

Handling Constrained Optimization in Factor Graphs for Autonomous Navigation

Barbara Bazzana Tiziano Guadagnino Giorgio Grisetti

Abstract—Factor graphs are graphical models used to represent a wide variety of problems across robotics, such as Structure from Motion (SfM), Simultaneous Localization and Mapping (SLAM) and calibration. Typically, at their core, they have an optimization problem whose terms only depend on a small subset of variables. Factor graph solvers exploit the locality of problems to drastically reduce the computational time of the Iterative Least-Squares (ILS) methodology. Although extremely powerful, their application is usually limited to unconstrained problems. In this paper, we model constraints over variables within factor graphs by introducing a factor graph version of the method of Lagrange Multipliers. We show the potential of our method by presenting a full navigation stack based on factor graphs. Differently from standard navigation stacks, we can model both optimal control for local planning and localization with factor graphs, and solve the two problems using the standard ILS methodology. We validate our approach in real-world autonomous navigation scenarios, comparing it with the de facto standard navigation stack implemented in ROS. Comparative experiments show that for the application at hand our system outperforms the standard nonlinear programming solver Interior-Point Optimizer (IPOPT) in runtime, while achieving similar solutions.

I. INTRODUCTION

Factor graphs are a general graphical formalism to model several problems. The robotics community extensively used factor graphs to approach estimation problems, such as SfM, SLAM and calibration [8], [14], [15], [22]. On the one hand, they provide a graphical representation that exposes the general structure of the problem which, if exploited, allows the design of efficient algorithms. On the other hand, factor graphs allow to naturally couple problems sharing common variables by simply joining them. The solution to the joined factor graph will be then the optimum of the coupled problem. Well-known methodologies to solve large factor graphs use variants of Iterative Least Squares (ILS) [13], [16], [17] that do not support *constrained* optimization in its original formulation.

Leveraging on the results of [25], in this paper we present an iterative version of the method of Lagrange Multipliers where constraints are modeled as a new class of factors. Being able to perform constrained optimization allows extending the domain of factor graphs to general optimization problems. As an application, we present here a factor graph-Model Predictive Control (MPC) framework for unicycle navigation. In particular, we address here localization and motion planning within a single framework as the same general purpose factor graph solver [13] is used for both problems. The aim of this paper is to present a unified formulation for estimation




Fig. 1: Left: a trajectory executed by our real robot avoiding an unforeseen obstacle, a schematic factor graph is superposed to it where some robot states are represented (pink circles). Both the localization problem (orange square) and the optimal control are depicted. States are linked by the motion model (yellow squares) and subject to obstacle avoidance and trajectory tracking (green squares). Right: our custom made unicycle.

and control problems. This unification can be very convenient as estimation and control could consistently share common information such as system dynamics, external disturbances, and uncertainty models. Thanks to the design of efficient algorithms this can be done while maintaining high computational performance.

Our system models the localization problem as scan matching with implicit data association where the likelihood of the current estimate is evaluated through the distance between the scan and the map. Experimental results show that, given the odometry as initial guess, this method has similar performances to the Adaptive Monte Carlo Localization ROS package `amcl`¹.

The optimal control problem we consider is that of generating a collision-free trajectory for a differential drive robot with actuation limits in a dynamic environment. To this extent, we introduce *constraint* factors and *obstacle avoidance* factors: the latter ones model the minimization of an artificial potential function [24] [3].

We validated our approach by running our system both in simulation and on a custom-made robot, where all software runs on a Raspberry PI 4. Besides showing the feasibility of our approach, real-world experiments highlight its moderate computational requirements. Fig. 1 shows a real-world trajectory and a simplistic representation of the related factor graph. We compare our system with the well-known ROS navigation stack, combining `amcl` for localization and the `teb_local_planner`², the Timed Elastic Band [21] ROS plugin of `move_base`³, for local planning. In particular, we provide a set of goals within the map of our Department and

Barbara Bazzana, Tiziano Guadagnino, Giorgio Grisetti are with the Department of Computer, Control, and Management Engineering “Antonio Ruberti”, Sapienza University of Rome, Rome, Italy, Email: {bazzana, guadagnino, grisetti}@diag.uniroma1.it.

¹<http://wiki.ros.org/amcl>

²http://wiki.ros.org/teb_local_planner

³http://wiki.ros.org/move_base

that of a factory-like environment, and consider the path length and associated duration for each segment.

Finally, experimental results confirm that when approaching constrained optimization problems on the models we analyzed, our generic solver outperforms a state-of-the-art nonlinear programming solver such as IPOPT in runtime while producing comparable trajectories and final cost value. This gives a glimpse of the computational performances of factor graphs motivating their use also for very large problems.

II. RELATED WORK

Factor graphs are a very powerful tool to model a great variety of unconstrained optimization problems, spacing from SLAM to SfM [15] [22]. In the recent literature, they have been applied to model optimal control problems as well [12], [18], [26], [28], [29]. To this extent, extensions of factor graph solvers have been proposed to handle constraints arising from optimal control modeling. Formalizing optimal control problems using factor graphs allows unifying the dual problems of estimation and control under the same representation. Thanks to this unification, common information and variables can be shared consistently in both estimation and control processes. Further, these two aspects can be addressed jointly taking advantage of efficient factor graph solvers [13], [16], [17].

To the best of our knowledge, Mukadam et al. *et al.* [18] were the first to optimize a unique factor graph for both trajectory estimation and motion planning, where trajectories are modeled using continuous Gaussian Processes [11]. However, they were only relying on soft constraints. Later on, the problem of addressing constrained optimization received increasing attention, starting from Ta *et al.* [26]. They present a factor graph version of Sequential Quadratic Programming (SQP) to handle nonlinear equality constraints and apply their approach to the development of an MPC framework on unmanned aerial vehicles. Our work can be viewed as an extension of this method, where we model both equality and inequality constraints using factor graphs. In place of SQP, we rely on the Method of Lagrange Multipliers because, differently from SQP, no quadratic programming solver is required for the internal iterations when addressing inequalities.

