COOPERATIVE LOCALIZATION: ON MOTION-INDUCED INITIALIZATION AND JOINT STATE ESTIMATION UNDER COMMUNICATION CONSTRAINTS

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

NIKOLAS TRAWNY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

STERGIOS ROUMELIOTIS, ADVISOR

AUGUST 2010

Acknowledgements

This thesis would not have been possible without the support of a number of people.

First of all, my thanks go to my adviser, Professor Stergios Roumeliotis, for his constant encouragement and guidance, for the long hours of passing along his knowledge and experience, for pushing me beyond my own limitations, and for his seemingly endless supply of interesting research problems. I am also thankful for the time and invaluable advice from my committee members, Professor Giannakis, Professor Saad, and Professor Banerjee, as well as from Professor Gebre-Egziabher, who upheld my connection to the aerospace world. Finally, I would like to thank Georganne Tolaas, who helped navigate the maze of rules and regulations of the Computer Science Ph.D. program. Thanks also go to my labmates, Tassos, Esha, Joel, Faraz, Paul, Sam, Ke, Gian-Luca, and Igor, with all of whom I have shared hours of discussion, work, laughter, food, and play.

Life in graduate school consisted not only of robotics. I am glad for the happy times spent with friends at the German Student Association. I would also like to thank the conductors and fellow musicians at campus band and our chamber group for the chance to play some great music after a day filled with math derivations. Additionally, a big thank you to the entire Peterson family, for their embodiment of "Minnesota Nice", and for their generous hospitality.

Most importantly, none of this would have been possible had it not been for the unwavering support and love of my family. Even though my studies sent me far away from home, my parents have always done everything so I could follow my dreams, and have never failed to believe in me. My wife, Miriam, has shared with me all the ups and downs of life as a grad student, and still loves me. I am fortunate to have you.

Last but not least, I gratefully acknowledge financial support from the National Science Foundation, IEEE, the NASA Jet Propulsion Laboratory, the University of Minnesota Department of Computer Science and Engineering, and the University of Minnesota Digital Technology Center.

Abstract

Teams of mobile robots are becoming increasingly popular to measure and estimate quantities of interest at spatially distributed locations. They have been used in tasks such as surveillance, search and rescue, and underwater- or space exploration. For these tasks, accurate localization, i.e., determining the position and orientation (pose) of each robot, is a fundamental requirement. Instead of localizing each robot in a team independently, Cooperative Localization (CL) incorporates robot-to-robot observations and *jointly* estimates all robots' poses, which improves localization accuracy for all team members. However, such joint estimation also creates significant challenges. In particular, initializing a joint estimation algorithm requires knowledge of all robots' poses with respect to a common frame of reference. This initialization is straightforward using GPS or manual measurements, but is difficult in the absence of external references. The second difficulty of CL is that it requires communicating large amounts of data, e.g., the robots' sensor measurements or state estimates. However, transmitting all these quantities is not always feasible, either due to bandwidth or power constraints.

This thesis offers novel solutions to the aforementioned problems. In the first part of the thesis, we investigate the problem of *CL initialization*, using robot-to-robot measurements acquired at different vantage points during robot motion. We focus on the most challenging case of distance-only measurements, and provide algorithms that compute the guaranteed global optimum of a nonlinear weighted Least Squares problem formulation. These techniques exploit recent advances in numeric algebraic geometry and optimization.

In the second part, we investigate the problem of *CL under communication constraints*. To reduce communication bandwidth, we propose using adaptively quantized measurements. We extend existing quantized filtering approaches to batch MAP estimators, and apply these techniques to multi-robot localization. We provide results on optimal threshold selection, as well as optimal bit allocation to efficiently utilize time-varying bandwidth. Our results are validated in simulation and experiments.

By providing solutions for two important problems in CL – motion-induced estimator initialization, and estimation under communication constraints – the research presented in this thesis aims to promote use of cooperative mobile robots in challenging real-world applications.

Contents

List of Tables

List of Figures

Nomenclature and Abbreviations

$\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{P})$	The multivariate Gaussian pdf in the variable $\mathbf{x}$, with mean $\boldsymbol{\mu}$ and covariance matrix $\mathbf{P}$.
$Q(x)$	The Gaussian tail probability, defined as $Q(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}} du$
$Q^{-1}(x)$	Inverse Gaussian tail probability
$\hat{\mathbf{x}}_{k\|l}$	Estimate of $\mathbf{x}$ at time step k, conditioned on all measurements up to time step l
$\det(\cdot)$	Determinant of a matrix
$\mathrm{tr}(\cdot)$	Trace of a matrix
${}^i\mathbf{p}$	A vector $\mathbf{p}$ expressed with respect to frame $\{i\}$
${}^i_j\mathbf{C}$	Rotation matrix that transforms vectors from frame $\{j\}$ to frame $\{i\}$

BQKF	Batch-quantized Kalman filter
cdf	Cumulative distribution function
CL	Cooperative Localization
C-SLAM	Cooperative simultaneous localization and mapping
EKF	Extended Kalman filter
GPS	Global positioning system
IMU	Inertial measurement unit
IQKF	Iteratively-quantized Kalman filter
IQMAP	Iteratively-quantized MAP estimator
i-WLS	Iterative weighted least squares
KF	Kalman filter
MAP	Maximum a posteriori
MLE	Maximum likelihood estimation
MMSE	Minimum mean square error
NEES	Normalized estimation error squared
pdf	Probability density function
PF	Particle filter

QMAP	(Single-bit) quantized MAP estimator
RMS	Root mean square (error)
SLAM	Simultaneous localization and mapping
SNR	Signal-to-noise ratio
SOI-KF	Sign-of-innovations Kalman filter
SVD	Singular value decomposition
UAV	Unmanned aerial vehicle
UKF	Unscented Kalman filter
WLS	Weighted least squares
WSN	Wireless sensor network

Chapter 1

Introduction

1.1 Multi-robot Cooperative Localization

Teams of mobile robots, or mobile wireless sensor networks (WSNs), are becoming increasingly popular to measure and estimate quantities of interest at spatially distributed locations. They have been used in tasks such as habitat construction [90, 129], inspection of industrial areas such as nuclear power plants or pipelines [21], aerial surveillance [9], search and rescue operations [114], and underwater- [17] or even space exploration [55]. Mobile sensor networks are not limited to robots. More and more portable devices, like cell phones, PDAs, and even cars are equipped with sensing and networking capabilities, enabling a totally new generation of location-based services, like monitoring of air-quality by pedestrians [93] or traffic flow from moving cars [26]. Compared to an individual robot or sensor, multi-robot systems have the advantages of faster and more extensive coverage, increased reliability to sensor failures, and higher estimation accuracy through sensor fusion.

Accurate *localization*, i.e., estimation of the robot's position and orientation or attitude (pose), is a key requirement for successful completion of any higher level task that involves motion in, sensing of, and interaction with the robot's environment. To make inferences from and act upon sensor observations (e.g., a rover detecting an interesting geologic formation on a planet's surface, or an unmanned aerial vehicle (UAV) spotting a victim in a search and rescue mission), it is of critical importance to know where this observation was made. In order to coherently fuse data in a distributed sensor network, the location and viewing direction of each sensor has to be known. As yet another example, navigation of mobile robots requires accurate knowledge of the robot's pose to successfully plan a path to reach the goal location. In most cases in practice, precise localization is achieved by fusing measurements from proprioceptive sensors (e.g., odometry or inertial measurement units [IMUs]) and exteroceptive sensors (e.g., GPS, cameras, or laser scanners). Localization based purely on proprioceptive sensors, also known as dead reckoning, suffers from unbounded error accumulation over time. This is due to the fact that dead reckoning continuously integrates noise and biases affecting the sensor measurements. For this reason,

exteroceptive sensor observations are often used to bound the error or reduce its growth rate. Localization, and simultaneous localization and mapping (SLAM), i.e., the process of localizing a robot while concurrently constructing a map of its surroundings, have been the subject of intense research over the last decades.

Teams of multiple robots or mobile sensors have the potential to improve their positioning accuracy by performing *Cooperative Localization* (CL). CL is a form of sensor sharing, and as such requires communication, either between the robots themselves, or between the robots and a fusion center. In CL, all robot poses are estimated jointly, using all available measurements from the entire team, including relative observations between the robots, for example of their mutual distance or bearing (see Fig. 1.1). Relative observations cannot bound the localization error, but significantly reduce the error growth rate [79]. CL is therefore particularly attractive in GPS-denied applications, such as underwater or indoors. Compared to (cooperative) SLAM, CL has the advantage of being independent of the environment, in the sense that it does not require solving the difficult perception problem of natural feature extraction and matching, nor does it in fact require the presence of stable features in the environment at all. Using robot-to-robot observations greatly simplifies the data association problem, since the robots or the signals used to obtain the measurements can be specifically designed for easy detection. CL allows to abstract much of the perception problem which is the cause of many difficulties in SLAM. At the same time, CL creates significant challenges of its own, some of which will be addressed in this thesis.

1.1.1 Initialization

In order to fuse geometric observations and relative measurements, e.g., of distance or bearing between the robots, the measurements have to be expressed with respect to a common frame of reference. This requires either estimating all robot poses with respect to the same reference frame, or estimating the relative pose between the robots. However, since the system and measurement models in CL are often nonlinear, the performance (in terms of convergence, consistency, and precision) of any linearized estimator such as the EKF greatly depends on the accuracy of the initial relative pose estimate. The process for determining this relative transformation between any two sensors (or a sensor and a given reference frame) is also known as extrinsic calibration. Accurate extrinsic calibration is a prerequisite not only for CL, but for any spatially distributed geometric estimation task, e.g., cooperative mapping, tracking, etc.

Often extrinsic calibration is performed using external references, such as manual measurements, GPS, or matched features in overlapping sensor observations. GPS in particular renders extrinsic calibration straightforward, since the measurements directly provide the position of all team members in a common global frame of reference. However, there exist many interesting application scenarios for multi-robot systems (e.g., indoors, underwater, or underground) where external references are either unavailable (e.g., in GPS-denied or

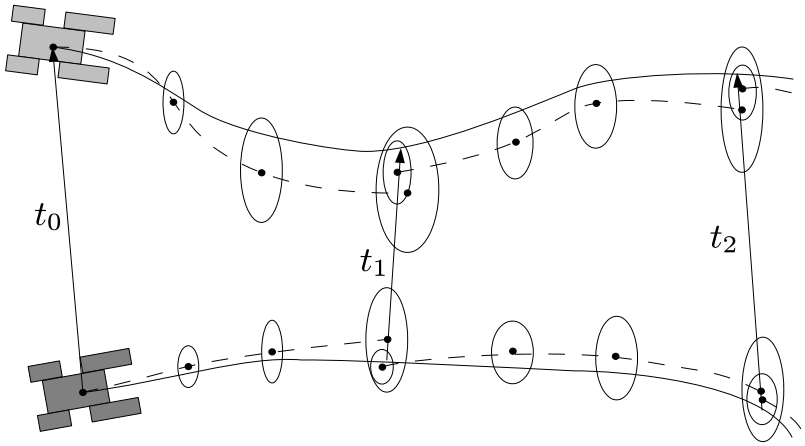

Figure 1.1: Schematic depiction of CL: Two robots localize using dead reckoning. Over time, their estimated poses (dashed lines) deviate from their true trajectories (solid lines), and their localization uncertainty, depicted by the 3σ error ellipses, grows. At times t_0, t_1, and t_2, the robots acquire observations, e.g., of their relative distance or bearing (arrows), which allow them to update their pose estimates and to reduce their localization uncertainty as shown by the smaller ellipses centered around their updated position estimates.

textureless areas), or do not yield results of sufficient accuracy. In these cases, alternative methods, referred to as motion-induced extrinsic calibration, can be used to determine the relative pose between robots based only on dead reckoning and robot-to-robot observations. More precisely, these methods seek to determine the relative pose by exploiting the geometric constraints arising from robot motion and interrobot measurements.

Motion-induced extrinsic calibration algorithms can be classified as minimal vs. overdetermined, and as deterministic vs. probabilistic. Roughly speaking, minimal problems have as many constraints as unknowns, and are often used in outlier rejection algorithms such as RANSAC [33]. Minimal problems described by systems of polynomials, such as the ones described in this thesis, may have multiple solutions. In contrast, overdetermined problems have more constraints (e.g., due to additional measurements) than necessary to solve the problem, which usually results in a unique solution and can sometimes be used to simplify the solution algorithm. Deterministic algorithms rely purely on system geometry, and do not consider measurement noise. If noise is present, overdetermined problems become inconsistent and no guarantees can be given for the accuracy of a deterministic, numerical solution. For this reason, it is preferable to use probabilistic algorithms (e.g., maximum likelihood estimators) that account for noise and uncertainty in a statistically sound fashion. Under certain assumptions, such as additive Gaussian measurement noise, these probabilistic estimators result in nonlinear weighted Least Squares (WLS) problems, which are typically solved iteratively, using methods like Gauss-Newton or Levenberg-Marquardt, initialized with the solution of a deterministic algorithm. However, any iterative, gradient-related optimization algorithm can at best provide guarantees about *local optimality* (more precisely, convergence to a stationary point), and can converge to a local minimum far from the true solution if poorly initialized. To overcome this problem, this thesis presents a new paradigm for motion-induced extrinsic calibration, where the the guaranteed *global optimum* of polynomial WLS problem formulations is computed analytically. Our results show that global optimality can be achieved at reasonable cost and provides considerable increase in accuracy over existing deterministic or stochastic approaches.

1.1.2 Communication

Once the initial transformation between any two robots has been determined, subsequent measurements can be processed in order to continuously track the current pose of all robots in the team, expressed with respect to a common frame of reference. For this purpose, a variety of estimation techniques have been proposed for CL, such as extended and unscented Kalman filters, particle filters, MAP estimators, etc. As previously discussed, CL treats localization as a joint estimation problem, entailing significant processing requirements whose reduction has been the subject of research over the past few years [53, 88, 83]. At the same time, joint estimation also imposes considerable communication requirements either between the robots themselves, or between the robots and a fusion center. Communication in CL

often encompasses exchange of individual motion measurements (i.e., linear and rotational velocity and acceleration), robot-to-robot relative distance and bearing measurements, and can be extended to also include GPS measurements, observations of point landmarks in the environment, or more comprehensive information for data association (e.g., camera images or laser scans).

In many cases, however, mobile sensor networks have to operate with only limited bandwidth that might preclude full communication of this large amount of data. Communication constraints can arise due to (i) the nature of the operational environment (e.g., limited bandwidth caused by strong signal attenuation and perturbation underwater or underground), (ii) the specific application domain (e.g., limited payload and power budgets, or stealth requirements in military operations), or (iii) economic motivations (e.g., in order to extend the operational life-span through lower power consumption resulting from reduced data transmissions). Exacerbating the problem of limited communication is the fact that the available communication resources will in general not be constant over time. For mobile robot applications, bandwidth will vary depending on a number of factors such as the distance between robots, the presence and nature of obstacles, available power, and current network load.

These consideration motivate the need for CL algorithms that can work efficiently under severe, time-varying communication constraints. Several approaches have been proposed to reduce communication in CL, e.g., based on approximate decoupling of the robot estimates, or based on sensor scheduling. In many of these cases, either correlations are not properly accounted for, or measurements are completely discarded, which may lead to suboptimal estimates. Instead, this thesis introduces a CL algorithm that reduces communication requirements through adaptive compression, i.e., quantization, of the transmitted measurements.[1] This approach requires novel estimation techniques, since standard methods such as the Kalman filter can only operate with real-valued measurements. Existing quantization techniques are mostly designed for signal reconstruction or parameter estimation, but not for estimation of dynamic random processes as in CL. Notable exceptions are the quantized filters of Ribeiro *et al.* [102] and Msechu *et al.* [81]. Inspired by these algorithms, this thesis develops novel batch MAP estimators for quantized measurements that are better suited to mitigate the strong nonlinearities arising not only from the robots' motion and measurement models, but also from quantization itself. Moreover, these algorithms allow using current surplus bandwidth to refine past quantized measurements and can thus make better use of time-varying bandwidth than filters.

[1] In most cases, real-valued measurements are available in digital representation, and are as such already "generically quantized", e.g., to 32 or 64 bits depending on the hardware. However, we are interested in "non-generic", adaptive quantization using $1 - N$ bits per measurement, optimized for the task of estimation. Hence, we will from now on refer to *real-valued* vs. *quantized* measurements when speaking about "generically" vs. task-specifically quantized measurements, respectively.

1.2 Research objectives

To summarize, this thesis investigates two of the key challenges in CL: relative robot pose estimation (initialization) and multi-robot pose tracking under communication constraints (bandwidth limitations).

1.2.1 Distance-based motion-induced extrinsic calibration

The first important contribution of this thesis is an algorithm for CL initialization in the absence of external references such as GPS. In particular, we formulate and solve the problem of estimating the relative transformation between two robots in 2-D and 3-D, using dead reckoning and robot-to-robot distance observations acquired during robot motion. By applying this algorithm pairwise to all members of the group, and selecting one of the robots' frames as the common frame of reference one can compute the team pose estimate required for CL initialization.

In contrast to traditional, only locally optimal approaches that first compute a coarse initial guess for the relative pose and then refine it using iterative nonlinear WLS, we develop and analyze algorithms that compute the guaranteed global optimum of the nonlinear WLS formulation. Specifically, we transform the WLS cost function into polynomial form, and then employ recent techniques from polynomial system solving and polynomial optimization to efficiently determine either all stationary points of the cost function, or directly its global minimum. We first analyze the WLS problem corresponding to the maximum likelihood estimate (MLE), show its poor scalability, and then provide relaxations that achieve virtually indistinguishable performance in near-real time.

In our analysis, we limit ourselves to the case of distance-only measurements between two robots, leaving other measurement modalities such as relative bearing or relative position as possible extensions for future work. Moreover, we assume that the robot-to-robot observations are outlier-free, and that all measurements are correctly time-stamped, i.e., that the robot clocks are synchronized. Finally, we assume that during initialization the robot localization errors are small, and discuss possible solutions when this assumption is violated.

Polynomial WLS problems arise often in robotics (kinematics) and computer vision, and the presented global optimization methods hold the promise of being applicable to a wide range of applications beyond the particular problem of CL initialization.

1.2.2 Cooperative Localization under communication constraints

The second important contribution of this thesis is an algorithm for CL that can adaptively and efficiently operate under severe, time-varying bandwidth limitations. More generally, we develop a batch MAP estimation algorithm for dynamic random processes using *quantized* measurements. In contrast to recently proposed minimum mean square error (MMSE) es-

timators (sign-of-innovation Kalman filter (SOI-KF) [102], or iteratively quantized Kalman filter (IQKF) [81]), which estimate the state at the current time step given all previous quantized measurements, we introduce novel maximum a posteriori (MAP) estimators, termed (single-bit) quantized MAP (QMAP) and iteratively quantized MAP (IQMAP), which provide state estimates over an entire time window. As in real-valued estimation, such batch estimation techniques enjoy certain advantages over filters in that (i) they act as a smoother, and can better mitigate linearization errors, and (ii) they can retroactively process past and time-delayed measurements. In the context of quantized estimation these advantages are very important, since quantization is a highly nonlinear process which violates the recurrent Gaussian approximation made by the SOI-KF and IQKF. Secondly, particularly in robotic applications, we are often faced with time-varying bandwidth availability. Batch estimators can efficiently make use of surplus bandwidth to refine the quantization of the most informative past measurements, whereas filters are limited to refining only the current observations. However, it is known from IQKF performance analysis that the accuracy gained by using additional bits to refine a single measurement quickly levels off, meaning that refining only the current observations becomes inefficient after some point. In this thesis, we introduce the QMAP and IQMAP formalisms, present optimal quantization thresholds for single- and multi-bit quantization of measurements, and analyze an efficient bit allocation scheme to address time-varying bandwidth. We then apply these algorithms to the problem of CL in simulation and experiments.

In this work, we operate under the assumption that a communication protocol is in place that ensures error-free, synchronous transmissions without packet loss throughout the network. Note that sufficiently powerful error-correction codes necessary to achieve this might require significant communication overhead. The reduced bandwidth requirements mentioned in this thesis refer to "goodput", i.e., the net amount of usable communication for the estimation task after subtracting the overhead. The estimation architecture itself assumes known system- and measurement models throughout the network, and operates in a "multi-centralized" fashion, so that every robot executes the exact same estimation algorithm using the same quantized measurements and quantization thresholds, resulting in an identical, common state estimate throughout the team.

Quantized estimation is very relevant in CL, since many of the most interesting application environments for multi-robot teams are subject to severe bandwidth limitations. Independently of the operating environment, modern robotic platforms are equipped with ever increasing numbers of sensors providing data at higher and higher rates and resolutions. Since it is unlikely that future communication links will be able to accommodate this growing information stream for too long, the use of adaptive, task-specific data compression approaches such as the ones presented in this thesis has the potential to significantly impact the operation of distributed robotics systems, by improving their ability to efficiently exchange information and adjust to varying resource availability.

1.3 Organization of the manuscript

Chapter 2 provides an overview of related work in CL, with particular focus on initialization and estimation using quantized measurements. Chapter 3 describes the distance-based relative pose estimation problem, and presents algorithms for global minimization of the WLS formulation and its relaxations. Their superior accuracy over traditional iterative minimization techniques is demonstrated through simulations and experiments. Chapter 4 describes the multi-centralized, quantized estimation framework for multi-robot CL under communication constraints. Results from simulation and experiments demonstrate and validate the performance and accuracy of MMSE and MAP estimators as a function of available bandwidth. Finally, chapter 5 provides concluding remarks and an outlook on future research directions.

Chapter 2

Related Work

2.1 Cooperative Localization

Robots that localize individually usually determine their pose by combining measurements from proprioceptive (motion) sensors (e.g., wheel encoders or inertial measurement units [IMUs]) and exteroceptive sensors such as sonars, laser scanners, or cameras. When multiple robots can communicate and measure relative distance and/or bearing to each other, they can combine these measurements and jointly estimate their poses, thus increasing the localization accuracy of the entire team [105, 79, 116]. This process is termed Cooperative Localization (CL).

Some CL algorithms determine the relative pose of the robot team at each time step independently, solely by combining all current (possibly redundant) exteroceptive measurements, e.g., via triangulation and LS [32, 109]. Under the assumption of bounded noise, this problem setup of CL has been solved by set intersection and cast as a convex optimization problem [110, 115]. These methods, often referred to as set membership approaches, have also been extended to cooperative SLAM [24]. Using only exteroceptive measurements to determine the pose of a robot team is closely related to the concept of static sensor network localization, since these methods make no use of measurements of the robots' motion. Dieudonné *et al.* [25] have shown that deterministic relative pose determination in a static 2-D sensor network based on relative distance and/or bearing information is NP-hard.

A more accurate way to solve CL is to consider relative measurements from multiple time steps as well as motion information from proprioceptive sensors in a filtering or batch estimation approach. One of the earliest forms of CL was presented in [64, 63, 61, 62], where only a subset of the robots moves while using the remaining ones as stationary landmarks in a triangulation-based localization scheme. Similar techniques for enhancing localization accuracy by careful coordination of robot motion were used in [36] and [99]. Later approaches removed the requirement for carefully coordinated robot motion, and applied different estimation techniques to address CL. Howard *et al.* [44], for example, combined proprioceptive and exteroceptive measurements of all robots and computed the MLE of the entire robot

team's pose history. In order to reduce the computational complexity ensuing from such batch estimators, recursive filtering has also been applied to CL [103, 73]. EKF-based CL in 2-D has recently been shown to be inconsistent [46], unless it is ensured that the estimator preserves the observability properties of the underlying nonlinear system (i.e., the three degrees of freedom corresponding to global translation and orientation are unobservable). If in addition a map of the environment is known, robots can perform global localization within the map, e.g., using Monte Carlo-based methods [22]. By exchanging their belief functions, cooperating robots can improve accuracy and speed of convergence of such algorithms [34]. As map-based localization is outside the scope of this work, we will not consider it further.

Despite the benefit of increased localization accuracy, CL is known to suffer from increased computational complexity ($O(N^4)$ for N robots observing each other) compared to individual localization. Several researchers have investigated methods to reduce this complexity, either by exactly distributing the computations among the robots [103, 83], or by approximately decoupling [45, 53, 88] or hierarchically restructuring the estimation problem [74].

Until now, however, little attention has been given to the problems of CL initialization and operation under communication constraints.

2.2 Distance-based motion-induced extrinsic calibration

The initialization problem of CL, in other words, estimating the pose of all team members with respect to a common frame of reference, is often assumed solved. Indeed, if GPS is available, this task is straightforward. Other techniques involve matching features in overlapping sensor observations. A common example is the use of laser scan matching to derive the relative pose between the origins of two scans [70, 87], e.g., using the iterative closest point (ICP) algorithm [13]. The relative pose can also be determined using camera images, for example from observations of matching points from two stereo cameras using Horn's method [42].

In this thesis, we focus on initialization methods that do not require external references, but instead rely on estimates of robot motion and robot-to-robot relative measurements. Our paper [118] gives an overview of deterministic algorithms for such motion-induced extrinsic calibration in 3-D, for the cases of distance and bearing measurements, bearing-only measurements, and distance-only measurements. In this thesis, we concentrate on the distance-only case, which is arguably the most difficult problem.

Previous research on leveraging sensor-to-sensor distance measurements to solve the relative localization problem has focussed primarily on *static* sensor networks. These approaches usually only determine the position of sensor nodes, not their relative orientation, using measurements to so-called anchor nodes with known global position. The position of the remaining sensors in the network can be uniquely determined if certain graph-rigidity

constraints are satisfied [4]. A number of algorithms for node localization using distance or squared distance measurements have been proposed, e.g., based on convex optimization [28, 10], centralized [108] or distributed [18] multidimensional scaling, or sum-of-squares relaxation [107, 84]. There also exist methods for simultaneous localization and mapping of a single mobile robot in 2-D using range measurements to static beacons [27], or to moving vehicles equipped with GPS, which are treated as moving beacons [7].

In 3-D, the (minimal) problem of extrinsic calibration based on six distance measurements is mathematically equivalent to the forward kinematics problem of the general Stewart-Gough mechanism [112]. This parallel manipulator consists of a base platform and a moving platform connected by six articulated legs of variable lengths. The forward kinematics problem is to determine the relative pose of the end effector with respect to the base, given the six leg lengths (distances) and the coordinates of the leg attachment points in the base and end effector frames, respectively. This deterministic problem has 40 (real and complex) solutions [97], which can be found based on nonlinear techniques from algebraic geometry [48, 67, 117]. In 2-D, similar analogies between robot pose estimation and the kinematics of planar linkage mechanisms can be drawn [131].

Recent work by Zhou and Roumeliotis [133] and by Trawny *et al.* [122, 117, 118] has addressed the problem of relative pose estimation based on distance measurements for two mobile robots moving in 2-D and 3-D, respectively, by providing deterministic and probabilistic methods for minimal and overdetermined problems. For the probabilistic approaches, the authors propose computing the MLE of the relative pose in a two-step process: First, they derive methods to compute a coarse estimate for the relative pose, based on solving a deterministic, overdetermined system of polynomial constraints using 4 or 5 distance measurements (2-D), or at least 10 distance measurements (3-D). In a second step, they use the result as an initial guess in i-WLS refinement.

Following a different paradigm, Shames *et al.* [107] present a probabilistic method designed for 2-D and 3-D network localization, which employs SOS relaxation to directly compute the global optimum of an (unweighted) LS cost function using squared distances, without the need for an initial guess. In an extension, they employ this strategy to indirectly infer the relative pose of two robots in 2-D by computing the positions of robot 2 expressed in robot 1's frame of reference. This directly yields the relative position, whereas the relative orientation has to be computed in a subsequent optimization step not further specified in the paper. Their proposed cost function contains components from n measurements of the squared distance between the robots, and $1/2n(n-1)$ distance constraints between the positions of robot 2 at different time steps as inferred from odometry.

Two-step methods, such as the ones proposed by [133, 122] can produce an initial guess that is far from the MLE. As a result, the i-WLS refinement may converge to a local minimum, or exceed the maximum number of iterations allowed to converge. In contrast, the approach of Shames *et al.* is a one-step method, but unfortunately, since the number

of unknowns grows with the number of measurements, it does not scale well. Moreover, this method effectively decouples position and orientation estimation, leading to suboptimal results. In chapter 3, we will investigate efficient, scalable algorithms to jointly estimate relative position and orientation by computing the global optimum of a polynomial LS cost function.

2.3 Estimation under communication constraints

2.3.1 Real-valued estimation

In estimation, we distinguish between classical estimation, where the object to estimate is considered unknown but deterministic and constant, and Bayesian estimation, where the estimated quantity is a random variable distributed according to some probability density function (pdf) called *prior*. Parameter estimation reasons about the pdf of the measurements *parameterized* by the unknown parameter, whereas Bayesian estimation uses the *posterior* pdf, i.e., the pdf of the estimated state *conditioned* on the measurements. Maximum Likelihood estimation (MLE) is a prominent parameter estimation technique, whereas minimum mean square error (MMSE) or maximum a posteriori (MAP) estimation are examples of Bayesian estimators [76, 54].

Bayesian estimation is particularly prevalent in estimation of random processes, in which the estimated quantities (the system state variables) change over time according to a specified dynamic system model. Estimation techniques for random processes are further categorized depending on the time-horizon of the measurements processed for estimating the state at a particular time instant (or over a certain time interval).

- Filters compute the best estimate of the current state at time step k based on all available measurements up to that time, by considering the posterior $p(\mathbf{x}_k|\mathbf{z}_{0:k})$. For linear problems with Gaussian noise, the optimal solution to the filtering problem is the well-known Kalman Filter [52]. For nonlinear problems, optimal filtering is intractable in general, and suboptimal approaches like the extended Kalman filter (EKF) [77], the unscented Kalman filter (UKF) [50], Gaussian Sum filters [2, 49], or particle filters (PF) [29] are used instead.

- Smoothers, on the other hand, compute the best estimate of the state at time step $k \leq K$ from the pdf $p(\mathbf{x}_k|\mathbf{z}_{0:K})$, i.e., based on all available measurements over the *entire* time horizon[1]. Smoothers are usually more accurate than filters, since intermediate states are estimated based on past, current, *and future* observations[2].

[1]The predictor, which is not considered in this work, is another type of estimator that estimates *future* states from past measurements, based on the pdf $p(\mathbf{x}_{k+l}|\mathbf{z}_{0:k})$, where $l > 1$.

[2]For example, for a time horizon of $K = 100$ time steps, the smoothed estimate of the state at time step $k = 10$ is based on measurements from time step $k = 1$ through time step $K = 100$, whereas the filter estimate of the state at time step $k = 10$ is only based on the measurements recorded between $k = 1$ and $k = 10$.

- Batch estimators only use past and current measurements, but estimate the *entire state history* concurrently, i.e., they reason about $p(\mathbf{x}_{0:K}|\mathbf{z}_{0:K})$. The MAP estimator[3] is a common representative of this class of estimators. Since batch estimators, as opposed to filters, estimate the entire state history (or a subset thereof), they can incorporate time-delayed measurements and retroactively correct data association errors. Their main benefit, however, is mitigation of linearization errors in nonlinear problems. For this reason, batch estimation is the method of choice in highly nonlinear computer vision problems like structure from motion or scene reconstruction [123]. Smoothing and batch LS have also found renewed interest in robotic applications (e.g., in CL [44, 83], and in SLAM for single [23, 51] and multiple robots [3, 56]).

A common assumption underlying the estimation techniques typically applied in robotics is that the measurements are *real-valued*. In our case, however, we seek to reduce communication requirements by quantizing the measurements.

2.3.2 Parameter estimation with quantized observations

If only *quantized* versions of the real-valued measurements are available, commonly used estimators (e.g., MLE or batch LS) are not applicable without modifications. Estimation with quantized observations has recently received growing interest in the research community, but compared to the wealth of research in real-valued estimation, our knowledge in this domain is still quite limited. Early work on distributed, quantized estimation was concerned mostly with estimation of scalar, deterministic parameters, based on noisy sensor observations that were quantized and transmitted to a fusion center.

Before we discuss the work on quantized parameter estimation, we want to emphasize that the design of an optimal quantizer (i.e., partitioning of the input space, design of the decoder, and choice of the distortion measure) is task-dependent [37]. In other words, quantizers for signal reconstruction, estimation, detection, or distributed processing will, in general, be different. For this reason, the well-known solution to signal reconstruction from quantized measurements, the Lloyd-Max quantizer [69, 75], is not necessarily the optimal solution to signal estimation.

To allow estimation of noisy parameters (noisy source coding) instead of pure signal reconstruction, the Lloyd-Max algorithm has been modified and generalized [31, 98]. The MSE distortion measure used in the original Lloyd-Max algorithm was replaced with the expected MSE of the estimated quantity, conditioned on the quantized, noisy observations. While in [98] the noisy measurements are encoded, transmitted, and used in an MMSE estimator at the decoder, it can be advantageous, under certain conditions, to first form an optimal estimate at the sensor, and then optimally quantize it using the usual MSE distortion metric [31].

[3] In the case of additive Gaussian noise, the MAP estimator can be written as a batch LS problem.

The above methods focus on estimation based on quantized observations from a *single source*. Lam and Reibman [65] extend the generalized Lloyd algorithm to estimation at a fusion center receiving quantized observations from *multiple sensors*. They investigate both MMSE and MAP distortion measures, and conclude that decoder (or estimator) design is more challenging for the MAP criterion, while the MMSE distortion entails a more difficult computation of partitions at the encoder; hence the choice between these two criteria depends on the task at hand. In either case, the optimal quantizer will require knowledge of the complete joint probability distribution of the parameter and all the sensor measurements. Gubner [39] alleviates this problem by constraining the structure of the estimator to find a *locally* optimal MMSE quantizer that requires only knowledge of the bivariate joint distribution of the parameter and the individual sensor measurements.

Papadopoulos *et al.* [89] introduce the concept of information loss as the ratio of the Cramer-Rao bounds achievable by quantized and real-valued estimators. In their analysis of a scalar signal corrupted by additive noise, they show that this loss is in general a function of the true parameter value, and that adding an extra signal strongly influences the shape of this function. Surprisingly, addition of a random signal (i.e., noise) with known covariance can greatly improve the worst-case information loss, while only slightly increasing the minimum information loss. This technique bears strong similarity to a process called nonsubtractive dithering [128] used in lossy compression of, e.g., audio or image data. In addition to random input, the authors evaluate the impact of adding a deterministic control feedback from the fusion center before quantization. Essentially, the signal estimate computed at the fusion center can be used to move the signal to be quantized to the region of minimal information loss.

Ribeiro and Giannakis [100] build on Papadopoulos' observations to study MLE of scalar parameters corrupted by Gaussian noise. In particular, they investigate optimal threshold design for different SNR scenarios, defined as the ratio of parameter dynamic range vs. measurement noise variance. Interestingly, they find that using carefully selected, different quantization thresholds at each sensor can generate the same effect as the control input proposed in [89]. Also, they show that the likelihood functions appearing in MLE for the Gaussian noise case are log-concave, a result we will find to carry over to our proposed MAP estimators (cf. Lemmas 1 and 3 in chapter 4.3).

The work of Luo [71, 72, 130] describes the design and analysis of algorithms for MMSE estimation of a quantized, scalar parameter corrupted by zero-mean noise with *unknown pdf* but known support. Initial work [72] where a pre-defined subset i of all nodes estimated the i-th most significant bit of the parameter, was later extended to an isotropic estimation scheme [71], in which each node executes the same estimation algorithm, and the choice of which significant bit to estimate was determined by the outcome of a random process. Finally, [130] analyzes even more general scenarios where the noise pdf is allowed to vary from sensor to sensor.

In [101], Ribeiro and Giannakis analyze the middle ground between quantized MLE of scalar parameters corrupted by noise with completely unknown pdf, noise with known pdf up to a finite set of parameters (e.g., the variance), and known noise distribution. Judicial choice of quantization thresholds allows estimation of the unknown parts of the noise pdf and leads to asymptotically efficient estimators with performance close to the unquantized estimator. The results of Luo [71, 72, 130] are thus extended in [101] in that the assumption of known support of the noise pdf is relaxed.

Finally, Shah *et al.* [106] analyze how to incorporate prior knowledge about the parameter within a MAP estimator based on binary observations. In particular, they consider linear functions of vector parameters corrupted by additive noise of known pdf. They establish that the posterior pdf is log-concave if the prior pdf and the cumulative noise distributions (cdfs) are log-concave. Vector observations corrupted by Gaussian noise are treated as a set of independent scalar observations, by appropriate pre-whitening of the measurements. We will use a similar approach for vector observations in multi-robot localization. For the case of Gaussian prior and noise, Shah *et al.* compute the approximate performance loss as the ratio of the corresponding diagonal elements of the Fisher information matrices of the quantized and real-valued estimators. We will extend this approach in our work to determine the performance loss of quantized MAP estimation for dynamic random processes.

Notice that all these approaches are designed to estimate a *constant* parameter, treated either as deterministic (e.g., [100, 71, 89]) or as a random variable (e.g., [106, 65, 39]). In the next section, we will consider estimating dynamic random processes, i.e., quantities that change over time. Estimation of dynamic quantities is of particular interest in robotic applications (e.g., CL [103], or tracking of a moving object [132]), and therefore the focus of our work.

2.3.3 Random process estimation with quantized observations

Despite the large number of publications on parameter estimation using quantized observations, only few approaches exist that use the same type of observations for estimating random processes.

Abdallah and Papadopoulos [1] extend their earlier work [89] on constant scalar estimation to estimating a time-varying scalar output of a linear random process with Gaussian system and measurement noise. They investigate the impact of control signals, noisy feedback, and channel corruption on the attainable estimation performance. The authors propose a two-tiered estimator, that first obtains a MAP estimate of the current signal based on the single-bit quantized sensor observations (without regard of the system dynamics), and then uses this estimate as a real-valued measurement update in a centralized EKF that considers the dynamics of the random process. However, this approach fails if the signal cannot be reconstructed completely from the quantized observations of a single time step. Instead, it would be desirable to have a filter that can directly incorporate quantized

observations.

This is exactly the subject of work by Ribeiro *et al.* [102] and Msechu *et al.* [81]. They develop a filtering scheme for scalar observations of a linear Gaussian random process, quantized with one (Sign-of-Innovation Kalman filter or SOI-KF) or multiple bits per measurement. For the multi-bit filtering approaches, they distinguish between batch-quantized (BQKF) [82] and iteratively-quantized (IQKF) [80] Kalman filters. Both single- and multi-bit approaches approximate the posterior pdf at each time step by a Gaussian, an approximation often employed in nonlinear estimation [58]. This approximation renders the estimation problem tractable, and, interestingly, results in a filter structure very similar to that of the regular Kalman filter. Even more surprising is the fact that the IQKF's performance (in terms of the trace of the covariance matrix) comes very close to that of the analog KF (i.e., the KF that processes real-valued, unquantized measurements), even if every measurement is quantized using only a few bits. In particular, the covariance reduction per measurement is 63.7% of that of the analog KF with only a single bit, and more than 98% with as little as 4 bits per measurement. This shows that quantization optimized for estimation can realize significant savings in communication with vanishingly small impact on estimation performance.

Despite these very promising results, the above quantized filtering methods still suffer from certain important limitations. First, even for linear systems with linear measurement models, quantization introduces nonlinearities. Moreover, realistic multi-robot applications that consider motion in 2-D or 3-D usually involve nonlinear models. Treating the posterior pdf as a Gaussian in this case is only approximate, and results in suboptimal estimation. As is the case for nonlinear estimation with real-valued measurements, we expect that *batch estimation* can mitigate errors resulting from linearization effects, and lead to improved estimation accuracy. Second, BQKF and IQKF cannot efficiently use time-varying bandwidth resources. Both filtering schemes discard the analog measurements after quantization. If additional bandwidth (i.e., bits) become available later on, they can only be used to refine the quantization of the current measurement. However, the performance analysis in [81] shows that the benefit of using more than 3 or 4 bits per measurement becomes marginal. Instead, the additional bits would be more useful if spent on *earlier* measurements to correct the state history, which is precisely what a batch algorithm allows. The development and evaluation of a quantized batch MAP estimator applied to the problem of CL is the subject of chapter 4 of this thesis.

2.3.4 Cooperative Localization with reduced communication

Several algorithms have been devised in an attempt to reduce computational complexity of CL. Unfortunately, the resulting savings in communication often come at the cost of optimality and estimation accuracy. We can broadly distinguish between different classes of approaches.

The first class reduces communication by communicating the condensed measurement information in form of state estimates and covariances, instead of communicating all measurements. Roumeliotis and Bekey [103] show that a centralized EKF estimator can be exactly distributed across the team. Each robot can propagate its own state separately, and information only needs to be broadcast across the network when robots meet and record relative measurements. Cross-correlations are fully accounted for by decomposing the corresponding submatrices of the covariance, and propagating the separate terms on each robot. Despite reduced computation and communication, this algorithm cannot cope with fluctuations in available bandwidth.

The following methods are approximate variants of the first class, that do not require broadcasting the measurements to all members of the team. Karam *et al.* [53] propose to reduce the communication load of a centralized EKF by replacing the exchange of measurements with intermittent exchange of state estimates and their covariances computed based on these measurements. The approach does not assume unique identification of other robots, and can handle intermittent communication with only a subset of the team. In their algorithm, each robot maintains a separate EKF-based pose estimate of the entire team, using proprioceptive sensors for egomotion estimation, and its own exteroceptive sensor observations (relative pose measurements) together with a tracking model to estimate the pose of the other vehicles. To achieve CL, all robots sporadically broadcast these estimates across the team. The states from other team members are then fused in a separate estimator (maintained by each robot) that identifies matching robot states in the different state vectors, and yields an improved group estimate by processing all available state estimates corresponding to the same robot as independent observations in a Kalman filter update. Without properly accounting for correlations arising from repeated fusion of the states and also from using a common tracking model, this algorithm can over time result in inconsistent estimates. Moreover, the method requires each robot to broadcast data quadratic in the number of observed robots per vehicle (state estimate plus covariance matrix), each time a fused estimate is desired. Hence, this method offers no means to adapt to variable bandwidth availability.

Bahr *et al.* [6] propose an algorithm where each robot maintains a bank of filters, each of which incorporating only the most recent relative measurements to a unique subset of other vehicles. The robots, assumed to be uniquely identifiable, periodically broadcast their set of own 3-D position estimates and corresponding covariance matrices, as well as bookkeeping information about the subset of vehicles whose measurements contributed to each estimate. This method is designed to explicitly avoid re-use of measurements and resulting inconsistency, but it has to discard a significant amount of information and is hence suboptimal.

