The unconstrained and inequality constrained moving horizon approach to robot localization

Gianluigi Pillonetto, Aleksander Aravkin, and Stefano Carpin

***Abstract*— We present a moving horizon approach for estimating the state of a nonlinear dynamic system possibly subject to inequality constraints. The method takes advantage of a recent algorithm proposed in the literature based on interior point methods. The approach exploits the same decomposition used for unconstrained Kalman-Bucy smoothers. Hence, the number of operations required by the algorithm scales linearly with the length of the horizon, making possible its use for online applications. We apply this method to the robot localization problem, showing that it is able to produce much more accurate results than the iterated Kalman filter with few additional computational effort.**

I. Introduction

Kalman filters (KF) are widely used to estimate states of dynamic stochastic systems in many different fields such as biomedicine, economy and robotics (see [13], [22] for detailed expositions of its properties). When the system under study is nonlinear, the simplest implementation is the extended or iterated Kalman filter (IKF) [9]. However, the estimates obtained by IKF may be quite distant from the minimum variance ones. In addition, this filter may be much sensitive to unknown initial conditions and local minima. Particle filters (PF) are an important alternative where optimization is replaced by propagation of a posterior density in sampled form and Monte Carlo integration [21], [24], [25], [11]. A problem is that these techniques call for delicate tuning of proposal densities to improve their convergence rates. In addition, robust statistical convergence criteria are still missing. We also notice that in many applications additional knowledge on the system state can be available, e.g. in the form of inequality constraints. Including this information may be important to improve the estimation process. However, while linear and nonlinear equality constraints on the state vector can be easily handled e.g. by augmenting the measurement model, see e.g. [7], [10], [20], [26], imposing affine or nonlinear inequality constraints is more difficult. This can further complicate the implementation of IKF and PF [4], [17], [20].

To deal with the above issues, in this paper we propose a new moving horizon approach [6], [16], [17], i.e. a filter where a moving window of previous measurements is processed. The approach can also efficiently handle inequality constraints on the state. In particular, we will use a recently proposed optimization algorithm that relies upon interior point methods [3]. The key feature of this method is that it takes advantage of the same decomposition used for unconstrained Kalman smoothers [2], [18]. In this way, the required operations scale linearly with the horizon length, making it feasible to simultaneously optimize with respect to the state value at all time indices in the moving window. This leads to a filter which is much less sensitive to unknown initial conditions and local minima.

In order to demonstrate the effectiveness of our technique, the estimator is applied to a robot localization problem. Due to its practical importance, this is one of the most widely studied topics in mobile robotics. Our focus is on mobile robots moving in planar environments where the goal is to estimate the position in the plane and the *yaw* orientation of the robot. More precisely, in the simulations presented in this paper the robot moves in an environment conditioned with landmarks located at known locations, and is equipped with a sensor identifying the landmarks and estimating their distance. The goal is to estimate the triple (x, y, θ), indicated as *pose* in the following. In this setting the inequality constraints can represent the known region where the robot is moving. The reader is referred to the book by Thrun et al. [23] for a thorough description of the most commonly used probabilistic techniques used to solve this problem. Most of the published literature in the field resorts to IKF for position tracking (where the initial pose of the robot is known) [12] and PF for global localization problems [11], [14], [21], [24], [25]. In this latter problem, the initial pose is unknown and it is crucial to track non Gaussian distributions. Remarkably, the moving horizon approach proposed in this paper will prove to be robust also for solving this kind of estimation tasks, with only a moderate additional computational effort in comparison to IKF.

The paper is organized as follows. Section II formally defines the problem and presents the theory supporting the presented moving horizon approach. Algorithmic details are presented in section III. The experimental setup and results are presented in Section IV. Finally, conclusions are offered in Section V.

II. Statement of the problem

A. *The moving horizon version of the inequality constrained Kalman smoothing problem*

We are given the noisy output measurements $z_k \in \mathbf{R}^m$ coming from the following dynamic system

$$x_k = g_k(x_{k-1}) + w_k \quad , \quad z_k = h_k(x_k) + v_k \tag{1}$$

G. Pillonetto is with the Department of Information Engineering, University of Padova, Padova Italy. A. Aravkin is with the Department of Mathematics, University of Washington, Seattle U.S.A. S. Carpin is with the School of Engineering, University of California, Merced U.S.A.