Yang *et al.* [29] include control inputs in a factor graph to solve a Linear Quadratic Regulator problem subject to auxiliary equality constraints. During the variable elimination process, when constrained variables are eliminated, a specialized solver is used for solving the constrained sub-problem separately. Differently from them, our method can handle both equality and inequality constraints and does not require a specialized solver, as it uses a standard ILS algorithm.

To the best of our knowledge, only Xie *et al.* [28] introduce both equality and inequality constraints in factor graphs. In particular, they present a factor graph version of a barrier-based approach to constrained optimization similar to the well-known Interior-Point Method. Our work can be viewed as complementary to Xie *et al.*, in that we present an implementation of the Method of Lagrange Multipliers, as an alternative to the Interior-Point Method.

Sodhi *et al.* [25] proposes an approach to refine the estimate of past trajectories: to accomplish collision-free trajectories, states are subject to inequality constraints. To the best of our knowledge they have been the first to introduce the Method of Lagrange Multipliers in the context of incremental smoothing. We rely on the same concepts to develop a factor-graph version of this method. We propose a different application than Sodhi *et al.*: instead of state estimation, we address optimal control problems to develop a factor graph-based MPC controller that can be used as an elegant and compact local planner. In addition, we address the problem of obstacle avoidance by designing a factor that relies on a distance cost function inspired by scan registration techniques. The resulting factor graph is solved by our generic factor graph solver [13].

III. FACTOR GRAPH OPTIMIZATION

In this paper, we model both localization and MPC using factor graphs, which are bipartite graphs with two kinds of nodes: variables and factors. Variables represent the state of our system, while factor nodes model measurements connecting the set of variables from which they depend.

More formally, let $\mathbf{x} = \mathbf{x}_{0:N}$ be the set of all variables which can span over arbitrary continuous domains. If the measurements $\mathbf{z} = \mathbf{z}_{0:N}$ are affected by Gaussian noise, we can represent the k^{th} factor as the tuple $\langle \mathbf{z}_k, \boldsymbol{\Omega}_k, \mathbf{h}_k(\cdot) \rangle$ comprising the mean \mathbf{z}_k , the information matrix $\boldsymbol{\Omega}_k$, and the prediction function $\mathbf{h}_k(\mathbf{x}_k)$, with \mathbf{x}_k the set of variables connected to the k^{th} factor. In this setting the negative log-likelihood of the factor graph becomes:

$$F(\mathbf{x}) = \sum_{k=0}^K \|\mathbf{h}_k(\mathbf{x}_k) - \mathbf{z}_k\|_{\boldsymbol{\Omega}_k} = \sum_{k=0}^K \|\mathbf{e}_k(\mathbf{x}_k)\|_{\boldsymbol{\Omega}_k} \quad (1)$$

Solving a factor graph means finding the \mathbf{x} which minimizes Eq. (1), i.e. the \mathbf{x} which is maximally consistent with the measurements. In the typical case of SLAM, variables can be robot states or landmark poses, while factors represent measurements correlating them. In optimal control, variables can be robot states or control inputs, whereas factors model terms of the objective function. We refer the reader to [10], [13] for more details. One of the aims of this paper is to address constrained optimization by extending classical ILS solvers.

ILS minimizes Eq. (1) by using Gauss-Newton (GN) algorithm that iteratively refines the current solution $\hat{\mathbf{x}}$ by solving the quadratic approximation of Eq. (1). The first-order Taylor expansion of the errors $\mathbf{e}_k(\mathbf{x}_k)$ around $\hat{\mathbf{x}}$ given the perturbation $\Delta \mathbf{x}$ is computed as:

$$\mathbf{e}_k(\hat{\mathbf{x}}_k + \Delta \mathbf{x}_k) \simeq \overbrace{\mathbf{e}_k(\hat{\mathbf{x}}_k)}^{\hat{\mathbf{e}}_k} + \mathbf{J}_k \Delta \mathbf{x}_k \quad (2)$$

By substituting Eq. (2) in Eq. (1), we obtain a quadratic form that approximates the cost function around $\hat{\mathbf{x}}$:

$$F(\hat{\mathbf{x}} + \Delta \mathbf{x}) \simeq c + 2\mathbf{b}^T \Delta \mathbf{x} + \Delta \mathbf{x}^T \mathbf{H} \Delta \mathbf{x} \quad (3)$$




Fig. 2: Factor graph modeling the localization problem. Each measurement contributes with a factor. The green square represents the odometry prior, while purple squares represent robot measurements $\mathbf{z} = \mathbf{z}_{0:K}$, e.g. laser endpoints.




Fig. 3: Factor graph modeling an optimal control problem: green squares represent prior factors imposing that the estimate is close to the desired values \mathbf{x}^{ref} and \mathbf{u}^{ref} , yellow squares factors modeling the motion model, blue squares upper and lower constraint factors, orange squares obstacle avoidance factors.

where

$$\mathbf{b} = \sum_{k=0}^K \mathbf{J}_k^T \Omega_k \hat{\mathbf{e}}_k \quad \mathbf{H} = \sum_{k=0}^K \mathbf{J}_k^T \Omega_k \mathbf{J}_k \quad (4)$$

The minimum $\Delta \mathbf{x}$ of the quadratic form of Eq. (3) is the solution of the linear system $\mathbf{H} \Delta \mathbf{x} = -\mathbf{b}$, while the next estimate is updated by applying the perturbation $\hat{\mathbf{x}} \leftarrow \hat{\mathbf{x}} + \Delta \mathbf{x}$.

IV. LOCALIZATION AS REGISTRATION ON A DISTANCE MATRIX

Localization aims at estimating the most likely robot pose in the world frame $\mathbf{X}_R^W \in SE(3)$, given the robot measurements $\mathbf{z} = \mathbf{z}_{0:K}$ and the odometry prior \mathbf{z}_O . The factor graph in Fig. 2 illustrates this problem. In the reminder, we will specialize this general schema in the case of laser-based localization. The localization module proposed leverages recent developments in photometric registration on a distance cost function $d(\cdot)$ for the laser endpoints.