Panzieri *et al.* [88] employ the *interlaced* EKF to perform CL without having to broadcast information to all robots. In their approach, each robot solves a decoupled estimation

problem, using the other robots as uncertain mapped landmarks. Specifically, when two robots are within measuring range, they exchange their current state estimates and covariances. In the relative measurement models, the robot positions are replaced by their respective estimates, and the measurement noise covariances are inflated by the uncertainty of the state estimates. A similar approach is proposed by Bahr *et al.* [7], where an autonomous underwater vehicle localizes using range measurements to GPS-equipped surface ships. Although such decoupling greatly reduces computational and communication requirements, it neglects correlations ensuing from robot-to-robot observations, which will eventually lead to inconsistent estimates.

The second class of methods reduce communication by restructuring the estimator. As an example, Martinelli [74] reduces communication and computation cost by introducing an estimation hierarchy. In particular, a team of robots is divided into subgroups, which perform CL only with their group leaders, who in turn perform CL among each other. This approach, however, is an approximation of the centralized EKF, since it neglects correlations arising across groups and can hence result in inconsistent estimates.

A third class of methods reduces communication by only selectively processing measurements. Mourikis and Roumeliotis [78], for example, consider CL for robot formations with limited communication bandwidth. The authors address this problem by optimally adjusting the measurement frequency of the different sensors within the given constraints. This is an example of optimal *sensor scheduling* to maximize steady-state localization accuracy under constrained communication. Measurements that do not contribute sufficiently to localization accuracy (as determined by the optimization process) are discarded; the remaining are transmitted as real-valued quantities. The proposed algorithm is only applicable to robots traveling in fixed formations, which renders the estimation problem approximately linear.

The above solutions attempt to reduce complexity and to decentralize the Cooperative Localization task, usually by transmitting less information (e.g., states and covariances instead of measurements, or only a subset of the measurements) to fewer nodes (e.g., only one-hop neighbors instead of broadcasting to the entire team). In many of these cases, either correlations are not properly accounted for (e.g., [88, 74]), or measurements have to be discarded (e.g., [78, 6]), resulting in suboptimal estimates compared to the centralized estimator that uses all measurements. Moreover, none of the existing approaches to CL is designed to efficiently exploit time-varying bandwidth resources.

In our work, we introduce a new approach to CL that reduces communication through compression, e.g., by means of quantization of sensor observations. If faced with communication restrictions, our algorithms will trade off bandwidth requirements for estimation accuracy, at the same time being able to refine the state estimate once more bandwidth becomes available. The details are described in chapter 4.

Note that the problem of asynchronous measurements and communication, i.e., not all

robots being able to communicate with all other robots at any given time, can be considered as another type of communication constraint, which, however, is not considered in this thesis. Only a few works address this problem, either exactly via the notion of check-points as in the work by Leung *et al.* [68], or approximately as in the work of Karam *et al.* [53] and Panzieri *et al.* [88]. Extending our proposed method to handle asynchronous communication would be an interesting subject of future work.

Chapter 3

Distance-based Motion-Induced Extrinsic Calibration

In this chapter, we describe a method for determining the relative transformation between two robots in the absence of GPS and other external references. In particular, we discuss how to estimate the relative pose between two robots based on robot-to-robot distance measurements and robot motion. By computing the pairwise relative pose between all robots in a team, this algorithm can be used to determine an estimate of the initial position and orientation of all robots with respect to a common reference frame, $\hat{\mathbf{x}}_{\text{ini}}$, and the corresponding covariance, $\mathbf{P}_0$. Having an accurate prior for the joint robot pose is an important prerequisite for CL, but is also necessary whenever data obtained by sensors at different locations are to be spatially correlated. Parts of this chapter have been published in a conference paper [120].

3.1 Introduction

Sensor fusion between multiple mobile robots critically relies on accurate extrinsic calibration, i.e., the knowledge of the robots' relative position and orientation (pose). In CL, knowledge of the relative pose is necessary to initialize the joint pose estimation problem. In other sensor fusion tasks, such as target tracking, geometric measurements have to be expressed with respect to a common frame of reference in order to allow matching and correlation, implying the need for knowing the relative pose. Common methods to determine the relative transformation between two robots make use of external references. Examples are direct manual measurements, absolute measurements with respect to a common frame (e.g., using GPS), or indirect correlation of overlapping sensor measurements (e.g., map or image matching). However, there exist scenarios, e.g., in GPS-denied or featureless areas, where these methods either provide insufficient accuracy or are even infeasible.

In the absence of external references, the robots can determine their relative pose based

on robot-to-robot sensor measurements and dead reckoning. Our recent paper [118] describes deterministic algorithms for computing the relative pose between pairs of robots based on motion and observations of robot-to-robot distance, bearing, or both. Arguably the most challenging case of such motion-induced extrinsic calibration occurs when the robots can only measure distance between each other. Loosely speaking, distance measurements offer the least amount of information in that they require the highest number of measurements from different positions (at least three in 2-D, six in 3-D) to determine a set of possible relative poses.[1] Mathematically, estimating relative pose from distance measurements requires solving systems of polynomial equations, which is to this day an active field of research. Robot-to-robot distance measurements can be acquired by various sensors such as sonar, radar, or laser, or indirectly as a function of the received communication signal (e.g., in terms of received signal strength or signal time of flight). Distance-based relative pose estimation in 2-D and 3-D will be the focus of this chapter (see Fig. 3.1).

Current approaches to this problem [133, 122] follow a two-step process. A deterministic algorithm is used to compute a coarse initial solution, followed by iterative WLS (i-WLS) refinement, e.g., using the Gauss-Newton algorithm [12]. Such a two-step procedure is a standard technique for solving nonlinear WLS problems. However, regardless of the iterative algorithm used (e.g., Gauss-Newton, Levenberg-Marquardt, Conjugate Gradient, etc.) or its numerical properties, algorithms using gradient-related search directions can only guarantee convergence to a stationary point, but cannot make any claims about global optimality (unless, of course, the cost function is convex). Instead, in this chapter we present a fundamentally different way to compute the maximum likelihood estimate (MLE), by solving the nonlinear WLS directly for the guaranteed *global* optimum. The approach is based on transforming the WLS cost function into polynomial form, and then employing recent techniques from polynomial system solving [16, 117, 60, 20] and polynomial optimization [92, 66, 126] to efficiently find either all stationary points of the cost function, or directly its global minimum.

The main contributions of this chapter are threefold:

1. We present the nonlinear WLS cost function for the relative pose problem, and a method to find its global minimum in 2-D and 3-D. Under the assumption of Gaussian noise, the optimal solution to the WLS will yield the MLE. We derive the corresponding first-order optimality conditions and show how to transform them into a system of polynomials, which can be solved for all stationary points, e.g., using homotopy continuation [124]. Our results show that this method is a feasible approach to find the MLE for a small number of measurements in 2-D, but scales poorly and is impractical in 3-D.

2. We present and analyze a relaxed WLS formulation based on squared distance mea-

[1] In general, i.e., for non-singular configurations, at least four (2-D) or seven (6-D) distance measurements are required in order to uniquely determine the relative pose.

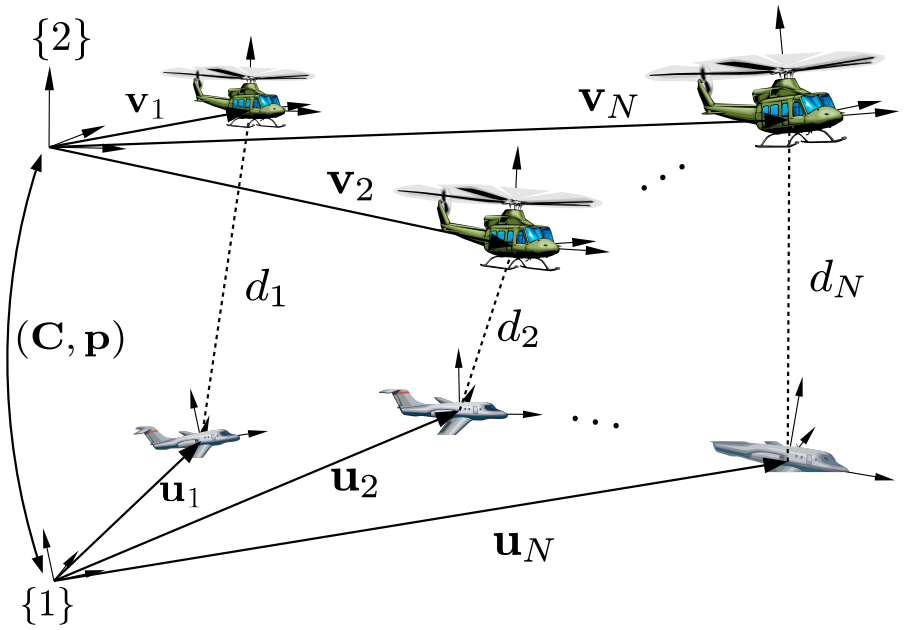

Figure 3.1: Robot 1 (airplane) and 2 (helicopter) take N distance measurements, $d_1, \ldots, d_N$, as they move. The robots know their positions $\mathbf{u}_i$, $\mathbf{v}_i$, in their respective global frames, $\{1\}$ and $\{2\}$, e.g., from odometry (2-D) or inertial navigation (3-D). The objective is to determine the translation, $\mathbf{p}$ and rotation, $\mathbf{C}$, between frames $\{1\}$ and $\{2\}$.

surements (SD-WLS). The resulting polynomial system is shown to have constant solution complexity, independent of the number of measurements.

3. We describe how to efficiently compute the global optimum of the SD-WLS, using a hybrid algebraic-numeric technique based on the eigendecomposition of a generalized companion matrix [117] in 2-D, and in 3-D using SOS relaxation [126]. In the latter case, we also present a fast technique for certifying global optimality of candidate stationary points that is based on Lagrange relaxation.

Results from simulations and experiments show that the SD-WLS estimator significantly outperforms the deterministic algorithms of [133] and [122], even if followed by i-WLS refinement, and yields performance almost indistinguishable from the original WLS estimator.

3.2 Geometric problem formulation

Consider two robots that acquire N robot-to-robot distance measurements $d_i,\ i = 1, \ldots, N$ while moving in 2-D or 3-D, as illustrated in Fig. 3.1. We assume that each robot can localize with respect to its own global frame of reference ($\{1\}$ and $\{2\}$), for example, using wheel odometry (2-D) or inertial measurements (3-D). Therefore, at the time the robots acquire the distance measurements, the coordinates of the first robot, $\mathbf{u}_i := {}^1\mathbf{u}_i$, and the second robot, $\mathbf{v}_i := {}^2\mathbf{v}_i$, are considered known. We further assume that the robots' clocks

are synchronized. The objective is to compute the 3 degree-of-freedom (2-D), or 6 degree-of-freedom (3-D) transformation between frames {1} and {2}, i.e., the translation vector $\mathbf{p} := {}^1\mathbf{p}_2$ and rotation matrix $\mathbf{C} := {}^1_2\mathbf{C}$.

The robot-to-robot distance d_i can be expressed as the length of vector $\mathbf{w}_i,\ i = 1, \ldots, N$, connecting the two robots at the time of measurement (see Fig. 3.1):

$$d_i = ||\mathbf{w}_i||_2 = \sqrt{\mathbf{w}_i^T \mathbf{w}_i}\,, \qquad \mathbf{w}_i := \mathbf{p} + \mathbf{C}\mathbf{v}_i - \mathbf{u}_i. \tag{3.1}$$

We parameterize rotation using a 4×1 unit quaternion

$$\bar{q} = \begin{bmatrix} q_1 & q_2 & q_3 & q_4 \end{bmatrix}^T = \begin{bmatrix} \mathbf{q}^T & q_4 \end{bmatrix}^T, \qquad \bar{q}^T \bar{q} = 1. \tag{3.2}$$

Employing the (non-Hamiltonian) convention for quaternions specified in [15], the rotation matrix $\mathbf{C}$ can be written as a nonsingular, quadratic expression in the elements of the quaternion $\bar{q}$

$$\mathbf{C}(\bar{q}) = (q_4^2 - \mathbf{q}^T\mathbf{q})\mathbf{I}_3 - 2q_4\lfloor \mathbf{q} \times \rfloor + 2\mathbf{q}\mathbf{q}^T \tag{3.3}$$

where $\mathbf{I}_3$ is the 3×3 identity matrix, and the symbol $\lfloor \mathbf{q} \times \rfloor$ denotes the skew-symmetric 3×3 matrix corresponding to the cross-product, so that

$$\mathbf{q} \times \mathbf{p} = \lfloor \mathbf{q} \times \rfloor \mathbf{p} = \begin{bmatrix} 0 & -q_3 & q_2 \\ q_3 & 0 & -q_1 \\ -q_2 & q_1 & 0 \end{bmatrix} \mathbf{p}. \tag{3.4}$$

More details about quaternion algebra can be found in the technical report [119].

In 2-D, the quaternion can be specified by $\bar{q} = \begin{bmatrix} 0 & 0 & \sin(\phi/2) & \cos(\phi/2) \end{bmatrix}^T$. For simplicity, we will instead (equivalently) define

$$\bar{q}_{(2-D)} := \begin{bmatrix} \sin\phi & \cos\phi \end{bmatrix}^T, \quad \bar{q}_{(2-D)}^T \bar{q}_{(2-D)} = 1 \tag{3.5}$$

$$\mathbf{C}(\bar{q})_{(2-D)} := \begin{bmatrix} \cos\phi & -\sin\phi \\ \sin\phi & \cos\phi \end{bmatrix} \tag{3.6}$$

and note that in the 2-D case, $\mathbf{C}(\bar{q})_{(2-D)}$ is linear in the elements of $\bar{q}_{(2-D)}$. To simplify notation, we will drop the subscript $(2-D)$ from now on as long as the correct expression is evident from context.

The approach of [133, 122] (and of all other deterministic formulations of this problem) is to assume noise-free measurements, and to stack the constraints provided by each observation into the following system of polynomial equations in the four (seven for 3-D) unknowns $\mathbf{p}$ and $\bar{q}$.

Deterministic System:

$$\mathbf{w}_i^T \mathbf{w}_i - d_i^2 = 0, \quad i = 1, \ldots, N \tag{3.7a}$$

$$\bar{q}^T \bar{q} - 1 = 0. \tag{3.7b}$$

In 2-D, Zhou and Roumeliotis [133] show that $N = 3$ measurements are necessary for this problem to have a discrete set of six possibly complex solutions (minimal problem). Moreover, $N > 3$ measurements uniquely determine the relative pose. Further, they present a linear algorithm to solve the overdetermined problem with $N = 5$ measurements for the unique solution.

For the 3-D case, we know from research on the forward kinematics of the general Stewart-Gough mechanism (e.g., [127, 48, 117]) that $N = 6$ measurements are necessary for this problem to have a discrete set of 40 possibly complex solutions[2] (minimal problem). In our recent publications [122, 118], we show that $N > 6$ measurements uniquely determine the relative pose. We also provide a linear algorithm to solve the overdetermined polynomial system using $N \geq 10$ measurements.

In practice, the robot-to-robot distance measurements are corrupted by noise

$$z_i = d_i + n_i, \quad i = 1, \ldots, N \tag{3.8}$$

$$\mathbf{n} = \begin{bmatrix} n_1 & \ldots & n_N \end{bmatrix}^T \sim \mathcal{N}(\mathbf{n}; \mathbf{0}, \mathbf{R}), \tag{3.9}$$

where we have modeled the measurement noise $\mathbf{n}$ as zero-mean Gaussian, with covariance matrix $\mathbf{R}$. In case of independent measurements, $\mathbf{R} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_\mathrm{N}^2)$, but in the presence of correlated noise, $\mathbf{R}$ can be full. Replacing d_i by the noisy measurements z_i, the overdetermined system (3.7) with $N > 3$ (2-D), or $N > 6$ (3-D), will be inconsistent, and no guarantees can be given for the solutions computed by the algorithms in [133] and [122]. In particular, these algorithms cannot account for the measurement noise in a statistically sound fashion. Therefore, their results are usually used as an initial guess for a subsequent i-WLS refinement [133, 122]. However, as shown in the simulation results in chapter 3.6, i-WLS refinement with poor initialization can result in estimates far from the global optimum.

As an alternative, we propose to combine the constraints (3.7a) arising from the measurements in a WLS fashion, thus correctly accounting for the measurement noise. In what follows, we present methods to solve the WLS problem directly for the *globally* optimal estimate of the relative pose. The approach is to formulate the first-order optimality conditions of the WLS cost function as a square system of polynomials, which can be solved directly to obtain all stationary points. A desirable property of this formulation is that without any modification it will yield the global optimum for the overdetermined system ($N > 3$,

[2]Since in 3-D $\mathbf{q}$ and $-\mathbf{q}$ represent the same rotation, system (3.7) has in fact 80 numerical, but only 40 physically distinct solutions.

or $N > 6$), as well as all solutions to the minimal problem (6 solutions for $N = 3$ in 2-D, or 40 solutions for $N = 6$ in 3-D), which will be global minima of the WLS cost function with cost function value identically equal to zero.

3.3 Weighted Least Squares (WLS)

Ideally, given (3.1) and (3.8), we would like to compute the MLE by solving the following nonlinear WLS problem.

WLS Cost Function:

$$\begin{aligned} &\min_{\mathbf{p},\bar{q}} \ \frac{1}{2}\,(\mathbf{h}_d - \mathbf{z})^T \mathbf{R}^{-1}(\mathbf{h}_d - \mathbf{z}) \\ &\text{s.t. } \bar{q}^T\bar{q} = 1 \end{aligned} \tag{3.10}$$

where the (distance) measurement function $\mathbf{h}_d$ is given by[3]

$$\mathbf{h}_d = \begin{bmatrix} \sqrt{\mathbf{w}_1^T\mathbf{w}_1} & \dots & \sqrt{\mathbf{w}_N^T\mathbf{w}_N} \end{bmatrix}^T. \tag{3.11}$$

Using the method of Lagrange multipliers [12], we write the Lagrangian

$$\mathcal{L}_d = \frac{1}{2}\,(\mathbf{h}_d - \mathbf{z})^T \mathbf{R}^{-1}(\mathbf{h}_d - \mathbf{z}) + \lambda(\bar{q}^T\bar{q} - 1) \tag{3.12}$$

and the corresponding KKT optimality conditions as

$$\nabla_{\mathbf{p}}\mathcal{L}_d = \begin{bmatrix} \frac{\mathbf{w}_1}{\sqrt{\mathbf{w}_1^T\mathbf{w}_1}} & \cdots & \frac{\mathbf{w}_N}{\sqrt{\mathbf{w}_N^T\mathbf{w}_N}} \end{bmatrix} \mathbf{R}^{-1}(\mathbf{h}_d - \mathbf{z}) = \mathbf{0} \tag{3.13a}$$

$$\nabla_{\bar{q}}\mathcal{L}_d = \begin{bmatrix} \frac{\nabla_{\bar{q}}(\mathbf{w}_1^T\mathbf{w}_1)}{2\sqrt{\mathbf{w}_1^T\mathbf{w}_1}} & \cdots & \frac{\nabla_{\bar{q}}(\mathbf{w}_N^T\mathbf{w}_N)}{2\sqrt{\mathbf{w}_N^T\mathbf{w}_N}} \end{bmatrix} \mathbf{R}^{-1}(\mathbf{h}_d - \mathbf{z}) + 2\lambda\bar{q} = \mathbf{0} \tag{3.13b}$$

$$\nabla_{\lambda}\mathcal{L}_d = \bar{q}^T\bar{q} - 1 = 0 \tag{3.13c}$$

where

$$\begin{aligned} \nabla_{\bar{q}}(\mathbf{w}_i^T\mathbf{w}_i) &= 2\nabla_{\bar{q}}\left((\mathbf{p} - \mathbf{u}_i)^T\mathbf{C}\mathbf{v}_i\right) \\ &= 2(\nabla_{\bar{q}}\mathrm{vec}(\mathbf{C}))\ \mathrm{vec}\left((\mathbf{p} - \mathbf{u}_\mathrm{i})\mathbf{v}_\mathrm{i}^\mathrm{T}\right) \end{aligned} \tag{3.14}$$

[3] In what follows, different subscripts for the measurement function will indicate different formulations and relaxations of the problem.

and

$$\nabla_{\bar{q}}\mathrm{vec}(\mathbf{C}_{(2-\mathrm{D})}) = \begin{bmatrix} 0 & 1 & -1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix} \tag{3.15}$$

$$\nabla_{\bar{q}}\mathrm{vec}(\mathbf{C}_{(3-\mathrm{D})}) = 2\begin{bmatrix} q_1 & q_2 & q_3 & q_2 & -q_1 & q_4 & q_3 & -q_4 & -q_1 \\ -q_2 & q_1 & -q_4 & q_1 & q_2 & q_3 & q_4 & q_3 & -q_2 \\ q_3 & q_4 & q_1 & -q_4 & -q_3 & q_2 & q_1 & q_2 & q_3 \\ q_4 & q_3 & -q_2 & -q_3 & q_4 & q_1 & q_2 & -q_1 & q_4 \end{bmatrix}. \tag{3.16}$$

Here, vec(.) denotes the column-wise vectorization operation.

In order to solve this system of equations and thus to determine all stationary points, we transform it into a system of polynomials by introducing auxiliary variables. Our objective is then to apply efficient polynomial system solving techniques that can find all roots simultaneously. In particular, choosing auxiliary variables

$$a_i := \sqrt{\mathbf{w}_i^T \mathbf{w}_i}, \quad i = 1, \ldots, N \tag{3.17}$$

we have a new measurement function (distance as a function of the auxiliary variable $\mathbf{a}$) $\mathbf{h}_{ad}$

$$\mathbf{h}_{ad} = \begin{bmatrix} a_1 & \ldots & a_N \end{bmatrix}^T \tag{3.18}$$

and we obtain the Lagrangian

$$\mathcal{L}_{ad} = \frac{1}{2}(\mathbf{h}_{ad} - \mathbf{z})^T \mathbf{R}^{-1}(\mathbf{h}_{ad} - \mathbf{z}) + \sum_{i=1}^{N} \mu_i(\mathbf{w}_i^T \mathbf{w}_i - a_i^2) + \lambda(\bar{q}^T \bar{q} - 1). \tag{3.19}$$

Taking the gradient with respect to the unknowns, we obtain the following system of $2N + 5$ ($2N + 8$ in 3-D) polynomials in the $2N + 5$ ($2N + 8$ in 3-D) unknowns $\mathbf{p}, \bar{q}, \lambda$, and $\mathbf{a} := \begin{bmatrix} a_1 & \ldots & a_N \end{bmatrix}^T$, $\boldsymbol{\mu} := \begin{bmatrix} \mu_1 & \ldots & \mu_N \end{bmatrix}^T$.

WLS First-Order Optimality Conditions:

$$\nabla_{\mathbf{p}} \mathcal{L}_{ad} = 2 \sum_{i=1}^{N} \mu_i \mathbf{w}_i = \mathbf{0} \tag{3.20a}$$

$$\nabla_{\bar{q}} \mathcal{L}_{ad} = \sum_{i=1}^{N} \mu_i \nabla_{\bar{q}} (\mathbf{w}_i^T \mathbf{w}_i) + 2\lambda \bar{q} = \mathbf{0} \tag{3.20b}$$

$$\nabla_{\mathbf{a}} \mathcal{L}_{ad} = \mathbf{R}^{-1}(\mathbf{h}_{ad} - \mathbf{z}) - 2 \begin{bmatrix} a_1 \mu_1 \\ \vdots \\ a_N \mu_N \end{bmatrix} = \mathbf{0} \tag{3.20c}$$

$$\nabla_{\mu_i} \mathcal{L}_{ad} = \mathbf{w}_i^T \mathbf{w}_i - a_i^2 = 0, \quad i = 1, \ldots, N \tag{3.20d}$$

$$\nabla_{\lambda} \mathcal{L}_{ad} = \bar{q}^T \bar{q} - 1 = 0. \tag{3.20e}$$

In 2-D, all polynomials are quadratic; in 3-D, (3.20a), (3.20b), and (3.20d) become cubic.

Notice that squaring the constraints on the a_i (cf. (3.17) vs. (3.20d)) may introduce spurious solutions with erroneous signs. Therefore, we will only accept real solutions that fulfill the original constraint (3.17) as candidate stationary points.

Once all the stationary points of (3.10) are found by directly solving system (3.20), the final step is to evaluate the cost function on the candidate points, and to choose the one with minimum value as the final, guaranteed globally optimal WLS estimate of the relative pose.

The numerical solution of (3.20) can be obtained, for example, using homotopy continuation [124]. In 2-D, system (3.20) has 70 solutions for $N = 3$ measurements, 240 solutions for $N = 4$, and 784 solutions for $N = 5$, as determined both through analysis of the Gröbner basis using Macaulay 2 [38] and numerically using PHCpack [124]. In our experiments, we solved example systems with $N = 4$ measurements in about 2 to 4 minutes (diagonal vs. dense covariance matrix), and $N = 5$ measurements in 12 to 20 minutes, using PHCpack on an Intel T9400 2.53 GHz laptop with 2GB of RAM. The reason for the increase in time is the introduction of additional monomials in the polynomial system due to cross-coupling induced by the off-diagonals of $\mathbf{R}$. In 3-D, unfortunately, it turns out that this process is computationally impracticable, as homotopy-based solutions require days of runtime, even for the minimal case of $N = 6$ measurements.

As evident from these results, this formulation scales quite poorly, because the number of unknowns (i.e., a_i, μ_i) grows with the number of measurements. The complexity of solving polynomial systems is exponential in the number of unknowns [30], rendering this approach feasible only for small-scale problems. This leads us to the introduction of an alternative formulation using squared distance measurements.

3.4 Squared distances WLS (SD-WLS)

A much preferable approach from the point of scalability is to solve the WLS based on noisy measurements of the *squared* distance

$$z'_i = d_i^2 + n'_i, \quad i = 1, \dots, N \tag{3.21}$$

$$\mathbf{n}' = \begin{bmatrix} n'_1 & \dots & n'_N \end{bmatrix} \sim \mathcal{N}(\mathbf{n}'; \mathbf{0}, \mathbf{R}') \tag{3.22}$$

where for simplicity we temporarily assume that the noise in the squared measurements is zero-mean Gaussian. Under this assumption, the SD-WLS will actually yield the MLE. We will relax this assumption in chapter 3.4.1.

The SD-WLS cost function is defined as

SD-WLS Cost Function:

$$\begin{aligned} &\min_{\mathbf{p},\bar{q}} \ \frac{1}{2} (\mathbf{h}_{sd} - \mathbf{z}')^T \mathbf{R}'^{-1} (\mathbf{h}_{sd} - \mathbf{z}') \\ &\text{s.t. } \bar{q}^T \bar{q} = 1 \end{aligned} \tag{3.23}$$

where

$$\mathbf{h}_{sd} := \begin{bmatrix} \mathbf{w}_1^T \mathbf{w}_1 & \dots & \mathbf{w}_N^T \mathbf{w}_N \end{bmatrix}^T. \tag{3.24}$$

Forming the Lagrangian $\mathcal{L}_{sd}$ analogously to (3.12), setting its gradient to zero and dividing by constant factors, we obtain the following polynomial system of 5 (8 in 3-D) equations in the 5 (8 in 3-D) unknowns $\mathbf{p}, \bar{q}$ and λ.

SD-WLS First-Order Optimality Conditions:

$$\nabla_{\mathbf{p}} \mathcal{L}_{sd} = \begin{bmatrix} \mathbf{w}_1 & \mathbf{w}_2 & \dots & \mathbf{w}_N \end{bmatrix} \mathbf{R}'^{-1} (\mathbf{h}_{sd} - \mathbf{z}') = \mathbf{0} \tag{3.25a}$$

$$\nabla_{\bar{q}} \mathcal{L}_{sd} = \begin{bmatrix} \nabla_{\bar{q}} (\mathbf{w}_1^T \mathbf{w}_1) & \dots & \nabla_{\bar{q}} (\mathbf{w}_N^T \mathbf{w}_N) \end{bmatrix} \mathbf{R}'^{-1} (\mathbf{h}_{sd} - \mathbf{z}') + 2\lambda \bar{q} = \mathbf{0} \tag{3.25b}$$

$$\nabla_{\lambda} \mathcal{L}_{sd} = \bar{q}^T \bar{q} - 1 = 0. \tag{3.25c}$$

In 2-D, the polynomials (3.25a) and (3.25b) are cubic[4], while in 3-D, they are quintic. The quaternion unit-norm constraint (3.25c) is quadratic in both cases.

Analysis of this system's Gröbner basis for different numerical instances using coefficients from a finite field shows that in 2-D, this system has 28, and in 3-D 1424 (possibly complex) solutions. This is a general result that holds for all cases of N and $\mathbf{R}'$. In particular, we have the following proposition:

[4]In practice, for the 2-D cases of both WLS and SD-WLS, we immediately eliminate λ and one equation by pre-multiplying (3.20b) and (3.25b) by $\begin{bmatrix} \cos\phi & -\sin\phi \end{bmatrix}$, thus increasing their degrees by one. In contrast, in 3-D similar elimination of λ requires six new equations, which would result in a non-square system of equations and is therefore avoided.

Proposition 1. *The number of solutions of the system* (3.25)*, and its solution complexity, is independent of the covariance matrix,* $\mathbf{R}'$*, and of the number of measurements,* N*, for* $N \geq 3$ *(*$N \geq 6$ *in 3-D).*

Proof. We base our proof on the following observation: If for a system of polynomials arising from a specific problem class, different (non-singular) instantiations of the problem differ only in the numeric coefficients but not in the structure of the polynomial system (i.e., the monomials comprising each polynomial remain the same), the leading monomials of the corresponding Gröbner basis will generally also be the same, with the consequence that the standard basis of the quotient ring and the number of solutions of the system is constant [19, 60]. We therefore have to show that the structure of system (3.25) and the monomials in each polynomial (3.25a)-(3.25c) are independent of N and $\mathbf{R}'$ generically. Clearly (but in contrast to the polynomial system (3.20)), the number of equations and the number of unknowns of (3.25) is constant and independent of N and $\mathbf{R}'$. Also, (3.25c) remains unchanged for different problem instantiations. To see that the monomials in (3.25a) and (3.25b) are independent of N and $\mathbf{R}'$, consider rewriting these equations as $\sum_{i,j=1}^{N} r'^{-1}_{i,j}(\mathbf{w}_i^T\mathbf{w}_i - z'_i)\mathbf{w}_j^T$ and $\sum_{i,j=1}^{N} r'^{-1}_{i,j}(\mathbf{w}_i^T\mathbf{w}_i - z'_i)\nabla_{\bar{q}}(\mathbf{w}_j^T\mathbf{w}_j)$, where $r'^{-1}_{i,j}$ denotes the (i,j)-th entry of $\mathbf{R}'^{-1}$. Generically, each polynomial $\mathbf{w}_i^T\mathbf{w}_i - z'_i$, $\mathbf{w}_i$, and $\nabla_{\bar{q}}(\mathbf{w}_i^T\mathbf{w}_i)$ contains the same monomials for every i. Analogously, each of the products $(\mathbf{w}_i^T\mathbf{w}_i - z'_i)\mathbf{w}_j^T$ and $(\mathbf{w}_i^T\mathbf{w}_i - z'_i)(\nabla_{\bar{q}}(\mathbf{w}_j^T\mathbf{w}_j))$ contains the same monomials generically, regardless of i and j. Finally, linear combinations of polynomials do not introduce new monomials, which concludes the proof. □

3.4.1 Gaussian approximation of squared noisy measurements

As discussed in the beginning of the section, the noise model (3.21) of SD-WLS is only an approximation if in fact only noisy measurements of the (non-squared) robot-to-robot distances are available. In particular, replacing (3.8) by (3.21) cannot be achieved by simply squaring the noisy measurements, i.e., $z'_i \neq z_i^2$. The reason is that $z_i^2 = d_i^2 + 2d_in_i + n_i^2$, and the corresponding noise term $s_i := 2d_in_i + n_i^2$ is *not* zero-mean Gaussian. Indeed, following the standard formulas to compute the pdf of functions of random variables [111], the pdf of the vector $\boldsymbol{\zeta} := \begin{bmatrix}\zeta_1 & \dots & \zeta_N\end{bmatrix}$ with $\zeta_i = z_i^2$ is given by

$$p(\boldsymbol{\zeta}) = \sum_{j=1}^{2^N} \frac{1}{2^N \prod_i \sqrt{\zeta_i}} \mathcal{N}(\boldsymbol{\gamma}_j; \mathbf{d}, \mathbf{R}) \tag{3.26}$$

where each $\boldsymbol{\gamma}_j,\ j = 1, \dots, 2^N$, is a vector of the form $\begin{bmatrix}\pm\sqrt{\zeta_1} & \dots & \pm\sqrt{\zeta_N}\end{bmatrix}^T$ with one of the 2^N possible different sign assignments for its individual elements, and $\mathbf{d} = \begin{bmatrix}d_1 & \dots & d_N\end{bmatrix}$.

However, the non-Gaussian pdf resulting from squaring a Gaussian random variable can be well approximated by a Gaussian pdf with matching first and second order moments.

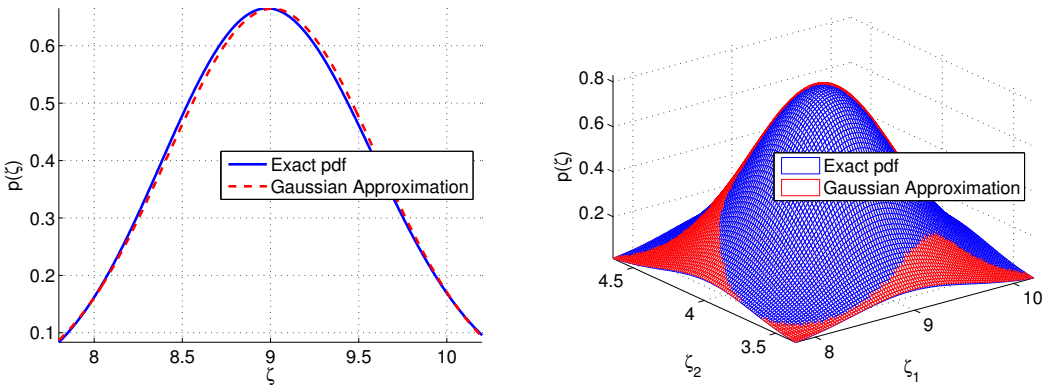

(a) (b)

Figure 3.2: Examples of approximating the pdf of squared distance measurements by a Gaussian. (a) Scalar measurement case, with $d_1 = 3\text{m}$ and $\sigma_1 = 0.1\text{m}$. (b) Two distance measurements with $\mathbf{d} = \begin{bmatrix}3\text{m} & 2\text{m}\end{bmatrix}^T$, and correlated noise with covariance $\mathbf{R} = \begin{bmatrix}(0.1\text{m})^2 & (0.01\text{m})^2 \\ (0.01\text{m})^2 & (0.08\text{m})^2\end{bmatrix}$. For small variances, the approximation is reasonably accurate: for the examples shown here, the KL divergence is $8.34 \cdot 10^{-4}$ and $2.0 \cdot 10^{-3}$, respectively.

Specifically, computing the mean, $\bar{s}_i$, and covariance, $\mathbf{\Sigma}$, of s_i yields

$$\bar{s}_i := E[s_i] = \mathbf{R}_{ii} \tag{3.27}$$

$$\mathbf{\Sigma}_{ii} := E[(s_i - \bar{s}_i)^2] = \mathbf{R}_{ii}(4d_i^2 + 2\mathbf{R}_{ii}) \tag{3.28}$$

$$\mathbf{\Sigma}_{ij} := E[(s_i - \bar{s}_i)(s_j - \bar{s}_j)] = \mathbf{R}_{ij}(4d_i d_j + 2\mathbf{R}_{ij})$$

We can now approximate (3.21) by setting

$$z_i' \simeq z_i^2 - \bar{s}_i \tag{3.29}$$

$$\mathbf{R}' \simeq \mathbf{\Sigma} \tag{3.30}$$

where we replace the distances d_i in the expression for $\mathbf{\Sigma}$ by their noisy measurements z_i.

As illustrated in Fig. 3.2, using (3.29), (3.30) to approximate the non-Gaussian pdf (3.26) by a Gaussian is reasonably accurate, particularly for high signal-to-noise ratios.

3.4.2 Accounting for uncertainty in robot positions

Our current noise models for WLS (3.8), and for SD-WLS (3.21) assume perfectly known robot positions $\mathbf{u}_i$, $\mathbf{v}_i$. If these positions are estimated using odometry or inertial measurements, this assumption will be violated. There are several possible approaches to address this issue:

- Inflate the measurement noise covariance matrix, $\mathbf{R}$. By doing so, uncertainty in the robots' positions is treated as an additional measurement error. Although this

method is only an approximation, its main advantage is that the estimate can still be obtained in a single step. It works particularly well when the position uncertainty is small and independent of the robot position itself, and was therefore used in our 3-D experiment, see chapter 3.7.2.

- Compute the Jacobians of the measurement function $\mathbf{h}$ with respect to the robots' positions, $\mathbf{H}_u := \frac{\partial \mathbf{h}}{\partial \mathbf{u}}$, and $\mathbf{H}_v := \frac{\partial \mathbf{h}}{\partial \mathbf{v}}$. This corresponds to the following linearization of the measurement model

$$\mathbf{z} = \mathbf{h}(\mathbf{p}, \bar{q}, \mathbf{u}, \mathbf{v}) + \mathbf{n} \simeq \mathbf{h}(\mathbf{p}, \bar{q}, \hat{\mathbf{u}}, \hat{\mathbf{v}}) + \mathbf{H}_u \tilde{\mathbf{u}} + \mathbf{H}_v \tilde{\mathbf{v}} + \mathbf{n} \tag{3.31}$$

  where $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{v}}$ denote the error in the robot positions $\mathbf{u}$ and $\mathbf{v}$, i.e., $\mathbf{u} \simeq \hat{\mathbf{u}} + \tilde{\mathbf{u}}$, and $\mathbf{v} \simeq \hat{\mathbf{v}} + \tilde{\mathbf{v}}$. Compute an augmented noise matrix $\tilde{\mathbf{R}} = \mathbf{R} + \mathbf{H}_u \mathbf{P}_{uu} \mathbf{H}_u^T + \mathbf{H}_v \mathbf{P}_{vv} \mathbf{H}_v^T$, where $\mathbf{P}_{uu}$ and $\mathbf{P}_{vv}$ denote the covariance of $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{v}}$, and use $\tilde{\mathbf{R}}$ as weighting matrix in the WLS cost function (see [122, 118] for details). This results in an iterative estimation process, since the weighting matrix will depend on the linearization points used for the Jacobians $\mathbf{H}_u$ and $\mathbf{H}_v$, e.g., $\mathbf{H}_u = \mathbf{H}_u(\hat{\mathbf{p}}, \hat{\bar{q}}, \hat{\mathbf{u}}, \hat{\mathbf{v}})$. The initial linearization point is obtained by solving the SD-WLS problem without accounting for position uncertainty. The globally optimal estimate of the relative pose computed for a given $\tilde{\mathbf{R}}$ is used to compute the next linearization point. Despite losing all theoretical claims of convergence and optimality, we find that in practice, this process converges after very few iterations and in most cases leads to consistent estimates. This method was employed in our 2-D experiment, see chapter 3.7.1.

- Augment the state vector to also include the robot positions, $\mathbf{u}$ and $\mathbf{v}$, and augment the measurement vector $\mathbf{z}$ with the odometry measurements (or any other measurements used for robot localization). Compute the MLE batch estimate using bundle adjustment [123, 56]. This is arguably the optimal treatment of position uncertainty, but given the large number of unknowns in the state vector, is not amenable to the global optimization procedures outlined in chapter 3.5.

In what follows, we will concentrate on the case where the robot positions are either perfectly known, or the measurement covariance has been appropriately inflated to account for each robot's localization uncertainty. Further, we assume that the measurements are outlier-free. If outliers are present, they can be removed using RANSAC [33] in combination with minimal problem solvers [133, 117].

3.4.3 Initializing CL

Up to now, we have focused on how to compute the initial relative pose between the robots. In order to use the results from the preceding methods to initialize the estimation process in CL (that tracks the pose of the robot team as it moves), we need to compute an estimate of

the robot team pose in a common frame of reference, as well as a corresponding covariance matrix.

For this purpose, we propose to use the relative pose estimated using the SD-WLS algorithm as an initial guess for a bundle adjustment scheme similarly to the one discussed in chapter 3.4.2 above, that fuses all odometric and relative observations, and computes the robot poses at all times a measurement was obtained in a common frame, for example, that of robot 1. It also provides a corresponding covariance in form of the inverse Hessian. Subsequently, we extract the estimates corresponding to the pose of the robot team at the last time instant, as well as their covariance, and use this as the Gaussian prior $\mathbf{x}_{\text{ini}}, \mathbf{P}_0$ commonly used in CL (see also chapter 4).

3.5 Solving the SD-WLS problem

In this section, we will address different algorithms to efficiently solve the minimization problem (3.23), or equivalently its KKT optimality conditions (3.25). Before, we will have to introduce some additional notation. We define a monomial in n variables, $\mathbf{x} = \begin{bmatrix} x_1 & \dots & x_n \end{bmatrix}^T$, as $\mathbf{x}^{\boldsymbol{\alpha}} := x_1^{\alpha_1} \cdot \ldots \cdot x_n^{\alpha_n}$, $\alpha_i \in \mathbb{N} \cup \{0\}$. The degree of a monomial is given by $\sum_i \alpha_i$. A real polynomial is a linear combination of monomials, $q(\mathbf{x}) = \sum_i q_i \mathbf{x}^{\boldsymbol{\alpha}_i}$, with coefficients $q_i \in \mathbb{R}$. The degree of a polynomial is given by the maximum degree of its monomials. Consider the vector $\mathbf{y}^{(d)}$ of monomials in n variables of degree up to d with length $\binom{n+d}{d}$

$$\mathbf{y}^{(d)} = \begin{bmatrix} 1 & x_1 & \dots & x_n & \dots & x_1^d & \dots & x_n^d \end{bmatrix}^T. \tag{3.32}$$

We can write any real polynomial of total degree up to d as

$$q(\mathbf{x}) = \mathbf{q}^T \mathbf{y}^{(d)} \tag{3.33}$$

where $\mathbf{q}$ denotes the (possibly very sparse) coefficient vector of $q(\mathbf{x})$. We will first discuss a technique to solve the multivariate polynomial system (3.25), followed by relaxations to directly solve (3.23).

3.5.1 Polynomial system solving via eigendecomposition of the multiplication matrix

Given the constant structure of the polynomial system (3.25), we can avoid using general purpose homotopy continuation solvers, such as PHCpack, and instead apply a faster, hybrid algebraic-numeric method based on the eigendecomposition of a generalized companion matrix [117, 20, 16, 5]. In particular, in the 2-D case we employ the normal set-based method, described in detail in [117, 16], which exploits the specific system structure to

provide all 28 solutions simultaneously, in a fast, non-iterative fashion.

