love, and support. You are my personal heroes. Last but not least, special thanks to my soulmate, Christine, with whom I have shared unbelievable moments of joy during the past two years. Thanks for never giving up on me. You are my rock.

¹Wikipedia, but without any error.

To my family.

Contents

1	Introduction	1
1.1	Motivation	3
1.2	Main Contributions	6
1.3	The Structure of the Thesis	7
1.4	Publications	7
1.5	General Notation	9
2	Simultaneous Localization and Mapping: A Survey	10
2.1	Classic Paradigms	10
2.2	Modern Paradigms	16
2.3	Nonlinear Least Squares Formulation	20
2.4	Summary	21
I	SLAM: Sparse Separable Nonlinear Least Squares	22
3	A Sparse Separable SLAM Back-End	23
3.1	Related Works	24
3.2	Problem Formulation	26
3.2.1	Graphical Representation of SLAM	26
3.2.2	Measurement Model	26
3.2.3	Point Estimation Criterion	29
3.3	Exploiting Separability in SLAM	29
3.3.1	Variable Projection	29
3.3.2	Kaufman’s Algorithm	31
3.4	Sparse Variable Projection	33
3.5	Projection Gain	37
3.6	Maximum A Posteriori Estimation	38
3.7	Implementation	41
3.8	Experimental Results	42
3.8.1	Convergence	42
3.8.2	Run Time	45
3.9	Discussion	48
3.10	Summary	49

II	SLAM: Estimation over Graphs	52
4	Estimation over Graphs	53
4.1	Estimation-over-Graph Problems	54
4.1.1	Diff-Net	54
4.1.2	Compass-SLAM	56
4.1.3	SLAM	57
4.2	Main Results	59
4.2.1	Average Degree \leftrightarrow Log-Likelihood Ratio	59
4.2.2	Algebraic Connectivity \leftrightarrow Worst-Case Error Variance	64
4.2.3	Tree Connectivity \leftrightarrow Volume of Uncertainty Ellipsoids	65
4.3	Experimental Results	70
4.4	Summary	74
5	Synthesis of Near-Tree-Optimal Graphs	76
5.1	Maximizing Tree-Connectivity	77
5.1.1	Characterizing t -Optimal Graphs in $\mathbb{G}_{n,m}$	78
5.2	The Edge Selection Problem	79
5.2.1	1-ESP ⁺	81
5.2.2	Exhaustive Search	82
5.3	Approximation Algorithms for k -ESP ⁺	83
5.3.1	Greedy Algorithm	83
5.3.2	Convex Relaxation	85
5.3.3	Certifying Near-Optimality	91
5.4	Dual Problem: δ -ESP*	92
5.4.1	Greedy Algorithm	92
5.4.2	Convex Relaxation	94
5.4.3	Certifying Near-Optimality	96
5.5	k -ESP ⁺ Over a Partition Matroid	97
5.5.1	Greedy Algorithm	99
5.5.2	Convex Relaxation	100
5.5.3	Certifying Near-Optimality	102
5.6	Improper k -ESP ⁺	103
5.6.1	Most Valuable Spanning Tree	104
5.6.2	Greedy Heuristics for Improper k -ESP ⁺	106
5.7	Applications	109
5.7.1	Near-D-Optimal Measurement Selection	109
5.7.2	D-Optimal Bootstrapping for SLAM	113
5.7.3	Network Reliability	114
5.7.4	Applications in Other Domains	116
5.8	Experimental Results	116
5.8.1	Random Graphs	117
5.8.2	Real Pose-Graph SLAM Dataset	117

5.9 Summary	118
6 Conclusion	122
6.1 SLAM: Sparse Separable Nonlinear Least Squares	122
6.2 SLAM: Estimation over Graphs	124
Appendices	129
A Preliminaries	130
A.1 Linear Algebra	130
A.2 Estimation Theory	131
A.3 Discrete Optimization	132
A.4 Graph Theory	134
B Proofs	139
B.1 Chapter 3	139
B.2 Chapter 4	141
B.3 Chapter 5	142
Bibliography	147

List of Figures

1.1	Different applications of SLAM	2
1.2	A motivating example	3
3.1	An illustration of the variable projection algorithm	39
3.2	Evolution of the projection gain	40
3.3	Average number of iterations	43
3.4	Overall run time as a function of edge density	47
3.5	An illustration of the run times listed in Table 3.4.	51
4.1	Impact of connectivity on estimation quality in SLAM	54
4.2	A toy example	60
4.3	Average of the log-likelihood ration (γ)	63
4.4	Evolution of the relative error as a function of δ	72
4.5	Evolution of the relative error as a function of δ	72
4.6	Log-determinant of the Fisher information matrix in terms of τ	73
5.1	Examples of k -ESP	80
5.2	Convex relaxation for k -ESP	85
5.3	Partition matroid constraints.	97
5.4	A counterexample	103
5.5	Illustration of the Greedy-Greedy algorithm.	107
5.6	Greedy-Greedy.	109
5.7	Illustration of the ϵ -trick.	110
5.8	The ϵ -trick: augmenting a spanning tree.	111
5.9	k -ESP ⁺ in random graphs	119
5.10	k -ESP ⁺ for pose-graph SLAM in the Intel Research Lab dataset.	120
5.11	Greedy design in the Intel Research Lab dataset	121

List of Tables

3.1	A summary of variable projections algorithms	37
3.2	GN vs. VP under different noise levels	44
3.3	LM vs. VP-LM under different noise levels	44
3.4	Total run time for 10 random walks with 10^5 poses	44
3.5	Real and synthetic SLAM benchmarks	49
4.1	Normalized tree-connectivity of publicly available datasets	71