Our robot is equipped with a 2D lidar and moves on a plane, hence $\mathbf{X}_R^W \in SE(2)$. Let us indicate as $\mathbf{z}_k \in \mathbb{R}^2$ a generic laser endpoint expressed in the robot frame. Let $d(\mathbf{p}) : \mathbb{R}^2 \rightarrow \mathbb{R}$ be the 2D function that returns the minimum distance between a point \mathbf{p} in world coordinates and the closest obstacle in the map. Knowing the robot pose \mathbf{X}_R^W we can estimate the distance between an endpoint and the closest obstacle as:

$$\mathbf{e}_k^l(\mathbf{X}_R^W) = d(\mathbf{X}_R^W \mathbf{z}_k). \quad (5)$$

Should the robot be perfectly localized in a static environment, all distances will be close to zero, therefore Eq. (5) is a perfect candidate to construct a factor. The localization graph will

result in the following cost function:

$$G(\mathbf{X}_R^W) = \sum_{k=0}^K \|\mathbf{e}_k^l(\mathbf{X}_R^W \mathbf{z}_k)\|_{\Omega_k^l} \quad (6)$$

$$\mathbf{X}_R^{W*} = \underset{\mathbf{X}_R^W}{\operatorname{argmin}} G(\mathbf{X}_R^W) \quad (7)$$

In our implementation the distance function is computed once, and stored as a 2D matrix having the same resolution of the map. Each cell of the matrix contains the value of $d(\mathbf{p})$ at that point. Subcell values are obtained through bilinear interpolation. The Jacobian of Eq. (5) is computed by applying the chain rule where the Jacobian term originated by the distance function is computed numerically. Besides being very intuitive, our approach combines very fast runtime with high accuracy and ease of implementation.

V. MPC WITH OBSTACLE AVOIDANCE FOR UNICYCLE

As for localization, we will now describe how an optimal control problem can be represented as a factor graph. MPC is a well-established technique for optimally controlling systems subject to constraints. At each time step, an MPC controller computes the best feasible trajectory which minimizes the assigned objective function, while avoiding obstacles and satisfying the constraints. Only the first control input is applied to the system; at the next time step, a new trajectory embedding information coming from the most recent measurements is computed. More in detail, let $\mathbf{u} = \mathbf{u}_{0:N-1}$ be the sequence of controls, $\mathbf{x} = \mathbf{x}_{0:N}$ the resulting chain of states, $\mathbf{u}^{ref} = \mathbf{u}_{0:N-1}^{ref}$ and $\mathbf{x}^{ref} = \mathbf{x}_{0:N}^{ref}$ the reference values for controls and states respectively, one possible log-likelihood is the Quadratic Regulator (QR) cost function $\Phi_{QR}(\mathbf{u}, \mathbf{x})$ [1] subject to the robot dynamics and the actuation limits:

$$\Phi_{QR}(\mathbf{u}, \mathbf{x}) = \sum_{n=0}^N \|\overbrace{\mathbf{x}_n - \mathbf{x}_n^{ref}}^{\mathbf{e}_n^x(\mathbf{x}_n)}\|_{\Omega_n^x} + \sum_{n=0}^{N-1} \|\overbrace{\mathbf{u}_n - \mathbf{u}_n^{ref}}^{\mathbf{e}_n^u(\mathbf{u}_n)}\|_{\Omega_n^u} \quad (8)$$

The link between Eq. (1) and Eq. (8) is evident when the control chain becomes part of the state. As the task of MPC is to compute both the trajectory *and* the sequence of controls, in the language of factor graphs controls will become variable nodes as well. Hence, each term in the above equation can be modeled as a prior factor of degree one, with reference value \mathbf{x}_n^{ref} or \mathbf{u}_n^{ref} . In contrast, obstacle avoidance and constrained optimization require more effort to be translated into the same formalism. Fig. 3 illustrates the factor graph formulation of an optimal control problem.

A feasible trajectory can be found by minimizing a cost function $g(\mathbf{x})$ that decreases with the distance from obstacles, being zero at safe locations. According to the theory of Artificial Potential Fields [24] [3], having denoted as k a positive real scalar, as μ a low-distance threshold, and as ρ

a high-distance threshold, we designed $g(\mathbf{x})$ as:

$$g(\mathbf{x}) = \begin{cases} k \left(\frac{1}{\mu} - \frac{1}{\rho} \right) & \text{if } d(\mathbf{x}) < \mu \\ k \left(\frac{1}{d(\mathbf{x})} - \frac{1}{\rho} \right) & \text{if } \mu < d(\mathbf{x}) < \rho \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

The above equation imposes a finite maximum value to $g(\mathbf{x})$ by clamping it to its maximum where $d(\mathbf{x}) < \mu$, and shapes a valley of zero-potential where $d(\mathbf{x}) > \rho$. Points in the region $d(\mathbf{x}) > \rho$ are safe and do not need to be pushed away by the optimization process. Therefore we model the obstacle avoidance through a factor whose error is:

$$\mathbf{e}_n^o(\mathbf{x}_n) = g(\mathbf{x}_n). \quad (10)$$

It is well known that Artificial Potential Fields are affected by local minima. In our implementation, we have modified the gradient of the repulsive field generated by Eq. (9) to overcome such minima. Using the idea of vortex fields [9], we add to the repulsive gradient a vector which is tangent to the equipotential contours of the repulsive field. By applying the obstacle avoidance factor to each robot state, the vortex field acts on the whole trajectory rather than on a single state [9], to deflect it from unexpected obstacles.