The idea behind the normal set-based method is to first expand the polynomial system (3.25), compactly written as

$$\mathbf{\Psi}(\mathbf{p}, \bar{q}) = \mathbf{0} \tag{3.34}$$

by adding new polynomials of the form

$$\psi_i'(\mathbf{p}, \bar{q}) = \psi_i(\mathbf{p}, \bar{q}) \cdot \mathbf{p}^{\boldsymbol{\alpha}_\mathbf{p}} \bar{q}^{\boldsymbol{\alpha}_{\bar{q}}} = 0 \tag{3.35}$$

The new polynomials, $\psi_i'(\mathbf{p}, \bar{q})$, are products of the original polynomials, $\psi_i(\mathbf{p}, \bar{q})$, with monomials in the unknowns $\mathbf{p}$ and $\bar{q}$. The new, expanded system of polynomials, consisting of $\psi_i(\mathbf{p}, \bar{q})$ and $\psi_i'(\mathbf{p}, \bar{q})$, involves monomials up to some degree ℓ. For the SD-WLS problem in 2-D, for example, we chose these monomials so as to create new polynomials up to degree $\ell = 8$. Note that adding these new polynomials does not change the solution set, since all $\psi_i'(\mathbf{p}, \bar{q})$ vanish on the roots of the original system, and do not add new solutions.

The next step is to write the expanded polynomial system in matrix form

$$\mathbf{C}_e \mathbf{y}^{(\ell)} = \mathbf{0}, \tag{3.36}$$

gathering the monomials in the vector $\mathbf{y}^{(\ell)}$, and the numeric coefficients in the expanded coefficient matrix $\mathbf{C}_e$. From this matrix one can extract a 28×28 generalized companion matrix that defines multiplication by a function (which we chose to be the first component of $\mathbf{p}$) within the so-called quotient ring. All 28 solutions can be extracted simultaneously from the left eigenvectors of this matrix [20, 117].

Our current Matlab implementation requires approximately 125 ms (as determined by Matlab's profiler) to solve an instance of the 2-D SD-WLS problem, of which approximately 16 ms are spent creating the expanded coefficient matrix $\mathbf{C}_e$, which for this problem is of dimension 532×495, with 4% non-zero entries. Extracting the multiplication matrix and all 28 solutions requires 109 ms.

Unfortunately, in the 3-D case, this approach to solving the SD-WLS problem is unwieldy for practical applications, since a solution using homotopy continuation (PHCpack) requires several hours, and the previously described multiplication matrix-based solvers require excessive memory (> 14 GB) for the internal matrix decompositions.

In the next two sections, we will therefore investigate two different relaxations of the SD-WLS problem, that can compute and certify the global optimum of SD-WLS within practicable time and resources.

3.5.2 Sum-of-Squares relaxation

SOS relaxation has been proposed as an efficient means to solve global polynomial optimization problems [91] without having to compute all critical points. SOS relaxation works by transforming real polynomial optimization problems into semidefinite programs (SDPs), which can be solved efficiently using interior point methods such as SeDuMi [113]. In the last decade, several toolboxes have been published to efficiently solve SOS problems, e.g., SOSTOOLS [94], GloptiPoly [40], or SparsePOP [126]. We have used SparsePOP to obtain solutions to (3.23). In the following, we will briefly outline the underlying theoretical principles, but refer the reader to [92, 66, 57, 125] for more details on SOS relaxation.

A polynomial $f(\mathbf{x}) = \mathbf{f}^T \mathbf{y}^{(2d)}$ of even degree $2d$ is a sum of squares, if it can be written as

$$f(\mathbf{x}) = \sum_i q_i(\mathbf{x})^2 = \sum_i (\mathbf{q}_i^T \mathbf{y}^{(d)})^2 = \mathbf{y}^{(d)^T} \left(\sum_i \mathbf{q}_i \mathbf{q}_i^T \right) \mathbf{y}^{(d)} = \mathbf{y}^{(d)^T} \mathbf{Q} \mathbf{y}^{(d)} \tag{3.37}$$

for some polynomials $q_i(\mathbf{x})$ so that $\mathbf{Q} \succcurlyeq 0$ is positive semidefinite. Note that this decomposition is not necessarily unique, and the set of matrices $\mathbf{Q}$ fulfilling (3.37) is an affine subspace [92]. If a polynomial $f(\mathbf{x})$ can be written as a sum of squares, it is necessarily non-negative for all values of $\mathbf{x}$. However, not all non-negative polynomials can be written as sums of squares.

One can determine whether a polynomial can be decomposed as a sum of squares by solving a convex semidefinite feasibility problem

$$\begin{aligned} &\text{find} \;\; \mathbf{Q} \\ &\text{s.t.} \; \langle \mathbf{A}_i, \mathbf{Q} \rangle = f_i \\ &\qquad \mathbf{Q} \succcurlyeq 0 \end{aligned} \tag{3.38}$$

where f_i are the $\binom{n+2d}{n}$ coefficients of the polynomial $f(\mathbf{x})$, $\langle \mathbf{A}_i, \mathbf{Q} \rangle = \mathrm{tr}(\mathbf{A}_i \mathbf{Q})$, and $\mathbf{A}_i$ are constant symmetric (selector or localizer) matrices, such that their (l, m)-th elements are given by

$$a_{i_{l,m}} = \begin{cases} 1 & \text{if } y_l^{(d)} y_m^{(d)} = y_i^{(2d)}, \\ 0 & \text{otherwise.} \end{cases} \tag{3.39}$$

$\mathbf{A}_i$ can be found directly from expanding $\mathbf{y}^{(d)^T} \mathbf{Q} \mathbf{y}^{(d)}$ and comparing coefficients with $f(\mathbf{x})$. Once a matrix $\mathbf{Q}$ is found, the polynomials $q_i(\mathbf{x})$ can be determined from an eigendecomposition $\mathbf{Q} = \sum_i \lambda_i \mathbf{v}_i \mathbf{v}_i^T$, so that $q_i(\mathbf{x}) = \sqrt{\lambda_i} \mathbf{v}_i^T \mathbf{y}^{(d)}$.

Consider now the unconstrained optimization problem

$$\min f(\mathbf{x}), \tag{3.40}$$

where $f(\mathbf{x})$ is a real polynomial with global minimum $f^* = \min f(\mathbf{x})$ that is achieved at $\mathbf{x}^*$, i.e., $f(\mathbf{x}^*) = f^*$.

To convert this to an SOS problem, we first introduce an auxiliary variable γ as lower bound for $f(\mathbf{x})$

$$\begin{aligned} &\max \gamma \\ &\text{s.t. } f(\mathbf{x}) - \gamma \geq 0 \ \forall \mathbf{x} \end{aligned} \tag{3.41}$$

and then replace the inequality constraint by

$$\begin{aligned} &\max \gamma \\ &\text{s.t. } f(\mathbf{x}) - \gamma \text{ is SOS.} \end{aligned} \tag{3.42}$$

Again, the relaxed problem can be written as an SDP in the variables γ and $\mathbf{Q}$ (not in $\mathbf{x}$) and solved efficiently. Let γ^* denote the optimal solution of (3.42). Since (3.42) is a relaxation, we always have that $\gamma^* \leq f^*$. However, in many cases we can verify that $\gamma^* = f^*$ and thus provide a certificate of global optimality. Moreover, if $\mathbf{x}^*$ is the unique[5] optimal point, the Lagrange multiplier (i.e., the dual variable of (3.42)) is exactly equal to the monomial vector $\mathbf{y}^{(d)}$ evaluated at $\mathbf{x}^*$, so that both the optimal function value, as well as the optimal solution point can be directly recovered from the solution of the SDP [92, 66]. Evaluating $f(\mathbf{x})$ on the point extracted from the dual variable at the solution of the SDP and comparing it to γ^* allows to verify if global optimality has been achieved.

As already stated above, not all nonnegative polynomials can be written as sums of squares, so that there exist cases where γ^* does not provide the exact global minimum of $f(\mathbf{x})$. In these cases, we can make use of the Positivstellensatz to form hierarchies of SOS relaxations with increasingly tighter lower bounds [66, 57]. The Positivstellensatz also helps to address constrained optimization, as needed for the SD-WLS problem.

The form of the Positivstellensatz used by Lasserre [66], due to Putinar [96], states that under mild assumptions[6], every polynomial $p(\mathbf{x})$, strictly positive on the compact set defined by polynomial inequalities $g_i(\mathbf{x}) \geq 0, i = 1, \ldots, r$, can be written as

$$p(\mathbf{x}) = q(\mathbf{x}) + \sum_{k=1}^{r} g_k(\mathbf{x}) t_k(\mathbf{x}) \tag{3.43}$$

where $q(\mathbf{x}), t_k(\mathbf{x})$ are SOS polynomials of sufficiently high degree.

For the unconstrained optimization problem, this means that even if solving (3.42) for γ does not provide the exact global minimum of $f(\mathbf{x})$, we can constrain the search within some

[5] If the global optimum is achieved by a discrete set of (multiple) points, other methods to recover their coordinates are available [41, 126].

[6] One way to fulfil the assumption is for the set of polynomial inequalities $g_i(\mathbf{x})$ to include the ball constraint $a^2 - \mathbf{x}^T\mathbf{x} \geq 0$, which can always be added with a chosen sufficiently large.

hypersphere and construct new SOS optimization problems to find progressively tighter bounds that will converge asymptotically to the global optimum f^* as the degree increases.[7] Specifically, replace $p(\mathbf{x})$ by $f(\mathbf{x}) - \gamma$, and introduce $g(\mathbf{x}) = a^2 - \mathbf{x}^T\mathbf{x} \geq 0$ for some a. Then solve

$$\begin{aligned} &\max \gamma \\ &\text{s.t. } f(\mathbf{x}) - \gamma - g(\mathbf{x})t(\mathbf{x}) \text{ is SOS.} \end{aligned} \tag{3.44}$$

for increasing degree of $t(\mathbf{x})$.

In the case of constrained optimization, i.e.,

$$\begin{aligned} &\min f(\mathbf{x}) \\ &\text{s.t. } h_i(\mathbf{x}) = 0, i = 1, \ldots, r \\ &\qquad g_i(\mathbf{x}) \geq 0, i = r+1, \ldots, m \end{aligned} \tag{3.45}$$

the inequality constraints can be used directly in the Positivstellensatz, and the equality constraints are added as two inequalities, $h_i(\mathbf{x}) \geq 0, -h_i(\mathbf{x}) \geq 0$. The resulting SOS problem will be

$$\begin{aligned} &\max \gamma \\ &\text{s.t. } f(\mathbf{x}) - \gamma - \sum_{i=1}^{r} h_i(\mathbf{x})\left(t_i(\mathbf{x}) - u_i(\mathbf{x})\right) - \sum_{j=r+1}^{m} g_j(\mathbf{x})v_j(\mathbf{x}) = q(\mathbf{x}) \\ &\qquad t_i(\mathbf{x}), u_i(\mathbf{x}), v_j(\mathbf{x}), q(\mathbf{x}) \text{ are SOS.} \end{aligned} \tag{3.46}$$

$$\begin{aligned} \Leftrightarrow &\max \gamma \\ &\text{s.t. } f(\mathbf{x}) - \gamma - \sum_{i=1}^{r} h_i(\mathbf{x})k_i(\mathbf{x}) - \sum_{j=r+1}^{m} g_j(\mathbf{x})v_j(\mathbf{x}) = q(\mathbf{x}) \\ &\qquad v_j(\mathbf{x}), q(\mathbf{x}) \text{ are SOS.} \end{aligned} \tag{3.47}$$

Notice that we can replace $(t_i(\mathbf{x}) - u_i(\mathbf{x}))$ by arbitrary (non-SOS) polynomials $k_i(\mathbf{x})$. Interestingly, one can show that the polynomials $k_i(\mathbf{x})$ and $v_j(\mathbf{x})$ play the role of generalized Lagrangian multipliers [57, 66].

To convert (3.46) into an SDP, one chooses the degree of $q(\mathbf{x})$ as an even number N_q with

$$N_q \geq \max\Big(\deg(f(\mathbf{x})), \deg(g_i(\mathbf{x})), \deg(h_i(\mathbf{x}))\Big),$$

which determines the degrees of $k_i(\mathbf{x})$ and $v_j(\mathbf{x})$. N_q is increased iteratively until γ can be

[7] In many cases in practice, convergence occurs already for low degrees of $q(\mathbf{x})$ and $t_k(\mathbf{x})$, but to our knowledge only few results exist for guaranteed finite convergence (see, e.g., [85]).

certified as the global optimum.

Having described the general process of constrained and unconstrained polynomial optimization using SOS relaxation, we now return to the specific application of minimizing the SD-WLS cost function (3.23) subject to the quaternion unit norm equality constraint $h(\mathbf{x}) = \bar{q}^T\bar{q} - 1$. To solve it, we chose the degree of $q(\mathbf{x})$ as $N_q = 6$, i.e., equal to the degree of the cost function. We used SparsePOP [126] in combination with SeDuMi [113] to formulate and solve the SOS program, since SparsePOP can exploit sparsity in the equation structure, and performs variable scaling to improve numerical stability. In addition to the quaternion unit norm constraint (3.25c), we also added the constraint $q_4 \geq 0$ in order to remove one of the two physically equivalent solutions $(\bar{q}, -\bar{q})$. Increasing the default precision used in SeDuMi to 10^{-11} helped to greatly reduce numerical errors of the solution. Bounding the unknowns (box constraints) helped in some instances to improve numerics as well, at the cost of increasing the size of the SDP problem. Increasing N_q, or adding a ball constraint was generally not found necessary to certify global optimality. In our simulations we have not encountered a case where global optimality could not be certified (up to numerical inaccuracies). A typical instance of a SparsePOP program of SD-WLS led to an SDP problem solved by SeDuMi of the form

$$\begin{aligned} \min_{\mathbf{x}} \ & \mathbf{c}^T\mathbf{x} \\ \text{s.t.} \ & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \in K \end{aligned}$$

To provide a notion of computational requirements, the dimensions of $\mathbf{A}$ were $\dim(\mathbf{A}) = 1582 \times 13731$, and the cone K had parameters (see SeDuMi user guide [113] for notation) $K.f = 330$ (number of free primal components), $K.l = 5$ (number of nonnegative components), $K.s = \begin{bmatrix}(110 \times 110) & (36 \times 36)\end{bmatrix}$ (dimensions of positive semi-definite constraints). A typical problem instance was solved in 50-70s on an Intel T9400 2.53 GHz laptop with 2GB of RAM.

3.5.3 Lagrange relaxation

In this section, we will introduce an alternative method that, contrary to SOS relaxation, does not find a global optimal solution of the 3-D SD-WLS problem, but can be used to certify global optimality of a given candidate stationary point. Its main advantage is that it is at least an order of magnitude faster than SOS relaxation. This process is motivated by the realization that particularly for the 3-D case, the process of finding a stationary point of the SD-WLS problem using iterative optimization methods (e.g., Gauss-Newton) is considerably faster than SOS relaxation. This is especially true if a good initialization point is available, for example from deterministic methods [117] used during a preliminary RANSAC [33] outlier rejection stage. The shortcoming of the iterative optimization methods, namely

their inability to certify global optimality, can be overcome if one can establish equality between the value of the cost function evaluated at the stationary point and a guaranteed lower bound of the cost function (which then provides a certificate for global optimality). In contrast to the previous section, where this lower bound was provided by SOS relaxation, in this section we will compute a lower bound using Lagrange relaxation.

In particular, we will introduce additional auxiliary variables to form a new Lagrangian, $\mathcal{L}_{rs}$, and consider its dual optimization problem. The key insights will be that the dual function, $Q(\boldsymbol{\mu}, \nu)$, can be evaluated very efficiently, and, more surprisingly, that despite the non-convexity of this problem, strong duality holds in a significant percentage of problem instances. A zero duality gap serves as certificate for global optimality of the SD-WLS solution.

The auxiliary variables are defined as $\mathbf{r} := \mathbf{C}^T\mathbf{p}$ and $s := \mathbf{p}^T\mathbf{p}$, yielding the optimization problem

$$\begin{aligned} \min_{\mathbf{p},\bar{q},\mathbf{r},s} \ & \frac{1}{2}(\mathbf{h}_{rs} - \mathbf{z}')^T\mathbf{R}'^{-1}(\mathbf{h}_{rs} - \mathbf{z}') \\ \text{s.t. } & \bar{q}^T\bar{q} - 1 = 0 \\ & \mathbf{C}^T\mathbf{p} - \mathbf{r} = \mathbf{0} \\ & \mathbf{p}^T\mathbf{p} - s = 0 \end{aligned} \tag{3.48}$$

with

$$\mathbf{h}_{rs} = \begin{bmatrix} \mathbf{w}_1^T\mathbf{w}_1 & \dots & \mathbf{w}_N^T\mathbf{w}_N \end{bmatrix}^T = \mathbf{H}\begin{bmatrix} \mathbf{p} \\ \mathbf{r} \\ s \end{bmatrix} + \mathbf{b} \tag{3.49}$$

$$\mathbf{H} = \begin{bmatrix} -2\mathbf{u}_1^T & 2\mathbf{v}_1^T & 1 \\ \vdots & \vdots & \vdots \\ -2\mathbf{u}_n^T & 2\mathbf{v}_n^T & 1 \end{bmatrix} \tag{3.50}$$

$$\mathbf{b} = \begin{bmatrix} \mathbf{u}_1^T\mathbf{u}_1 + \mathbf{v}_1^T\mathbf{v}_1 - 2\mathbf{u}_1^T\mathbf{C}\mathbf{v}_1 \\ \vdots \\ \mathbf{u}_n^T\mathbf{u}_n + \mathbf{v}_n^T\mathbf{v}_n - 2\mathbf{u}_n^T\mathbf{C}\mathbf{v}_n \end{bmatrix}. \tag{3.51}$$

The corresponding Lagrangian is

$$\mathcal{L}_{rs} = \frac{1}{2}(\mathbf{h}_{rs} - \mathbf{z}')^T\mathbf{R}'^{-1}(\mathbf{h}_{rs} - \mathbf{z}') + \lambda(\bar{q}^T\bar{q} - 1) + \boldsymbol{\mu}^T(\mathbf{C}^T\mathbf{p} - \mathbf{r}) + \nu(\mathbf{p}^T\mathbf{p} - s) \tag{3.52}$$

The stationary points can be found by solving the

SD-WLS r/s First-Order Optimality Conditions:

$$\nabla_{\begin{bmatrix}\mathbf{p}^T & \mathbf{r}^T & s\end{bmatrix}^T}\mathcal{L}_{rs} = \mathbf{H}^T\mathbf{R}'^{-1}\left(\mathbf{H}\begin{bmatrix}\mathbf{p}\\ \mathbf{r}\\ s\end{bmatrix} + \mathbf{b} - \mathbf{z}'\right) + \begin{bmatrix}\mathbf{C}\boldsymbol{\mu} + 2\nu\mathbf{p}\\ -\boldsymbol{\mu}\\ -\nu\end{bmatrix} = \mathbf{0} \tag{3.53a}$$

$$\nabla_{\bar{q}}\mathcal{L}_{rs} = (\nabla_{\bar{q}}\mathbf{h}_{rs})\mathbf{R}'^{-1}\left(\mathbf{H}\begin{bmatrix}\mathbf{p}\\ \mathbf{r}\\ s\end{bmatrix} + \mathbf{b} - \mathbf{z}'\right) + (\nabla_{\bar{q}}\boldsymbol{\mu}\mathbf{C}^T\mathbf{p}) + 2\lambda\bar{q} = \mathbf{0} \tag{3.53b}$$

$$\nabla_{\lambda}\mathcal{L}_{rs} = \bar{q}^T\bar{q} - 1 = 0 \tag{3.53c}$$

$$\nabla_{\boldsymbol{\mu}}\mathcal{L}_{rs} = \mathbf{C}^T\mathbf{p} - \mathbf{r} = \mathbf{0} \tag{3.53d}$$

$$\nabla_{\nu}\mathcal{L}_{rs} = \mathbf{p}^T\mathbf{p} - s = 0 \tag{3.53e}$$

where

$$\nabla_{\bar{q}}\mathbf{h}_{rs} = \nabla_{\bar{q}}\mathbf{b} = -2\begin{bmatrix}\nabla_{\bar{q}}\mathbf{u}_1^T\mathbf{C}\mathbf{v}_1 & \dots & \nabla_{\bar{q}}\mathbf{u}_N^T\mathbf{C}\mathbf{v}_N\end{bmatrix} \tag{3.54}$$

$$\nabla_{\bar{q}}\mathbf{u}_i^T\mathbf{C}\mathbf{v}_i = (\nabla_{\bar{q}}\mathrm{vec}(\mathbf{C}))\ \mathrm{vec}\left(\mathbf{u}_\mathrm{i}\mathbf{v}_\mathrm{i}^\mathrm{T}\right) \tag{3.55}$$

$$\nabla_{\bar{q}}\boldsymbol{\mu}\mathbf{C}^T\mathbf{p} = (\nabla_{\bar{q}}\mathrm{vec}(\mathbf{C}))\ \mathrm{vec}\left(\mathbf{p}\boldsymbol{\mu}^\mathrm{T}\right) \tag{3.56}$$

Notice that a stationary point $\mathbf{p}^*, \bar{q}^*, \lambda^*$ fulfilling the KKT conditions of the SD-WLS problem (3.25) is also a stationary point for this formulation with auxiliary variables $\mathbf{r}$ and s, since it will (exactly) fulfill the optimality conditions (3.53) above, after computing $\mathbf{r}^* = \mathbf{C}(\bar{q}^*)^T\mathbf{p}^*$ and $s^* = \mathbf{p}^{*T}\mathbf{p}^*$, as well as the corresponding Lagrange multipliers $\boldsymbol{\mu}^*, \nu^*$ using the bottom part of (3.53a)

$$\begin{bmatrix}\boldsymbol{\mu}^*\\ \nu^*\end{bmatrix} = \begin{bmatrix}\mathbf{0}_{4\times 3} & \mathbf{I}_{4\times 4}\end{bmatrix}\mathbf{H}^T\mathbf{R}'^{-1}\left(\mathbf{H}\begin{bmatrix}\mathbf{p}^*\\ \mathbf{r}^*\\ s^*\end{bmatrix} + \mathbf{b} - \mathbf{z}'\right). \tag{3.57}$$

To see this, compare the expressions for (3.25) and (3.53) evaluated at $\mathbf{p}^*, \bar{q}^*, \lambda^*, \mathbf{r}^*, s^*, \boldsymbol{\mu}^*$ and ν^*.

To certify global optimality, we will now consider the Lagrangian dual [12] of this optimization problem, after only dualizing the constraints in $\mathbf{r}$ and s. For this, consider the Lagrangian

$$\mathcal{L}_{\mathrm{Lag}} = \frac{1}{2}(\mathbf{h}_{rs} - \mathbf{z}')^T\mathbf{R}'^{-1}(\mathbf{h}_{rs} - \mathbf{z}') + \boldsymbol{\mu}^T(\mathbf{C}^T\mathbf{p} - \mathbf{r}) + \nu(\mathbf{p}^T\mathbf{p} - s) \tag{3.58}$$

The dual function is given by

$$Q(\boldsymbol{\mu},\nu) = \min_{\mathbf{p},\bar{q},\mathbf{r},s} \mathcal{L}_{\text{Lag}}(\mathbf{p},\bar{q},\mathbf{r},s,\boldsymbol{\mu},\nu) \quad \text{s.t. } \bar{q}^T\bar{q} = 1 \tag{3.59}$$

It is important to observe that in the minimization problem of computing the dual function, the equality constraints on $\mathbf{r}$ (3.53d) and s (3.53e) are not explicitly enforced (however, they hold for an optimal choice of $\boldsymbol{\mu}$ and ν). For this reason, the dual function $Q(\boldsymbol{\mu},\nu)$ always provides a lower bound to the original minimization problem (3.48).

Next we will describe how to compute the dual function for arbitrary, but known values of $\boldsymbol{\mu}$ and ν. Computing the dual function (3.59) requires solving a minimization problem, which will again be accomplished by computing its stationary points by solving a polynomial system. As before, we use the method of Lagrange multipliers to obtain a new Lagrangian, $\mathcal{L}_{\text{Lag}'}$, which is formally equal to (3.52), but differs in that $\boldsymbol{\mu},\nu$ are now considered fixed (known) parameters, not variable Lagrange multipliers.

$$\mathcal{L}_{\text{Lag}'} = \frac{1}{2}(\mathbf{h}_{rs}-\mathbf{z}')^T\mathbf{R}'^{-1}(\mathbf{h}_{rs}-\mathbf{z}') + \lambda(\bar{q}^T\bar{q}-1) + \boldsymbol{\mu}^T(\mathbf{C}^T\mathbf{p}-\mathbf{r}) + \nu(\mathbf{p}^T\mathbf{p}-s) \tag{3.60}$$

The corresponding first-order optimality conditions consequently closely resemble those of (3.53), except for the equality constraints.

SD-WLS Lagrangian Dual First-Order Optimality Conditions:

$$\nabla_{\begin{bmatrix}\mathbf{p}^T & \mathbf{r}^T & s\end{bmatrix}^T}\mathcal{L}_{\text{Lag}'} = \mathbf{H}^T\mathbf{R}'^{-1}\left(\mathbf{H}\begin{bmatrix}\mathbf{p}\\ \mathbf{r}\\ s\end{bmatrix} + \mathbf{b} - \mathbf{z}'\right) + \begin{bmatrix}\mathbf{C}\boldsymbol{\mu} + 2\nu\mathbf{p}\\ -\boldsymbol{\mu}\\ -\nu\end{bmatrix} = \mathbf{0} \tag{3.61a}$$

$$\nabla_{\bar{q}}\mathcal{L}_{\text{Lag}'} = (\nabla_{\bar{q}}\mathbf{h}_{rs})\mathbf{R}'^{-1}\left(\mathbf{H}\begin{bmatrix}\mathbf{p}\\ \mathbf{r}\\ s\end{bmatrix} + \mathbf{b} - \mathbf{z}'\right) + (\nabla_{\bar{q}}\boldsymbol{\mu}\mathbf{C}^T\mathbf{p}) + 2\lambda\bar{q} = \mathbf{0} \tag{3.61b}$$

$$\nabla_{\lambda}\mathcal{L}_{\text{Lag}'} = \bar{q}^T\bar{q} - 1 = 0 \tag{3.61c}$$

The fact that the measurement function $\mathbf{h}_{rs}$ is linear in $\mathbf{p},\mathbf{r}$, and s results in a decoupled problem structure and allows to first solve for $\mathbf{p},\mathbf{r}$, and s in terms of $\bar{q}$ from (3.61a) to obtain the following quadratic expression in $\bar{q}$

$$\begin{bmatrix}\mathbf{p}\\ \mathbf{r}\\ s\end{bmatrix} = \left(\mathbf{H}^T\mathbf{R}'^{-1}\mathbf{H} + \begin{bmatrix}2\nu\mathbf{I}_{3\times3} & \mathbf{0}_{3\times4}\\ \mathbf{0}_{4\times3} & \mathbf{0}_{4\times4}\end{bmatrix}\right)^{-1}\left(\mathbf{H}^T\mathbf{R}'^{-1}(\mathbf{z}'-\mathbf{b}) + \begin{bmatrix}-\mathbf{C}\boldsymbol{\mu}\\ \boldsymbol{\mu}\\ \nu\end{bmatrix}\right). \tag{3.62}$$

We substitute this expression for $\mathbf{p},\mathbf{r},s$ in (3.53b), and then proceed to solve the system [(3.61b), (3.61c)] (recall that the values of $\boldsymbol{\mu}$ and ν are considered known), which consists

of one quadratic and four cubic polynomials in only five unknowns $\bar{q}$ and λ. Analysis of its Groebner basis reveals that this system has 80 (possibly complex) solutions and its BKK bound is tight (the BKK bound is a formula due to Kushnirenko, Bernstein, and Khovanskiisee that provides an upper bound to the number of solutions of a polynomial system as a function of its monomial structure; see e.g., [19]). Importantly, it can be solved very fast (in approx. 3-5s using PHCpack, with room for considerable speed-up by using multiplication matrix-based methods as described in chapter 3.4). This system's real solutions, $\bar{q}^*$, are backsubstituted into (3.62) to find all real critical points of (3.59). We evaluate (3.59) on all critical points and retain the solution with minimum cost function value as the dual function value $Q(\boldsymbol{\mu}, \nu)$. Weak duality states that the duality gap, $\frac{1}{2}\ (\mathbf{h}_{sd} - \mathbf{z}')^T \mathbf{R}'^{-1} (\mathbf{h}_{sd} - \mathbf{z}') - Q(\boldsymbol{\mu}, \nu)$, is always nonnegative, i.e., $Q(\boldsymbol{\mu}, \nu)$ is always a lower bound to the global minimum f^*.

Despite the optimization problem(3.48) being non-convex, we have found empirically that there exist a significant number of cases where the duality gap is zero.[8] To exploit this fact, we consider a stationary point of the SD-WLS problem, which is also a stationary point of (3.48). We then compute the Lagrange multipliers $\boldsymbol{\mu}, \nu$ corresponding to this stationary point using (3.57), and evaluate the dual function Q for this particular choice of $\boldsymbol{\mu}, \nu$ by following the process described above. Note that at the solution of this process, the optimal points $\mathbf{p}^*, \bar{q}^*, \mathbf{r}^*, s^*$ that achieve $Q(\boldsymbol{\mu}, \nu)$ (i.e., the global minimum of $\mathcal{L}_{\text{Lag}}$) do not have to be the same as the values of these variables at the stationary point of the SD-WLS problem under consideration[9]. If, however, they do coincide, and the cost function value of (3.48) equals that of $Q(\boldsymbol{\mu}, \nu)$, we have a certificate of global optimality for the stationary point under consideration. The total cost of computing the certificate comprises the cost for finding the stationary point (e.g., by solving minimal problems and subsequent iterative SD-WLS refinement using Gauss-Newton), the cost to compute corresponding values of $\mathbf{r}, s, \boldsymbol{\mu}$, and ν (see (3.57)), as well as the cost for solving the minimization problem (3.59) to compute $Q(\boldsymbol{\mu}, \nu)$. The total cost will be at least one order of magnitude lower than solving the SOS relaxation problem. Hence, this process can be used as a fast "pre-screening" step before SOS relaxation. Notice that a candidate stationary point that fails the above process cannot be ruled out as the global optimum. In the event that the stationary point cannot be certified using Lagrange relaxation, one will have to fall back to the SOS relaxation method described in chapter 3.5.2.

Alternatively, one could solve the actual dual problem, that is, to maximize $Q(\boldsymbol{\mu}, \nu)$ with respect to the Lagrangian multipliers, which is a concave (but usually discontinuous) optimization problem, without need of a candidate stationary point. To this end, we have investigated the subgradient and cutting plane methods [12]. The cutting plane method usually converges within about 50 iterations, but may not produce a usable approximation

[8] The conditions under which the duality gap is zero are subject of ongoing research.

[9] The original stationary point fulfilling (3.25) will be among the solution set, since it also fulfills (3.53). However, there may be other stationary points in the solution set resulting in lower values of $\mathcal{L}_{\text{Lag}}$.

to the primal optimal solution if the duality gap is non-zero. The total cost, however, starts to be comparable to that of SOS relaxation, which in our experiments has always provided the global optimum and is hence the preferred approach. Notice that the process of computing $\boldsymbol{\mu}^*$ and ν^* at the candidate SD-WLS stationary point described above serves to provide a good initial guess for the optimal solution of the dual problem. In principle, if the duality gap is indeed zero, maximizing the (concave) dual function $Q(\boldsymbol{\mu}, \nu)$ with respect to the Lagrange multipliers $\boldsymbol{\mu}$ and ν will converge to the solution $\boldsymbol{\mu}^*$ and ν^* at the global SD-WLS optimum even without special initialization. However, this process takes about as long as solving the SOS relaxation, but at the same time runs the risk of encountering a case with nonzero duality gap and hence not being able to provide a certificate of global optimality.

Algorithm 1 Certify global optimality of stationary point using Lagrange relaxation.

Require: Stationary point $\mathbf{p}^*, \bar{q}^*, \lambda^*$ of SD-WLS problem (3.23)
Ensure: Certificate of global optimality or FAIL
  $\mathbf{r}^* \leftarrow \mathbf{C}(\bar{q}^*)^T \mathbf{p}^*$
  $s^* \leftarrow \mathbf{p}^{*T} \mathbf{p}^*$
  Compute $\boldsymbol{\mu}^*, \nu^*$ using (3.57)
  Compute $\begin{bmatrix} \mathbf{p}^T & \mathbf{r}^T & s \end{bmatrix}$ as a function of $\bar{q}$ using (3.62)
  Substitute above into [(3.61b), (3.61c)] and solve for $\bar{q}, \lambda$ (80 solutions)
  Backsubstitute solutions for $\bar{q}, \lambda$ into expression for $\begin{bmatrix} \mathbf{p}^T & \mathbf{r}^T & s \end{bmatrix}$
  Evaluate (3.60) on all 80 stationary points, retain the one with lowest cost function value as dual function value $Q(\boldsymbol{\mu}^*, \nu^*)$.
  if $Q(\boldsymbol{\mu}^*, \nu^*) = \mathcal{L}_{rs}(\mathbf{p}^*, \bar{q}^*, \lambda^*, \mathbf{r}^*, s^*, \boldsymbol{\mu}^*, \nu^*)$ **then**
    return Global optimality certified
  else
    return FAIL (proceed with SOS relaxation)
  end if

3.6 Simulation results

We compared the performance of the WLS, the SD-WLS, and the deterministic algorithms of [133, 122] for the overdetermined case where $N = 5$ ($N = 10$) noisy (but outlier-free) measurements are available. Specifically, we conducted Monte Carlo simulations using measurement noise with diagonal covariance matrix $\mathbf{R} = \sigma_d^2 \mathbf{I}$ for different values of σ_d, and 1000 trials per setting. The trajectories were chosen so that the robot-to-robot distances varied randomly between 1-2 m, and the displacement between measurements varied between 3-6 m. The parameters z_i' and $\mathbf{R}'$ were determined based on the approximations (3.29) and (3.30).

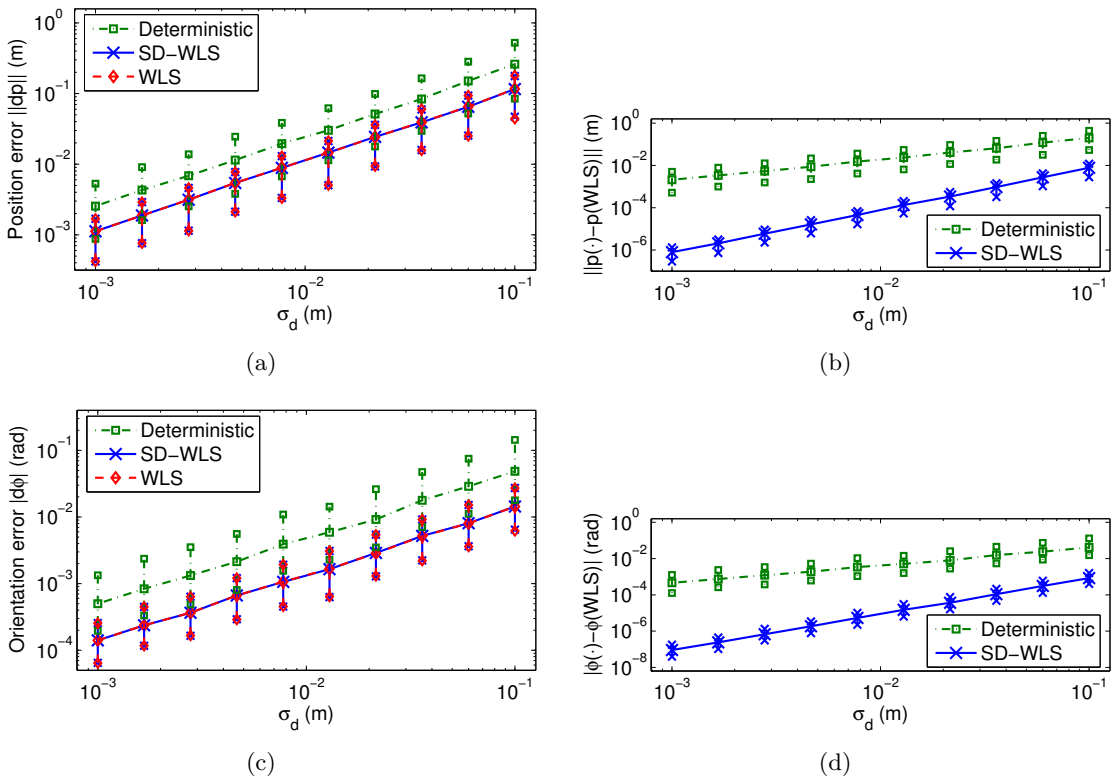

Figure 3.3: 2-D simulation results from 1000 trials for each noise setting, comparing WLS, SD-WLS and the linear algorithm of [133] (stand-alone), using 5 distance measurements corrupted by zero-mean Gaussian measurement noise with covariance $\mathbf{R} = \sigma_d^2\mathbf{I}$. Plotted are the median and the 25% and 75% quantiles. The accuracy of the WLS and the SD-WLS formulation is almost identical, and higher than that of the linear method by a nearly constant factor. The SD-WLS solution is close to that of WLS particularly for small noise values. (a) Norm of position error with respect to ground truth. (b) Difference in position with respect to WLS estimate. (c) Orientation error with respect to ground truth. (d) Difference in orientation with respect to WLS estimate.

3.6.1 2-D case

For 2-D, the results are shown in Fig. 3.3. As evident, SD-WLS consistently and significantly outperforms the deterministic algorithm (stand-alone). From the doubly logarithmic plot one can deduce that the median position error of the linear algorithm in [133] with respect to the true solution is approximately twice as large as the WLS and the SD-WLS error, and the orientation error approximately 3.5 times as large, independent of the measurement noise standard deviation. Even more remarkable is the fact that the performance of the SD-WLS is virtually indistinguishable from that of WLS, despite it not being optimal in the maximum likelihood sense (due to the approximation of chapter 3.4.1). The errors for all algorithms grow approximately linearly with increasing noise standard deviation, with a higher growth rate for the deterministic method than for WLS and SD-WLS. SD-WLS and WLS achieve approximately the same percentage of inconsistent estimates (i.e., estimation error $> 3\sigma$, see Fig. 3.4(b)). Note that WLS here refers to the global optimum of (3.20), which we were able to compute exactly due to the small number of measurements.

In addition to the stand-alone output of the linear and SD-WLS algorithms, we have also investigated the performance of Levenberg-Marquardt-based i-WLS using the outputs of the deterministic and the SD-WLS algorithms, and even ground truth as initialization points. The latter is of course not available in real applications, but serves as comparison for best possible initialization. The results shown in Fig. 3.4 confirm that SD-WLS-based initialization outperforms deterministic initialization in terms of required iterations and number of inconsistent estimates. Furthermore, initializing i-WLS with the result of the deterministic method can indeed lead to convergence to a local minimum, as illustrated in Fig. 3.4(c). A more surprising result is the fact that SD-WLS initialization even outperforms initialization with ground truth in terms of required iterations and in terms of achieving the global WLS optimum. In fact, in the presence of large noise, the ground truth might not reside within the basin of attraction of the global WLS optimum, that is, even initialization of i-WLS with the ground truth might result in a local minimum. However, Fig. 3.4(b) also shows that achieving the global optimum does not necessarily imply a consistent estimate.

3.6.2 3-D case

The results shown in Fig. 3.5 for the 3-D case using 10 measurements are very similar. The improvement over the linear method is even more pronounced than in 2-D, with the error of the deterministic method being up to 6 times larger than that of SD-WLS for position, and up to 50 times for orientation. We attribute that to the increased number of variables and increased degree of the polynomial system. We have observed that increasing the number of measurements improves the performance of the deterministic method noticeably (see Fig. 3.6), but it remains greatly inferior to SD-WLS.

Note that unlike the 2-D case, we were unable to compute the certified global optimum of WLS for this set of simulations. What is referred to as WLS here is therefore based on

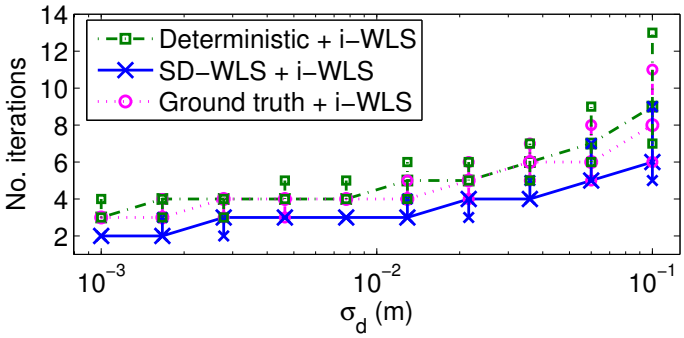

(a)

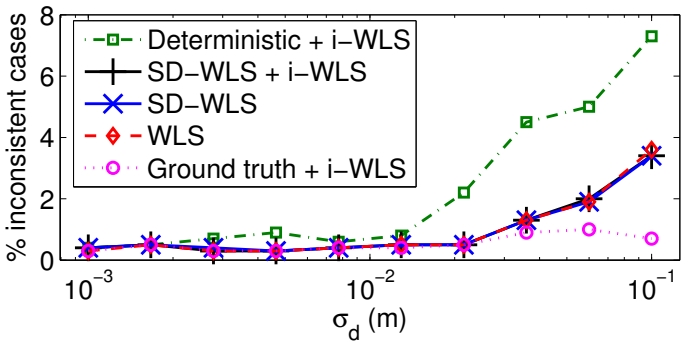

(b)

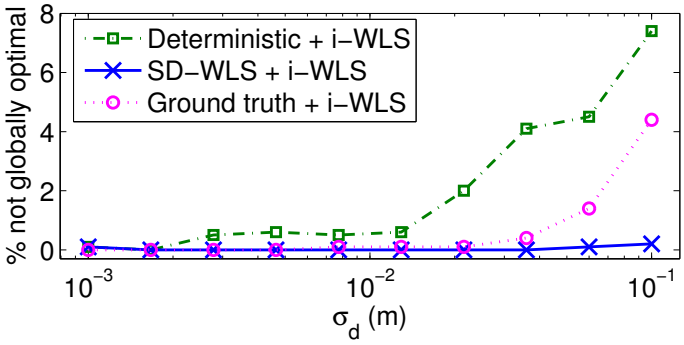

(c)

Figure 3.4: 2-D: Comparison of Levenberg-Marquardt-based i-WLS initialized with the result of the linear algorithm of [133], with the result of SD-WLS, and with ground truth. (a) Number of iterations to convergence (median and 25% / 75% quantiles). (b) Percentage of inconsistent cases (error $> 3\sigma$). (c) Percentage of cases where i-WLS does not reach the global optimum, depending on the method of initialization.

the relative pose estimate with lowest WLS cost function value, chosen from the results of i-WLS initialized with ground truth, linear algorithm, or SD-WLS, respectively, which we hope is a reasonably good approximation of the true global optimum.