Let us now concentrate on the kinematic model $\mathbf{f}(\cdot)$ mapping current state \mathbf{x}_n and control \mathbf{u}_n to next state \mathbf{x}_{n+1} . The error can be modeled as a soft constraint:

$$\mathbf{e}_n^p(\mathbf{x}_n) = \mathbf{x}_{n+1} - \mathbf{f}(\mathbf{x}_n, \mathbf{u}_n). \quad (11)$$

The corresponding factor has an information matrix Ω_n^p giving the stiffness of the soft constraint. In an MPC fashion, at each time step a graph is constructed starting from the current localization estimate which becomes the first (fixed) variable \mathbf{x}_0 . Our robot is a unicycle, controlled in translational and rotational velocities, denoted respectively with v_n and ω_n . Using Runge-Kutta integration [5] with T_s as integration interval and being $(x_n, y_n)^T$ the position and θ_n the orientation of the n^{th} pose \mathbf{x}_n , the kinematics is as follows:

$$x_{n+1} = x_n + v_n T_s \cos \left(\theta_n + \frac{\omega_n T_s}{2} \right) \quad (12)$$

$$y_{n+1} = y_n + v_n T_s \sin \left(\theta_n + \frac{\omega_n T_s}{2} \right) \quad (13)$$

$$\theta_{n+1} = \theta_n + \omega_n T_s \quad (14)$$

Applying Eq. (12) in the solution scheme of Sec. III might result in controls $\mathbf{u} = (v, \omega)_{0:N-1}^T$ that exceed the actuation limits. Hence we need to enforce *inequality* constraints over the variables \mathbf{u}_n , which are not addressed by regular ILS. In the next section, we introduce *constraint factors* and *constrained variables* to model both inequality and equality constraints.

Being described by the factors $\mathbf{e}_n^x(\mathbf{x}_n)$, $\mathbf{e}_n^u(\mathbf{u}_n)$, $\mathbf{e}_n^o(\mathbf{x}_n)$, and $\mathbf{e}_n^p(\mathbf{x}_n)$, the problem has a block diagonal structure. In our implementation, we impose the variable ordering $\mathbf{x}_0, \mathbf{u}_0, \mathbf{x}_1, \dots$ to guarantee a maximally sparse structure of the linear system.

VI. EXTENDING ITERATIVE LEAST SQUARES SOLVERS FOR CONSTRAINED OPTIMIZATION

Lagrange Multipliers [2] is perhaps the most known method to deal with constrained optimization. Although the method tackles nonlinear constraint functions as well, we will focus on the linear case to give a simple overview. Consider the problem:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{k=0}^K \overbrace{\|\mathbf{e}_k(\mathbf{x}_k)\|_{\Omega_k}}^{F(\mathbf{x})}, \quad \text{s.t. } \mathbf{C}\mathbf{x} + \mathbf{c} = 0 \quad (15)$$

where $\mathbf{C} \in \mathbb{R}^{c \times n}$. The augmented Lagrangian of Eq. (15) is defined as:

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}^c; \rho^c) = F(\mathbf{x}) + (\boldsymbol{\lambda}^c)^T (\mathbf{C}\mathbf{x} + \mathbf{c}) + \frac{\rho^c}{2} \|\mathbf{C}\mathbf{x} + \mathbf{c}\|_2^2. \quad (16)$$

Here, $\boldsymbol{\lambda}^c \in \mathbb{R}^c$ is a vector of Lagrange Multipliers and the term $\frac{\rho^c}{2} \|\cdot\|_2^2$, with $\rho^c \in \mathbb{R}$, is a penalty term. It can be shown [4] that the minimum of Eq. (15) can be found by iteratively computing:

$$\mathbf{x} \leftarrow \underset{\mathbf{x}}{\operatorname{argmin}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}^c; \rho^c) \quad (17)$$

$$\boldsymbol{\lambda}^c \leftarrow \boldsymbol{\lambda}^c + \rho^c (\mathbf{C}\mathbf{x} + \mathbf{c}). \quad (18)$$

This method can be generalized to deal also with inequality constraints $\mathbf{D}\mathbf{x} + \mathbf{d} \leq 0$ with $\mathbf{D} \in \mathbb{R}^{d \times n}$ by adding an additional multiplier $\boldsymbol{\lambda}^d \in \mathbb{R}^d$ and an additional penalty term weighted by $\rho^d \in \mathbb{R}$. The new augmented Lagrangian then becomes:

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}^c, \boldsymbol{\lambda}^d; \rho^c, \rho^d) = & F(\mathbf{x}) \\ & + (\boldsymbol{\lambda}^c)^T (\mathbf{C}\mathbf{x} + \mathbf{c}) + \frac{\rho^c}{2} \|\mathbf{C}\mathbf{x} + \mathbf{c}\|_2^2 \\ & + (\boldsymbol{\lambda}^d)^T (\mathbf{D}\mathbf{x} + \mathbf{d}) + \frac{\rho^d}{2} \|\mathbf{D}\mathbf{x} + \mathbf{d}\|_2^2 \end{aligned} \quad (19)$$

The current state estimate is the \mathbf{x} which minimizes Eq. (19). Further, the update-law of the inequality multiplier $\boldsymbol{\lambda}^d$ is the following:

$$\boldsymbol{\lambda}^d \leftarrow \max(0, \boldsymbol{\lambda}^d + \rho^d (\mathbf{D}\mathbf{x} + \mathbf{d})) \quad (20)$$

In our application, control inputs are subject to inequality constraints, namely: $|v| < v_{max}$ and $|\omega| < \omega_{max}$.

The Method of Lagrange Multipliers belongs to the class of primal-dual methods which split the optimization in two steps: a primal step (17) and a dual step (18, 20). Implementing the dual update requires extending the *constrained variables* with the multipliers. An iteration proceeds by updating the constrained variable \mathbf{x} while keeping the multipliers fixed through the primal step; once \mathbf{x} is refined, the multipliers $\boldsymbol{\lambda}^c$ and $\boldsymbol{\lambda}^d$ are updated with \mathbf{x} fixed through the dual step.

The primal step Eq. (17) is a minimization that can be solved through ILS since the multipliers are fixed. Let $\hat{\mathbf{x}}$ be the current solution, $\hat{\mathbf{c}} = \mathbf{C}\hat{\mathbf{x}} + \mathbf{c}$ and $\hat{\mathbf{d}} = \mathbf{D}\hat{\mathbf{x}} + \mathbf{d}$ be the constraint values computed at $\hat{\mathbf{x}}$. During the linearization step

of ILS, the quadratic form approximating Eq. (19) around $\hat{\mathbf{x}}$ then becomes:

$$\mathcal{L}(\hat{\mathbf{x}} + \Delta\mathbf{x}, \lambda^c, \lambda^d; \rho^c, \rho^d) \simeq c + 2\mathbf{b}^T \Delta\mathbf{x} + \Delta\mathbf{x}^T \mathbf{H} \Delta\mathbf{x} \quad (21)$$

where

$$\mathbf{b} = \sum_{k=0}^K \mathbf{J}_k^T \Omega_k \hat{\mathbf{e}}_k + \underbrace{\frac{\rho^c}{2} \mathbf{C}^T \hat{\mathbf{c}} + \frac{1}{2} \mathbf{C}^T \lambda^c}_{\mathbf{b}^c} + \underbrace{\frac{\rho^d}{2} \mathbf{D}^T \hat{\mathbf{d}} + \frac{1}{2} \mathbf{D}^T \lambda^d}_{\mathbf{b}^d} \quad (22)$$

$$\mathbf{H} = \sum_{k=0}^K \mathbf{J}_k^T \Omega_k \mathbf{J}_k + \underbrace{\frac{\rho^c}{2} \mathbf{C}^T \mathbf{C}}_{\mathbf{H}^c} + \underbrace{\frac{\rho^d}{2} \mathbf{D}^T \mathbf{D}}_{\mathbf{H}^d}.$$

By comparing Eq. (22) and Eq. (4), it turns out that the primal update can be implemented in ILS solvers by introducing *constraint factors* that contribute to the \mathbf{H} and the \mathbf{b} matrices respectively through \mathbf{H}^c and \mathbf{b}^c or \mathbf{H}^d and \mathbf{b}^d , depending on the type of constraint. Our derivation of constraint factors is made possible by the Lagrangian being a simple extension of the usual cost function. As a result, the primal update is similar to the usual update of ILS, which then triggers the update of the multipliers through the dual update.

The addition of constraints factors $|v| < v_{max}$ and $|\omega| < \omega_{max}$ to the problem formulation presented in Sec. V does not change the block diagonal structure of the problem.

VII. EXPERIMENTS

We tested our system both in simulation and on a real custom made unicycle robot⁴ equipped with a 2D lidar (InnoMaker-LD06), running on a Raspberry PI4 board at our Department, whose map is shown in Fig. 4. Comparative experiments have been conducted in simulation with a similar robot configuration, on a laptop Intel(R) Core(TM) i7-10750H CPU running at 2.60GHz with 16GB of RAM. We evaluate separately our factor graph-based localization (Sec. VII-A) and our MPC controller (Sec. VII-C). For a better understanding of the behavior of MPC optimization, we report the internals of an MPC iteration in Sec. VII-B. In all simulations, we set velocity limits to $|v| < 1$ m/s and $|\omega| < 1$ rad/s.

Our navigation stack relies on 2 components: a global planner and a local planner. Whenever a new goal is selected the global planner calculates the obstacle-free path by Dijkstra's algorithm. This path is then fed to the MPC controller that calculates the optimal control, and considers potential dynamic obstacles detected through the lidar. The robot estimates its position by using the factor graph-based localizer of Sec. IV. We compared our method with the most recent version of the ROS navigation stack consisting of ROS `amcl` package for localization and ROS `move_base` for planning and control. `amcl` implements the Adaptive Particle Filter localization [19]. As local planner for `move_base`, we have been using the ROS package implementation of the Timed Elastic Band method [21] because the base local planner was not always able to perform the assigned navigation tasks.




Fig. 4: Map of our Department.




Fig. 5: Map of factory-like environment.

In Sec. VII-D we compare our method with the well-known state-of-the-art general purpose nonlinear programming solver IPOPT [20], which is based on Interior-Point Method. IPOPT is reliable and robust in constraint handling, with cheaper iterations than those of SQP. Our solver is on average 75 times faster while reaching a similar value of the cost function.

A. Localization

We evaluated the performance of our optimization-based localization system and of `amcl` by measuring the Absolute Pose Error (APE) [23] between the estimated pose and the ground truth provided by ROS `stage` environment⁵. In our implementation of the proposed localization method:

- the cost of laser endpoints \mathbf{z}_k where $d(\mathbf{X}_R^W \mathbf{z}_k)$ is smaller than the resolution of the map are scaled down by 0.25;
- laser endpoints \mathbf{z}_k where $d(\mathbf{X}_R^W \mathbf{z}_k) \geq 0.3$ m are discarded in the optimization to disregard the effects of unexpected obstacles;
- laser endpoints \mathbf{z}_k further than 10.0 m from the robot are dropped to focus the localization effort on the local neighborhood and to reject the effect of potential distortions in the map.

Fig. 6 plots the evolution of the rotational and translational error over time of both approaches. With an average translational error of 0.087 ± 0.023 m and rotational error of 0.022 ± 0.027 rad our approach has similar performances to `amcl` which has an error of 0.090 ± 0.023 m, 0.014 ± 0.019 rad.