Figure 3.7 shows the comparison of i-WLS initialized with the solution from SD-WLS, the linear method, and with ground truth. Again, initialization with SD-WLS requires the smallest number of iterations (with lowest variability) in the Levenberg-Marquardt i-WLS optimization, and also results in considerably fewer inconsistent estimates than i-WLS initialized with the linear method.

Figure 3.7(c) shows the percentage of cases where i-WLS does not converge to the (approximate) global optimum. As in 2-D, we observe cases where i-WLS initialized with ground truth fails to reach the global optimum, especially for large noise variance. Initializing i-WLS with the results from the linear algorithm results in a local optimum in a significant percentage of cases (up to 10%, and more often than in 2-D), even for low measurement noise standard deviation. If a solution has converged to a local optimum, the median difference of the cost function value to the (approximate) global optimum is 46.3 (2,756%) if initialized using the linear method, 0.92 (55%) if initialized with ground truth, and 0.05 (4.2%) if initialized with SD-WLS, which again shows that the linear method can sometimes yield unusable estimates far from the optimal solution. The median distance between a locally optimal and (approximately) globally optimal relative position estimate is 0.99 m (linear initialization), 0.74 m (initialization with ground truth) and 0.24 m (initialization with SD-WLS). Given that the relative distance between the robots was only between 1 m and 2 m, these errors are significant for all methods, and indicate that the local optima can be physically far from each other and from the global optimum.

Figure 3.8 depicts the percentage of problem instances for which SD-WLS optimality can be certified using the method of Lagrange relaxation described in chapter 3.5.3. Notice that for 12 measurements, 99% of cases can be certified efficiently, but even for 10 measurements, Lagrange relaxation certifies 42% of problem instances. Therefore, in particular with growing number of measurements, checking optimality using Lagrange relaxation before resorting to SOS relaxation can lead to significant savings of time and computation. Note that Figure 3.8 shows the percentage for very small distance measurement noise. We have observed that the curve flattens slightly with increasing measurement noise, with a minimum of approximately 65% of cases certified for 12 measurements with $\sigma_d = 0.1$ m, which is still considerable.

The simulation results demonstrate the superiority of SD-WLS over the deterministic algorithms, both if used stand-alone and with i-WLS refinement. However, we also note that the computational complexity of the deterministic solvers and Levenberg-Marquardt i-WLS is significantly lower than the SD-WLS solution methods presented in this chapter. Our simulation results are intended to provide decision guidelines for the tradeoffs between accuracy and computation requirements, which are usually application dependent.

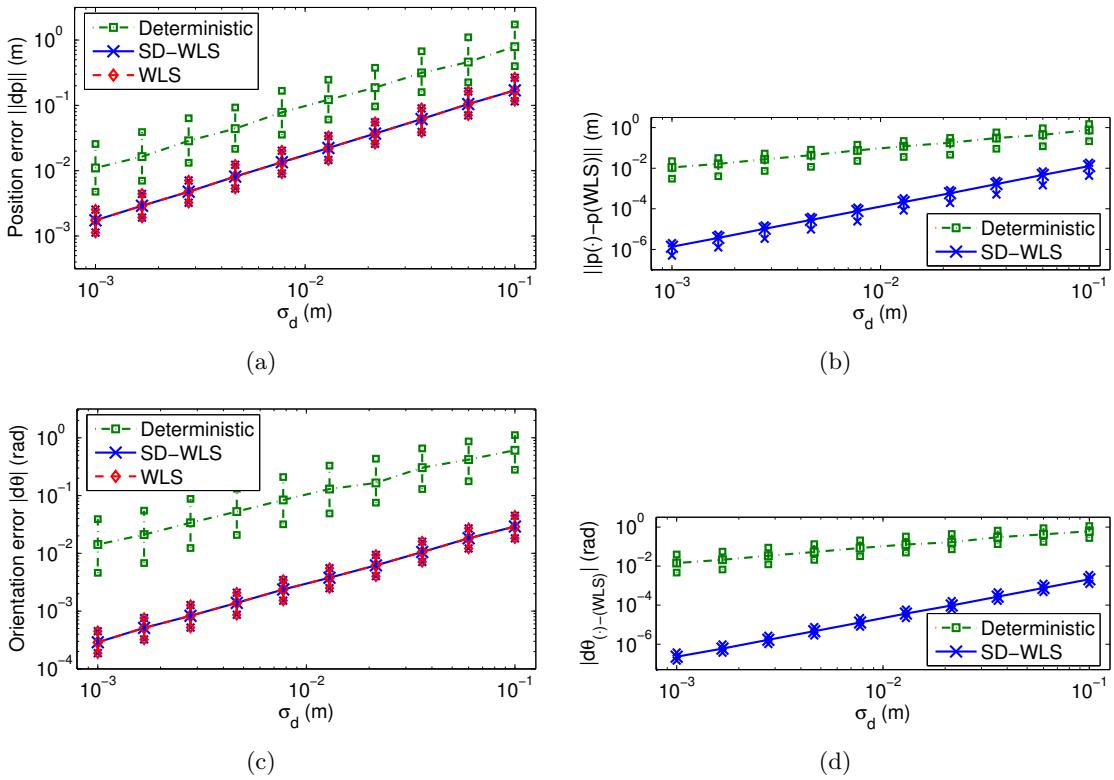

Figure 3.5: 3-D Simulation results from 1000 trials for each noise setting, comparing WLS, SD-WLS and the linear algorithm of [122], using 10 distance measurements corrupted by zero-mean Gaussian measurement noise with covariance $\mathbf{R} = \sigma_d^2 \mathbf{I}$. Plotted are the median and the 25% and 75% quantiles. The accuracy of the WLS and the SD-WLS formulation is almost identical, and higher than that of the linear method by a constant factor. The SD-WLS solution is close to that of WLS particularly for small noise values. (a) Norm of position error with respect to ground truth. (b) Difference in position with respect to WLS. (c) Orientation error with respect to ground truth. (d) Difference in orientation with respect to WLS.

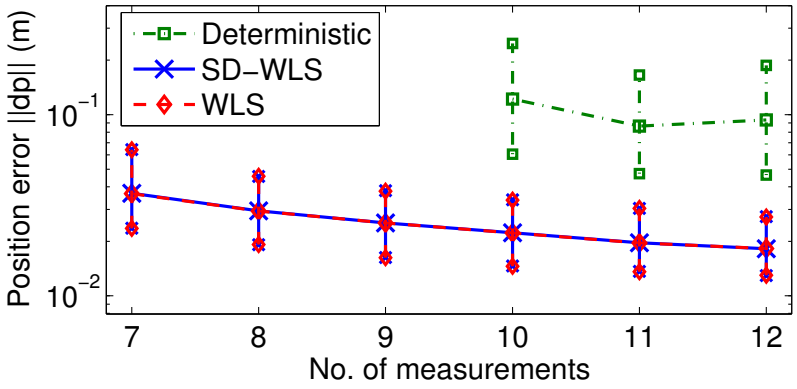

Figure 3.6: 3-D: Median position error and 25% and 75% quantiles for $\sigma_d = 0.013$ m as a function of the number of measurements. WLS (dashed line) and SD-WLS (solid line) show indistinguishable performance that decreases monotonically with increasing number of measurements. The deterministic algorithm of [122] (dash-dotted line) also benefits from an increased number of measurements, but exhibits considerably larger errors and variability.

3.7 Experimental results

In addition to simulations, we have tested the distance-based relative pose estimation algorithms in 2-D and 3-D in two real-world experiments. The 2-D experiment was carried out using two Pioneer-II robots, with ground truth established by a calibrated overhead camera. In the 3-D experiment we performed extrinsic calibration of two Scorbot ER-VII manipulator arms instead of mobile robots, in order to obtain reliable ground truth. The results clearly demonstrate the ability of the proposed algorithms to provide accurate estimates in realistic applications.

3.7.1 2-D case: Extrinsic calibration of two Pioneer-II robots

The 2-D experiment was carried out using two Pioneer-II robots moving randomly in an area of 4 m $\times$ 5 m (see Fig. 3.9(a)). Each robot estimated its position from 10 Hz wheel odometry, using a differential drive kinematic model with noisy wheel velocity measurements having standard deviation of $\sigma_v = 8$ mm/s. The ground truth was established from observations using a calibrated ceiling-mounted camera. These data also provided synthetic relative distance measurements by adding white, zero-mean Gaussian noise with standard deviation $\sigma_d = 0.05$ m.

We compared the estimates of the initial relative pose obtained using the SD-WLS, the linear algorithm with 5 measurements, and the i-WLS initialized with SD-WLS. For the SD-WLS solution we processed two iterations. The first accounted only for noise in the distance measurements, and its optimum was used as linearization point for the Jacobians with respect to the robot positions, in order to correctly account for uncertainty due to odometry in the second iteration (see chapter 3.4.2). The results in Fig. 3.9(b) show that SD-WLS

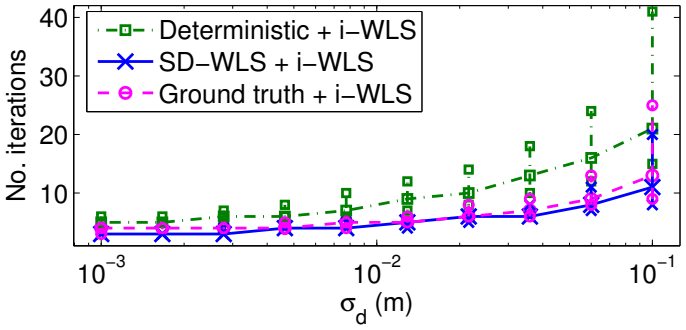

(a)

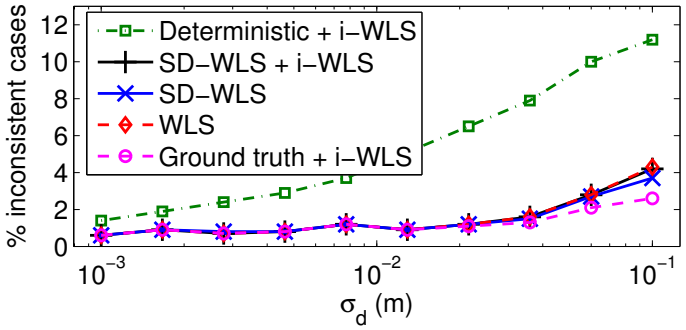

(b)

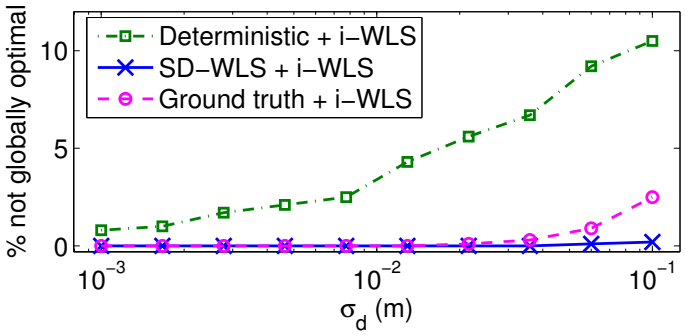

(c)

Figure 3.7: 3-D: Comparison of i-WLS initialized with the result of the linear algorithm of [118], that of SD-WLS, and with ground truth. (a) Number of iterations to convergence (median and 25% / 75% quantiles). (b) Percentage of inconsistent cases (error $> 3\sigma$). (c) Percentage of cases where i-WLS does not reach the global optimum, depending on the method of initialization.

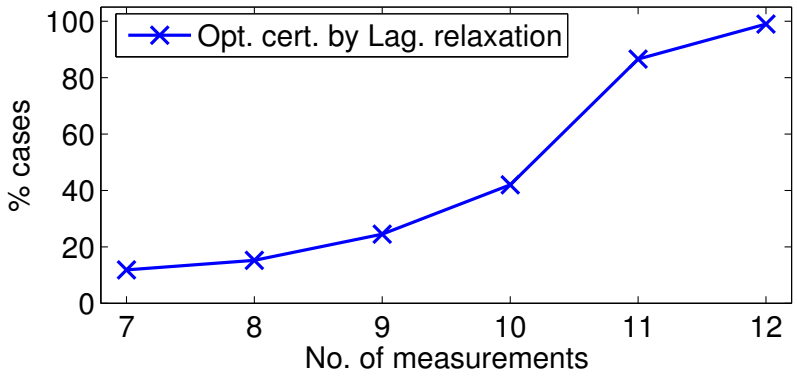

Figure 3.8: 3-D: Percentage of cases where Lagrange relaxation certifies global optimality as a function of the number of measurements, for $\sigma_d = 0.001m$.

is significantly more accurate than the linear method. Notice that the accuracy is almost independent of the number of measurements, since the growing position uncertainty due to integrated odometry errors cancels gains from processing additional distance measurements. The errors after i-WLS refinement are in the order of 10 cm for relative position and 3 degrees for relative orientation.

3.7.2 3-D case: Extrinsic calibration of two Scorbot ER-VII manipulators

For the 3-D experiment, we computed the relative pose between the base frames of two Scorbot ER-VII manipulator arms mounted on two tables (see Fig. 3.10). For this purpose, we moved the manipulator tips to 12 random positions, and manually measured the distance between them using a measuring tape. The position of the tips, i.e., the values of $\mathbf{u}_i$ and $\mathbf{v}_i$, were determined from the encoder readings and the robot forward kinematics. We assumed the measurement uncertainty as $\sigma_d = 3$ mm, which accounted for measurement errors, as well as for inaccuracies in the manipulator intrinsic calibration.

To establish ground truth, we operated under the assumption that the robots' z-axes were aligned, and perpendicular to the table surfaces. As a result, the true relative roll and pitch angles are assumed to be zero. Also, the true difference in height can be determined by measuring the vertical distance between the table surfaces, and by reading off the z-coordinate of the table surface from the gripper position when it is touching the table top. The remaining degrees of freedom, i.e., the true relative position in x and y, as well as the yaw angle, were measured using a 2-D Sick-LMS laser scanner. The laser scan plane intersected the cylindrical robot base, as well as two flat panels attached to the robots' arms when they were moved to their home positions aligned with the robots' x-axes. The vector connecting the centers of two circles fitted to the reflections of the robots' bases provided the relative position $\mathbf{p}$ between the robots' base frames, expressed in the laser frame. Two lines fitted to the panel reflections provided the direction of the robots' x-axes in the laser frame. Together with the assumption that both robots' z-axes were aligned with the laser

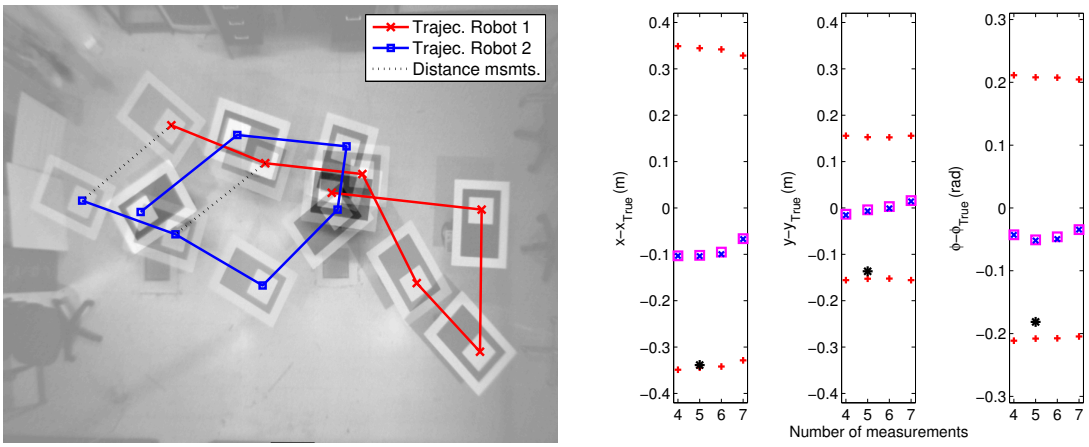

(a) (b)

Figure 3.9: Experiment with two Pioneer robots equipped with markers for optical tracking using an overhead camera. (a) The camera images provided ground truth for relative pose, and served to create 7 synthetic relative distance measurements at the marked points (for clarity, only the first two measurements are drawn as dotted lines). (b) Error in relative pose as a function of the number of measurements. Shown are i-WLS refinement of SD-WLS (blue x) and corresponding 3σ-bounds (red +), as well as the output of pure SD-WLS (magenta square) and the linear algorithm (black star, for 5 measurements only).

z-axis, this completely determined the orientation of the robots' base frames with respect to the laser frame. This in turn allowed computing the true yaw angle, and the ground truth value of $\mathbf{p}$ expressed in either manipulator's frame.[10]

We have compared the ground truth with the results of the linear algorithm [118] with 10-12 measurements, SD-WLS, and of i-WLS initialized with SD-WLS (see Fig. 3.11). Clearly, the results from SD-WLS and i-WLS are almost indistinguishable, which was expected due to the small measurement noise variance. The large errors incurred by the linear algorithm using 10 and 11 measurements demonstrate the possibly large errors of deterministic algorithms when faced with noisy measurements, even in relatively low noise settings. However, notice that using 12 measurements, the linear algorithm yields a relatively accurate solution, even though it is still clearly outperformed by the least squares approaches.

The accuracy achieved with 12 distance measurements is in the order of 3 cm for position, and 2.1 deg for orientation (3σ). The residuals between measured and expected distances for the least squares solution with 12 measurements are less than 3.5 mm, indicating a very precise fit. Contrasting this to the relatively large uncertainty in the final calibration confirms the relatively low information content of distance measurements for the purpose of relative pose estimation, as compared, e.g., to relative distance and bearing. An important

[10] As evident from the description, the process to obtain ground truth is rather tedious, and does not provide an easy mechanism to accurately compute all rotational degrees of freedom for general configurations. This should be contrasted to the easy process of obtaining distance measurements between the tool tips, which lends itself to automation.

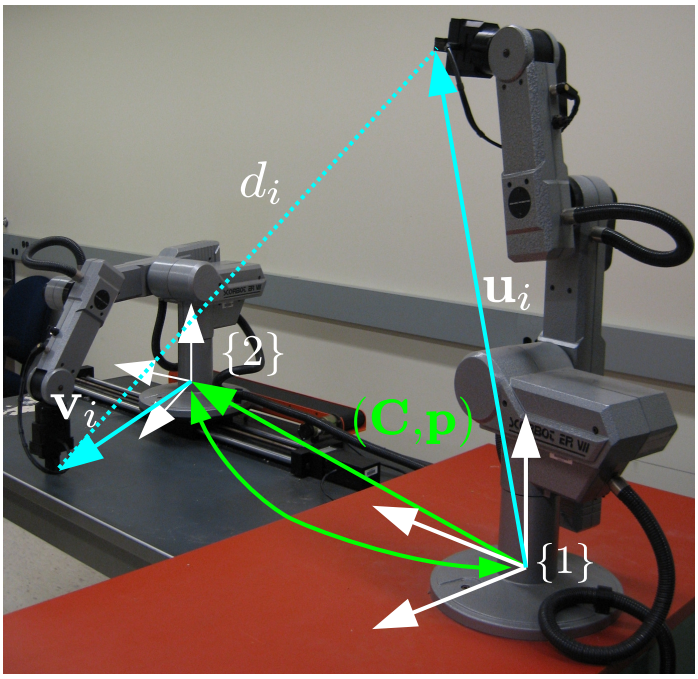

Figure 3.10: Experiment to determine the relative pose between the base frames of two Scorbot ER-VII manipulators. Distance between the gripper tips was measured manually using a measuring tape.

topic of ongoing research is the impact of measurement geometry on the estimation accuracy (similar to geometric dilution of precision (GDOP) in GPS), and the problem of determining the optimal gripper locations to obtain the most informative distance measurements.

3.8 Summary

In this chapter, we have presented methods to estimate the relative pose between two robots based on robot-to-robot distance measurements and knowledge of the robots' egomotion, both in 2-D and in 3-D. In particular, we have proposed to directly determine the global optimum of the corresponding WLS problem by finding all stationary points as the roots of a square multivariate polynomial system. We have shown how to construct this polynomial system for the WLS of the original problem, which is also the MLE for distance measurements corrupted by Gaussian noise. The complexity of solving this system was shown to scale very poorly, since the number of variables grows with the number of measurements. On the other hand, the alternative SD-WLS formulation, using squared distance measurements, was shown to have constant solution complexity, independent of the number of measurements or the structure of the covariance matrix. In 2-D, solving the problem is carried out very efficiently using recent hybrid algebraic-numeric techniques to solve multivariate polynomial systems based on the eigendecomposition of a generalized companion matrix [117]. In 3-D, we have presented SOS- and Lagrange relaxations to efficiently solve

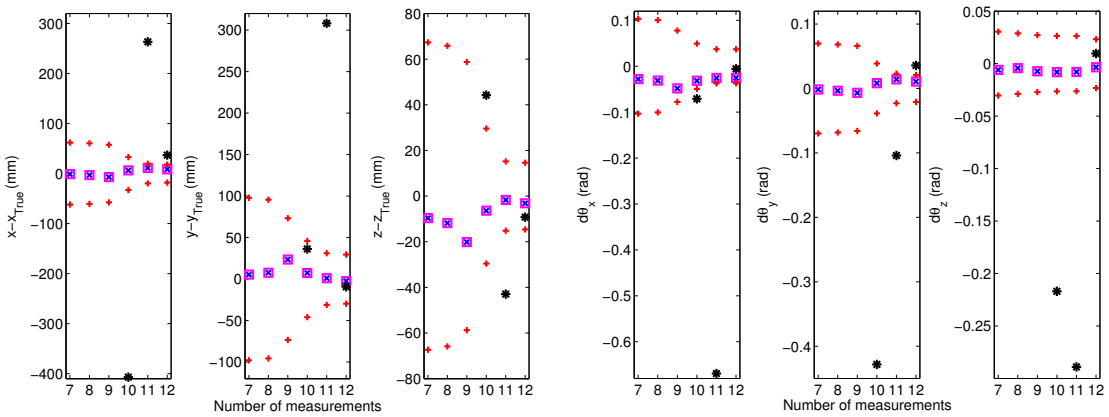

(a) (b)

Figure 3.11: Extrinsic calibration of two Scorbot ER-VII manipulators. (a) Error in relative position as a function of the number of measurements. Shown are i-WLS refinement of SD-WLS (blue x) and corresponding 3σ-bounds (red +), as well as the output of pure SD-WLS (magenta square) and the deterministic algorithm [118] (black stars, for 10-12 measurements only). (b) Error in relative orientation.

the SD-WLS problem. Despite its name, SOS relaxation is able to find and certify the globally optimal solution to the SD-WLS problem. Lagrange relaxation was presented as a fast, alternative method to certify global SD-WLS optimality of a candidate stationary point. Simulation results showed zero duality gap for a significant percentage of problem instances, rapidly growing with increasing number of measurements. Despite losing optimality in the MLE sense, the SD-WLS method performed almost indistinguishably from the WLS estimator in simulation, and was shown to be far more accurate than the deterministic methods presented in [133] and [118], over a wide range of noise levels.

In a broader context, the research presented in this chapter demonstrates that the newly available tools for polynomial-system solving allow moving towards global optimization, if the cost function and the first-order optimality conditions can be transformed into polynomial form. The latter is often the case for geometric problems in robotics (kinematics) and computer vision. Moreover, we have provided a concrete application example (i.e., robot-to-robot extrinsic calibration) for which we showed that for overdetermined systems, combining polynomial constraints in a WLS formulation is far more accurate than solving the non-square system of polynomial constraints deterministically.

Chapter 4

Cooperative Localization under Communication Constraints

While the previous chapter has presented a method to initialize CL, this chapter will focus on the joint estimation task of CL itself once initialization is complete. In particular, we will concentrate on CL when communication bandwidth is severely limited and fluctuating. To adapt to these communication constraints, we propose to use quantized measurements in the estimation process. To this end, this chapter will describe existing MMSE and novel MAP estimators to estimate dynamic random processes using quantized measurements, and apply them to the problem of CL. Part of this chapter has been published in a conference paper [121].

4.1 Introduction

Networks of mobile robots have recently been proposed for tasks such as aerial surveillance [9], search and rescue operations [114], underwater- [17] or even space exploration [55]. For any robotic task, accurate localization, i.e., determining the position and orientation (pose) of each robot, is a fundamental requirement. In contrast to the simple solution of localizing each robot in a team independently, Cooperative Localization (CL) [99, 44, 103] incorporates robot-to-robot observations and *jointly* estimates all robots' poses, which significantly improves localization accuracy for all team members [79]. Unfortunately, this benefit comes at the cost of increased computation and communication.

Recent research has focussed primarily on reducing the *computational requirements* of CL (e.g., [88, 53, 74, 83]). However, also the *communication requirements* of CL are substantial, since, depending on the approach, either measurements from both proprioceptive and exteroceptive sensors, or state estimates and covariance matrices have to be exchanged. Communicating all these quantities might be infeasible if the robots are subject to communication constraints due to (i) the nature of the *operational environment* (e.g., limited

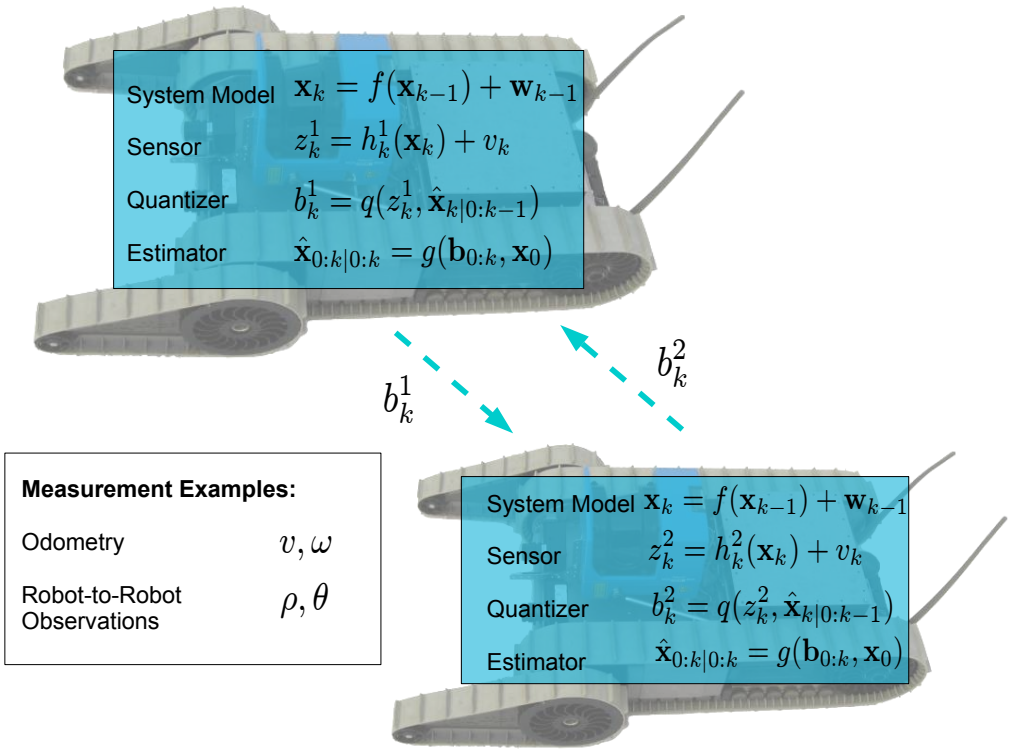

Figure 4.1: Multi-centralized, Cooperative Localization using quantized measurements with two robots. At every time step k, each robot i quantizes its measurements (e.g., linear and rotational velocity, robot-to-robot distance and bearing) and broadcasts the resulting bit sequences, $\mathbf{b}_k^i$, to its partner. Every robot computes the same estimate, $\hat{\mathbf{x}}_{0:k|0:k}$, of the pose of the entire team, based on all quantized measurements available up to time step k, as well as on shared knowledge of system- and measurement models and their associated noise characteristics.

bandwidth caused by strong signal attenuation and distortion underwater or underground), (ii) the specific *application domain* (e.g., stealth requirements in military operations), or (iii) *economic motivations* (e.g., in order to extend the operational life-span through lower power consumption resulting from reduced data transmissions).

In this chapter, we investigate algorithms that can perform CL with significantly reduced communication, by transmitting (severely) quantized instead of real-valued sensor observations. For this purpose, we present a minimum mean square error (MMSE) filter [81] and derive a new maximum a posteriori (MAP) batch estimator specifically designed for dynamic random process estimation with quantized observations. We show that, despite significantly reduced communication requirements, their performance in CL comes very close to that of real-valued estimators. In particular, we propose a "multi-centralized" system (cf. Fig. 4.1), where each robot executes the exact same estimation algorithm, and broadcasts its quantized observations to all other robots. Combined with the shared knowledge of system- and sensor models, all robots are able to compute identical quantization thresholds and identical

joint state estimates.

The contributions of this chapter are threefold: (i) we develop for the first time single- and multi-bit MAP estimators for estimation of random processes with quantized observations, (ii) we develop optimal threshold selection and bit allocation schemes for time-varying bandwidth, and (iii) we apply quantized filter and MAP algorithms as a novel solution to the problem of multi-robot Cooperative Localization under communication constraints.

4.2 MMSE estimation

Let us assume a discrete-time, linear dynamic system with M_k scalar measurements per time step

$$\mathbf{x}_k = \mathbf{F}_{k-1}\mathbf{x}_{k-1} + \mathbf{G}_{k-1}\mathbf{w}_{k-1}, \quad \mathbf{x}_0 \sim \mathcal{N}(\mathbf{x}_0; \mathbf{x}_{\text{ini}}, \mathbf{P}_0) \tag{4.1}$$

$$z_{km} = \mathbf{h}_{km}^T \mathbf{x}_k + v_{km}, \quad m = 1, \ldots, M_k \tag{4.2}$$

with zero-mean, white Gaussian, uncorrelated system and measurement noise with covariance

$$E[\mathbf{w}_k \mathbf{w}_l^T] = \delta_{kl}\mathbf{Q}_k, \; E[v_{km} v_{ln}] = \delta_{km,ln}\sigma_{km}^2 \tag{4.3}$$

where δ_{kl} is the Kronecker delta ($\delta_{kl} = 1$ if $k = l$, else $\delta_{kl} = 0$). Following the discussion in chapter 3, we will assume that initialization has been completed, and the initial estimate, $\mathbf{x}_{\text{ini}}$, and its covariance, $\mathbf{P}_0$, are known (see chapter 3.4.3).

Notice that in this formulation the only control input is system noise. Usually in mobile robotics, the system dynamics are given by a kinematic model, with linear and rotational velocity and/or acceleration measurements acting as control input. Here, we assume instead a statistical motion model (e.g., zero acceleration driven by white noise [8]), allowing us to include the control variables (e.g., linear and rotational velocities) in the state vector. As a consequence, measurements from proprioceptive and exteroceptive sensors can be treated identically. Notice further that this formulation also allows vector-valued measurements, since these can be decomposed into several scalar measurements after appropriate pre-whitening. Finally, we note that the linear formulation (4.1)-(4.3) is used for ease of presentation. Later on, we will use linearization to accommodate the *nonlinear* system and measurement models prevalent in robotics.

4.2.1 Kalman filter

It is well known that the optimal MMSE estimate of the state $\mathbf{x}_k$ given all measurements up to time step k (denoted as $\hat{\mathbf{x}}_{k|k}$ and $\mathbf{z}_{0:k}$, respectively) is the conditional mean of the

posterior pdf $p[\mathbf{x}_k|\mathbf{z}_{0:k}]$

$$\hat{\mathbf{x}}_{k|k} := E[\mathbf{x}_k|\mathbf{z}_{0:k}] = \int_{-\infty}^{\infty} \mathbf{x}_k p[\mathbf{x}_k|\mathbf{z}_{0:k}] d\mathbf{x}_k \tag{4.4}$$

For a linear system with Gaussian noise, both the prior and the posterior pdf are Gaussian, and hence are completely characterized by their mean and covariance. Both moments of the pdf are computed by the regular Kalman filter, described in Algorithm 2. Notice that the state update requires knowledge of the (real-valued) measurement z_{km}.

Algorithm 2 Kalman filter

1: KF Propagation

$$\begin{aligned}\hat{\mathbf{x}}_{k|k-1} &= \mathbf{F}_{k-1}\hat{\mathbf{x}}_{k-1|k-1} \\ \mathbf{P}_{k|k-1} &= \mathbf{F}_{k-1}\mathbf{P}_{k-1|k-1}\mathbf{F}_{k-1}^T + \mathbf{G}_{k-1}\mathbf{Q}_{k-1}\mathbf{G}_{k-1}^T\end{aligned}$$

2: KF Measurement Update, $m = 1, \ldots, M_k$

$$\begin{aligned}\hat{\mathbf{x}}_{k|k,m} &= \hat{\mathbf{x}}_{k|k,m-1} + \frac{\mathbf{P}_{k|k,m-1}\mathbf{h}_{km}}{\mathbf{h}_{km}^T\mathbf{P}_{k|k,m-1}\mathbf{h}_{km} + \sigma_{km}^2}(z_{km} - \mathbf{h}_{km}^T\hat{\mathbf{x}}_{k|k,m-1}) \\ \mathbf{P}_{k|k,m} &= \mathbf{P}_{k|k,m-1} - \frac{\mathbf{P}_{k|k,m-1}\mathbf{h}_{km}\mathbf{h}_{km}^T\mathbf{P}_{k|k,m-1}}{\mathbf{h}_{km}^T\mathbf{P}_{k|k,m-1}\mathbf{h}_{km} + \sigma_{km}^2} \\ \hat{\mathbf{x}}_{k|k,0} &:= \hat{\mathbf{x}}_{k|k-1},\ \hat{\mathbf{x}}_{k|k} := \hat{\mathbf{x}}_{k|k,M_k},\ \mathbf{P}_{k|k,0} := \mathbf{P}_{k|k-1},\ \mathbf{P}_{k|k} := \mathbf{P}_{k|k,M_k}\end{aligned}$$

4.2.2 Single- and multi-bit Kalman filter (SOI-KF and IQKF)

Assume now, that due to communication constraints we cannot afford to transmit the entire measurement z_{km}, but only a quantized version b_{km}. Further, assume that each robot can broadcast its quantized measurements to all other robots (i.e., by transmitting at a fixed frequency, or in a round-robin fashion), and that error-correcting codes ensure error-free transmission[1]. Finally, assume known system-, measurement-, and noise models, and a commonly shared state estimate throughout the entire network. The last assumption is guaranteed (after appropriate initialization) by requiring each member of the network to execute the *same* estimation algorithm using the *same* quantized observations (a process we refer to as "multi-centralized" estimation). As a result, every robot is able to reproduce the quantization thresholds used by the other robots, since they are uniquely determined based on the shared estimate [102].

The quantization of the measurements introduces nonlinearity into the system, causing the Gaussianity of the posterior pdf to break down, even when the prior pdf is Gaussian

[1]The specific design of the scheduling and routing algorithms is beyond the scope of this work.

Algorithm 3 SOI-KF

1: SOI-KF Propagation ≡ KF Propagation (cf. Alg. 2)

2: SOI-KF Quantization Rule

$$b_{km} = \begin{cases} +1 & \text{if } z_{km} - \mathbf{h}_{km}^T \hat{\mathbf{x}}_{k|k,m-1} > 0 \\ -1 & \text{otherwise} \end{cases}$$

3: SOI-KF Measurement Update, $m = 1, \ldots, M_k$

$$\hat{\mathbf{x}}_{k|k,m} = \hat{\mathbf{x}}_{k|k,m-1} + \sqrt{\frac{2}{\pi}} \frac{\mathbf{P}_{k|k,m-1}\mathbf{h}_{km}}{\sqrt{\mathbf{h}_{km}^T \mathbf{P}_{k|k,m-1}\mathbf{h}_{km} + \sigma_{km}^2}} b_{km}$$

$$\mathbf{P}_{k|k,m} = \mathbf{P}_{k|k,m-1} - (2/\pi) \frac{\mathbf{P}_{k|k,m-1}\mathbf{h}_{km}\mathbf{h}_{km}^T \mathbf{P}_{k|k,m-1}}{\mathbf{h}_{km}^T \mathbf{P}_{k|k,m-1}\mathbf{h}_{km} + \sigma_{km}^2}$$

$$\hat{\mathbf{x}}_{k|k,0} := \hat{\mathbf{x}}_{k|k-1},\ \hat{\mathbf{x}}_{k|k} := \hat{\mathbf{x}}_{k|k,M_k},\ \mathbf{P}_{k|k,0} := \mathbf{P}_{k|k-1},\ \mathbf{P}_{k|k} := \mathbf{P}_{k|k,M_k}$$

and both system- and measurement models are linear. As a result, exact MMSE estimation becomes, in general, intractable. However, if the quantized measurement is chosen as the sign of the innovation, and at each time step the posterior pdf $p[\mathbf{x}_k|\mathbf{b}_{0:k}]$ is approximated as Gaussian with mean $\hat{\mathbf{x}}_{k|k}$ and covariance $\mathbf{P}_{k|k}$, then, as shown in [102], the resulting SOI-KF estimator (cf. Algorithm 3) has a form very similar to the KF. Remarkably, one can show that in a continuous-time formulation, the performance (in terms of MSE) of the SOI-KF is the same as that of a real-valued KF with $\pi/2$ times larger measurement noise variance [102]. In other words, by transmitting only a single bit instead of a real-valued scalar, one still achieves about 64% of the performance of the regular KF.

More recent work [81] has shown that similar results hold for quantization with more than one bit per measurement. Specifically, the SOI-KF measurement update shown in Algorithm 3 can be repeated for each bit in an iterative fashion, leading to the iteratively quantized KF, or IQKF. The performance is shown to improve with the number of bits, and even for as little as 4 bits, it reaches 98% of the performance of the real-valued KF.[2]

4.3 MAP estimation

Notice that neither SOI-KF nor IQKF are exact, since they approximate the (non-Gaussian) posterior pdf by a Gaussian at every step. For nonlinear systems, the additional linearization of system- and/or measurement model in the EKF (or, analogously, in the SOI-EKF [102] or IQEKF [81]) will only exacerbate this problem, and may lead over time to filter inconsistency. To overcome these limitations, in this section we investigate real-valued and

[2] In order to properly account for correlations, the state has to be augmented with the measurement noise in this case. For further details, we refer the interested reader to [81].

quantized batch MAP estimation. We expect the smoothing effect of batch estimation to mitigate errors arising from linearization and Gaussian approximation.

A batch MAP estimate is the best estimate of the *entire* state history (denoted as $\mathbf{x}_{0:K}$) given all measurements up to time K (denoted as $\mathbf{z}_{0:K}$). Contrary to the MMSE, which is the *mean* of the posterior pdf, the MAP estimate is given by its *mode*. Using Bayes' rule and the Markov property of the dynamic system, we can write the batch MAP estimate as

$$\begin{aligned}
&\hat{\mathbf{x}}_{0:K|0:K} = \arg\max p(\mathbf{x}_{0:K}|\mathbf{z}_{0:K}) \\
&= \arg\max \frac{1}{p(\mathbf{z}_{0:K})} p(\mathbf{z}_{0:K}|\mathbf{x}_{0:K}) p(\mathbf{x}_{0:K}) \\
&= \arg\max \prod_{k=0}^{K} \prod_{m=1}^{M_k} p(z_{km}|\mathbf{x}_k) \prod_{k=0}^{K-1} p(\mathbf{x}_{k+1}|\mathbf{x}_k) \cdot p(\mathbf{x}_0) \\
&= \arg\min \Big(- \sum_{k=0}^{K} \sum_{m=1}^{M_k} \log p(z_{km}|\mathbf{x}_k) - \sum_{k=0}^{K-1} \log p(\mathbf{x}_{k+1}|\mathbf{x}_k) - \log p(\mathbf{x}_0) \Big)
\end{aligned} \tag{4.5}$$

For the linear system with Gaussian noise (4.1)-(4.3), the conditional pdfs in (4.5) are also Gaussian, and their logarithm is a quadratic form (up to a constant), so that the MAP estimate can be formulated as a standard weighted Least Squares (WLS) problem (see Sec. 4.3.3).

4.3.1 Single-bit quantized MAP estimation (QMAP)

The situation becomes significantly more challenging when replacing real-valued by quantized measurements. For illustrative purposes, we first analyze the case where each measurement is quantized with exactly one bit. Similarly to the SOI-KF, we will compute the quantized measurement using as threshold, τ_{km}, the expected measurement based on the current MAP estimate, which is the optimal choice as shown later in Lemma 2.

QMAP Quantization Rule:

$$\begin{aligned}
b_{km} &= \begin{cases} +1 & \text{if } z_{km} - \tau_{km} > 0 \\ -1 & \text{otherwise} \end{cases} \\
\text{where } \tau_{km} &= \mathbf{h}_{km}^T \hat{\mathbf{x}}_{k|0:k,m-1}
\end{aligned} \tag{4.6}$$

For computing the posterior pdf and the MAP estimate as in (4.5), we will have to determine the quantized measurement likelihoods $p(b_{km}|\mathbf{x}_k)$ instead of $p(z_{km}|\mathbf{x}_k)$. Due to the assumption of Gaussian measurement noise, we can compute these measurement

likelihoods in terms of the Gaussian tail probability or Q-function (see Appendix A.1) as

$$p(b_{km}|\mathbf{x}_k) = Q\left(\frac{b_{km}(\tau_{km} - \mathbf{h}_{km}^T\mathbf{x}_k)}{\sigma_{km}}\right) \tag{4.7}$$

$$\begin{aligned}\text{since } p(b_{km} = 1\,|\mathbf{x}_k) &= \Pr\left\{z_{km} - \tau_{km} > 0\,|\mathbf{x}_k\right\} \\ &= \Pr\left\{v_{km} > \tau_{km} - \mathbf{h}_{km}^T\mathbf{x}_k\,|\mathbf{x}_k\right\}\end{aligned}$$

and similarly for $p(b_{km} = -1\,|\mathbf{x}_k)$. We are now ready to state the main result for the single-bit quantized MAP in the following proposition:

Proposition 2 (Single-Bit Quantized MAP (QMAP))**.** *Assume the linear model of* (4.1)-(4.3)*. If a single bit is allocated per measurement, with a quantization rule as in* (4.6)*, the posterior pdf is given by*

$$\begin{aligned}&p(\mathbf{x}_{0:K}|\mathbf{b}_{0:K}) \\ &= \frac{1}{p(\mathbf{b}_{0:K})}\prod_{k=0}^{K}\prod_{m=1}^{M_k} p(b_{km}|\mathbf{x}_k)\prod_{k=0}^{K-1} p(\mathbf{x}_{k+1}|\mathbf{x}_k)\cdot p(\mathbf{x}_0) \\ &\propto \prod_{k=0}^{K}\prod_{m=1}^{M_k} Q\left(\frac{b_{km}(\tau_{km} - \mathbf{h}_{km}^T\mathbf{x}_k)}{\sigma_{km}}\right)\prod_{k=0}^{K-1}\mathcal{N}(\mathbf{x}_{k+1};\mathbf{F}_k\mathbf{x}_k, \mathbf{G}_k\mathbf{Q}_k\mathbf{G}_k^T)\cdot\mathcal{N}(\mathbf{x}_0;\mathbf{x}_{\text{ini}},\mathbf{P}_0)\end{aligned} \tag{4.8}$$

The proposition makes again use of Bayes' rule and the Markov property of the system, just as in the derivation of (4.5), with the important difference of using the quantized instead of the analog (Gaussian) measurement likelihood.