⁴<https://www.marrtino.org>

⁵<http://wiki.ros.org/stage>

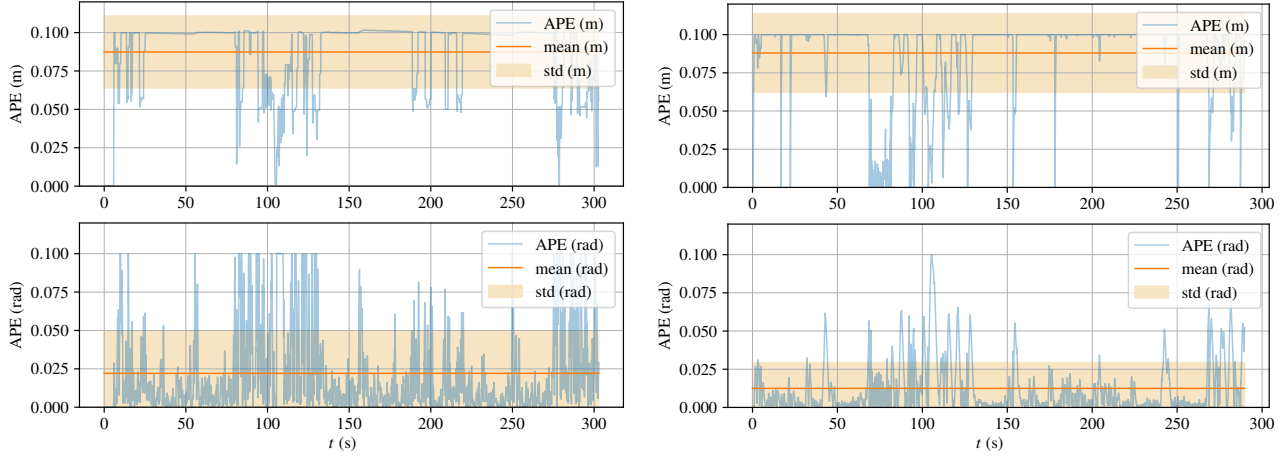


Fig. 6: APE with respect to linear and rotational part: on the left our factor graph based approach, on the right ROS `amcl`

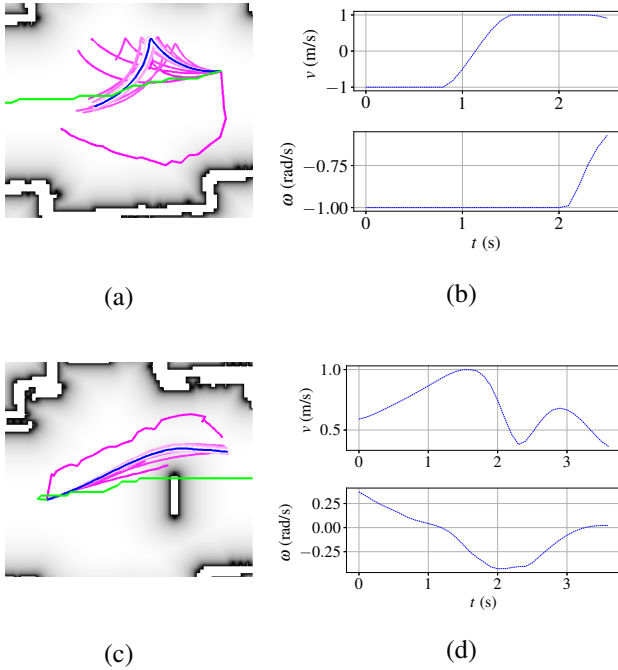


Fig. 7: Backward maneuver from right to left: (a) evolution of optimization from green curve to blue one passing through whitening pink; (b) corresponding control inputs; Obstacle avoidance maneuver from left to right: (c) evolution of optimization from green curve to blue one passing through whitening pink; (d) corresponding control inputs.

B. Maneuvers through MPC

Our approach is capable of automatically generating control inputs that allow the robot to perform complex maneuvers. In the reminder, we will focus on two situations where the trajectories generated by our MPC highly differ from the initial guess: backward motion and obstacle avoidance. Complex maneuvers result from satisfying the requirements expressed in the objective function while fulfilling the constraints.

Suppose that the robot is oriented so that the target is behind its back. In this case, many local planning algorithms switch to a re-orientation maneuver based on performing some kind of

backward motion until the angle between the robot orientation and the target is below a given threshold. Being based on optimization, our approach will automatically generate backward maneuvers for re-orientation, simply by minimizing the error between the desired and the actual path. Fig. 7(a) illustrates the evolution of the optimization of our MPC while solving for a time horizon. Thanks to the obstacle avoidance factor, the system has the knowledge of the map through the $g(\cdot)$ function of Sec. V, represented as the background. In our implementation of Eq. (9) we use $k = 0.075$, $\rho = 0.8$ m, and $\mu = 0.05$ m. The optimization starts from the initial guess illustrated by the green curve in Fig. 7(a) and ends at the blue curve. Intermediate steps are shown in pink, whitening as the iterations proceed. The temporal evolution of the control inputs at the optimum is shown in Fig. 7(b), where the cusp corresponds to the zero-velocity point and the limits are always met.

As we have mentioned in Sec. V, our optimization factor graph outputs a feasible trajectory by imposing that it minimizes the function $g(\cdot)$ over all poses. Fig. 7(c) highlights the effect of the obstacle avoidance factors, using the color scheme defined in the previous paragraph. When the initial guess coming from the static planner crosses an unforeseen obstacle, our MPC local planner deflects the initial guess to overcome the obstacle while reaching the goal. In correspondence to the inflection point of the trajectory, the angular velocity inverts its sign while the linear velocity reaches its maximum Fig. 7(d).

C. MPC based navigation

We report here the results obtained by running our navigation stack and the ROS one to travel across 8 goals 30 times within the map of our Department shown in Fig. 4. To compare the approaches we evaluated the path length and the duration to reach each of the goals. Our approach safely drove the robot during the whole 12 hours of simulation producing highly repeatable results that are summarized Tab. I. For each path we store the ground truth poses published by ROS `stage` and the duration in seconds. The path length is computed by summing the distance between subsequent ground truth poses.

(a) Path length: mean \pm standard deviation (m).

Path	Ours	ROS Navigation Stack
AC	175.42 \pm 0.68 (0.39%)	171.39 \pm 2.99 (1.74%)
CB	171.65 \pm 0.54 (0.31%)	169.37 \pm 1.09 (0.64%)
BD	169.52 \pm 0.81 (0.48%)	165.91 \pm 0.26 (0.16%)
DE	91.08 \pm 0.37 (0.41%)	89.66 \pm 0.19 (0.21%)
EA	93.12 \pm 0.33 (0.35%)	91.66 \pm 0.95 (1.03%)
AB	18.11 \pm 0.26 (1.43%)	18.07 \pm 0.21 (1.16%)
BC	172.76 \pm 0.80 (0.46%)	169.43 \pm 0.21 (0.12%)
CD	18.71 \pm 0.19 (1.01%)	19.07 \pm 0.33 (1.73%)

(b) Path duration: mean \pm standard deviation (s)

Path	Ours	ROS Navigation Stack
AC	197.663 \pm 1.63 (0.82%)	186.13 \pm 4.77 (2.56%)
CB	193.80 \pm 4.35 (2.24%)	190.78 \pm 13.21 (6.92%)
BD	189.84 \pm 1.07 (0.56%)	176.45 \pm 1.06 (0.60%)
DE	104.43 \pm 1.24 (1.19%)	97.32 \pm 0.87 (0.89%)
EA	108.74 \pm 1.37 (1.26%)	108.46 \pm 13.00 (11.99%)
AB	31.59 \pm 1.18 (3.74%)	25.73 \pm 0.78 (3.03%)
BC	194.11 \pm 3.57 (1.84%)	182.08 \pm 1.46 (0.80%)
CD	33.67 \pm 1.14 (3.39%)	27.31 \pm 0.47 (1.72%)

TABLE I: Statistical analysis of path length and path duration on DIAG map.

(a) Path length: mean \pm standard deviation (m)

Path	Ours	ROS Navigation Stack
AB	154.30 \pm 1.03 (0.67%)	148.47 \pm 0.27 (0.18%)
BC	86.52 \pm 0.31 (0.35%)	83.99 \pm 0.18 (0.21%)
CD	142.06 \pm 0.76 (0.53%)	136.97 \pm 0.26 (0.19%)
DE	130.88 \pm 1.31 (1.00%)	124.99 \pm 0.19 (0.15%)

(b) Path duration: mean \pm standard deviation (s)

Path	Ours	ROS Navigation Stack
AB	182.06 \pm 1.63 (0.89%)	197.25 \pm 11.78 (5.97%)
BC	158.96 \pm 1.73 (0.11%)	172.09 \pm 9.16 (5.32%)
CD	91.22 \pm 0.65 (0.71%)	97.99 \pm 5.59 (5.70%)
DE	147.81 \pm 3.00 (2.03%)	161.27 \pm 10.99 (6.81%)

TABLE II: Statistical analysis of path length and path duration on a factory-like map.

Using our MPC controller, the standard deviation of path length and duration never exceed 1.43% and 3.74% respectively. Corresponding values obtained using the ROS navigation system are 1.74% and 11.99%. The mean value of the percentage standard deviation on path duration is 1.88% using our system and 3.56% using ROS, the same value on path length is 0.60% using our system and 0.85% using ROS. On long paths, where narrow passages account only for a small portion of the whole path, ROS duration is on average 95% of ours, while on paths mostly consisting of narrow passages ROS duration is 81% of ours. We can infer that the difference in path duration is due to our approach begin more conservative in narrow passages.

To verify this and further validate our approach, we have tested the two navigation stacks also on the map of the factory-like environment in Fig. 5. We report in Tab. II the results obtained by running our navigation stack to reach 5 goals for 30 times. In this case, our approach performs better than the ROS controller in terms of path duration, still satisfying the control limits. As for the DIAG map, path duration results are more repeatable, showing a smaller standard deviation. The percentage standard deviation on path length is unchanged compared to the DIAG map results.

Solver	Initial Cost	Final Cost	Iterations	Time[s]
Ours		53.612	118	0.011
IPOPT	7938.200	50.423	445	1.459
Ours		68.682	125	0.013
IPOPT	3367.190	63.744	809	2.497
Ours		64.588	158	0.013
IPOPT	6199.650	60.097	931	2.943
Ours		447.180	189	0.031
IPOPT	22011.100	443.779	479	1.872
Ours		568.386	366	0.055
IPOPT	27744.090	563.406	549	1.951
Ours		484.154	294	0.043
IPOPT	21341.260	480.467	439	1.682

TABLE III: Comparison with IPOPT: convergence analysis.




Fig. 8: Comparison with IPOPT: control inputs corresponding to the backward maneuver of last row of Tab. III.

D. Comparison with IPOPT

We evaluated the performance of our solver compared to the well-known solver IPOPT. In particular, we fed an instance of our optimal control problem to the Interface For nonlinear Optimizers (IPOPT) [27], having set IPOPT as a nonlinear programming solver. Four different goals were considered starting from a given initial pose, resulting in both forward and backward trajectories. Tab. III shows the results obtained, while Fig. 8 plots control inputs optimized with the two methods corresponding to the last row of Tab. III. As a termination criterion for our solver, we consider the norm of the perturbation vector and that of the constraint violations. Optimization is stopped when: $\|\Delta \mathbf{x}\|_2 < \epsilon_x$, $\|\mathbf{C}\mathbf{x} + \mathbf{c}\|_\infty < \epsilon_c$, $\|\max(\mathbf{D}\mathbf{x} + \mathbf{d}, 0)\|_\infty < \epsilon_d$, where $\mathbf{0}$ is the vector of all zeros, and $\epsilon_x, \epsilon_c, \epsilon_d = 1e-4$.

Our solver is around 2 orders of magnitude faster than IPOPT at converging. This makes our solver preferable for online onboard applications. Two aspects make this possible. The first is the Method of Lagrange Multipliers, whose iterations are more compact than those of Interior-Point. The second is that factor graph optimization is highly efficient for sparse optimization problems, such as the one we are solving. The final cost of our solution is slightly higher, however both a quantitative comparison with the initial value reported in Tab. III and a graphical comparison based on Fig. 8 show that the two solutions are very similar. In summary, our evaluation suggests that the proposed framework is considerably more efficient than standard non-linear programming solvers. At the same time, the proposed framework allows combining estimation and control problems under a unified methodology.

VIII. CONCLUSIONS

We presented an approach to embedding constraints on variables as factors of a graph, based on the Method of Lagrange Multipliers. Having introduced an Iterative Least Squares solver for constrained optimization, we have contributed to extending the variety of problems across robotics that can be modeled using factor graphs.