The following Lemma is an important consequence of this proposition:

Lemma 1. *The posterior pdf of the QMAP given in* (4.8) *is log-concave in* $\mathbf{x}$*.*

Proof. The lemma follows from the facts that the Gaussian pdf is log-concave [14], integrals of log-concave pdfs over convex sets are log-concave [95], hence making $p(b_{km}|\mathbf{x}_k)$ log-concave, and log-concavity is closed under multiplication [14]. □

Log-concavity ensures that the MAP estimate is unique for every choice of τ_{km} and can be found using efficient convex optimization techniques [14], which are guaranteed to converge to the global optimum.

Fig. 4.2 illustrates the difference between real-valued and QMAP estimation for a simple case of a scalar prior, and a single, scalar measurement. Note that while the real-valued posterior is Gaussian, the quantized posterior is skewed and heavy-tailed. However, its log-concavity ensures that the MAP estimate is unique and can be computed efficiently.

The following lemma states that for the case of a Gaussian prior and scalar measurement, quantized by a single bit, choosing the quantization threshold τ_{km} as the expected value of the measurement is not only optimal for MMSE [102], but also for MAP estimation.

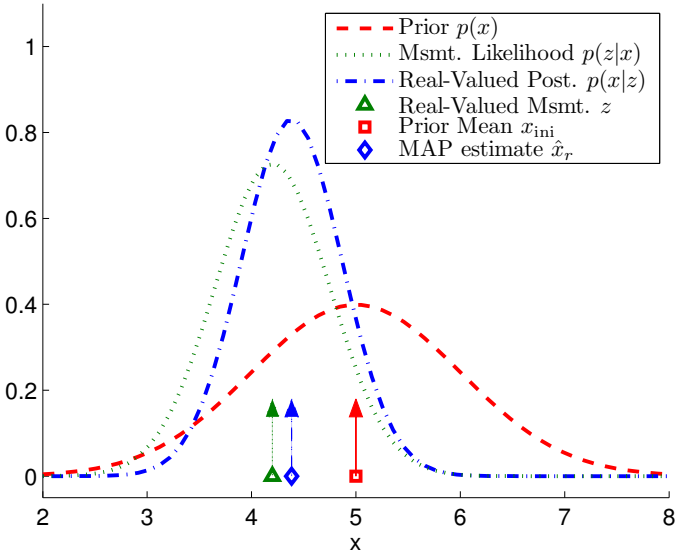

(a) Pdfs for real-valued estimation, and the corresponding MAP estimate $\hat{x}_r$.

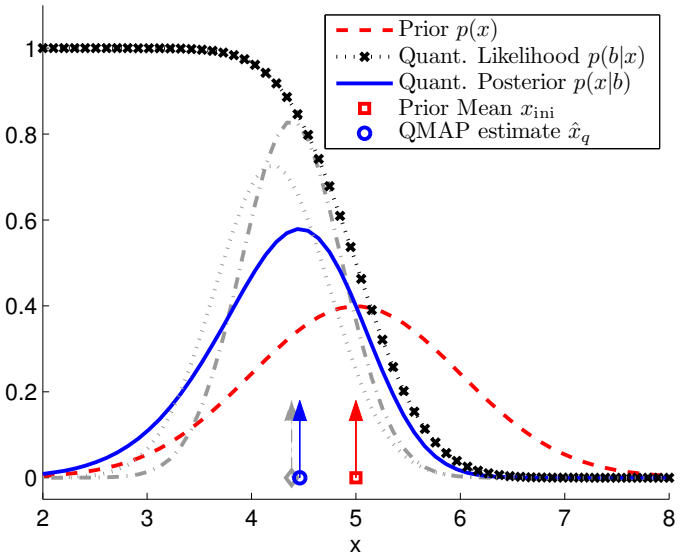

(b) Pdfs for quantized estimation, and the corresponding MAP estimate $\hat{x}_q$ (pdfs for real-valued are shown shaded for comparison).

Figure 4.2: Comparison of real-valued and quantized measurement-based posterior pdfs for scalar prior and one scalar measurement $z = x + v$, $x \sim \mathcal{N}(x; 5, 1)$, $v \sim \mathcal{N}(v; 0, 0.3)$, for the realization $z = 4.2$ quantized with a single bit at the threshold $\tau = x_{\text{ini}} = 5$. The posterior for the quantized measurement is skewed and heavy-tailed, compared to the real-valued (Gaussian) posterior.

Lemma 2. *Assume the linear model of* (4.1)-(4.3)*. If a single bit is allocated per measurement, with a quantization rule as in* (4.6)*, this choice of quantization threshold minimizes the Bayesian "hit-or-miss" risk function used in MAP estimation.*

Proof. See Appendix A.2 □

4.3.2 Multi-bit iteratively quantized MAP estimation (IQMAP)

In case *multiple bits* are available for quantizing a particular measurement, we require an iterative, multi-bit quantization scheme with *adaptive* quantization thresholds. For this purpose, we define the following quantization rule to determine the ith bit allocated to measurement z_{km}.

IQMAP Quantization rule:

$$b_{km,i} = \begin{cases} +1 & \text{if } z_{km} - \tau_{km,i} > 0,\ i = 1, \ldots, I \\ -1 & \text{otherwise} \end{cases} \tag{4.9}$$

We will address how to choose $\tau_{km,i}$ momentarily.

The idea behind this approach is to define an interval containing z_{km} and iteratively reduce its size. Before receiving the first bit, the measurement z_{km} is known trivially to lie in the interval $-\infty < z_{km} < \infty$. When the first bit arrives, e.g., $b_{km,1} = 1$, we can update this interval to be $\tau_{km,1} < z_{km} < \infty$. After receiving a sequence of bits $\mathbf{b}_{km} = \begin{bmatrix} b_{km,1} & \ldots & b_{km,I} \end{bmatrix}$, we can establish successively tighter lower and upper bounds, i.e., $\tau_{km,\ell} < z_{km} \leq \tau_{km,u}$, depending on the specific quantization outcomes and thresholds. In this case, the measurement likelihood of the bit sequence $\mathbf{b}_{km}$ can be computed as $p(\mathbf{b}_{km}|\mathbf{x}_k) = \Pr\{\tau_{km,\ell} < z_{km} \leq \tau_{km,u}|\mathbf{x}_k\} = Q\left(\frac{\tau_{km,\ell} - \mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right) - Q\left(\frac{\tau_{km,u} - \mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right)$, leading to the following proposition:

Proposition 3 (Iteratively quantized MAP (IQMAP))**.** *Consider the linear system* (4.1)-(4.3)*. If multiple bits are allocated per measurement according to the quantization rule* (4.9)*, the posterior pdf is given by*

$$\begin{aligned} & p(\mathbf{x}_{0:K}|\mathbf{b}_{0:K}) \\ & = \frac{1}{p(\mathbf{b}_{0:K})} \prod_{k=0}^{K} \prod_{m=1}^{M_k} p(\mathbf{b}_{km}|\mathbf{x}_k) \prod_{k=0}^{K-1} p(\mathbf{x}_{k+1}|\mathbf{x}_k) \cdot p(\mathbf{x}_0) \qquad (4.10) \\ & \propto \prod_{k=0}^{K} \prod_{m=1}^{M_k} \left[Q\left(\frac{\tau_{km,\ell} - \mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right) - Q\left(\frac{\tau_{km,u} - \mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right) \right] \\ & \quad \cdot \prod_{k=0}^{K-1} \mathcal{N}(\mathbf{x}_{k+1}; \mathbf{F}_k\mathbf{x}_k, \mathbf{G}_k\mathbf{Q}_k\mathbf{G}_k^T) \cdot \mathcal{N}(\mathbf{x}_0; \mathbf{x}_{\text{ini}}, \mathbf{P}_0) \qquad (4.11) \end{aligned}$$

Notice that if the measurements are quantized with a single bit, the IQMAP becomes identical to the QMAP algorithm (see Proposition 2). Following similar arguments as in the QMAP case, we also have the following lemma:

Lemma 3. *The posterior pdf of the IQMAP, given by* (4.11)*, is log-concave.*

Proof. We first show that $p(\mathbf{b}_{km}|\mathbf{x}_k)$ is log-concave. We have

$$p(\mathbf{b}_{km}|\mathbf{x}_k) = \int_{\tau_{km,l}}^{\tau_{km,u}} p(z_{km}|\mathbf{x}_k)dz_{km} \quad \text{, where } p(z_{km}|\mathbf{x}_k) = \mathcal{N}(z_{km}; \mathbf{h}_{km}^T\mathbf{x}_k, \sigma_{km}^2) \tag{4.12}$$

This is an integration over a convex set of a conditional pdf that is log-concave in both arguments, z_{km} and $\mathbf{x}_k$. According to [95, p.12] it follows that $p(\mathbf{b}_{km}|\mathbf{x}_k)$ is also log-concave. To show that the entire posterior is log-concave, we note that the Gaussian prior is log-concave, and log-concavity is closed under multiplication. □

As for the QMAP, the consequence is that the IQMAP estimate is unique and can be found efficiently. Recall that these results hold for linear system and measurement models. Unfortunately, as in real-valued estimation, global log-concavity cannot be guaranteed in general if IQMAP is applied to systems with nonlinear dynamic- and measurement models.

Fig. 4.3 shows an illustrative example of the effects of iterative quantization of a scalar measurement with a scalar prior, using one to four bits. The example demonstrates that both measurement likelihood and posterior pdf using quantized measurements approximate their counterparts using real-valued measurements increasingly well with growing number of bits. With as little as four bits, the curves in the two cases are visually almost indistinguishable, resulting in nearly identical estimates.

We now address the question of how to successively select new quantization thresholds. We propose to choose $\tau_{km,i}$ so that, given all previous quantized observations, the measurement z_{km} could lie on each side of the new threshold with *equal probability*, and hence both possible realizations of the next bit, $b_{km,i}$, are equiprobable.

Proposition 4 (IQMAP Threshold Selection)**.** *Assume measurement z_{km} has been quantized using $i-1$ bits, yielding lower and upper thresholds $\tau_{km,\ell}$ and $\tau_{km,u}$. Further, approximate the conditional pdf $p(z_{km}|\mathbf{b}_{0:k})$ by a Gaussian with mean $\mathbf{h}_{km}^T\hat{\mathbf{x}}_{k|0:k}$ and covariance $\mathbf{h}_{km}^T\mathbf{P}_{k|0:k}\mathbf{h}_{km} + \sigma_{km}^2$, where $\hat{\mathbf{x}}_{k|0:k}$ denotes the current MAP estimate of the state at time step k conditioned on all quantized observations $\mathbf{b}_{0:k}$, and $\mathbf{P}_{k|0:k}$ its covariance. The new threshold $\tau_{km,i}$ that divides the interval between $\tau_{km,\ell}$ and $\tau_{km,u}$ into two equiprobable regions can then be computed as*

$$\tau_{km,i} = \mathbf{h}_{km}^T\hat{\mathbf{x}}_{k|0:k} + \sqrt{\mathbf{h}_{km}^T\mathbf{P}_{k|0:k}\mathbf{h}_{km} + \sigma_{km}^2}\,Q^{-1}\left(\frac{1}{2}\left[Q\left(\tau'_{km,\ell}\right) + Q\left(\tau'_{km,u}\right)\right]\right) \tag{4.13}$$

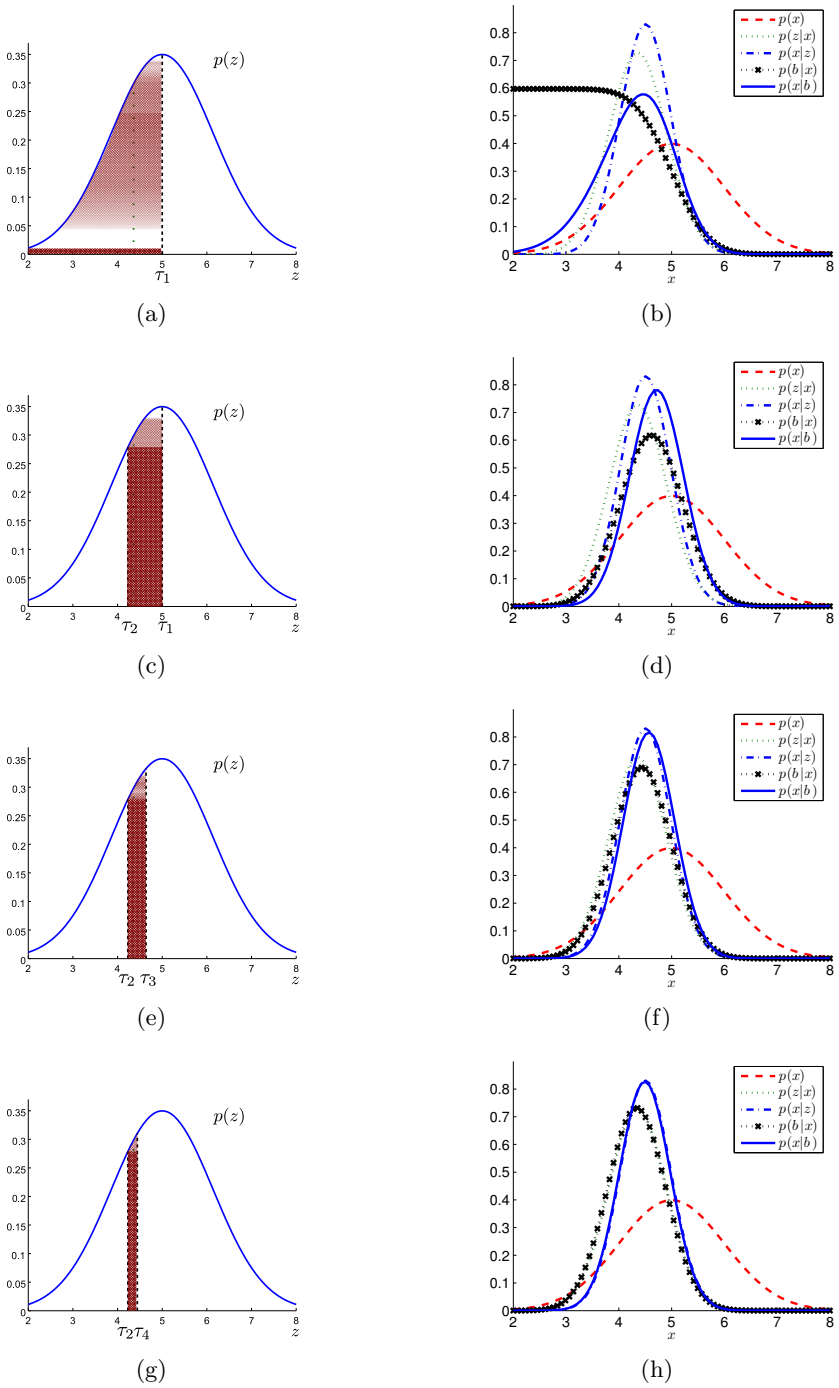

Figure 4.3: Comparison of real-valued and quantized measurement-based posterior pdfs for scalar prior and one scalar measurement $z = x + v$, $x \sim \mathcal{N}(x; 5, 1)$, $v \sim \mathcal{N}(v; 0, 0.3)$, for the realization $z = 4.36$ iteratively quantized with one [(a), (b)] up to four bits [(g), (h)]. The figures on the left show the quantization thresholds along with the measurement probability. The figures on the right show the prior, the measurement likelihoods, and the posterior pdfs for quantized and real-valued measurements. Quantization with four bits results in visually nearly indistinguishable posteriors.

where $\tau'_{km,*} := \frac{\tau_{km,*} - \mathbf{h}_{km}^T \hat{\mathbf{x}}_{k|0:k}}{\sqrt{\mathbf{h}_{km}^T \mathbf{P}_{k|0:k} \mathbf{h}_{km} + \sigma_{km}^2}}$

To see this, we utilize the Gaussian approximation of $p(z_{km}|\mathbf{b}_{0:k})$ to express the probability that z_{km} lies on either side of the new threshold in terms of the Q-function. The new threshold $\tau_{km,i}$ should then fulfill

$$\begin{aligned} &\Pr\left\{\tau_{km,\ell} < z_{km} \leq \tau_{km,i} \middle| \mathbf{b}_{0:k}\right\} = \Pr\left\{\tau_{km,i} < z_{km} \leq \tau_{km,u} \middle| \mathbf{b}_{0:k}\right\} \\ \Leftrightarrow &Q\left(\tau'_{km,\ell}\right) - Q\left(\tau'_{km,i}\right) = Q\left(\tau'_{km,i}\right) - Q\left(\tau'_{km,u}\right) \end{aligned}$$

From this, the proposition follows.

The choice of equiprobable quantization is optimal with respect to maximizing the entropy of the next bit. More formally:

Lemma 4. *Consider the linear system* (4.1)-(4.3)*. If multiple bits are allocated per measurement according to the quantization rule* (4.9)*, the choice of quantization threshold* $\tau_{km,i}$ *according to* (4.13) *maximizes the entropy of bit* $b_{km,i}$*.*

Proof. Consider the definition for the entropy of bit $b_{km,i}$

$$H_{b_{km,i}} = -\Pr\{b_{km,i} = 1\} \log(\Pr\{b_{km,i} = 1\}) - \Pr\{b_{km,i} = -1\} \log(\Pr\{b_{km,i} = -1\}) \tag{4.14}$$

$$= -\Pr\{b_{km,i} = 1\} \log(\Pr\{b_{km,i} = 1\}) - (1 - \Pr\{b_{km,i} = 1\} \log(1 - \Pr\{b_{km,i} = 1\}) \tag{4.15}$$

Setting the first derivative with respect to $\Pr\{b_{km,i} = 1\}$ to zero yields

$$\frac{dH_{b_{km,i}}}{d\Pr\{b_{km,i} = 1\}} = -\log(\Pr\{b_{km,i} = 1\}) + \log(1 - \Pr\{b_{km,i} = 1\}) = 0 \tag{4.16}$$

$$\Rightarrow \Pr\{b_{km,i} = 1\} = 1 - \Pr\{b_{km,i} = 1\} \Rightarrow \Pr\{b_{km,i} = 1\} = \frac{1}{2} \tag{4.17}$$

The second derivative evaluated at this solution is

$$\left.\frac{d^2 H_{b_{km,i}}}{d\Pr\{b_{km,i} = 1\}^2}\right|_{\Pr\left\{b_{km,i}=1\right\}=\frac{1}{2}} = -\frac{1}{\frac{1}{2}} - \frac{1}{1 - \frac{1}{2}} < 0 \tag{4.18}$$

proving that equiprobable thresholds indeed maximize the entropy. □

Entropy-maximizing thresholds are a robust choice for a quantizer, however, they do not necessarily guarantee optimality with respect to the MAP Bayes risk function as in the case of single-bit quantization. Note that neither the QMAP nor the IQMAP threshold selection schemes (nor MAP estimation in general) are designed to minimize the MSE (as is the case for the IQEKF). The MAP estimator rather minimizes the "hit-and-miss" Bayes risk (see Appendix A.2). However, a typical criterion for estimator performance evaluation

is the Root Mean Square (RMS) error, which is closely related to the MSE. Therefore, using the RMS as comparative performance metric confers an inherent advantage to the IQEKF over the IQMAP estimator. Still, the IQMAP has the capability of relinearization and smoothing, giving it an edge over the filters that can potentially compensate this shortcoming. If RMS error was the design metric for an application, it is conceivable to combine the advantages of both estimators, by using the thresholds computed by the IQEKF, and smooth the estimates over a time window using the IQMAP algorithm, possibly at a lower frequency to amortize the computations to solve the MAP optimization problem over several time steps.

4.3.3 Solving the (IQ-)MAP estimation problem

We will now describe how to actually solve the IQMAP optimization problem. To start, we review how to solve the real-valued MAP estimation problem. We first write (4.5) as a standard WLS problem

$$\hat{\mathbf{x}}_{0:K|0:K} = \arg\min \Bigg(\sum_{k=0}^{K} \sum_{m=1}^{M_k} \frac{1}{2} \| z_{km} - \mathbf{h}_{km}^T \mathbf{x}_k \|^2_{\sigma_k^{-2}} + \sum_{k=0}^{K-1} \frac{1}{2} \| \mathbf{x}_{k+1} - \mathbf{F}_k \mathbf{x}_k \|^2_{(\mathbf{G}_k \mathbf{Q}_k \mathbf{G}_k^T)^{-1}} + \frac{1}{2} \| \mathbf{x}_0 - \mathbf{x}_{\text{ini}} \|^2_{\mathbf{P}_0^{-1}} \Bigg). \tag{4.19}$$

Here, $\mathbf{x}_{\text{ini}}$ and $\mathbf{P}_0$ are the initial or prior state estimate for the robot team and its covariance. The notation $\|\mathbf{x}\|_{\mathbf{W}}$ denotes the weighted L_2-norm, which can be converted to a regular L_2-norm by pre-whitening

$$\|\mathbf{x}\|_{\mathbf{W}} = \sqrt{\mathbf{x}^T \mathbf{W} \mathbf{x}} = \|\mathbf{W}^{1/2} \mathbf{x}\|_2 \tag{4.20}$$

where $\mathbf{W}^{1/2}$ denotes the matrix square-root so that $(\mathbf{W}^{1/2})^T \mathbf{W}^{1/2} = \mathbf{W}$.

Following the appropriate pre-whitening step, we can stack the summands of (4.19) to form a large regular Least Squares system with the closed-form solution given by the normal equations

$$\hat{\mathbf{x}}_{0:K|0:K} = \arg\min \frac{1}{2} \| \mathbf{A} \mathbf{x}_{0:K} - \mathbf{b} \|_2^2 \tag{4.21}$$

$$= (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b} \tag{4.22}$$

where $\mathbf{A}$, $\mathbf{b}$ are functions of the quantities $\mathbf{F}_k, \mathbf{G}_k, \mathbf{Q}_k, z_{km}, \mathbf{h}_{km}, \sigma_k^2, \mathbf{x}_{\text{ini}}$, and $\mathbf{P}_0$. The covariance of this estimate is given by $\mathbf{P}_{0:K|0:K} = (\mathbf{A}^T \mathbf{A})^{-1}$.

We will solve the WLS problem using QR-decomposition, which according to Golub and van Loan [35] is numerically more accurate than the method of normal equations for problems with small residual $\mathbf{A}\hat{\mathbf{x}}_{0:K|0:K} - \mathbf{b}$. Let $\mathbf{A} = \mathbf{QR}$ denote the thin QR decompo-

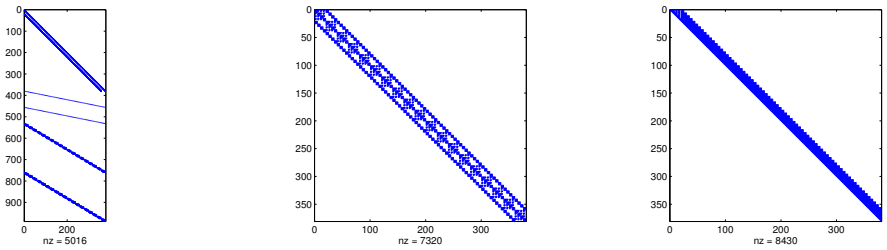

(a) Measurement matrix $\mathbf{A}$ (b) Information Matrix $\mathbf{A}^T\mathbf{A}$ (c) $\mathbf{R}$-Matrix from $\mathbf{A} = \mathbf{QR}$

Figure 4.4: Measurement matrix, information matrix, and Cholesky factor for 2-D Cooperative Localization with 4 robots and 20 time steps. The five banded structures in $\mathbf{A}$ correspond to contributions from initial conditions and stochastic motion model, linear and rotational velocity measurements, and relative distance and bearing measurements, respectively. All matrices are extremely sparse (1.3%, 5.1%, 5.8% non-zero elements respectively)

sition [35] of $\mathbf{A}$, such that $\mathbf{Q}$ is rectangular and has the same dimension as $\mathbf{A}$ and $\mathbf{R}$ is a square, upper triangular matrix. Then, $\|\mathbf{A}\mathbf{x}_{0:K} - \mathbf{b}\|_2^2 = \|\mathbf{R}\mathbf{x}_{0:K} - \mathbf{Q}^T\mathbf{b}\|_2^2 + \|\mathbf{d}\|_2^2$, where $\mathbf{d}$ is a constant independent of $\mathbf{x}_{0:K}$ ($\mathbf{d}$ is the projection of $\mathbf{b}$ onto the left Nullspace of $\mathbf{A}$), and we can efficiently compute the WLS estimate by solving $\mathbf{R}\hat{\mathbf{x}}_{0:K|0:K} = \mathbf{Q}^T\mathbf{b}$ via back-substitution. The matrix $\mathbf{R}$ is the Cholesky factor of the information matrix (the inverse covariance), hence $\mathbf{P} = (\mathbf{R}^T\mathbf{R})^{-1}$. Figure 4.4 shows the structure and sparsity patterns of the measurement matrix $\mathbf{A}$, the information matrix $\mathbf{A}^T\mathbf{A}$, and the upper triangular matrix $\mathbf{R}$ for a 2-D localization example.

For nonlinear system and/or measurement models, we obtain a nonlinear version of 4.19

$$\hat{\mathbf{x}}_{0:K|0:K} = \arg\min\Bigg(\sum_{k=0}^{K}\sum_{m=1}^{M_k}\frac{1}{2}\|z_{km} - h_{km}(\mathbf{x}_k)\|^2_{\sigma_k^{-2}} + \sum_{k=0}^{K-1}\frac{1}{2}\|\mathbf{x}_{k+1} - f_k(\mathbf{x}_k)\|^2_{(\mathbf{G}_k\mathbf{Q}_k\mathbf{G}_k^T)^{-1}} + \frac{1}{2}\|\mathbf{x}_0 - \mathbf{x}_{\text{ini}}\|^2_{\mathbf{P}_0^{-1}}\Bigg). \tag{4.23}$$

which we linearize around the current estimate.

The optimization problem is then solved using i-WLS, with the p-th iteration given by the linearized subproblem

$$\boldsymbol{\delta}\mathbf{x}^* = \arg\min\frac{1}{2}\|\mathbf{A}(\hat{\mathbf{x}}^{(p)}_{0:K|0:K})\boldsymbol{\delta}\mathbf{x} - \mathbf{b}(\hat{\mathbf{x}}^{(p)}_{0:K|0:K})\|_2^2 \tag{4.24}$$

$$\hat{\mathbf{x}}^{(p+1)}_{0:K|0:K} = \hat{\mathbf{x}}^{(p)}_{0:K|0:K} + \boldsymbol{\delta}\mathbf{x}^* \tag{4.25}$$

The notation is chosen to explicitly indicate the dependence of $\mathbf{A}$ and $\mathbf{b}$ on the current linearization point $\hat{\mathbf{x}}^{(p)}_{0:K|0:K}$.

In QMAP, (4.8), and IQMAP, (4.11), the measurement likelihoods are not Gaussian, and hence the above process does not apply directly. The optimization problem to solve is

(recall (4.8))

$$\hat{\mathbf{x}}_{0:K|0:K} = \arg\min\Bigg(\sum_{k=0}^{K}\sum_{m=1}^{M_k} -\log Q\left(\frac{b_{km}(\tau_{km}-\mathbf{h}_{km}^T\mathbf{x}_k)}{\sigma_{km}}\right) + \sum_{k=0}^{K-1}\frac{1}{2}\|\mathbf{x}_{k+1}-\mathbf{F}_k\mathbf{x}_k\|^2_{(\mathbf{G}_k\mathbf{Q}_k\mathbf{G}_k^T)^{-1}} + \frac{1}{2}\|\mathbf{x}_0-\mathbf{x}_{\text{ini}}\|^2_{\mathbf{P}_0^{-1}}\Bigg). \tag{4.26}$$

Following a sequential quadratic programming strategy, the Newton optimization step can be cast into a form analogous to i-WLS.

In particular, let $u_{km} = \frac{b_{km}(\tau_{km}-\mathbf{h}_{km}^T\hat{\mathbf{x}}_{k|0:K})}{\sigma_{km}}$. Then the second-order Taylor series expansion of $-\log p(b_{km}|\mathbf{x}_k)$ for the QMAP is given by

$$\begin{aligned} -\log Q\left(\frac{b_{km}(\tau_{km}-\mathbf{h}_{km}^T\mathbf{x}_k)}{\sigma_{km}}\right) &\approx -\log Q(u_{km}) - \frac{\mathcal{N}(u_{km};0,1)}{Q(u_{km})}\cdot\frac{b_{km}}{\sigma_{km}}\mathbf{h}_{km}^T\boldsymbol{\delta}\mathbf{x}_k \\ &+ \frac{1}{2\sigma_{km}^2}\frac{-u_{km}\mathcal{N}(u_{km};0,1)\cdot Q(u_{km}) + \big(\mathcal{N}(u_{km};0,1)\big)^2}{Q^2(u_{km})}\boldsymbol{\delta}\mathbf{x}_k^T\mathbf{h}_{km}\mathbf{h}_{km}^T\boldsymbol{\delta}\mathbf{x}_k \qquad (4.27) \\ &= \frac{1}{2}\|c_{km}\frac{\mathbf{h}_{km}^T}{\sigma_{km}}\boldsymbol{\delta}\mathbf{x}_k - d_{km}\|_2^2 + e_{km} \qquad (4.28) \end{aligned}$$

where we have used the fact that $b_{km}^2 = 1$ and denoted

$$c_{km} = \frac{\sqrt{-u_{km}\mathcal{N}(u_{km};0,1)\cdot Q(u_{km}) + \big(\mathcal{N}(u_{km};0,1)\big)^2}}{Q(u_{km})} \tag{4.29}$$

$$d_{km} = \frac{b_{km}\cdot\mathcal{N}(u_{km};0,1)}{\sqrt{-u_{km}\mathcal{N}(u_{km};0,1)\cdot Q(u_{km}) + \big(\mathcal{N}(u_{km};0,1)\big)^2}} \tag{4.30}$$

and e_{km} is an additive constant irrelevant for optimization.

Analogously, for the IQMAP the optimization problem becomes (see (4.11))

$$\hat{\mathbf{x}}_{0:K|0:K} = \arg\min\Bigg(\sum_{k=0}^{K}\sum_{m=1}^{M_k} -\log\left[Q\left(\frac{\tau_{km,\ell}-\mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right) - Q\left(\frac{\tau_{km,u}-\mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right)\right] + \sum_{k=0}^{K-1}\frac{1}{2}\|\mathbf{x}_{k+1}-\mathbf{F}_k\mathbf{x}_k\|^2_{(\mathbf{G}_k\mathbf{Q}_k\mathbf{G}_k^T)^{-1}} + \frac{1}{2}\|\mathbf{x}_0-\mathbf{x}_{\text{ini}}\|^2_{\mathbf{P}_0^{-1}}\Bigg). \tag{4.31}$$

We define $u_{km} = \frac{\tau_{km,u}-\mathbf{h}_{km}^T\hat{\mathbf{x}}_{k:0:K}}{\sigma_{km}}$ and $l_{km} = \frac{\tau_{km,l}-\mathbf{h}_{km}^T\hat{\mathbf{x}}_{k:0:K}}{\sigma_{km}}$, and obtain the following

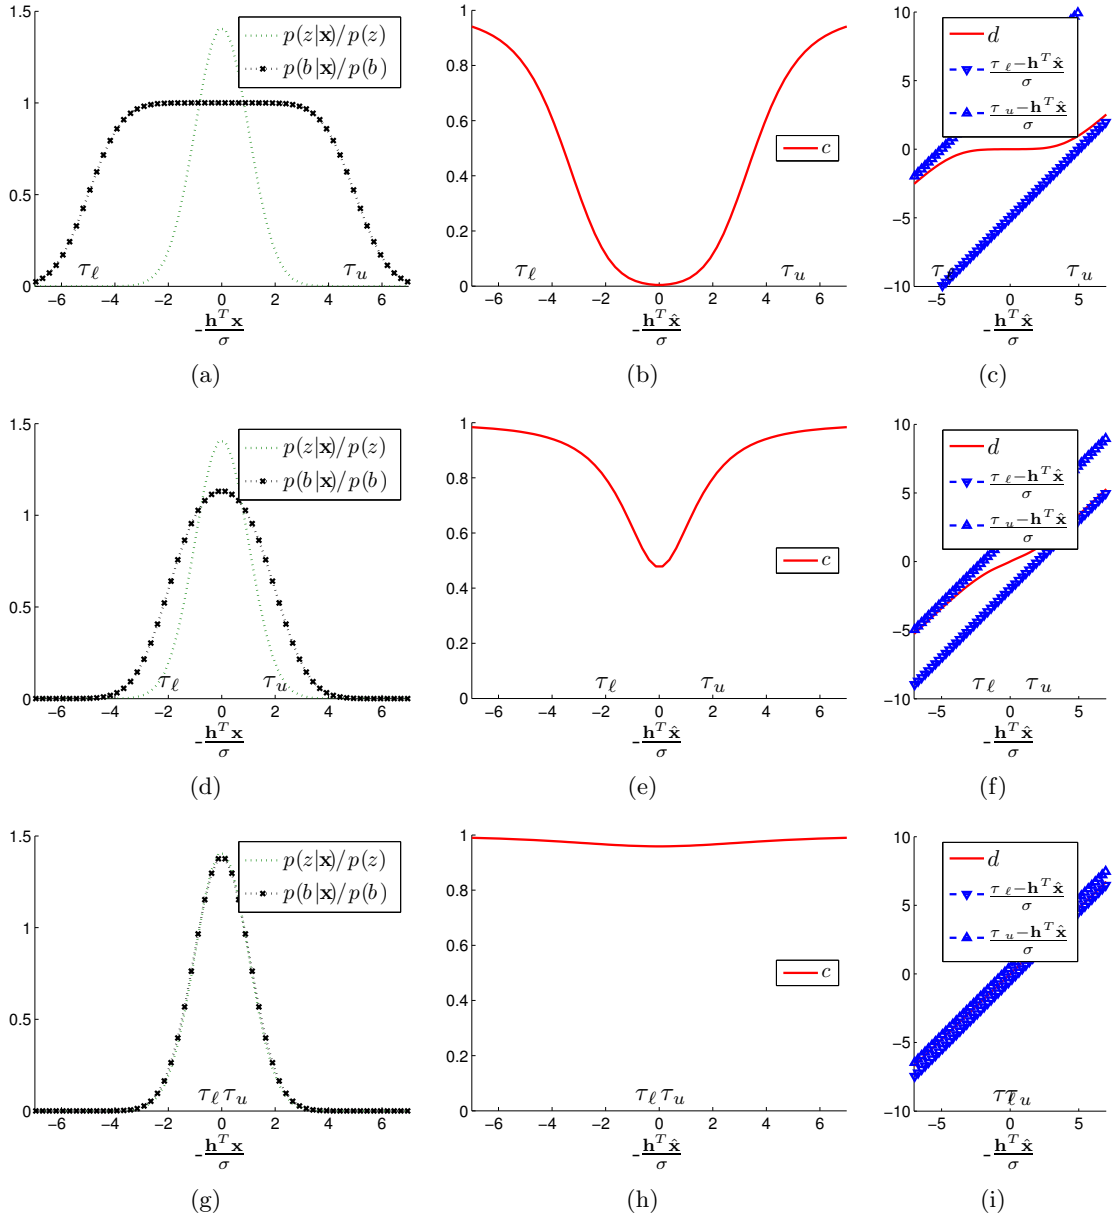

Figure 4.5: Comparison of quantized and real-valued measurement likelihood [(a),(d),(g)], measurement vector scaling factor c [(b),(e),(h)], and residual d [(c),(f),(i)], for $\tau_u - \tau_\ell = 10$ [(a)-(c)], $\tau_u - \tau_\ell = 4$ [(d)-(f)], and $\tau_u - \tau_\ell = 1$ [(g)-(i)], as a function of the normalized measurement residual $\frac{z-\mathbf{h}^T\mathbf{x}}{\sigma}$. The graphs x-axis is the normalized residual $\frac{z-\mathbf{h}^T\mathbf{x}}{\sigma}$, with $z = 0$ assumed for centering, and $\sigma = 1$. Quantized likelihood, measurement vector, and residual approximate their real-valued counterparts increasingly well as $\tau_u - \tau_\ell \to 0$.

second-order Taylor series expansion of $-\log p(\mathbf{b}_{km}|\mathbf{x}_k)$ for the IQMAP

$$
\begin{aligned}
&-\log\left[Q\left(\frac{\tau_{km,\ell}-\mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right)-Q\left(\frac{\tau_{km,u}-\mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right)\right]\approx-\log\Big[Q(l_{km})-Q(u_{km})\Big]\\
&+\frac{1}{\sigma_{km}}\frac{\mathcal{N}(u_{km};0,1)-\mathcal{N}(l_{km};0,1)}{Q(l_{km})-Q(u_{km})}\mathbf{h}^T\boldsymbol{\delta}\mathbf{x}_k\\
&+\frac{1}{2\sigma_{km}^2}\Bigg(\Big(u_{km}\mathcal{N}(u_{km};0,1)-l_{km}\mathcal{N}\left(l_{km};0,1\right)\Big)\cdot\Big(Q\left(l_{km}\right)-Q(u_{km})\Big)\\
&+\Big(\mathcal{N}\left(u_{km};0,1\right)-\mathcal{N}(l_{km};0,1)\Big)^2\Bigg)\frac{1}{\Big(Q\left(l_{km}\right)-Q(u_{km})\Big)^2}\boldsymbol{\delta}\mathbf{x}_k^T\mathbf{h}_{km}\mathbf{h}_{km}^T\boldsymbol{\delta}\mathbf{x}_k \qquad (4.32)\\
&=\frac{1}{2}\|c_{km}\frac{\mathbf{h}_{km}^T}{\sigma_{km}}\boldsymbol{\delta}\mathbf{x}_k-d_{km}\|_2^2+e_{km} \qquad (4.33)
\end{aligned}
$$

where

$$
\begin{aligned}
c_{km}=&\Bigg(\Big(u_{km}\mathcal{N}(u_{km};0,1)-l_{km}\mathcal{N}(l_{km};0,1)\Big)\Big(Q(l_{km})-Q(u_{km})\Big)\\
&+\Big(\mathcal{N}(u_{km};0,1)-\mathcal{N}(l_{km};0,1)\Big)^2\Bigg)^{\frac{1}{2}}\cdot\frac{1}{Q(l_{km})-Q(u_{km})} \qquad (4.34)\\
d_{km}=&-(\mathcal{N}(u_{km};0,1)-\mathcal{N}(l_{km};0,1))\\
&\cdot\Bigg(\Big(u_{km}\mathcal{N}(u_{km};0,1)-l_{km}\mathcal{N}(l_{km};0,1)\Big)\Big(Q(l_{km})-Q(u_{km})\Big)\\
&+\Big(\mathcal{N}(u_{km};0,1)-\mathcal{N}(l_{km};0,1)\Big)^2\Bigg)^{-\frac{1}{2}} \qquad (4.35)
\end{aligned}
$$

and again e_{km} is an additive constant irrelevant for optimization.

In comparison, in the real-valued case the negative log likelihood is given by

$$\frac{1}{2}\|z_{km}-\mathbf{h}_{km}^T\mathbf{x}_k\|_{\sigma_{km}^2}^2 \qquad (4.36)$$

which, after pre-whitening and second-order Taylor expansion around $\hat{\mathbf{x}}_{k:0:K}$ can be written as

$$\frac{1}{2}\left\|\frac{\mathbf{h}_{km}^T}{\sigma_{km}}\boldsymbol{\delta}\mathbf{x}_k-\frac{z_{km}-\mathbf{h}_{km}^T\hat{\mathbf{x}}_{k|0:K}}{\sigma_{km}}\right\|_2^2 \qquad (4.37)$$

Note that since this is a linear measurement model, the Taylor series expansion is unnecessary, but will help to illustrate the relationship to the quantized measurement case. From these expressions, we see that the quantized equals the real-valued measurement vector up to a scale factor c_{km} (cf. (4.28) and (4.33) vs. (4.37)).

To better understand the meaning of the expressions for c_{km} and d_{km}, consider the graphical depiction for different values of $u_{km} - l_{km}$ shown in Fig. 4.5, for $z = 0$ and $\sigma = 1$. Observe that the measurement likelihood exhibits a flat center region for large values of $u_{km} - l_{km}$. Intuitively, as long as a residual falls within the lower and upper bound, its likelihood is approximately constant, since we do not have any more precise information about its exact location. Mathematically, this flat region is explained by the difference of the two shifted Q-functions appearing in the likelihood, which themselves are relatively flat outside the sigmoidal region around zero.

The flat region translates into a near-zero value of c_{km}, and hence near-zero scaling of the measurement vector around the origin. Note that loosely speaking, due to the connection between the measurement matrix $\mathbf{A}$ and the information matrix $\mathbf{A}^T\mathbf{A}$, this also means less information coming from this measurement within the region around zero. Consider now the value for d_{km}, corresponding to the normalized residual in the real-valued case. We see that it is asymptotically bounded by u_{km} and l_{km}, which are shifted versions of the 45 degree bisecting line through the origin. Notice that in the real-valued case, this line and the graphic depiction of d_{km} should coincide. For large values of $u_{km} - l_{km}$, the value for d_{km} is reduced to almost zero, as long as the predicted measurement is well within the thresholds. As the difference between upper and lower thresholds decreases, we see that the flat region slowly vanishes, the dip in the curve of c_{km} diminishes, and the curve of b_{km} approaches a straight line.