In this paper, we focus on optimal control for MPC as an application field of constrained optimization to support the validity of our approach. In particular, we develop a full navigation stack based on factor graphs: both localization and MPC are modeled here through factors. Experiments show that our navigation stack compares favorably to the ROS navigation stack. Moreover, experiments conducted with IPOPT confirm that our solver gives the advantage of both an improvement in computational time and a unified formulation for estimation and control problems.

While our navigation stack is a working proof of concept, our method is general and can be easily used to model other classes of problems. As future work in the field of optimal control, different kinematic and dynamic models can be considered. Active SLAM [6], [7] is an example application where both estimation and control are coupled. The two problems share the knowledge of the map and the robot position estimate and contribute to the goal of having a robot that autonomously builds an accurate map in the shortest possible time. In the near future we envision future applications of our findings in this domain.

REFERENCES

- [1] Alberto Bemporad, Manfred Morari, Vivek Dua, and Efstratios N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, 2002.
- [2] Dimitri P Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.
- [3] Farid Bounini, Denis Gingras, Herve Pollart, and Dominique Gruyer. Bounded artificial potential field method for online path planning applications. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 180–185, 2017.
- [4] Stephen P. Boyd, Neal Parikh, Eric King wah Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3:1–122, 2011.
- [5] J.C. Butcher. A history of runge-kutta methods. *Applied Numerical Mathematics*, 20(3):247–260, 1996.
- [6] H. Carrillo and J.A. Castellanos. *Navigation Under Uncertainty Based on Active SLAM Concepts*, pages 209–235. Springer International Publishing, 2015.
- [7] Yongbo Chen, Shoudong Huang, and Robert Fitch. Active SLAM for mobile robots with area coverage and obstacle avoidance. *IEEE/ASME Transactions on Mechatronics*, 25(3):1182–1192, 2020.
- [8] M.D. Cicco, B.D. Corte, and G. Grisetti. Unsupervised Calibration of Wheeled Mobile Platforms. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2016.
- [9] Alessandro De Luca and Giuseppe Oriolo. Local incremental planning for nonholonomic mobile robots. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pages 104–110 vol.1, 1994.
- [10] Frank Dellaert. Factor graphs: Exploiting structure in robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):141–166, 2021.
- [11] Jing Dong, Mustafa Mukadam, Byron Boots, and Frank Dellaert. Sparse gaussian processes on matrix lie groups: A unified framework for optimizing continuous-time trajectories. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6497–6504, 2018.
- [12] Jing Dong, Mustafa Mukadam, Frank Dellaert, and Byron Boots. Motion planning as probabilistic inference using gaussian processes and factor graphs. In *Robotics: Science and Systems XII*, 2016.
- [13] Giorgio Grisetti, Tiziano Guadagnino, Irvin Aloise, Mirco Colosi, Bartolomeo Della Corte, and Dominik Schlegel. Least squares optimization: from theory to practice. *Robotics*, 9(3):51, July 2020.
- [14] Giorgio Grisetti, Rainer Kümmerle, and Kai Ni. Robust optimization of factor graphs by using condensed measurements. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 581–588, 2012.
- [15] Viorela Ila, Lukas Polok, Marek Solony, and Pavel Svoboda. SLAM++: a highly efficient and temporally scalable incremental slam framework. *The International Journal of Robotics Research*, 36(2):210–230, 2017.
- [16] Michael Kaess, Ananth Ranganathan, and Frank Dellaert. iSAM: Incremental smoothing and mapping. *IEEE Transactions on Robotics*, 24(6):1365–1378, 2008.
- [17] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. G2o: A general framework for graph optimization. In *2011 IEEE International Conference on Robotics and Automation*, pages 3607–3613, 2011.
- [18] Mustafa Mukadam, Jing Dong, Frank Dellaert, and Byron Boots. STEAP: simultaneous trajectory estimation and planning. *Autonomous Robots*, 43:415–434, 2019.
- [19] Patrick Pfaff, Wolfram Burgard, and Dieter Fox. Robust monte-carlo localization using adaptive likelihood models. In Henrik I. Christensen, editor, *European Robotics Symposium*, volume 22 of *Springer Tracts in Advanced Robotics*, pages 181–194. Springer, 2006.
- [20] Florian A. Potra and Stephen J. Wright. Interior-point methods. *Journal of Computational and Applied Mathematics*, 124(1):281–302, 2000.
- [21] Christoph Rösmann, Wendelin Feiten, Thomas Wösch, Frank Hoffmann, and Torsten Bertram. Efficient trajectory optimization using a sparse model. In *2013 European Conference on Mobile Robots*, pages 138–143, 2013.
- [22] Johannes L. Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [23] David Schubert, Thore Goll, Nikolaus Demmel, Vladyslav Usenko, Jorg Stuckler, and Daniel Cremers. The TUM vi benchmark for evaluating visual-inertial odometry. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2018.
- [24] Joe Sfeir, Maarouf Saad, and Hamadou Saliah-Hassane. An improved artificial potential field approach to real-time mobile robot path planning in an unknown environment. In *2011 IEEE International Symposium on Robotic and Sensors Environments (ROSE)*, pages 208–213, 2011.
- [25] Paloma Sodhi, Sanjiban Choudhury, Joshua G. Mangelson, and Michael Kaess. ICS: Incremental constrained smoothing for state estimation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 279–285, 2020.
- [26] Duy-Nguyen Ta, Marin Kobilarov, and Frank Dellaert. A factor graph approach to estimation and model predictive control on unmanned aerial vehicles. In *2014 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 181–188, 2014.
- [27] Alexander W Winkler. Ifopt - A modern, light-weight, Eigen-based C++ interface to Nonlinear Programming solvers Ipopt and Snopt., 2018.
- [28] Mandy Xie, Alejandro Escontrela, and Frank Dellaert. A factor-graph approach for optimization problems with dynamics constraints. *CoRR*, abs/2011.06194, 2020.
- [29] Shuo Yang, Gerry Chen, Yetong Zhang, Howie Choset, and Frank Dellaert. Equality constrained linear optimal control with factor graphs. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9717–9723, 2021.