We further observe from the graphs, as expected, that in the limit as the upper and lower thresholds converge to the actual measurement and the difference between u_{km} and l_{km} tends to zero, the expressions for c_{km} and d_{km} tend to 1 and $\frac{z_{km} - \mathbf{h}^T\hat{\mathbf{x}}_{k|0:K}}{\sigma_{km}} \left(= u_{km} = l_{km} \right)$ respectively, which are the values known from real-valued iterative MAP estimation.

Implementation issues

As a note for implementation in software, the numeric expressions for c_{km} and d_{km} are not very stable in regions far away from the origin. However, Fig. 4.5 provides the graphical asymptotes, which can be used in place of the numerical expressions in these regions. Further, we note that

$$Q(l_{km}) - Q(u_{km}) = 1 - Q(-l_{km}) - (1 - Q(-u_{km})) = Q(-u_{km}) - Q(-l_{km}). \tag{4.38}$$

In a computing system with finite precision, this can be used to convert expressions numerically close to $1 - 1$ (negative arguments of the Q-function) into expressions numerically close to $0 - 0$ (positive arguments of the Q-function), which can be computed with greater accuracy.

Further, in order to compute the quantization thresholds for the IQMAP, we further require knowledge of the covariance matrix $\mathbf{P}_{k|0:k}$. If the state k corresponds to the final m

elements in the state vector, this covariance can be obtained efficiently by inverting only a small, $m \times m$ portion of the upper triangular Cholesky factor $\mathbf{R}$, which can be done using back-substitution. In particular, note that $\mathbf{P} = (\mathbf{A}^T\mathbf{A})^{-1} = (\mathbf{R}^T\mathbf{R})^{-1} = \mathbf{R}^{-1}\mathbf{R}^{-T}$, and recall the fact that the inverse of an upper triangular matrix is upper triangular. Assume the matrix $\mathbf{P}$ is of dimension $n \times n$. We can compute the lower right submatrix $\mathbf{P}_{mm} := \mathbf{P}(n-m+1:n, n-m+1:n)$ (in Matlab notation) by computing the upper triangular matrix $\mathbf{R}_{mm} = \mathbf{R}(n-m+1:n, n-m+1:n)$, followed by the product $\mathbf{P}_{mm} = \mathbf{R}_{mm}^{-1}\mathbf{R}_{mm}^{-T}$. Especially for constant-rate quantization, having to compute the covariance only for the final states in the state vector in order to quantize the current measurements is a very typical operation. If other parts of the covariance matrix are needed, it might be necessary to compute the full covariance matrix, and then extract the required submatrix.

Finally, in order to prevent unbounded growth of the state vector (and hence of matrix $\mathbf{A}$) and the associated computational cost, it is advisable to marginalize out past states (see, e.g., [83] for a discussion of marginalization in MAP-based CL). Marginalization allows to maintain a constant-size state vector commensurate with the available computing resources. Depending on the choice of states to marginalize, this corresponds to an estimate of a sliding time window (as opposed to the entire state history) with uniform or non-uniform spacing in time.

4.3.4 Optimal bit allocation

MAP estimation, contrary to filtering, allows to include information from asynchronous, i.e., time-delayed measurements. In the context of quantized distributed estimation, it also allows to utilize unexpectedly available bandwidth resources to transmit additional information about earlier measurements. This section will address the question of how to allocate those additional bits. The performance criterion will be the expected reduction in uncertainty, as measured by the trace of the posterior covariance matrix.

A filter is forced to spend additionally available bits to refine only the measurements at the current time step. However, we know from the performance analysis in [81], that the gain in performance after spending more and more bits to quantize a particular measurement follows a law of diminishing returns, which will be confirmed by the simulation results later. A batch MAP estimator can spend the bits to refine previously quantized measurements, and there exist a variety of methods to determine which particular measurements to refine.

The optimal allocation would be the one that yields the global minimum of the performance criterion for the entire batch of bits to be allocated. This is a resource allocation problem and known to be hard, since it in principle requires to evaluate all possible combinations of measurements to refine. In order to avoid the associated combinatorial complexity, we will resort to heuristic allocation approaches.

The first (naïve) heuristic would be to uniformly allocate the bits over all available measurements, i.e., to evenly spread the available bits over the time horizon, if necessary

including tie-breaking rules, such as starting with oldest or most crudely quantized bits first.[3] This approach is clearly suboptimal, as not all measurements carry the same amount of information. Intuitively, spending bits on a measurement with large measurement noise variance, these bits will mostly quantize noise and hence be essentially wasted.

Instead, we propose a (greedy) per-bit optimal approach that allocates each bit individually by determining the one measurement to refine (using the iterative threshold selection scheme described in Sec. 4.3.2) which yields the maximum expected reduction in the trace of the posterior covariance matrix. This process is iterated until all bits are allocated. To this end, we compute for each measurement the new threshold $\tau_{km,\oplus}$ according to the threshold selection rule described in Proposition 4. Depending on the outcome of the quantization, there are two possible new likelihoods for this measurement, either $\left[Q\left(\frac{\tau_{km,l}-\mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_k m}\right)-Q\left(\frac{\tau_{km,\oplus}-\mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right)\right]$ or $\left[Q\left(\frac{\tau_{km,\oplus}-\mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right)-Q\left(\frac{\tau_{km,u}-\mathbf{h}_{km}^T\mathbf{x}_k}{\sigma_{km}}\right)\right]$.

For both cases, we can compute the new vector $c_{km,\oplus}\frac{\mathbf{h}_{km}^T}{\sigma_{km}}$ from Eq. (4.34). Notice that the particular measurement in question corresponds to exactly one line containing the row vector $c_{km}\frac{\check{\mathbf{h}}_{km}^T}{\sigma_{km}}$ in the measurement matrix, where $\check{\mathbf{h}}_{km}^T=\begin{bmatrix}0 & \dots & 0 & \mathbf{h}_{km}^T & 0 & \dots & 0\end{bmatrix}$ is a vector of zeros with the measurement jacobian inserted at the column position corresponding to the state at time k. After allocating the bit to refine this measurement, this row vector would now be replaced by the new vector $\mathbf{c}_{km,\oplus}\frac{\check{\mathbf{h}}_{km}^T}{\sigma_{km}}$.

Let us assume that at the current estimate, the least squares problem can be written as

$$\min\|\mathbf{A}\cdot\boldsymbol{\delta}\mathbf{x}-\mathbf{b}\|_2^2 \tag{4.39}$$

and from solving this system we also have the current QR-factorization

$$\mathbf{A}=\begin{bmatrix}\mathbf{A}_1\\ c_{km}\frac{\check{\mathbf{h}}_{km}^T}{\sigma_{km}}\\ \mathbf{A}_2\end{bmatrix}=\mathbf{Q}\cdot\mathbf{R} \tag{4.40}$$

Further, the covariance of the current estimation error is given by

$$\mathbf{P}=(\mathbf{A}^T\mathbf{A})^{-1}=(\mathbf{R}^T\mathbf{R})^{-1}=\mathbf{R}^{-1}\mathbf{R}^{-T} \tag{4.41}$$

To help decide whether to choose measurement z_{km} as the candidate measurement for which to allocate the bit, we have to quantify the impact that the corresponding change in the measurement vector has on the trace of the posterior covariance matrix. If we choose

[3] Notice that in the (I)QMAP framework, only whole bits can be allocated per measurement.

to refine the quantization of z_{km}, the new measurement matrix will be given by

$$\mathbf{A}_{\oplus} = \begin{bmatrix} \mathbf{A}_1 \\ c_{km,\oplus} \frac{\breve{\mathbf{h}}_{km}^T}{\sigma_{km}} \\ \mathbf{A}_2 \end{bmatrix} \tag{4.42}$$

Our goal is to compute the new posterior covariance matrix, whose trace we seek to minimize

$$\mathbf{P}_{\oplus} = (\mathbf{A}_{\oplus}^T \mathbf{A}_{\oplus})^{-1} \tag{4.43}$$

$$= \left(\mathbf{A}_1^T \mathbf{A}_1 + \mathbf{A}_2^T \mathbf{A}_2 + c_{km,\oplus}^2 \frac{\breve{\mathbf{h}}_{km} \breve{\mathbf{h}}_{km}^T}{\sigma_{km}^2} \right)^{-1} \tag{4.44}$$

$$= \left(\mathbf{A}_1^T \mathbf{A}_1 + \mathbf{A}_2^T \mathbf{A}_2 + c_{km}^2 \frac{\breve{\mathbf{h}}_{km} \breve{\mathbf{h}}_{km}^T}{\sigma_{km}^2} - c_{km}^2 \frac{\breve{\mathbf{h}}_{km} \breve{\mathbf{h}}_{km}^T}{\sigma_{km}^2} + c_{km,\oplus}^2 \frac{\breve{\mathbf{h}}_{km} \breve{\mathbf{h}}_{km}^T}{\sigma_{km}^2} \right)^{-1} \tag{4.45}$$

$$= \left(\mathbf{A}^T \mathbf{A} + (c_{km,\oplus}^2 - c_{km}^2) \frac{\breve{\mathbf{h}}_{km} \breve{\mathbf{h}}_{km}^T}{\sigma_{km}^2} \right)^{-1} \tag{4.46}$$

Notice that this is a rank one update on $\mathbf{P}$, so that we can compute $\mathbf{P}_{\oplus}$ using the Sherman-Morrison formula as a special case of the matrix inversion lemma [43].

$$\mathbf{P}_{\oplus} = \left(\mathbf{A}^T \mathbf{A} + \breve{\mathbf{h}}_{km} \frac{c_{km,\oplus}^2 - c_{km}^2}{\sigma_{km}^2} \breve{\mathbf{h}}_{km}^T \right)^{-1} \tag{4.47}$$

$$= \mathbf{P} - \mathbf{P} \breve{\mathbf{h}}_{km} \left(\breve{\mathbf{h}}_{km}^T \mathbf{P} \breve{\mathbf{h}}_{km} + \frac{\sigma_{km}^2}{c_{km,\oplus}^2 - c_{km}^2} \right)^{-1} \breve{\mathbf{h}}_{km}^T \mathbf{P} \tag{4.48}$$

We are interested in minimizing the trace of the updated covariance matrix

$$\operatorname{tr} \mathbf{P}_{\oplus} = \operatorname{tr} \left(\mathbf{P} - \mathbf{P} \breve{\mathbf{h}}_{km} \left(\breve{\mathbf{h}}_{km}^T \mathbf{P} \breve{\mathbf{h}}_{km} + \frac{\sigma_{km}^2}{c_{km,\oplus}^2 - c_{km}^2} \right)^{-1} \breve{\mathbf{h}}_{km}^T \mathbf{P} \right) \tag{4.49}$$

$$= \operatorname{tr} \mathbf{P} - \operatorname{tr} \left(\mathbf{P} \breve{\mathbf{h}}_{km} \left(\breve{\mathbf{h}}_{km}^T \mathbf{P} \breve{\mathbf{h}}_{km} + \frac{\sigma_{km}^2}{c_{km,\oplus}^2 - c_{km}^2} \right)^{-1} \breve{\mathbf{h}}_{km}^T \mathbf{P} \right) \tag{4.50}$$

$$= \operatorname{tr} \mathbf{P} - \operatorname{tr} \left(\breve{\mathbf{h}}_{km}^T \mathbf{P} \mathbf{P} \breve{\mathbf{h}}_{km} \left(\breve{\mathbf{h}}_{km}^T \mathbf{P} \breve{\mathbf{h}}_{km} + \frac{\sigma_{km}^2}{c_{km,\oplus}^2 - c_{km}^2} \right)^{-1} \right) \tag{4.51}$$

$$= \operatorname{tr} \mathbf{P} - \operatorname{tr} \left(\breve{\mathbf{h}}_{km}^T \mathbf{R}^{-1} \mathbf{R}^{-T} \mathbf{R}^{-1} \mathbf{R}^{-T} \breve{\mathbf{h}}_{km} \left(\breve{\mathbf{h}}_{km}^T \mathbf{R}^{-1} \mathbf{R}^{-T} \breve{\mathbf{h}}_{km} + \frac{\sigma_{km}^2}{c_{km,\oplus}^2 - c_{km}^2} \right)^{-1} \right) \tag{4.52}$$

Minimizing the trace of the updated covariance corresponds to

$$\max \operatorname{tr}\left(\breve{\mathbf{h}}_{km}^T\mathbf{R}^{-1}\mathbf{R}^{-T}\mathbf{R}^{-1}\mathbf{R}^{-T}\breve{\mathbf{h}}_{km}\left(\breve{\mathbf{h}}_{km}^T\mathbf{R}^{-1}\mathbf{R}^{-T}\breve{\mathbf{h}}_{km}+\frac{\sigma_{km}^2}{c_{km,\oplus}^2-c_{km}^2}\right)^{-1}\right) \tag{4.53}$$

which is a scalar. This quantity can be computed efficiently using Alg. 4.

Algorithm 4 Computing the reduction in $\operatorname{tr}\mathbf{P}$ due to a new bit

Require: $\breve{\mathbf{h}}_{km}$, σ_{km}, $c_{km,\oplus}$, c_{km}, and $\mathbf{R}^{-1}$, the latter computed using backsubstitution

Ensure: $\Delta_{tr}=\left(\breve{\mathbf{h}}_{km}^T\mathbf{R}^{-1}\mathbf{R}^{-T}\breve{\mathbf{h}}_{km}+\frac{\sigma_{km}^2}{c_{km,\oplus}^2-c_{km}^2}\right)^{-1}\breve{\mathbf{h}}_{km}^T\mathbf{R}^{-1}\mathbf{R}^{-T}\mathbf{R}^{-1}\mathbf{R}^{-T}\breve{\mathbf{h}}_{km}$

$\mathbf{v}_1 \leftarrow \mathbf{R}^{-T}\breve{\mathbf{h}}_{km}$ ($n\times 1$ vector)
$v_2 \leftarrow \breve{\mathbf{h}}_{km}^T\mathbf{R}^{-1}\mathbf{R}^{-T}\breve{\mathbf{h}}_{km}=\mathbf{v}_1^T\mathbf{v}_1$ (scalar)
$\mathbf{v}_3 \leftarrow \mathbf{R}^{-1}\mathbf{R}^{-T}\breve{\mathbf{h}}_{km}=\mathbf{R}^{-1}\mathbf{v}_1$ ($n\times 1$ vector)
$v_4 \leftarrow \breve{\mathbf{h}}_{km}^T\mathbf{R}^{-1}\mathbf{R}^{-T}\mathbf{R}^{-1}\mathbf{R}^{-T}\breve{\mathbf{h}}_{km}=\mathbf{v}_3^T\mathbf{v}_3$ (scalar)
$\Delta_{tr} \leftarrow v_4\left(v_2+\frac{\sigma_{km}^2}{c_{km,\oplus}^2-c_{km}^2}\right)^{-1}$

Notice that the most costly operation, the inversion of $\mathbf{R}$, needs to be done only once. The cost for inverting a triangular matrix is generally quadratic in the size of the state vector, but here, since $\mathbf{R}$ is generally banded, the cost is linear. The remaining operations from Alg. 4 have to be carried out once for each measurement, with the exception of the final step, which has to be performed twice for both of the candidate $c_{km,\oplus}$ from the two possible outcomes of the quantization. Based on the threshold selection rule described in Proposition 4, we assume that the outcome of the quantization with the additional bit will occur with equal probability. Therefore, we can compute the expected reduction of the trace of the covariance by simply averaging the results obtained from the two different values of $c_{km,\oplus}$. We repeat this process for all possible measurements, and finally choose for refinement the measurement that results in the maximum expected covariance reduction.

Notice that this approach is approximate, since computation of the final covariance matrix would require iteration to convergence which we disregard due to excessive computational complexity. The simulation results confirm that the approach outlined above yields sufficient results for practical purposes.

4.4 Simulation results

We tested the IQEKF and the IQMAP for CL with two robots moving in 2D, using the real-valued EKF and MAP as benchmark. The system model for each robot was given by a constant-velocity motion model [8]. The continuous-time dynamics for each robot were

given by

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{G}_c \begin{bmatrix} w_{\dot{V}} \\ w_{\dot{\omega}} \end{bmatrix} \tag{4.54}$$

with the state defined as $\mathbf{x} = \begin{bmatrix} x & y & \phi & V & \omega \end{bmatrix}^T$, $\mathbf{f}(\mathbf{x}) = \begin{bmatrix} V\cos\phi & V\sin\phi & \omega & 0 & 0 \end{bmatrix}^T$, and $\mathbf{G}_c = \begin{bmatrix} \mathbf{0}_{3\times 2}^T & \mathbf{I}_{2\times 2} \end{bmatrix}^T$. Further, we chose $\sigma_{\dot{V}} = 0.6325$ m/s$^2/\sqrt{\text{Hz}}$, and $\sigma_{\dot{\omega}} = 0.4967$ rad/s^2 $/\sqrt{\text{Hz}}$ as the continuous-time noise standard deviations for the motion model. The simulated trajectories followed these characteristics.

After first-order discretization with time step δt we obtain

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \delta t \mathbf{f}(\mathbf{x}_k) + \mathbf{w}_d \tag{4.55}$$

with system Jacobian

$$\mathbf{\Phi}_k = \begin{bmatrix} 1 & 0 & -V_k \delta t \sin\phi_k & \delta t \cos\phi_k & 0 \\ 0 & 1 & V_k \delta t \cos\phi_k & \delta t \sin\phi_k & 0 \\ 0 & 0 & 1 & 0 & \delta t \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

and discrete noise covariance $E[\mathbf{w}_d \mathbf{w}_d^T] = \mathbf{Q}_d$ where

$$\mathbf{Q}_d = \int_{t_k}^{t_{k+1}} \mathbf{\Phi}(t_{k+1}, \tau) \mathbf{G}_c \mathbf{Q}_c \mathbf{G}_c^T \mathbf{\Phi}(t_{k+1}, \tau)^T d\tau \tag{4.56}$$

The odometry measurements of linear and rotational velocity, corrupted by zero-mean Gaussian noise with $\sigma_{Vm} = 0.07$ m/s and $\sigma_{\omega m} = 0.28$ rad/s, were recorded at 10 Hz and treated as regular measurement updates.

We assumed robot-to-robot distance and bearing measurements at a frequency of 1 Hz, corrupted by additive Gaussian measurement noise with $\sigma_d = 0.05$ m and $\sigma_\theta = 0.09$ rad. Notice that none of the robots received absolute position measurements, rendering the system unobservable, with pose estimation errors increasing over time.

We compared the performance of IQEKF and IQMAP with all measurements quantized with 1-4 bits (constant per run) against that of real-valued EKF and MAP in Monte Carlo simulation. The RMS errors are shown in Fig. 4.6. We can see that the error decreases with increasing number of quantization bits, and that the performance comes close to that of the real-valued estimators with as little as 4 bits per measurements. Moreover, the performance gain per additional bit becomes increasingly small. Both quantized and real-valued MAP estimators perform consistently and considerably better than their filtering counterparts (up to 30% improvement in orientation error), as was expected due to their ability to mitigate

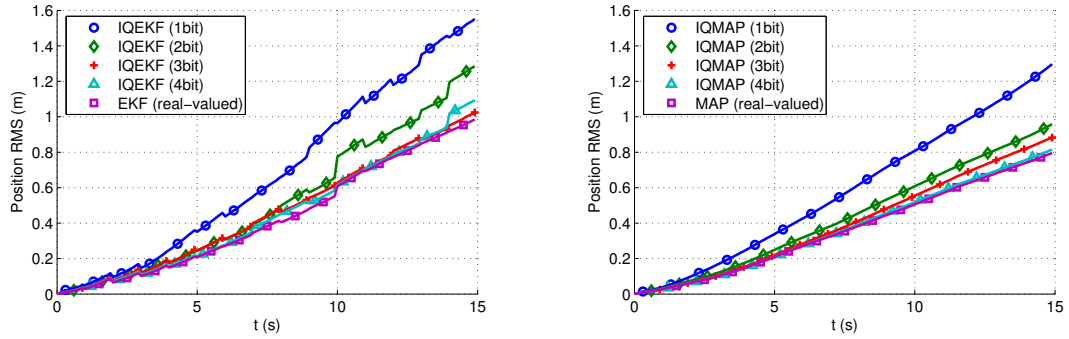

(a) Position (Filters)

(b) Position (MAP)

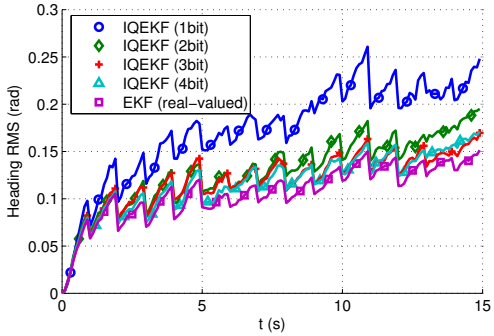

(c) Heading (Filters)

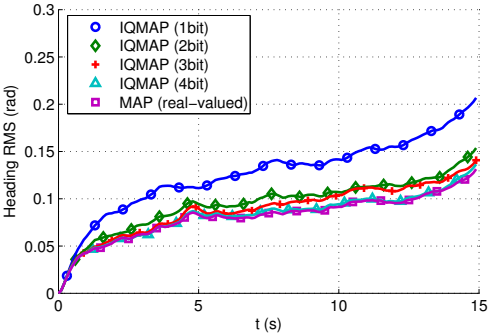

(d) Heading (MAP)

Figure 4.6: RMS error for $N_R = 2$ robots and $N_{MC} = 60$ Monte Carlo trials, for EKF, MAP, IQMAP (1-4 bits) and IQEKF (1-4 bits), computed as $\mathrm{RMS(t)} = \sqrt{\frac{1}{\mathrm{N_R}}\frac{1}{\mathrm{N_{MC}}}\sum_{\mathrm{i=1}}^{\mathrm{N_R}}\sum_{\mathrm{j=1}}^{\mathrm{N_{MC}}}(\hat{\mathrm{x}}_{\mathrm{i,j}}(\mathrm{t}) - \mathrm{x}_{\mathrm{i,j}}(\mathrm{t}))^2}$.

linearization errors in nonlinear problems such as CL. Also, the noticeable sawtooth pattern in the filters, resulting from the 1 Hz robot-to-robot measurements, is effectively smoothed out by the MAP estimators.

A typical result for the estimation error and the corresponding 3σ-bounds for the case of single-bit quantization is shown in Fig. 4.7. As expected, real-valued MAP outperforms QMAP, which in turn is more accurate than SOI-EKF. However, note the behavior of the quantized filter for the velocity estimate: Since both system and measurement function for this quantity are considered linear by the SOI-EKF, it (wrongly) converges and reaches steady state. The QMAP estimator, however, correctly accounts for the nonlinearity introduced by quantization, and more accurately depicts the evolution of uncertainty, as indicated by the tighter but variable 3σ bounds.

In a next simulation, we tested time-varying bit allocation. Using the same simulation setup as before, we quantized all measurements using two bits, and halted the simulation after 50 time steps. We then simulated the availability of additional 50 bits, roughly 10% of the 424 bits allocated up to that point. We computed the trace of the covariance matrix

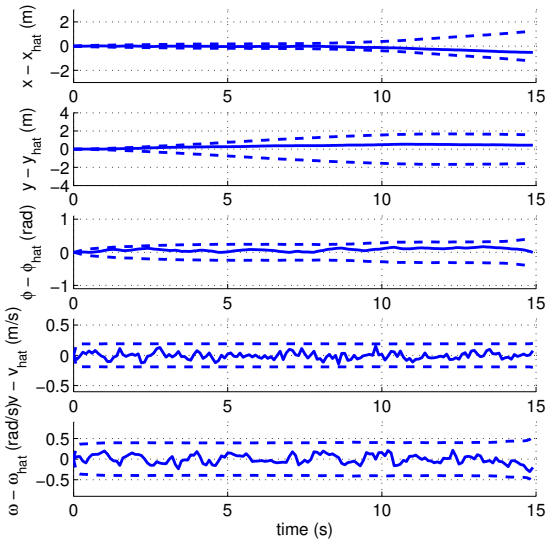

(a) MAP

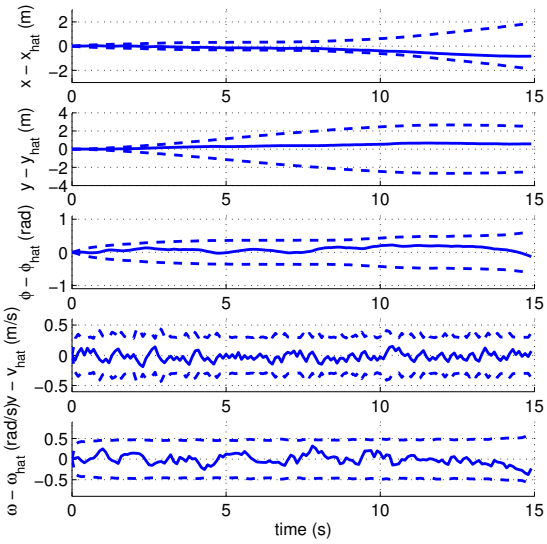

(b) QMAP

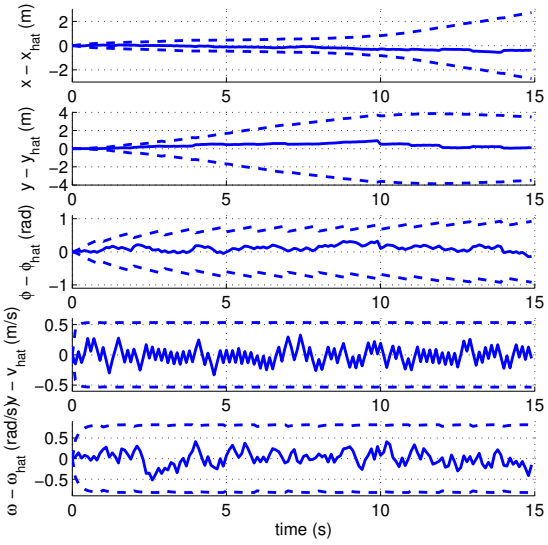

(c) SOI-EKF

Figure 4.7: Typical errors and 3σ-bounds (dashed) for the state estimates of one of the robots, for real-valued MAP, QMAP, and SOI-EKF (one bit). The quantization incurs performance loss (as well as considerable savings in communication). QMAP has better performance than SOI-EKF, showing its ability to better cope with linearization errors. The velocity covariance estimates of the SOI-EKF (but not of the QMAP) wrongly converge to steady state.

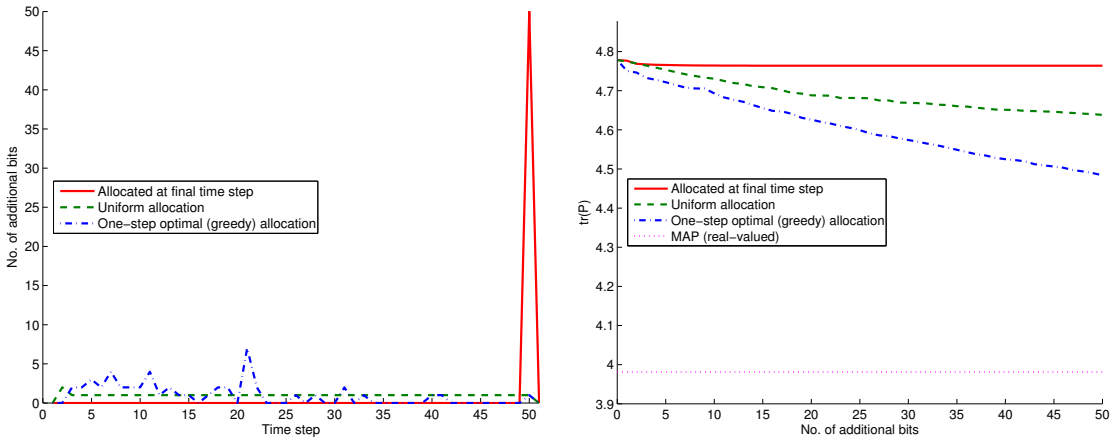

(a) (b)

Figure 4.8: Time-varying bit allocation. After two robots perform CL for 50 time steps and quantize each measurement using two bits, we simulate availability of additional 50 bits (roughly 10% of the 424 bits spent so far). Comparison of three strategies: Allocation only at final time step (solid line), uniform allocation (dashed line), and the proposed optimal strategy (dash-dotted line). The trace of the covariance achieved by real-valued MAP is shown for comparison (dotted line). (a) Final allocation of bits per time step. (b) Evolution of the trace of the covariance matrix.

of the entire state history after each of the bits was allocated.

We compared three strategies, with results shown in Fig. 4.8. The first strategy spent all bits evenly on measurements acquired only at the last time step. This mimics the behavior of a filter, which only maintains an estimate of the current state and is therefore incapable of refining past measurements. As shown in [81] for the MMSE estimator, the gain in precision drops rapidly when increasing the number of bits spent per measurement (law of diminishing returns). In accordance with this finding, we see the reduction in uncertainty leveling off after about 15 additional bits, and the overall reduction in uncertainty being rather minimal. The second strategy allocated bits sequentially and uniformly one per time step. This appears to be a reasonable heuristic, and simulation results show that it performs better than the previous strategy. Finally, we evaluated the one-step optimal (greedy) strategy described in Sec. 4.3.4 designed to maximize the expected reduction in uncertainty per bit, as measured by the reduction in the trace of the covariance. As evident from the plots, this strategy is the most efficient of all three methods at reducing uncertainty. For this particular simulation, after spending only an additional 10% of the previously allocated bits, it reduces the difference in $\operatorname{tr} \mathbf{P}$ between IQMAP (2bit) and real-valued MAP by about 32%, compared to 16% for the uniform bit allocation, and 2% when refining only the observations at the last time step.

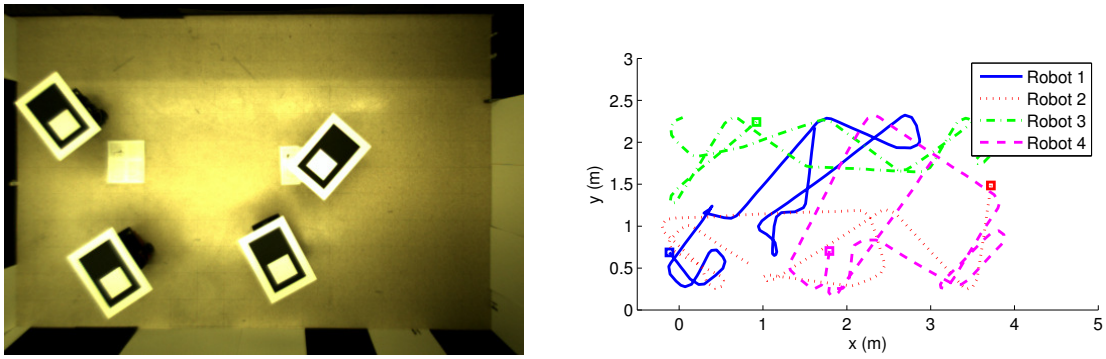

(a) (b)

Figure 4.9: Experimental setup: (a) Four Pioneer I robots with mounted calibration targets as viewed by the overhead camera, moving in a 2.5×4.5 m arena. (b) True robot trajectories. Starting points are marked with a square.

4.5 Experimental results

The quantized MAP estimation algorithms were also tested in a real world experiment. For this purpose, we employed four robots moving randomly in a 2.5×4.5 m arena, having odometry and relative position estimates available at 1 Hz (see Fig. 4.9). The robots moved at 0.1 m/s while avoiding collisions and obstacles. A calibrated overhead camera was used to establish ground truth as well as to provide synthetic measurements of relative distance and bearing.

Despite being of the same model, the robots exhibited heterogeneous noise characteristics in their odometry, with the noise standard deviations varying between $\sigma_{Vm} = 0.3$ cm/s to $\sigma_{Vm} = 0.6$ cm/s for linear velocity, and $\sigma_{\omega m} = 0.0078$ rad/s to $\sigma_{\omega m} = 0.02$ rad/s for rotational velocity. These different values can be attributed to different tire pressures, or manufacturing imperfections. The standard deviation for the constant velocity statistical motion model were set to $\sigma_V = 0.12$ m/s and $\sigma_\omega = 0.45$ rad/s, which were tuned after the experiment based on the maximum observed change per time step in measured velocity and turn rate for all robots. The observations of relative distance and bearing were extracted from the overhead camera images and corrupted by Gaussian noise with standard deviation of $\sigma_d = 0.05$ m for distance and $\sigma_\theta = 2$ deg for relative bearing.

The error plots (see Figs. 4.10-4.13) show that the IQEKF and the IQMAP perform very well in a real-world CL application. Moreover, with as little as four bits, the quantized estimators closely approach the performance of their real-valued counterparts. Again, we see that the IQEKF estimators incorrectly converge to steady state, whereas the IQMAP algorithms bound the error more dynamically, which helps to avoid overconfidence. Spikes in errors and corresponding 3σ-bounds, e.g., robot 4 at $t = 120$ s, coincide with an abrupt change in robot velocity or turn rate. This behavior is expected, since these points represent outliers with respect to the constant-velocity statistical motion model. The IQMAP

Table 4.1: Experimental results for quantized MAP CL with four robots.

RMS Position error (m)				
	Robot 1	Robot 2	Robot 3	Robot 4
IQMAP (1bit):	0.1637	0.2680	0.2748	0.2346
IQMAP (2bit):	0.1285	0.1608	0.1312	0.1596
IQMAP (3bit):	0.1593	0.1805	0.1522	0.1761
IQMAP (4bit):	0.1423	0.1546	0.1454	0.1570
MAP (real-valued):	0.1247	0.1616	0.1185	0.1453
RMS Orientation error (rad)				
	Robot 1	Robot 2	Robot 3	Robot 4
IQMAP (1bit):	0.2046	0.2330	0.2133	0.1971
IQMAP (2bit):	0.0564	0.0702	0.0623	0.0835
IQMAP (3bit):	0.0435	0.0470	0.0452	0.0509
IQMAP (4bit):	0.0384	0.0388	0.0424	0.0435
MAP (real-valued):	0.0465	0.0481	0.0474	0.0477
NEES				
	Robot 1	Robot 2	Robot 3	Robot 4
IQMAP (1bit):	3.7119	3.8164	3.7352	4.6111
IQMAP (2bit):	5.5733	6.3689	4.7954	6.8632
IQMAP (3bit):	7.7805	9.8302	7.5528	13.9581
IQMAP (4bit):	8.0233	8.2656	7.2866	12.6239
MAP (real-valued):	9.1525	9.9432	6.6416	8.1227

estimators almost always achieve lower RMS over the MMSE estimators, particularly for orientation (see Tables 4.1 and 4.2). We observe that in this particular experiment, the RMS error of the IQMAP using 4 bits per measurement is even slightly lower than the analog MAP. This is not a general result, an on average we would expect the analog estimators to consistently outperform their quantized counterparts (as is the case for the filter). Another interesting observation is the fact that the NEES for the IQMAP is very low for very low quantization rates, and increases with increasing number of bits. We attribute that to the fact that the covariance estimates tighten up faster than the estimation error decreases, or, in other words, that the 1-bit QMAP is very conservative in estimating uncertainty, whereas the higher-rate IQMAP estimators rather tend toward overconfidence, reflecting in a higher NEES value. Again, these results are for a single experiment, and may not hold generally.

4.6 Extension to hybrid estimation architectures

One of the shortcomings of the proposed multi-centralized estimation architecture is that it forces each robot to use only quantized measurements, even though it has access to its own, unquantized measurements. This way, each robot voluntarily relinquishes informa-

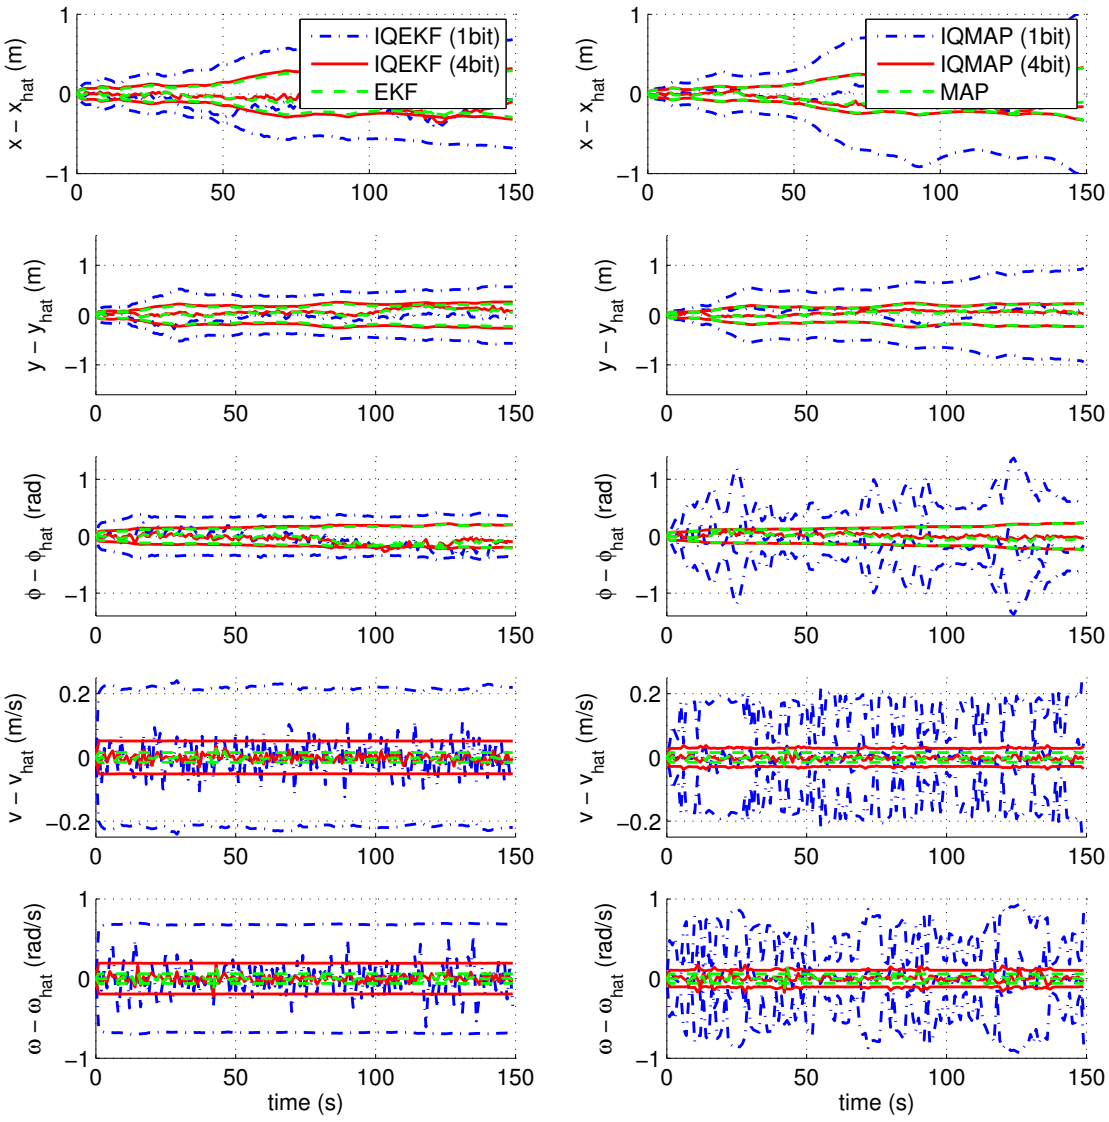

Figure 4.10: Experimental results. Estimation errors and corresponding 3σ-bounds for robot 1. (a) Filters. (b) MAP estimators.

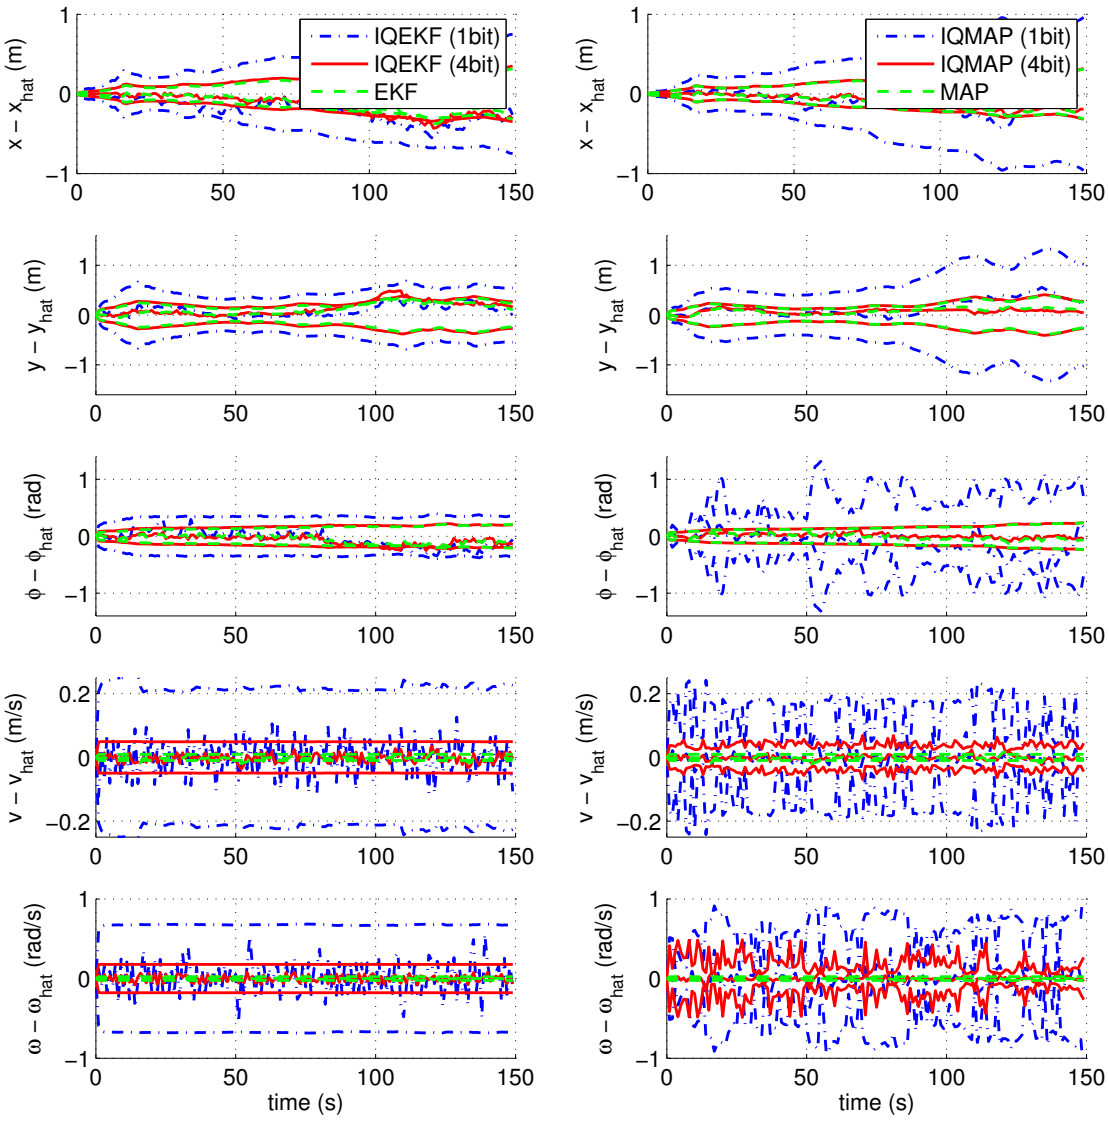

Figure 4.11: Experimental results. Estimation errors and corresponding 3σ-bounds for robot 2. (a) Filters. (b) MAP estimators.

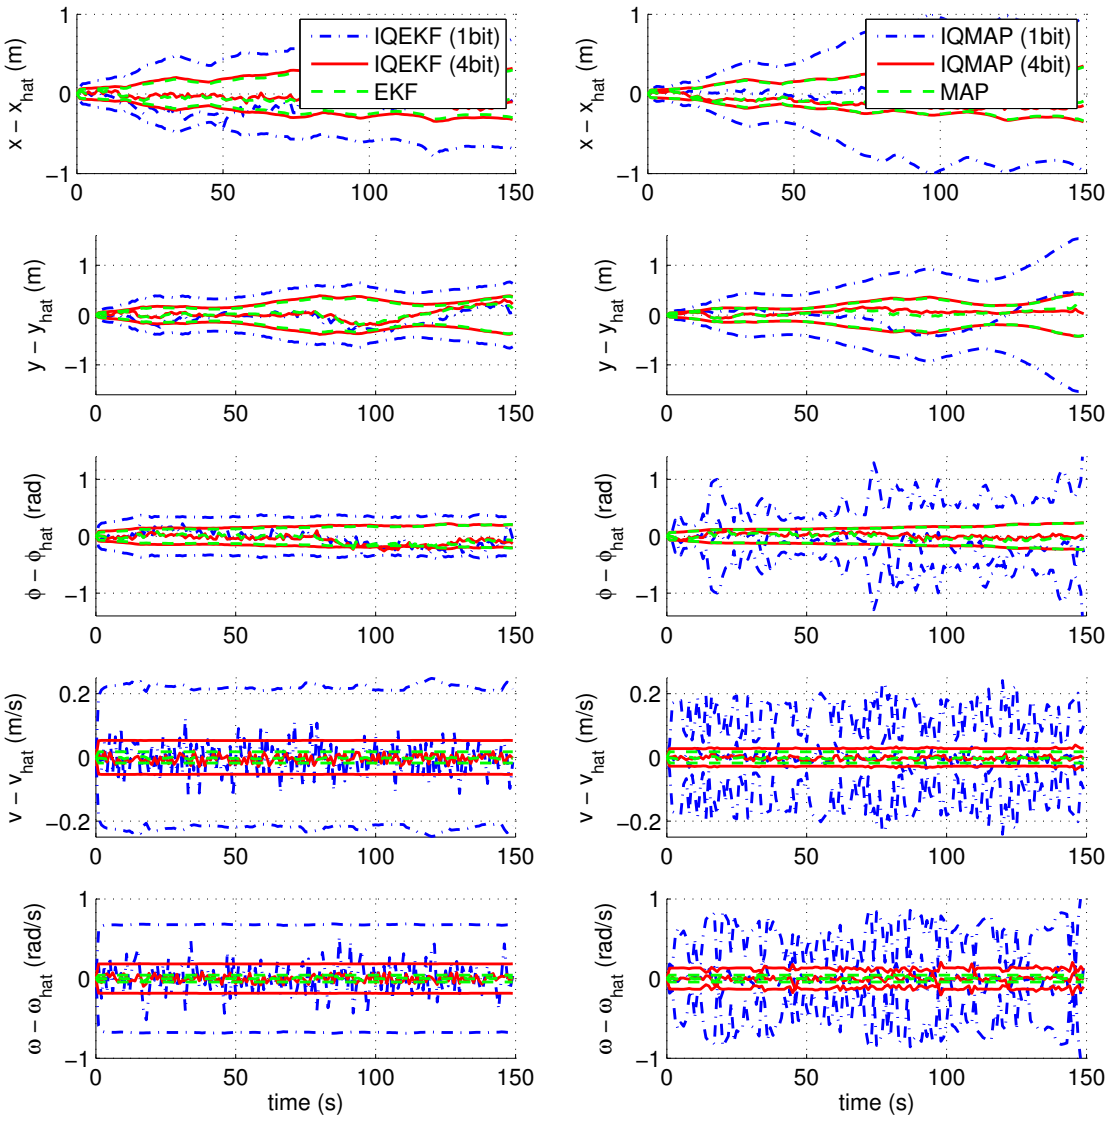

Figure 4.12: Experimental results. Estimation errors and corresponding 3σ-bounds for robot 3. (a) Filters. (b) MAP estimators.

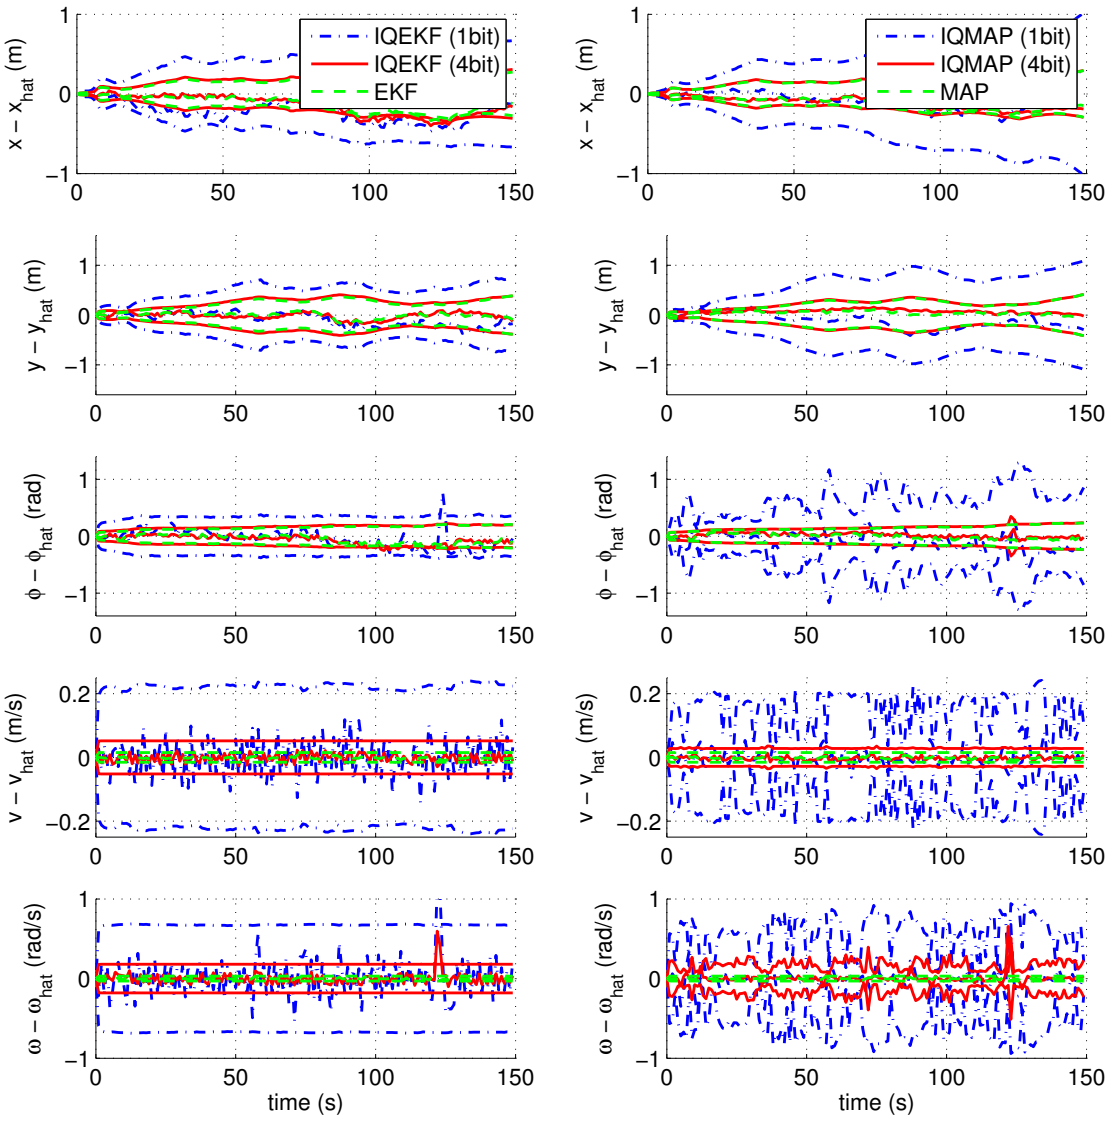

Figure 4.13: Experimental results. Estimation errors and corresponding 3σ-bounds for robot 4. (a) Filters. (b) MAP estimators.

Table 4.2: Experimental results for quantized MMSE CL with four robots.

RMS Position error (m)				
	Robot 1	Robot 2	Robot 3	Robot 4
IQEKF (1bit):	0.2035	0.2841	0.2211	0.3050
IQEKF (2bit):	0.1426	0.2434	0.1512	0.2200
IQEKF (3bit):	0.1407	0.2775	0.1730	0.2186
IQEKF (4bit):	0.1594	0.2702	0.1354	0.2030
EKF (real-valued):	0.1302	0.2218	0.1165	0.1718
RMS Orientation error (rad)				
	Robot 1	Robot 2	Robot 3	Robot 4
IQEKF (1bit):	0.1500	0.1573	0.1400	0.1601
IQEKF (2bit):	0.1110	0.1147	0.1089	0.1183
IQEKF (3bit):	0.1428	0.1448	0.1413	0.1398
IQEKF (4bit):	0.1018	0.1053	0.1045	0.1043
EKF (real-valued):	0.0874	0.0900	0.0881	0.0880
NEES				
	Robot 1	Robot 2	Robot 3	Robot 4
IQEKF (1bit):	5.1338	5.7128	4.8443	6.3003
IQEKF (2bit):	4.5735	4.5177	4.0932	5.2185
IQEKF (3bit):	8.4986	8.4313	7.3373	9.0999
IQEKF (4bit):	9.0996	9.0798	7.9321	10.4721
EKF (real-valued):	10.4850	11.3549	7.3406	9.0899

tion it could use to obtain more accurate estimates[4]. The reason this is necessary is that multi-centralized quantized estimation only works if all robots agree on the quantization thresholds, which are computed directly from the current estimate. For this to be equal among all robots, the estimators must only use information available to all team members, a symmetry which would be violated if one robot were to use its own unquantized measurements.

To overcome this limitation, we propose an extended, hybrid estimation architecture, which comprises two estimators per robot (see Fig. 4.14). For robot i, the quantized estimation module works as before, by processing the quantized measurements from the entire robot team, and using the resulting estimate $\hat{\mathbf{x}}_{0:K|\mathbf{b}^i_{0:K},\mathbf{b}^{j\neq i}_{0:K}}$ to compute the quantization thresholds. Note that the notation $\mathbf{b}^{j\neq i}_{0:K}$ refers to quantized measurements from all robots except for robot i, hence it can be thought of as a set of vectors. An additional, hybrid estimation module processes its own, unquantized measurements, together with the quantized measurements from the other team members $j \neq i$, to obtain a hybrid estimate $\hat{\mathbf{x}}_{0:K|\mathbf{z}^i_{0:K},\mathbf{b}^{j\neq i}_{0:K}}$. These hybrid estimates will differ between robots, but will be more accurate than the quantized estimates since they process more information. The performance of this estimator will be between that of the multi-centralized IQMAP, and that of the real-valued

[4]This holds under the assumption that the quantized measurements communicated from the other robots provide more information compared to the case where the robot localizes on its own

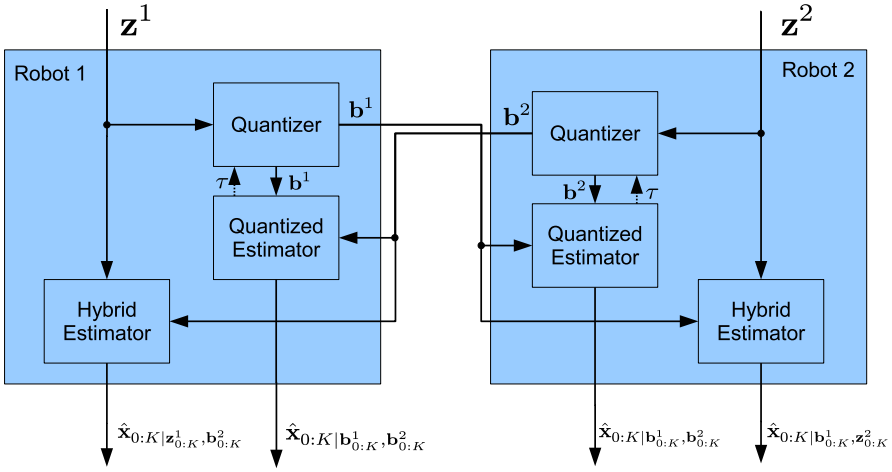

Figure 4.14: Hybrid estimator structure. Each robot i has an additional estimation module that uses its own, unquantized measurements $\mathbf{z}^i$, together with the bit sequences $\mathbf{b}^j$ from all other team members $j \neq i$, to produce an estimate $\hat{\mathbf{x}}_{0:K|\mathbf{z}^i_{0:K},\mathbf{b}^{j\neq i}_{0:K}}$. Note that unlike the quantized estimates $\hat{\mathbf{x}}_{0:K|\mathbf{b}^i_{0:K},\mathbf{b}^{j\neq i}_{0:K}}$, these hybrid estimates will differ from robot to robot.

MAP estimator. Especially in cases of relatively large bit rates, where the performance of IQMAP closely approaches that of the real-valued MAP, the gain in accuracy achieved by a hybrid estimator might only be small. However, we expect that this hybrid estimator will outperform the IQMAP under extreme bandwidth limitations, where the own analog measurements carry more information than the crudely quantized observations from the rest of the team.

A subject of ongoing research is how to use this more accurate estimate for the benefit of the remaining team members. Any such attempt of course has to comply with the strict bandwidth constraints, so it is of great importance to minimize any ensuing communication overhead. Moreover, the two estimates $\hat{\mathbf{x}}_{0:K|\mathbf{b}^i_{0:K},\mathbf{b}^{j\neq i}_{0:K}}$ and $\hat{\mathbf{x}}_{0:K|\mathbf{z}^i_{0:K},\mathbf{b}^{j\neq i}_{0:K}}$ are not independent, and so any further information sharing between the two corresponding estimators must properly account for correlations. For these reasons, exact feedback approaches might turn out computationally demanding and instead call for approximate methods. Possible ideas involve creating new, inferred measurements (e.g., a function of the difference between the two state estimates), or biasing the quantization outcome in order to minimize the difference between the two estimators.

4.7 Summary

In this chapter, we have derived a new MAP estimator (IQMAP) for estimation of random processes with quantized observations. We have shown its application, as well as that of a quantized filter (IQEKF) [81], to the problem of multi-robot Cooperative Localization under communication constraints. Compared to the regular encoding using 32 or 64 bits per real scalar measurement, both algorithms can significantly reduce communication requirements by allowing to encode scalar measurements using as little as 1 bit (or more, depending on the available resources).

We have further shown that the IQMAP estimator offers increased accuracy compared to the IQEKF, due to mitigation of linearization errors. For linear (or linearized) systems, finding the IQMAP estimate was shown to be a convex optimization problem, which guarantees a unique solution and allows the use of highly-efficient solvers. We have provided optimal threshold selection schemes for the case of single- and multi-bit quantization of scalar measurements. Additionally, the IQMAP offers the ability to efficiently utilize time-varying bandwidth, by using surplus bits to refine current and past measurements. To this end, we have proposed a one-step optimal (greedy) bit allocation scheme and shown in simulation that it outperforms two other bit allocation strategies. In particular, it was shown to use additionally available bits to greater benefit (i.e., resulting in larger reduction of uncertainty) than the IQEKF, which can only refine the current measurements and therefore suffers from rapidly diminishing gain in accuracy per additional bit spent.

Chapter 5

Concluding Remarks

5.1 Summary of contributions

The work presented in the preceding chapters has focussed on providing new techniques for challenging problems in Cooperative Localization, particularly the issues of *CL initialization* in the absence of GPS, and of *CL under severe communication constraints*. The main contributions of this work can be summarized as follows:

- **Globally optimal range-based relative pose estimation**
  Chapter 3 provided a novel method for relative pose estimation between robots in 2-D and 3-D, based on robot-to-robot distance measurements and estimates of the robots' motion. Knowledge of the relative pose between robots is necessary to initialize the CL estimation problem. It can be obtained, e.g., from external references (manual measurements, GPS, correlating overlapping sensor measurements, etc.), but these may not always be available or may not yield results of sufficient accuracy. Our proposed method exploits the mobility of the robots by determining the relative pose from distance measurements acquired at different vantage points during robot motion.

  The key contribution of this chapter, however, lies in the guaranteed globally optimal solution of the corresponding nonlinear WLS problem formulation, which by far transcends the problem of relative pose estimation. In our approach, we first transform the problem into the minimization of a polynomial cost function and then exploit novel tools from algebraic geometry and optimization theory for an efficient solution. This approach is potentially applicable to many geometric problems in robotics (kinematics) and computer vision. Additionally, it allows to go beyond the traditional approach of coarse initialization (e.g., using deterministic methods) followed by iterative optimization, which can at best guarantee local optimality. Moreover, it provides a benchmark for suboptimal, but potentially more computationally efficient approaches, and also offers tantalizing insights into nonlinear least squares problems in general, e.g., the observation that for low SNR, a local optimum might be closer to

the ground truth than the global optimum.

- **Quantized MAP estimator**
  Chapter 4 presented a novel MAP estimation scheme for dynamic random processes using severely and adaptively quantized measurements. Estimators usually assume access to real-valued measurements. Existing quantization schemes have until now either focussed on signal reconstruction or on parameter estimation. Only recently have there been (MMSE) filters for linear, dynamic Markov random processes [102, 81] with quantized observations. In this thesis, we have extended this approach to quantized MAP estimation with a single (QMAP) and multiple bits (IQMAP) per measurement. As batch estimation approaches, QMAP and IQMAP have two distinct advantages over the corresponding filter variants: (i) they are better suited to handle nonlinearities due to quantization and nonlinear measurement models prevalent in robotics, since they allow relinearization and utilize all measurements over an entire time horizon, and (ii) they can exploit time-varying bandwidth, by retroactively refining past measurements if more bits become available. These characteristics are important in order to apply quantized estimation in mobile robotics.

  Chapter 4 derived the optimization problem for the QMAP and IQMAP estimators and established their log-concavity for linear models. Moreover, we provided threshold selection rules for both cases, and established their optimality with respect to the hit-and-miss Bayes risk function used in real-valued MAP for the single bit case, and with respect to entropy for the IQMAP case. Finally, we provided a one-step optimal (greedy) method to allocate additional bits that minimized the expected reduction in uncertainty per bit.

With this work, we provide new tools to the robotics community that will enable and foster cooperative navigation in mobile sensor networks even in bandlimited and GPS-denied application domains.

5.2 Future research directions

The work on globally optimal WLS estimation for motion-induced relative pose estimation in chapter 3 opens a number of avenues for future research. The most straightforward idea is to exploit a similar technique for other sensing modalities, i.e., relative bearing or range-and bearing measurements [118], either in homogeneous or heterogeneous robot teams. Additional constraints could come in form of relative velocity measurements (e.g., from Doppler radar), or other external information such as observations of a common gravity vector. Similarly, it would be desirable to extend this approach to not only estimate the relative pose between two robots at a time, but between the entire robot team concurrently. Finally, multi-robot systems will always face the issue of precise clock synchronization,

which might be mitigated by including the unknown clock offset in the state vector.

More generally, the techniques for global optimization of polynomial cost functions can be applied in a variety of robotics and computer vision problems. One potential application we are currently investigating is inconsistency of nonlinear estimation problems. It is well-known that estimators such as the I-EKF can compute a single local minimum of a particular least squares cost function [11], which is a serious limitation for multimodal cost functions. The presented techniques allow to compute all stationary points, which in turn permits to track all local minima (or a subset thereof) in an efficiently steered multi-hypothesis estimation approach. We have already successfully applied this approach to the problem of target tracking using range-only [47] and bearing-only measurements. In computer vision, polynomial techniques have been predominantly applied to solve minimal problems (e.g., the 5-point algorithm [86]), which have important application in outlier rejection methods such as RANSAC [33]. However, given the results of chapter 3, it would be natural to apply global WLS optimization to the overdetermined variants of these problems. This is the subject of ongoing research.

Many of these extensions may increase the number of unknowns to levels infeasible to tackle in reasonable time with currently available processing and memory resources. However, the demonstrated success in smaller scale problems, together with the ongoing trend to faster and more capable computing hardware, gives reason to believe in a great future for polynomial global optimization methods.

Concerning quantized estimation, even though the results of chapter 4 have demonstrated applicability of the quantized MAP estimation methodology in CL, a number of capabilities should be added in order to find widespread use in mobile robotics applications. Most obvious among these are the extension to cooperative simultaneous localization and mapping (C-SLAM), and the extension to 3-D. The first will require techniques for landmark initialization, in turn necessitating additional prior assumptions on landmark location (e.g., uniformly within a limited region defined by maximum sensing range), and extension of the estimation framework to non-Gaussian pdfs. The latter, i.e., the extension of CL to 3-D, is a subject of intense current interest also in the classical real-valued estimation community, since it is seen as a major stepping stone to widespread practical use of mobile sensor networks. In this context we should also mention research efforts to reduce computational complexity. Currently, the quantization thresholds are computed based on the best MAP estimate, requiring very frequent solution of the entire MAP optimization problem. Other approaches might be less costly, e.g., employing a quantized filter variant (SOI-KF or IQKF) to propagate the state and determine the thresholds, and to solve the MAP problem less frequently, allowing amortization of the computational cost over several time steps. Moreover, the current multi-centralized architecture requires duplication of computations in all robots. It would be interesting to investigate distributed estimation schemes such as the one proposed by Nerurkar *et al.* [83] and modify them to work with

limited communication. Additionally, especially for robot networks that communicate and perform relative measurements very infrequently, it would be desirable to allow local integration of unquantized motion information measurements between relative observations, and to appropriately quantize this integrated information prior to transmission instead of sending each quantized proprioceptive measurement individually.

The optimality criteria used for threshold selection and bit allocation could also be further investigated. It might be of value to replace the maximum entropy criterion for multi-bit thresholds by other objective functions, such as the MAP Bayes risk as for the single bit case. Simulations for linear, 1-D CL toy problems suggest that quantization threshold selection has significant impact on estimator performance. Also the method for bit allocation is currently designed to optimally allocate a single bit. Finding the best subset of measurements to allocate several bits simultaneously is known to be a hard problem, but the theory of submodularity might offer a route for efficient suboptimal selection schemes with constant factor performance guarantees [59]. Moreover, bit allocation bears a relationship to sensor scheduling, in that the available resources are directed towards the most informative measurements. In this sense, investigating synergies between the techniques used by Mourikis and Roumeliotis [78] for sensor scheduling in robot formations and bit allocation in quantized CL as outlined in chapter 4.3.4 might be worthwhile. Another assumption we would like to relax is that of synchronous measurements. In the current framework, all robots have to have access to all measurements at the same time so as to ensure equal quantization thresholds. In practice, however, this assumption may not hold in multi-robot systems communicating over asynchronous, multi-hop networks. One avenue to pursue would be the concept of check points that has been introduced to address asynchronous communication for regular CL [68]. Finally, we have mentioned extended hybrid estimation architectures. Since it is only natural for each robot to use as much information as is available (i.e., to process its own unquantized measurements in addition to the quantized measurements received from the other team members), the question of how to reconcile the hybrid and a purely quantized estimate deserves further research.

As we finish this work, we are excited to find multi-robot systems and mobile sensor networks standing at the brink of widespread field deployment in real-world applications. We hope that our work helps to bring mobile robotics one step closer to this goal.

Bibliography

[1] M. M. Abdallah and H. C. Papadopoulos, "Sequential signal encoding and estimation for distributed sensor networks," in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Salt Lake City, UT, May 7–11, 2001, pp. 2577–2580.

[2] D. Alspach and H. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *IEEE Transactions on Automatic Control*, vol. 17, no. 4, pp. 439–448, Aug. 1972.

[3] L. A. A. Andersson and P. E. J. Nygards, "C-SAM: Multi-robot SLAM using square root information smoothing," in *Proc. of the IEEE International Conference on Robotics and Automation*, Pasadena, CA, May 19–23, 2008, pp. 2798–2805.

[4] J. Aspnes, T. Eren, D. K. Goldenberg, A. S. Morse, W. Whiteley, Y. R. Yang, B. D. O. Anderson, and P. N. Belhumeur, "A theory of network localization," *IEEE Transactions on Mobile Computing*, vol. 5, no. 12, pp. 1663–1678, Dec. 2006.

[5] W. Auzinger and H. J. Stetter, "An elimination algorithm for the computation of all zeros of a system of multivariate polynomial equations," in *Numerical Mathematics (Singapore 1988), International Series of Numerical Mathematics.* Basel, Switzerland: Birkhäuser, 1988, vol. 86, pp. 11–30.

[6] A. Bahr, M. Walter, and J. Leonard, "Consistent cooperative localization," in *Proc. of the IEEE International Conference on Robotics and Automation*, Kobe, Japan, May 12–17, 2009, pp. 3415–3422.

[7] A. Bahr, J. J. Leonard, and M. F. Fallon, "Cooperative Localization for Autonomous Underwater Vehicles," *The International Journal of Robotics Research*, vol. 28, no. 6, pp. 714–728, 2009.

[8] Y. Bar-Shalom and X.-R. Li, *Estimation and tracking: Principles, techniques, and software.* Norwood, MA: Artech House, Inc., 1993.

[9] R. W. Beard, T. W. McLain, D. B. Nelson, D. Kingston, and D. Johanson, "Decentralized cooperative aerial surveillance using fixed-wing miniature UAVs," *Proc. of the IEEE*, vol. 94, no. 7, pp. 1306–1324, Jul. 2006.

[10] A. Beck, P. Stoica, and J. Li, "Exact and approximate solutions of source localization problems," *IEEE Transactions on Signal Processing*, vol. 56, no. 5, pp. 1770–1778, May 2008.

[11] B. Bell and F. Cathey, "The iterated Kalman filter update as a Gauss-Newton method," *IEEE Transactions on Automatic Control*, vol. 38, no. 2, pp. 294–297, Feb. 1993.

[12] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Athena Scientific, 2003.

[13] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.

[14] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[15] W. G. Breckenridge, "Quaternions proposed standard conventions," Jet Propulsion Laboratory, Pasadena, CA, Interoffice Memorandum IOM 343-79-1199, 1999.

[16] M. Byröd, Z. Kukelova, K. Josephson, T. Pajdla, and K. Åström, "Fast and robust numerical solutions to minimal problems for cameras with radial distortion," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Anchorage, AK, Jun. 23–28, 2008, pp. 1–8.

[17] C. J. Cannell and D. J. Stilwell, "A comparison of two approaches for adaptive sampling of environmental processes using autonomous underwater vehicles," in *Proc. of MTS/IEEE OCEANS*, Washington, DC, Sep. 19–23, 2005, pp. 1514–1521.

[18] J. A. Costa, N. Patwari, and I. Alfred O. Hero, "Distributed weighted-multidimensional scaling for node localization in sensor networks," *ACM Transactions on Sensor Networks*, vol. 2, no. 1, pp. 39–64, Feb. 2006.

[19] D. Cox, J. Little, and D. O'Shea, *Using Algebraic Geometry*, 2nd ed. Springer, 2005.

[20] ——, *Ideals, Varieties, and Algorithms*, 3rd ed. Springer, 2008.

[21] A. J. Davison and N. Kita, "Active visual localisation for cooperating inspection robots," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Takamatsu, Japan, Oct. 30 – Nov. 5, 2000, pp. 1709–1715.

[22] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte Carlo localization for mobile robots," in *Proc. of the IEEE International Conference on Robotics and Automation*, Detroit, MI, May 10–15, 1999, pp. 1322–1328.

[23] F. Dellaert, "Square root SAM," in *Proc. of Robotics: Science and Systems*, Cambridge, MA, Jun. 8–11, 2005, pp. 177–184.

[24] M. Di Marco, A. Garulli, A. Giannitrapani, and A. Vicino, "Simultaneous localization and map building for a team of cooperating robots: a set membership approach," *IEEE Transactions on Robotics and Automation*, vol. 19, no. 2, pp. 238–249, Apr. 2003.

[25] Y. Dieudonne, O. Labbani-Igbida, and F. Petit, "Deterministic robot-network localization is hard," *IEEE Transactions on Robotics*, vol. 26, no. 2, pp. 331–339, Apr. 2010.

[26] M. D. Dikaiakos, A. Florides, T. Nadeem, and L. Iftode, "Location-aware services over vehicular ad-hoc networks using car-to-car communication," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 8, pp. 1590–1602, Oct. 2007.

[27] J. Djugash and S. Singh, "A robust method of localization and mapping using only range," in *International Symposium on Experimental Robotics*, Athens, Greece, Jul. 14–17, 2008.

[28] L. Doherty, K. S. J. Pister, and L. E. Ghaoui, "Convex position estimation in wireless sensor networks," in *Proc. of the Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE INFOCOM*, Anchorage, AK, Apr. 2001, pp. 1655–1663.

[29] A. Doucet, N. de Freitas, and N. Gordon, Eds., *Sequential Monte Carlo Methods in Practice*, ser. Statistics for Engineering and Information Science. New York: Springer-Verlag, May 2001.

[30] I. Z. Emiris and A. Rege, "Monomial bases and polynomial system solving," in *Proc. of the International Symposium on Symbolic and Algebraic Computation*. New York, NY, USA: ACM, 1994, pp. 114–122.

[31] Y. Ephraim and R. M. Gray, "A unified approach for encoding clean and noisy sources by means of waveform and autoregressive model vector quantization," *IEEE Transactions on Information Theory*, vol. 34, no. 4, pp. 826–834, Jul. 1988.

[32] Y. Feng, Z. Zhu, and J. Xiao, "Heterogeneous multi-robot localization in unknown 3D space," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China, Oct. 9–15, 2006, pp. 4533–4538.

[33] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[34] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization," *Autonomous Robots*, vol. 8, no. 3, pp. 325–344, 2000.

[35] G. H. Golub and C. F. van Loan, *Matrix Computations*, 3rd ed. Baltimore and London: Johns Hopkins University Press, 1996.

[36] R. Grabowski and P. Khosla, "Localization techniques for a team of small robots," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Maui, HI, Oct. 29 – Nov. 3, 2001, pp. 1067–1072.

[37] R. M. Gray, "Quantization in task-driven sensing and distributed processing," in *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing*, Toulouse, France, May 14–19, 2006, pp. V1049–V1052.

[38] D. R. Grayson and M. E. Stillman, "Macaulay 2, a software system for research in algebraic geometry," Available at http://www.math.uiuc.edu/Macaulay2/.

[39] J. A. Gubner, "Distributed estimation and quantization," *IEEE Transactions on Information Theory*, vol. 39, no. 4, pp. 1456–1459, Jul. 1993.

[40] D. Henrion, J. B. Lasserre, and J. Lofberg, "GloptiPoly 3: moments, optimization and semidefinite programming," *Optimization Methods and Software*, vol. 24, no. 4–5, pp. 761–779, 2009.

[41] D. Henrion and J.-B. Lasserre, "Detecting global optimality and extracting solutions in GloptiPoly," in *Positive Polynomials in Control*, ser. Lecture Notes in Control and Information Sciences. Springer, 2005, vol. 312, pp. 293–310.

[42] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *Journal of the Optical Society of America. A, Optics and Image Science*, vol. 4, no. 4, pp. 629–642, Apr. 1987.

[43] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.

[44] A. Howard, M. J. Mataric, and G. S. Sukhatme, "Localization for mobile robot teams using maximum likelihood estimation," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, Sep. 30 – Oct. 4, 2002, pp. 434–439.

[45] ——, "Putting the 'I' in 'team': an ego-centric approach to cooperative localization," in *Proc. of the IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, Sep. 14–19, 2003, pp. 868–874.

[46] G. P. Huang, N. Trawny, A. I. Mourikis, and S. I. Roumeliotis, "On the consistency of multi-robot cooperative localization," in *Proc. of Robotics: Science and Systems*, Seattle, WA, Jun.28–Jul.1, 2009, pp. 65–72. [Online]. Available: http://www.roboticsproceedings.org/rss05/p9.html

[47] G. P. Huang, K. X. Zhou, N. Trawny, and S. I. Roumeliotis, "A bank of MAP estimators for single-sensor range-only target tracking," in *Proc. of the American Control Conference*, Baltimore, MD, Jun. 30 – Jul. 2, 2010, pp. 6974–6980.

[48] M. L. Husty, "An algorithm for solving the direct kinematics of general Stewart-Gough platforms," *Mechanism and Machine Theory*, vol. 31, no. 4, pp. 365–380, 1996.

[49] K. Ito and K. Xiong, "Gaussian filters for nonlinear filtering problems," *IEEE Transactions on Automatic Control*, vol. 45, no. 5, pp. 910–927, May 2000.

[50] S. Julier, J. Uhlmann, and H. Durrant-Whyte, "A new method for the nonlinear transformation of means and covariances in filters and estimators," *IEEE Transactions on Automatic Control*, vol. 45, no. 3, pp. 477–482, 2000.

[51] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Fast incremental smoothing and mapping with efficient data association," in *Proc. of the IEEE International Conference on Robotics and Automation*, Rome, Italy, Apr. 10–14, 2007, pp. 1670–1677.

[52] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME–Journal of Basic Engineering*, vol. 82, pp. 35–45, 1960.

[53] N. Karam, F. Chausse, R. Aufrère, and R. Chapuis, "Localization of a group of communicating vehicles by state exchange," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China, Oct. 9–15, 2006, pp. 519–524.

[54] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory.* Upper Saddle River, New Jersey: Prentice Hall, 1993.

[55] S. B. Kesner, J. S. Plante, P. J. Boston, T. Fabian, and S. Dubowsky, "Mobility and power feasibility of a microbot team system for extraterrestrial cave exploration," in *Proc. of the IEEE International Conference on Robotics and Automation*, Rome, Italy, Apr. 10–14, 2007, pp. 4893–4898.

[56] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller, "Multiple relative pose graphs for robust cooperative mapping," in *Proc. of the IEEE International Conference on Robotics and Automation*, Anchorage, AK, May 3–8, 2010, pp. 3185–3192.

[57] S. Kim, M. Kojima, and H. Waki, "Generalized lagrangian duals and sums of squares relaxations of sparse polynomial optimization problems," *SIAM Journal on Optimization*, vol. 15, no. 3, pp. 697–719, 2005.

[58] J. H. Kotecha and P. M. Djuric, "Gaussian particle filtering," *IEEE Transactions on Signal Processing*, vol. 51, no. 10, pp. 2592–2601, Oct. 2003.

[59] A. Krause and C. Guestrin, "Near-optimal nonmyopic value of information in graphical models," in *Proc. of the Twenty-First Annual Conference on Uncertainty in Artificial Intelligence (UAI-05)*. Arlington, Virginia: AUAI Press, 2005, pp. 324–331.

[60] Z. Kukelova and T. Pajdla, "A minimal solution to the autocalibration of radial distortion," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Minneapolis, MN, Jun. 18–23, 2007, pp. 1–7.

[61] R. Kurazume and S. Hirose, "Study on cooperative positioning system: optimum moving strategies for CPS-III," in *Proc. of the IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May 16–20, 1998, pp. 2896–2903.

[62] ——, "An experimental study of a cooperative positioning system," *Autonomous Robots*, vol. 8, no. 1, pp. 43–52, Jan. 2000.

[63] R. Kurazume, S. Hirose, S. Nagata, and N. Sashida, "Study on cooperative positioning system (basic principle and measurement experiment)," in *Proc. of the IEEE International Conference on Robotics and Automation*, Minneapolis, MN, Apr. 22–28, 1996, pp. 1421–1426.

[64] R. Kurazume, S. Nagata, and S. Hirose, "Cooperative positioning with multiple robots," in *Proc. of the IEEE International Conference on Robotics and Automation*, San Diego, CA, May 8–13, 1994, pp. 1250–1257.

[65] W.-M. Lam and A. R. Reibman, "Design of quantizers for decentralized estimation systems," *IEEE Transactions on Communications*, vol. 41, no. 11, pp. 1602–1605, Nov. 1993.

[66] J. B. Lasserre, "Global optimization with polynomials and the problem of moments," *SIAM Journal on Optimization*, vol. 11, no. 3, pp. 796–817, 2001.

[67] T.-Y. Lee and J.-K. Shim, "Improved dialytic elimination algorithm for the forward kinematics of the general Stewart-Gough platform," *Mechanism and Machine Theory*, vol. 38, pp. 563–577, Jun. 2003.

[68] K. Leung, T. Barfoot, and H. Liu, "Decentralized localization of sparsely-communicating robot networks: A centralized-equivalent approach," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 62–77, Feb. 2010.

[69] S. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, Mar. 1982.

[70] F. Lu and E. Milios, "Robot pose estimation in unknown environments by matching 2D range scans," *Journal of Intelligent and Robotic Systems*, vol. 18, no. 3, pp. 249–275, Mar. 1997.

[71] Z.-Q. Luo, "An isotropic universal decentralized estimation scheme for a bandwidth constrained ad hoc sensor network," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 4, pp. 735–744, Apr. 2005.

[72] ——, "Universal decentralized estimation in a bandwidth constrained sensor network," *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 2210–2219, Jun. 2005.

[73] R. Madhavan, K. Fregene, and L. E. Parker, "Distributed heterogeneous outdoor multi-robot localization," in *Proc. of the IEEE International Conference on Robotics and Automation*, Washington, DC, May 11–15, 2002, pp. 374–381.

[74] A. Martinelli, "Improving the precision on multi robot localization by using a series of filters hierarchically distributed," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Diego, CA, Oct. 29 – Nov. 2, 2007, pp. 1053–1058.

[75] J. Max, "Quantizing for minimum distortion," *IEEE Transactions on Information Theory*, vol. 6, no. 1, pp. 7–12, Mar. 1960.

[76] P. S. Maybeck, *Stochastic Models, Estimation and Control*. New York: Academic Press, 1979, vol. 1.

[77] ——, *Stochastic Models, Estimation and Control*. New York: Academic Press, 1979, vol. 2.

[78] A. I. Mourikis and S. I. Roumeliotis, "Optimal sensor scheduling for resource-constrained localization of mobile robot formations," *IEEE Transactions on Robotics*, vol. 22, no. 5, pp. 917–931, Oct. 2006.

[79] ——, "Performance analysis of multirobot cooperative localization," *IEEE Transactions on Robotics*, vol. 22, no. 4, pp. 666–681, Aug. 2006.

[80] E. J. Msechu, S. I. Roumeliotis, A. Ribeiro, and G. B. Giannakis, "Distributed iteratively quantized Kalman filtering for wireless sensor networks," in *Conference Record of the 41st Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, Nov. 4–7, 2007, pp. 646–650.

[81] ——, "Decentralized quantized Kalman filtering with scalable communication cost," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3727–3741, Aug. 2008.

[82] E. J. Msechu, A. Ribeiro, S. I. Roumeliotis, and G. B. Giannakis, "Distributed Kalman filtering based on quantized innovations," in *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing*, Las Vegas, NV, Mar. 31 – Apr. 4, 2008, pp. 3293–3296.

[83] E. Nerurkar, S. I. Roumeliotis, and A. Martinelli, "Distributed maximum a posteriori estimation for multi-robot cooperative localization," in *Proc. of the IEEE International Conference on Robotics and Automation*, Kobe, Japan, May 12–17, 2009, pp. 1402 – 1409.

[84] J. Nie, "Sum of squares method for sensor network localization," *Computational Optimization and Applications*, vol. 43, no. 2, pp. 151–179, 2009.

[85] J. Nie, J. Demmel, and B. Sturmfels, "Minimizing polynomials via sum of squares over the gradient ideal," *Mathematical Programming*, vol. 106, no. 3, pp. 587–606, 2006.

[86] D. Nister, "An efficient solution to the five-point relative pose problem," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 6, pp. 756–770, Jun. 2004.

[87] A. Nüchter, K. Lingemann, J. Hertzberg, and H. Surmann, "6D SLAM-3D mapping outdoor environments," *Journal of Field Robotics*, vol. 24, no. 8-9, pp. 699–722, 2007.

[88] S. Panzieri, F. Pascucci, and R. Setola, "Multirobot localisation using interlaced extended Kalman filter," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China, Oct. 9–15, 2006, pp. 2816–2821.

[89] H. C. Papadopoulos, G. W. Wornell, and A. V. Oppenheim, "Sequential signal encoding from noisy measurements using quantizers with dynamic bias control," *IEEE Transactions on Information Theory*, vol. 47, no. 3, pp. 978–1002, Mar. 2001.

[90] C. A. C. Parker, H. Zhang, and C. R. Kube, "Blind bulldozing: multiple robot nest construction," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Las Vegas, NV, Oct. 27–31, 2003, pp. 2010–2015.

[91] P. Parrilo and B. Sturmfels, "Minimizing polynomial functions," in *Algorithmic and quantitative real algebraic geometry, DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. AMS, 2003, vol. 60, pp. 83–99.

[92] P. A. Parrilo, "Semidefinite programming relaxations for semialgebraic problems," *Mathematical Programming*, vol. 96, no. 2, pp. 293–320, May 2003.

[93] S. Patterson, B. Bamieh, and A. El Abbadi, "Environmental tomography: Ubiquitous sensing with mobile devices," in *Proc. of the IEEE 24th International Conference on Data Engineering*, Cancun, Mexico, Apr. 7–12, 2008, pp. 1560–1563.

[94] S. Prajna, A. Papachristodoulou, P. Seiler, and P. A. Parrilo, *SOSTOOLS: Sum of squares optimization toolbox for MATLAB*, Available from http://www.cds.caltech.edu/sostools and http://www.mit.edu/~parrilo/sostools, 2004.

[95] A. Prékopa, "Logarithmic concave measures and related topics," in *Stochastic Programming*, M. A. H. Dempster, Ed. Academic Press, 1980, pp. 63–82.

[96] M. Putinar, "Positive polynomials on compact semi-algebraic sets," *Indiana University Mathematics Journal*, vol. 42, pp. 969–984, 1993.

[97] M. Raghavan, "The Stewart platform of general geometry has 40 configurations," *Journal of Mechanical Design*, vol. 115, pp. 277–282, 1993.

[98] A. Rao, D. Miller, K. Rose, and A. Gersho, "Generalized vector quantization: jointly optimal quantization and estimation," in *Proc. of the 1995 IEEE International Symposium on Information Theory*, Whistler, BC, Sep. 17–22, 1995, p. 432.

[99] I. M. Rekleitis, G. Dudek, and E. E. Milios, "Multi-robot collaboration for robust exploration," in *Proc. of the IEEE International Conference on Robotics and Automation*, San Francisco, CA, Apr. 24–28, 2000, pp. 3164–3169.

[100] A. Ribeiro and G. B. Giannakis, "Bandwidth-constrained distributed estimation for wireless sensor networks – Part I: Gaussian case," *IEEE Transactions on Signal Processing*, vol. 54, no. 3, pp. 1131–1143, Mar. 2006.

[101] ——, "Bandwidth-constrained distributed estimation for wireless sensor networks – Part II: unknown probability density function," *IEEE Transactions on Signal Processing*, vol. 54, no. 7, pp. 2784–2796, Jul. 2006.

[102] A. Ribeiro, G. B. Giannakis, and S. Roumeliotis, "SOI-KF: Distributed Kalman filtering with low-cost communications using the sign of innovations," *IEEE Transactions on Signal Processing*, vol. 54, no. 12, pp. 4782–4795, Dec. 2006.

[103] S. I. Roumeliotis and G. A. Bekey, "Distributed multirobot localization," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 781–795, Oct. 2002.

[104] M. R. Sampford, "Some inequalities on Mill's ratio and related functions," *The Annals of Mathematical Statistics*, vol. 24, no. 1, pp. 130–132, 1953.

[105] A. C. Sanderson, "A distributed algorithm for cooperative navigation among multiple mobile robots," *Advanced Robotics*, vol. 12, no. 4, pp. 335–349, 1998.

[106] F. A. Shah, A. Ribeiro, and G. B. Giannakis, "Bandwidth-constrained MAP estimation for wireless sensor networks," in *Proc. of the 39th Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, Oct. 30 – Nov. 2, 2005, pp. 215–219.

[107] I. Shames, P. Bibalan, B. Fidan, and B. Anderson, "Polynomial methods in noisy network localization," in *Proc. of the 17th Mediterranean Conference on Control and Automation*, Thessaloniki, Greece, Jun. 24–26, 2009, pp. 1307–1312.

[108] Y. Shang, W. Ruml, Y. Zhang, and M. P. J. Fromherz, "Localization from mere connectivity," in *Proc. of the 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Annapolis, MD, Jun. 1–3, 2003, pp. 201–212.

[109] J. Spletzer, A. K. Das, R. Fierro, C. J. Taylor, V. Kumar, and J. P. Ostrowski, "Cooperative localization and control for multi-robot manipulation," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Maui, HI, Oct. 29 – Nov. 3, 2001, pp. 631–636.

[110] J. R. Spletzer and C. J. Taylor, "A bounded uncertainty approach to multi-robot localization," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Las Vegas, NV, Oct. 27–31, 2003, pp. 1258–1265.

[111] H. Stark and J. W. Woods, *Probability and Random Processes with Applications to Signal Processing*. Prentice Hall, 2001.

[112] D. Stewart, "A platform with six degrees of freedom," in *Proc. of the Inst. of Mechanical Engineers*, vol. 180, no. 15, 1965, pp. 371–376.

[113] J. F. Sturm, "Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11, no. 1–4, pp. 625–653, 1999.

[114] H. Sugiyama, T. Tsujioka, and M. Murata, "Collaborative movement of rescue robots for reliable and effective networking in disaster area," in *Proc. of the International Conference on Collaborative Computing: Networking, Applications and Worksharing*, San Jose, CA, Dec. 19–21, 2005.

[115] C. J. Taylor and J. R. Spletzer, "A bounded uncertainty approach to cooperative localization using relative bearing constraints," in *Proc. of the IEEE/RSJ International*

Conference on Intelligent Robots and Systems, San Diego, CA, Oct. 27 – Nov. 3, 2007, pp. 2500–2506.

[116] N. Trawny and T. Barfoot, "Optimized motion strategies for cooperative localization of mobile robots," in *Proc. of the IEEE International Conference on Robotics and Automation*, New Orleans, LA, Apr. 26 – May 1, 2004, pp. 1027–1032.

[117] N. Trawny, X. S. Zhou, and S. I. Roumeliotis, "3D relative pose estimation from six distances," in *Proc. of Robotics: Science and Systems*, Seattle, WA, Jun. 28 – Jul. 1, 2009.

[118] N. Trawny, X. S. Zhou, K. Zhou, and S. I. Roumeliotis, "Interrobot transformations in 3-D," *IEEE Transactions on Robotics*, vol. 26, no. 2, pp. 226–243, Apr. 2010.

[119] N. Trawny and S. I. Roumeliotis, "Indirect Kalman filter for 3-D pose estimation," University of Minnesota, Dept. of Comp. Sci. & Eng., MARS Lab, Tech. Rep. 2005-002, Jan. 2005.

[120] ——, "On the global optimum of planar, range-based robot-to-robot relative pose estimation," in *Proc. of the IEEE International Conference on Robotics and Automation*, Anchorage, AK, May 3–8, 2010, pp. 3200–3206.

[121] N. Trawny, S. I. Roumeliotis, and G. B. Giannakis, "Cooperative multi-robot localization under communication constraints," in *Proc. of the IEEE International Conference on Robotics and Automation*, Kobe, Japan, May 12–17, 2009, pp. 4394–4400.

[122] N. Trawny, X. S. Zhou, K. X. Zhou, and S. I. Roumeliotis, "3-D relative pose estimation from distance-only measurements," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Diego, CA, Oct. 29 – Nov. 2, 2007, pp. 1071–1078.

[123] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon, "Bundle adjustment – a modern synthesis," in *Vision Algorithms: Theory and Practice*, ser. Lecture Notes in Computer Science, B. Triggs, A. Zisserman, and R. Szeliski, Eds., vol. 1883. Springer-Verlag, 2000, pp. 298–372.

[124] J. Verschelde, "Algorithm 795: PHCpack: A general-purpose solver for polynomial systems by homotopy continuation," *ACM Transactions on Mathematical Software*, vol. 25, no. 2, pp. 251–276, 1999.

[125] H. Waki, S. Kim, M. Kojima, and M. Muramatsu, "Sums of squares and semidefinite program relaxations for polynomial optimization problems with structured sparsity," *SIAM Journal on Optimization*, vol. 17, no. 1, pp. 218–242, 2006.

[126] H. Waki, S. Kim, M. Kojima, M. Muramatsu, and H. Sugimoto, "Algorithm 883: SparsePOP—a sparse semidefinite programming relaxation of polynomial optimization problems," *ACM Transactions on Mathematical Software*, vol. 35, no. 2, pp. 1–13, 2008.

[127] C. W. Wampler, "Forward displacement analysis of general six-in-parallel sps (Stewart) platform manipulators using soma coordinates," *Mechanism and Machine Theory*, vol. 31, no. 3, pp. 331–337, 1996.

[128] R. Wannamaker, S. Lipshitz, J. Vanderkooy, and J. Wright, "A theory of nonsubtractive dither," *IEEE Transactions on Signal Processing*, vol. 48, no. 2, pp. 499–516, Feb. 2000.

[129] J. Wawerla, G. S. Sukhatme, and M. J. Mataric, "Collective construction with multiple robots," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, Sep. 30 – Oct. 4, 2002, pp. 2696–2701.

[130] J. J. Xiao and Z. Q. Luo, "Decentralized estimation in an inhomogeneous sensing environment," *IEEE Transactions on Information Theory*, vol. 51, no. 10, pp. 3564–3575, Oct. 2005.

[131] F. Zhang, G. Pereira, and V. Kumar, "Necessary and sufficient conditions for localization of robot networks," in *ASME International Design Engineering Technical Conferences. Mechanics and Robotics Conference*, Salt Lake City, UT, Sep. 28 – Oct. 2, 2004, pp. 13–22.

[132] K. Zhou and S. I. Roumeliotis, "Optimal motion strategies for range-only constrained multi-sensor target tracking," *IEEE Transactions on Robotics*, vol. 24, no. 5, pp. 1168–1185, Oct. 2008.

[133] X. S. Zhou and S. I. Roumeliotis, "Robot-to-robot relative pose estimation from range measurements," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1379–1393, Dec. 2008.

Appendix A

Appendices for Chapter 4

A.1 Properties of the Q-Function (Gaussian tail function)

Definition

$$Q(x) := \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}} du \tag{A.1}$$

$$= \Pr\{X > x\} \quad \text{if } X \sim \mathcal{N}(X; 0, 1) \tag{A.2}$$

The Q function is monotonically decreasing and log-concave (cf. Fig. A.1) [14, Ex. 3.39].

Example values

$$Q(-\infty) = 1 \qquad Q(0) = \frac{1}{2} \qquad Q(\infty) = 0 \tag{A.3}$$

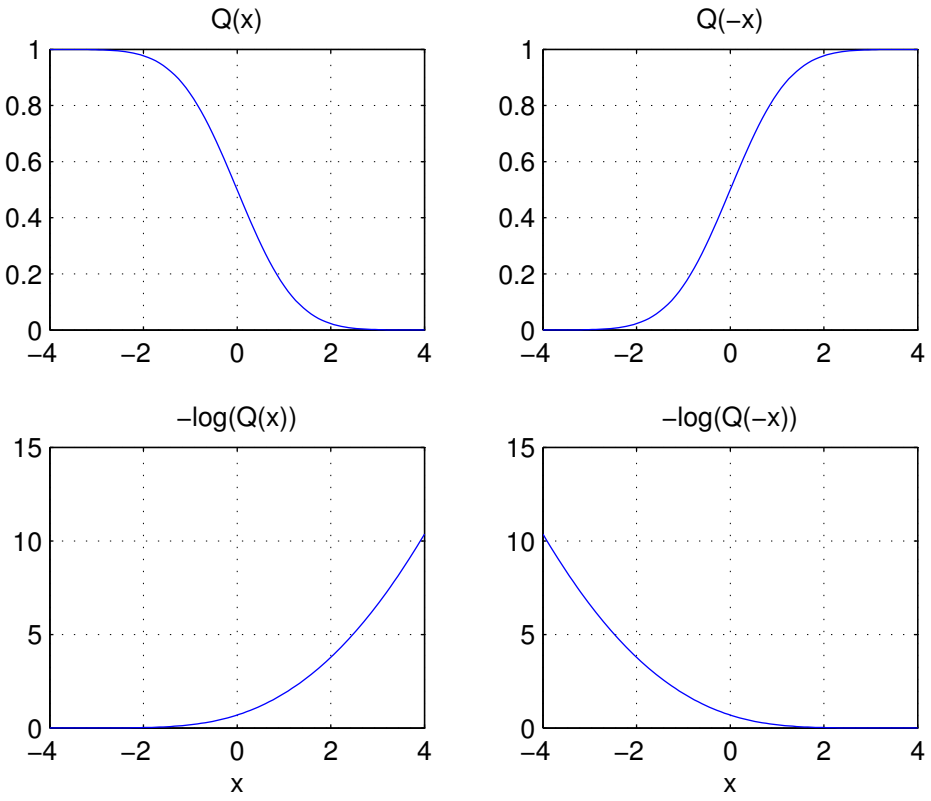

Figure A.1: The Q-function and its negative logarithm. The Q-function is monotonically decreasing and log-concave.

Q(-x)

$$Q(-x) = \int_{-x}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}} du \tag{A.4}$$

$$= 1 - \int_{-\infty}^{-x} \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}} du \qquad |v = -u \tag{A.5}$$

$$= 1 + \int_{\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{-\frac{v^2}{2}} dv \tag{A.6}$$

$$= 1 - \int_{x}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{v^2}{2}} dv \tag{A.7}$$

$$= 1 - Q(x) \tag{A.8}$$

$$= \Pr\{X \leq x\} \quad \text{if } X \sim \mathcal{N}(X; 0, 1) \tag{A.9}$$

Relation to erf **and** erfc

$$\operatorname{erf}(x) := \frac{2}{\sqrt{\pi}} \int_{0}^{x} e^{-t^2} dt \tag{A.10}$$

$$= 1 - 2Q(\sqrt{2}\, x) \tag{A.11}$$

$$\operatorname{erfc}(x) := \frac{2}{\sqrt{\pi}} \int_{x}^{\infty} e^{-t^2} dt \tag{A.12}$$

$$= 1 - \operatorname{erf}(x) \tag{A.13}$$

$$= 2Q(\sqrt{2}\, x) \tag{A.14}$$

$$Q(x) = \frac{\operatorname{erfc}\left(\frac{x}{\sqrt{2}}\right)}{2} \tag{A.15}$$

$$= \frac{1}{2}\left(1 - \operatorname{erf}\left(\frac{x}{\sqrt{2}}\right)\right) \tag{A.16}$$

General normal distribution If $X \sim \mathcal{N}(X; \mu, \sigma^2)$ then

$$\Pr\{X > x\} = \frac{1}{\sqrt{2\pi}\,\sigma} \int_x^\infty e^{-\frac{1}{2}\frac{(v-\mu)^2}{\sigma^2}} dv \quad |u = \frac{v-\mu}{\sigma},\ du = \frac{1}{\sigma} dv \tag{A.17}$$

$$= \frac{1}{\sqrt{2\pi}} \int_{\frac{x-\mu}{\sigma}}^\infty e^{-\frac{u^2}{2}} du \tag{A.18}$$

$$= Q\left(\frac{x-\mu}{\sigma}\right) \tag{A.19}$$

$$\Pr\{X \leq x\} = Q\left(-\frac{x-\mu}{\sigma}\right) \tag{A.20}$$

Inverse Function

$$Q^{-1}(y) = \sqrt{(}2)\,\mathrm{erfinv}(1-2y) \tag{A.21}$$

Derivatives

$$\frac{dQ(x)}{dx} = \frac{d}{dx} \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}} du \tag{A.22}$$

$$= -\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \tag{A.23}$$

$$= -\mathcal{N}(x; 0, 1) \tag{A.24}$$

where we made use of the Leibniz-rule for parameter integrals.

$$\frac{d^2Q(x)}{dx^2} = -(-x)\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \tag{A.25}$$

$$= x\mathcal{N}(x; 0, 1) \tag{A.26}$$

A.2 Optimal quantization threshold for single-bit quantization

In this appendix, we will prove Lemma 2 which states that the optimal quantization threshold in terms of minimizing the hit-or-miss Bayes cost for a single bit quantization of a linear measurement with Gaussian noise and Gaussian prior is the expected measurement.

In a MAP framework, we minimize the Bayes risk with respect to the vector "hit-or-

miss" cost function [54, p. 372]

$$\mathcal{C}(\mathbf{x}-\hat{\mathbf{x}}) = \begin{cases} 0 & ||\mathbf{x}-\hat{\mathbf{x}}||_2 < \epsilon \\ 1 & \text{else} \end{cases} \tag{A.27}$$

for some small ϵ. The Bayes Risk $\mathcal{R}$ is then

$$\mathcal{R} = E[\mathcal{C}(\mathbf{x}-\hat{\mathbf{x}})] \tag{A.28}$$

$$= \int \mathcal{C}(\mathbf{x}-\hat{\mathbf{x}}^+)p(\mathbf{x}, b=+1)d\mathbf{x} + \int \mathcal{C}(\mathbf{x}-\hat{\mathbf{x}}^-)p(\mathbf{x}, b=-1)d\mathbf{x} \tag{A.29}$$

where $\hat{\mathbf{x}}^+$ and $\hat{\mathbf{x}}^-$ refer to the estimates depending on the realization of the quantized measurement b.

When substituting the cost function into the first term, it can be decomposed as

$$\int \mathcal{C}(\mathbf{x}-\hat{\mathbf{x}}^+)p(\mathbf{x}, b=+1)d\mathbf{x} = \int\limits_{||\mathbf{x}-\hat{\mathbf{x}}^+||\geq\epsilon} 1 \cdot p(\mathbf{x}, b=+1)d\mathbf{x} + \int\limits_{||\mathbf{x}-\hat{\mathbf{x}}^+||<\epsilon} 0 \cdot p(\mathbf{x}, b=+1)d\mathbf{x} \tag{A.30}$$

where the first part represents almost the complete marginalization of $\mathbf{x}$, except for the contribution of the region within the epsilon-hypersphere, that is,

$$\int \mathcal{C}(\mathbf{x}-\hat{\mathbf{x}}^+)p(\mathbf{x}, b=+1)d\mathbf{x} = \int p(\mathbf{x}, b=+1)d\mathbf{x} - \int\limits_{||\mathbf{x}-\hat{\mathbf{x}}^+||<\epsilon} p(\mathbf{x}, b=+1)d\mathbf{x} \tag{A.31}$$

where the first part now equals $p(b=+1)$.

Proceeding analogously with the contribution of the negative quantization outcome, we can transform the Bayes Risk into

$$\mathcal{R} = p(b=+1) - \int\limits_{||\mathbf{x}-\hat{\mathbf{x}}^+||<\epsilon} p(\mathbf{x}, b=+1)d\mathbf{x} + p(b=-1) - \int\limits_{||\mathbf{x}-\hat{\mathbf{x}}^-||<\epsilon} p(\mathbf{x}, b=-1)d\mathbf{x} \tag{A.32}$$

$$\approx 1 - \left(p(b=+1|\mathbf{x})p(\mathbf{x})\Big|_{\mathbf{x}=\hat{\mathbf{x}}^+} + p(b=-1|\mathbf{x})p(\mathbf{x})\Big|_{\mathbf{x}=\hat{\mathbf{x}}^-} \right) \int\limits_{||\mathbf{x}||<\epsilon} d\mathbf{x} \tag{A.33}$$

$$= 1 - \left(Q\left(\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma}\right)\mathcal{N}(\hat{\mathbf{x}}^+; \boldsymbol{\mu}, \mathbf{P}) + Q\left(-\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma}\right)\mathcal{N}(\hat{\mathbf{x}}^-; \boldsymbol{\mu}, \mathbf{P}) \right) \int\limits_{||\mathbf{x}||<\epsilon} d\mathbf{x} \tag{A.34}$$

where we assume that ϵ is sufficiently small so that the pdf is approximately constant within the ϵ-hypersphere, and we use the fact that $p(b=+1) + p(b=-1) = 1$.

The goal is now to minimize $\mathcal{R}$ with respect to the quantization threshold τ and the

estimates $\hat{\mathbf{x}}^+$ and $\hat{\mathbf{x}}^-$. To this end, we compute the partial derivatives

$$\begin{aligned}\frac{\partial \mathcal{R}}{\partial \hat{\mathbf{x}}^+} = -\Bigg[&\frac{1}{\sigma}\mathcal{N}\left(\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma}; 0, 1\right)\mathcal{N}(\hat{\mathbf{x}}^+; \boldsymbol{\mu}, \mathbf{P})\mathbf{h} \\ &- Q\left(\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma}\right)\mathcal{N}(\hat{\mathbf{x}}^+; \boldsymbol{\mu}, \mathbf{P})\mathbf{P}^{-1}(\hat{\mathbf{x}}^+ - \boldsymbol{\mu})\Bigg]\int\limits_{||\mathbf{x}||<\epsilon} d\mathbf{x} \end{aligned} \tag{A.35}$$

$$\begin{aligned}\frac{\partial \mathcal{R}}{\partial \hat{\mathbf{x}}^-} = -\Bigg(&-\frac{1}{\sigma}\mathcal{N}\left(-\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma}; 0, 1\right)\mathcal{N}(\hat{\mathbf{x}}^-; \boldsymbol{\mu}, \mathbf{P})\mathbf{h} \\ &- Q\left(-\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma}\right)\mathcal{N}(\hat{\mathbf{x}}^-; \boldsymbol{\mu}, \mathbf{P})\mathbf{P}^{-1}(\hat{\mathbf{x}}^- - \boldsymbol{\mu})\Bigg)\int\limits_{||\mathbf{x}||<\epsilon} d\mathbf{x} \end{aligned} \tag{A.36}$$

$$\begin{aligned}\frac{\partial \mathcal{R}}{\partial \tau} = -\frac{1}{\sigma}\Bigg(&\mathcal{N}\left(-\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma}; 0, 1\right)\mathcal{N}(\hat{\mathbf{x}}^-; \boldsymbol{\mu}, \mathbf{P}) \\ &- \mathcal{N}\left(\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma}; 0, 1\right)\mathcal{N}(\hat{\mathbf{x}}^+; \boldsymbol{\mu}, \mathbf{P})\Bigg)\int\limits_{||\mathbf{x}||<\epsilon} d\mathbf{x} \end{aligned} \tag{A.37}$$

Setting the derivatives to zero, we obtain the following optimality conditions

$$\hat{\mathbf{x}}^+ = \boldsymbol{\mu} + \lambda^+\mathbf{P}\mathbf{h} \tag{A.38}$$

$$\hat{\mathbf{x}}^- = \boldsymbol{\mu} - \lambda^-\mathbf{P}\mathbf{h} \tag{A.39}$$

$$\mathcal{N}\left(-\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma}; 0, 1\right)\mathcal{N}(\hat{\mathbf{x}}^-; \boldsymbol{\mu}, \mathbf{P}) - \mathcal{N}\left(\frac{\tau - \mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma}; 0, 1\right)\mathcal{N}(\hat{\mathbf{x}}^+; \boldsymbol{\mu}, \mathbf{P}) = 0 \tag{A.40}$$

where the defining equations for λ^+, λ^- are given by

$$\lambda^+ = \frac{1}{\sigma}\frac{\mathcal{N}\left(\frac{\tau - \mathbf{h}^T\boldsymbol{\mu} - \mathbf{h}^T\lambda^+\mathbf{P}\mathbf{h}}{\sigma}; 0, 1\right)}{Q\left(\frac{\tau - \mathbf{h}^T\boldsymbol{\mu} - \mathbf{h}^T\lambda^+\mathbf{P}\mathbf{h}}{\sigma}\right)} \tag{A.41}$$

$$\lambda^- = \frac{1}{\sigma}\frac{\mathcal{N}\left(-\frac{\tau - \mathbf{h}^T\boldsymbol{\mu} + \mathbf{h}^T\lambda^-\mathbf{P}\mathbf{h}}{\sigma}; 0, 1\right)}{Q\left(-\frac{\tau - \mathbf{h}^T\boldsymbol{\mu} + \mathbf{h}^T\lambda^-\mathbf{P}\mathbf{h}}{\sigma}\right)} \tag{A.42}$$

To demonstrate that $\tau = \mathbf{h}^T\boldsymbol{\mu}$ is a stationary point, we first observe that in this case $\lambda^+ = \lambda^- = \lambda$, since for both the defining equations reduce to (see Fig. A.2)

$$\lambda = \frac{1}{\sigma}\frac{\mathcal{N}\left(\frac{-\mathbf{h}^T\lambda\mathbf{P}\mathbf{h}}{\sigma}; 0, 1\right)}{Q\left(\frac{-\mathbf{h}^T\lambda\mathbf{P}\mathbf{h}}{\sigma}\right)} \tag{A.43}$$

and, moreover, the first order optimality condition for τ, (A.40) holds, since

$$\mathcal{N}\left(-\frac{\tau-\mathbf{h}^T\boldsymbol{\mu}+\mathbf{h}^T\lambda^-\mathbf{P}\mathbf{h}}{\sigma};0,1\right)\mathcal{N}(\boldsymbol{\mu}-\lambda^-\mathbf{P}\mathbf{h};\boldsymbol{\mu},\mathbf{P})$$
$$-\mathcal{N}\left(\frac{\tau-\mathbf{h}^T\boldsymbol{\mu}-\mathbf{h}^T\lambda^+\mathbf{P}\mathbf{h}}{\sigma};0,1\right)\mathcal{N}(\boldsymbol{\mu}+\lambda^+\mathbf{P}\mathbf{h};\boldsymbol{\mu},\mathbf{P}) \tag{A.44}$$
$$=\mathcal{N}\left(-\frac{\mathbf{h}^T\lambda\mathbf{P}\mathbf{h}}{\sigma};0,1\right)\mathcal{N}(-\lambda\mathbf{P}\mathbf{h};\mathbf{0},\mathbf{P})-\mathcal{N}\left(-\frac{\mathbf{h}^T\lambda\mathbf{P}\mathbf{h}}{\sigma};0,1\right)\mathcal{N}(\lambda\mathbf{P}\mathbf{h};\mathbf{0},\mathbf{P}) \tag{A.45}$$
$$=0 \tag{A.46}$$

where the last equality exploits the fact that a zero-mean Gaussian pdf is symmetric.

To show that this stationary point is in fact a minimum, we consider the Hessian of the Bayes risk function. We determine the following second derivatives:

$$\begin{aligned}\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^{+^2}}=-\Bigg[&\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma}\mathcal{N}(\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^+;\boldsymbol{\mu},\mathbf{P})\frac{1}{\sigma^2}\mathbf{h}\mathbf{h}^T\\&-\mathcal{N}(\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^+;\boldsymbol{\mu},\mathbf{P})\mathbf{P}^{-1}(\hat{\mathbf{x}}^+-\boldsymbol{\mu})\frac{\mathbf{h}^T}{\sigma}\\&-\mathcal{N}(\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^+;\boldsymbol{\mu},\mathbf{P})\frac{\mathbf{h}}{\sigma}(\hat{\mathbf{x}}^+-\boldsymbol{\mu})^T\mathbf{P}^{-1}\\&+Q(\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma})\mathcal{N}(\hat{\mathbf{x}}^+;\boldsymbol{\mu},\mathbf{P})\mathbf{P}^{-1}(\hat{\mathbf{x}}^+-\boldsymbol{\mu})(\hat{\mathbf{x}}^+-\boldsymbol{\mu})^T\mathbf{P}^{-1}\\&-Q(\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma})\mathcal{N}(\hat{\mathbf{x}}^+;\boldsymbol{\mu},\mathbf{P})\mathbf{P}^{-1}\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x}\end{aligned} \tag{A.47}$$

$$\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^+\partial\hat{\mathbf{x}}^-}=\mathbf{0} \tag{A.48}$$

$$\begin{aligned}\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^+\partial\tau}=-\Bigg[&-\frac{1}{\sigma^2}\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma}\mathcal{N}(\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^+;\boldsymbol{\mu},\mathbf{P})\mathbf{h}\\&+\frac{1}{\sigma}\mathcal{N}(\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^+;\boldsymbol{\mu},\mathbf{P})\mathbf{P}^{-1}(\hat{\mathbf{x}}^+-\boldsymbol{\mu})\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x}\end{aligned} \tag{A.49}$$

$$\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^{-^2}} = -\Bigg[-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma}\mathcal{N}(-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^-;\boldsymbol{\mu},\mathbf{P})\frac{1}{\sigma^2}\mathbf{h}\mathbf{h}^T$$
$$+\mathcal{N}(-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^-;\boldsymbol{\mu},\mathbf{P})\mathbf{P}^{-1}(\hat{\mathbf{x}}^--\boldsymbol{\mu})\frac{\mathbf{h}^T}{\sigma}$$
$$+\mathcal{N}(-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^-;\boldsymbol{\mu},\mathbf{P})\frac{\mathbf{h}}{\sigma}(\hat{\mathbf{x}}^--\boldsymbol{\mu})^T\mathbf{P}^{-1}$$
$$+Q(-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma})\mathcal{N}(\hat{\mathbf{x}}^-;\boldsymbol{\mu},\mathbf{P})\mathbf{P}^{-1}(\hat{\mathbf{x}}^--\boldsymbol{\mu})(\hat{\mathbf{x}}^--\boldsymbol{\mu})^T\mathbf{P}^{-1}$$
$$-Q(-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma})\mathcal{N}(\hat{\mathbf{x}}^-;\boldsymbol{\mu},\mathbf{P})\mathbf{P}^{-1}\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x} \tag{A.50}$$

$$\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^-\partial\tau} = -\Bigg[\frac{1}{\sigma^2}\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma}\mathcal{N}(-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^-;\boldsymbol{\mu},\mathbf{P})\mathbf{h}$$
$$-\frac{1}{\sigma}\mathcal{N}(-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma};0,1)\mathcal{N}(\hat{\mathbf{x}}^-;\boldsymbol{\mu},\mathbf{P})\mathbf{P}^{-1}(\hat{\mathbf{x}}^--\boldsymbol{\mu})\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x} \tag{A.51}$$

$$\frac{\partial^2\mathcal{R}}{\partial\tau^2} = -\frac{1}{\sigma^2}\Bigg[-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma}\mathcal{N}\left(-\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^-}{\sigma};0,1\right)\mathcal{N}(\hat{\mathbf{x}}^-;\boldsymbol{\mu},\mathbf{P})$$
$$+\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma}\mathcal{N}\left(\frac{\tau-\mathbf{h}^T\hat{\mathbf{x}}^+}{\sigma};0,1\right)\mathcal{N}(\hat{\mathbf{x}}^+;\boldsymbol{\mu},\mathbf{P})\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x} \tag{A.52}$$

Evaluating these at $\tau=\mathbf{h}^T\boldsymbol{\mu}$ and the corresponding optimal solutions for $\hat{\mathbf{x}}^+$ and $\hat{\mathbf{x}}^-$ (see (A.38),(A.39)), and employing the substitution $\mathcal{N}\left(\frac{-\mathbf{h}^T\lambda\mathbf{P}\mathbf{h}}{\sigma};0,1\right)=\lambda\sigma Q\left(\frac{-\mathbf{h}^T\lambda\mathbf{P}\mathbf{h}}{\sigma}\right)$ from (A.43), we have

$$\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^{+^2}} = -\mathcal{N}(\lambda\mathbf{P}\mathbf{h};\mathbf{0},\mathbf{P})Q(-\frac{\lambda\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma})\Bigg[-\big(\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}+1\big)\lambda^2\mathbf{h}\mathbf{h}^T-\mathbf{P}^{-1}\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x} \tag{A.53}$$

$$\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^+\partial\hat{\mathbf{x}}^-} = \mathbf{0} \tag{A.54}$$

$$\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^+\partial\tau} = -\mathcal{N}(\lambda\mathbf{P}\mathbf{h};\mathbf{0},\mathbf{P})Q(-\frac{\lambda\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma})\Bigg[\big(\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}+1\big)\lambda^2\mathbf{h}\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x} \tag{A.55}$$

$$\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^{-^2}} = -\mathcal{N}(\lambda\mathbf{P}\mathbf{h};\mathbf{0},\mathbf{P})Q(-\frac{\lambda\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma})\Bigg[-\big(\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}+1\big)\lambda^2\mathbf{h}\mathbf{h}^T-\mathbf{P}^{-1}\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x} \tag{A.56}$$

$$\frac{\partial^2\mathcal{R}}{\partial\hat{\mathbf{x}}^-\partial\tau} = -\mathcal{N}(\lambda\mathbf{P}\mathbf{h};\mathbf{0},\mathbf{P})Q(-\frac{\lambda\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma})\Bigg[\big(\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}+1\big)\lambda^2\mathbf{h}\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x} \tag{A.57}$$

$$\frac{\partial^2\mathcal{R}}{\partial\tau^2} = -\mathcal{N}(\lambda\mathbf{P}\mathbf{h};\mathbf{0},\mathbf{P})Q(-\frac{\lambda\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma})\Bigg[-2\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}\lambda^2\Bigg]\int\limits_{||\mathbf{x}||<\epsilon}d\mathbf{x} \tag{A.58}$$

We can now write the Hessian, evaluated at the critical point, as

$$\mathbf{H} = \mathcal{N}(\lambda\mathbf{Ph}; \mathbf{0}, \mathbf{P})Q(-\frac{\lambda\mathbf{h}^T\mathbf{Ph}}{\sigma}) \int\limits_{||\mathbf{x}||<\epsilon} d\mathbf{x}$$
$$\cdot \begin{bmatrix} \mathbf{P}^{-1} + (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}\mathbf{h}^T & \mathbf{0} & -(1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h} \\ \mathbf{0} & \mathbf{P}^{-1} + (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}\mathbf{h}^T & -(1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h} \\ -(1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}^T & -(1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}^T & 2\frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2}\lambda^2 \end{bmatrix} \quad \text{(A.59)}$$

It remains to show that it is positive definite. To this end, we first observe that the leading scalar factors are positive. Also, the two matrices on the block-diagonal are each positive definite, since they are inverse covariance matrices with positive, additive rank-one updates. This further implies that all leading principal submatrices are positive definite, and hence their determinants are positive. This, together with a positive determinant of $\mathbf{H}$ would imply that $\mathbf{H}$ is positive definite [43].

The determinant of $\mathbf{H}$ can be computed using the Schur complement, as

$$\det(\mathbf{H}) = (\mathcal{N}(\lambda\mathbf{Ph}; \mathbf{0}, \mathbf{P})Q(-\frac{\lambda\mathbf{h}^T\mathbf{Ph}}{\sigma}) \int\limits_{||\mathbf{x}||<\epsilon} d\mathbf{x})^k$$
$$\cdot \det\left(\begin{bmatrix} \mathbf{P}^{-1} + (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}\mathbf{h}^T & \mathbf{0} \\ \mathbf{0} & \mathbf{P}^{-1} + (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}\mathbf{h}^T \end{bmatrix}\right)$$
$$\cdot \left(2\frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2}\lambda^2 - (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})^2\lambda^4\right.$$
$$\left.\cdot \begin{bmatrix}\mathbf{h}^T & \mathbf{h}^T\end{bmatrix} \begin{bmatrix} \mathbf{P}^{-1} + (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}\mathbf{h}^T & \mathbf{0} \\ \mathbf{0} & \mathbf{P}^{-1} + (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}\mathbf{h}^T \end{bmatrix}^{-1} \begin{bmatrix}\mathbf{h} \\ \mathbf{h}\end{bmatrix}\right) \quad \text{(A.60)}$$

where k denotes the number of rows/columns of $\mathbf{H}$. Since the scalar factors are positive and the block-diagonal matrix is positive definite, we only need to concern ourselves with the final scalar expression. After invoking the matrix inversion lemma, this factor can be written as

$$2\frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2}\lambda^2 - (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})^2\lambda^4$$
$$\cdot \begin{bmatrix}\mathbf{h}^T & \mathbf{h}^T\end{bmatrix} \begin{bmatrix} \mathbf{P}^{-1} + (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}\mathbf{h}^T & \mathbf{0} \\ \mathbf{0} & \mathbf{P}^{-1} + (1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2\mathbf{h}\mathbf{h}^T \end{bmatrix}^{-1} \begin{bmatrix}\mathbf{h} \\ \mathbf{h}\end{bmatrix} \quad \text{(A.61)}$$
$$= 2\frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2}\lambda^2 - 2(1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})^2\lambda^4\frac{\mathbf{h}^T\mathbf{Ph}}{1 + \mathbf{h}^T\mathbf{Ph}(1 + \frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2})\lambda^2} \quad \text{(A.62)}$$
$$= 2\frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2}\lambda^2\frac{1 - \lambda^2\sigma^2(\frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2} + 1)}{1 + \mathbf{h}^T\mathbf{Ph}(\frac{\mathbf{h}^T\mathbf{Ph}}{\sigma^2} + 1)\lambda^2} \quad \text{(A.63)}$$

What remains to be shown is therefore

$$1 - \lambda^2\sigma^2(\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2} + 1) \geq 0 \tag{A.64}$$

Our strategy will be to find an upper bound for $\lambda\sigma$ and show that the above inequality is satisfied for this upper bound (see Fig. A.2). Recall the definition (A.43) of $\lambda\sigma$

$$\lambda\sigma = \frac{\mathcal{N}\left(\frac{-\lambda\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma}; 0, 1\right)}{Q\left(\frac{-\lambda\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma}\right)} \tag{A.65}$$

We perform a change of variables to obtain

$$x = -\frac{\lambda\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma} \tag{A.66}$$

$$\lambda\sigma = -x\frac{\sigma^2}{\mathbf{h}^T\mathbf{P}\mathbf{h}} \tag{A.67}$$

$$-x\frac{\sigma^2}{\mathbf{h}^T\mathbf{P}\mathbf{h}} = \frac{\mathcal{N}(x; 0, 1)}{Q(x)} \tag{A.68}$$

Let us further introduce the functions

$$f(x) := -x\frac{\sigma^2}{\mathbf{h}^T\mathbf{P}\mathbf{h}} \tag{A.69}$$

$$g(x) := \frac{\mathcal{N}(x; 0, 1)}{Q(x)} \tag{A.70}$$

We remark that $g(x)$ is well-studied in the statistics community where its reciprocal is known as the Mill's ratio. Further, it is known [104] that $g(x)$ is upper bounded by the function

$$h(x) := \frac{2}{\sqrt{x^2 + 4} - x} \geq g(x) \quad \forall x \tag{A.71}$$

Among the properties of $g(x)$ and $h(x)$ are that both are strictly positive and monotonically increasing [104], which can be verified by checking global non-negativity of their derivative. In contrast, $f(x)$ is linear and monotonically decreasing, since $-\frac{\sigma^2}{\mathbf{h}^T\mathbf{P}\mathbf{h}} < 0$.

Let us denote the solution of (A.68) as x_1, so that $f(x_1) = g(x_1)$. Also, define as x_2 the solution of the modified system $f(x) = h(x)$, where $g(x)$ has been replaced by its upper bound. We now show that $x_2 \leq x_1$ by contradiction. Assume that $x_2 > x_1$. Then, using the defining equations, the bounding property of $h(x)$, and function monotonicity, we have

the following chain of inequalities

$$\text{Ass. } x_2 > x_1 \tag{A.72}$$

$$f(x_1) = g(x_1) \leq h(x_1) \leq h(x_2) = f(x_2) \tag{A.73}$$

$$\Rightarrow f(x_1) \leq f(x_2) \quad \text{Contradiction!} \tag{A.74}$$

which is a contradiction since $f(x)$ is monotonically decreasing.

Notice further, that from the definition of $f(x)$, we have

$$(\lambda\sigma) = f(x_1) \tag{A.75}$$

$$(\lambda\sigma)' := f(x_2) \geq (\lambda\sigma) \tag{A.76}$$

which follows from $x_2 \leq x_1$. We therefore now have an upper bound for $\lambda\sigma$ which can be computed analytically. In particular, after some algebraic manipulations,

$$f(x_2) = h(x_2) \tag{A.77}$$

$$x_2 = \pm \frac{\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}}{\sqrt{1 + \frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}}} \tag{A.78}$$

of which only the negative solution is valid. We further have

$$(\lambda\sigma)' = h(x_2) = \frac{1}{\sqrt{1 + \frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}}} \tag{A.79}$$

Inserting this into our original positivity constraints from the determinant, we have

$$1 - \lambda^2\sigma^2(\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2} + 1) \geq 1 - (\lambda\sigma)'^2(\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2} + 1) = 1 - \frac{1}{1 + \frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2}}(\frac{\mathbf{h}^T\mathbf{P}\mathbf{h}}{\sigma^2} + 1) = 0 \tag{A.80}$$

From this, we conclude that indeed the Hessian is positive definite at the stationary point, and that therefore the choice of threshold $\tau = \mathbf{h}^T\boldsymbol{\mu}$ is a minimizer of the MAP Bayes risk.

As a side note for the interested reader, we also have an additional bound, $(\lambda\sigma)'' = \sqrt{\frac{2}{\pi}}$, which follows from $g(0) = \sqrt{\frac{2}{\pi}}$ and the monotonicity of $g(x)$ (see Fig. A.2). We can use these bounds on $\lambda\sigma$ to compare the corresponding bounds on the solutions for the estimates $\hat{\mathbf{x}}$ to the estimate from the SOI-KF. In particular, we have that

$$\hat{\mathbf{x}} = \boldsymbol{\mu} + b\lambda\mathbf{P}\mathbf{h} \tag{A.81}$$

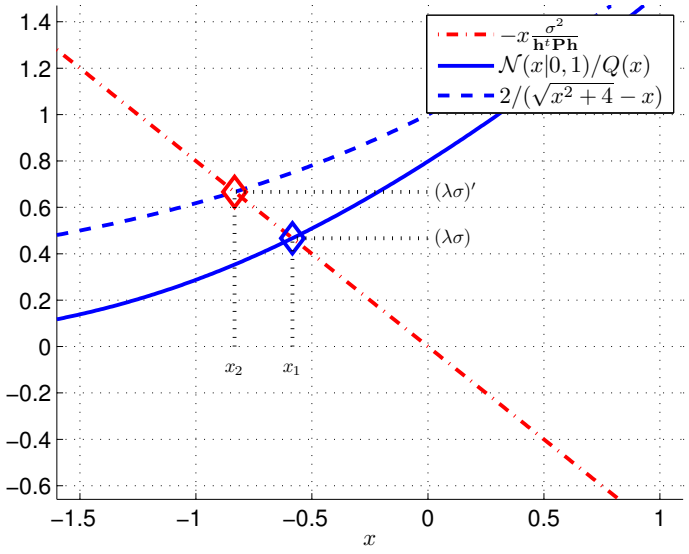

Figure A.2: Graphical depiction of the computation and upper bounding of $\lambda\sigma$.

where

$$\lambda_{\text{SOI-KF}} = \sqrt{\frac{2}{\pi}}\frac{1}{\sqrt{\mathbf{h}^T\mathbf{P}\mathbf{h}+\sigma^2}} \tag{A.82}$$

$$\lambda' = \frac{1}{\sqrt{\mathbf{h}^T\mathbf{P}\mathbf{h}+\sigma^2}} \tag{A.83}$$

$$\lambda'' = \sqrt{\frac{2}{\pi}}\frac{1}{\sqrt{\sigma^2}} \tag{A.84}$$

so that

$$\lambda' > \lambda_{\text{SOI-KF}}, \quad \lambda'' \geq \lambda_{\text{SOI-KF}} \tag{A.85}$$

In simulations we have further observed that $\lambda_{\text{SOI-KF}} \geq \lambda_{\text{MAP}}$. Although not proven, this observation appears reasonable in view of the mean-median-mode inequality, since the SOI-KF as an MMSE estimator determines the mean of the posterior pdf whereas the QMAP determines the mode, and the quantized posterior is unimodal and heavy-tailed.