where $x_k \in \mathbf{R}^n$. The noise vectors $w_k \in \mathbf{R}^n$ and $v_k \in \mathbf{R}^m$ are all mutually independent and

$$w_k \sim \mathbf{N}(0, Q_k) \quad , \quad v_k \sim \mathbf{N}(0, R_k) \tag{2}$$

where $Q_k \in \mathbf{R}^{n\times n}$ and $R_k \in \mathbf{R}^{m\times m}$ are known autocovariance matrices. We also indicate with $f_k : \mathbf{R}^n \to \mathbf{R}^\ell$ those known functions that model the constraints given by

$$f_k(x_k) \leq 0 \quad \text{for} \quad k = 1, \ldots, N \tag{3}$$

Let

$$\begin{aligned} S_k(x_k, x_{k-1}) = &\frac{1}{2}[z_k - h_k(x_k)]^{\mathrm{T}} R_k^{-1} [z_k - h_k(x_k)] \\ &+ \frac{1}{2}[x_k - g_k(x_{k-1})]^{\mathrm{T}} Q_k^{-1} [x_k - g_k(x_{k-1})] \end{aligned}$$

denote the residual sum of squares at time index k. In the spirit of moving horizon methods, let K be the moving index and M the length of the horizon. In addition, define

$$\begin{aligned} &x^K = \{x_q, \ldots, x_K\} \qquad z^K = \{z_q, \ldots, z_K\} \\ &q = \max\{1, K - M + 1\} \end{aligned}$$

Letting

$$S^K(x^K) = \sum_{k=q}^{K} S_k(x_k, x_{k-1})$$

the corresponding moving horizon problem is

$$\begin{array}{lll} \text{minimize} & S^K(x^K) & \text{w.r.t.} \quad x^K \\ & \text{s.t.} \quad f_k(x_k) \leq 0, & k = q, q+1, \ldots, K. \end{array} \tag{4}$$

In the above equations, x_{q-1} is treated as a known parameter and the initial conditions are given by the following special inizializations

$$g_q(x_{q-1}) = x_{q-1}, \qquad w_q \sim \mathbf{N}(0, V_q)$$

that model the state vector x_q as Gaussian with mean x_{q-1} and autocovariance V_q. Notice that no matter how large N is, problem (4) is bounded in size by M.

The following result is easily obtained.

Theorem 1: Assume that $x_q \sim \mathbf{N}(x_{q-1}, V_q)$ where x_{q-1} and V_q are known quantities. Then, the maximum a posteriori estimate of the process x^K, conditional on z^K and the constraints (3) acting at instants $k = q, q+1, \ldots, K$, is the solution of the problem 4.

In the scenario of robot localization, the constraints specified in (3) may e.g. represent the region where the robot is moving. Notice that another interpretation of our model is that the event $f(x_k) \leq 0$ accounts for the physical constraints of the environment by making the distribution of x_k conditional on x_{k-1} a truncated Gaussian.

III. The Numerical Algorithm

A. The Quadratic Programming Sub-problem

The algorithm presented in [3] is now briefly recalled and easily adapted to our framework. In particular, we will show that for any moving horizon index K, obtaining the estimate of x^K requires only $O(Mn^3)$ operations. In particular, each moving horizon problem is solved by introducing affine approximations that are first-order accurate for a state sequence $y^K = \{y_q, \ldots, y_K\}$ near a fixed state sequence x^K. Define the affine approximations $\tilde{f}_k$, $\tilde{g}_k$, and $\tilde{h}_k$ by

$$\begin{aligned} \tilde{f}_k(x_k; y_k) &= f_k(x_k) + f_k^{(1)}(x_k)(y_k - x_k) \\ \tilde{g}_k(x_k; y_k) &= g_k(x_k) + g_k^{(1)}(x_k)(y_k - x_k) \\ \tilde{h}_k(x_k; y_k) &= h_k(x_k) + h_k^{(1)}(x_k)(y_k - x_k) \end{aligned}$$

Then, the residual sum of squares function associated with the K-th moving horizon problem and with the above affine approximations is denoted by

$$\tilde{S}^K(x^K; y^K) = \sum_{k=q}^{K} \tilde{S}_k(x_k, x_{k-1}; y_k, y_{k-1}) \tag{5}$$

where

$$\begin{aligned} &\tilde{S}_k(x_k, x_{k-1}; y_k, y_{k-1}) = \\ &\quad (1/2)[y_k - \tilde{g}_k(x_{k-1}; y_{k-1})]^{\mathrm{T}} Q_k^{-1} [y_k - \tilde{g}_k(x_{k-1}; y_{k-1})] \\ &\quad +(1/2)[z_k - \tilde{h}_k(x_k; y_k)]^{\mathrm{T}} R_k^{-1} [z_k - \tilde{h}_k(x_k; y_k)] \end{aligned}$$

Following [3], the nonlinear problem (4) is solved by solving quadratic programming (QP) subproblems given by:

$$\begin{array}{lll} \text{minimize} & \tilde{S}^K(x^K; y^K) \quad \text{w.r.t.} \quad y^K & \\ \text{subject to} & \tilde{f}_k(x_k; y_k) \leq 0, & k = q, q+1, \ldots, K \end{array} \tag{6}$$

Define $A_k \in \mathbf{R}^{n\times n}$ and $C_k \in \mathbf{R}^{n\times n}$ by

$$\begin{aligned} A_k &= -M_k^{-1} g_k^{(1)}(x_{k-1}) \\ C_k &= \left\{ \begin{array}{c} M_k^{-1} + h_k^{(1)}(x_k)^{\mathrm{T}} R_k^{-1} h_k^{(1)}(x_k) \\ +g_{k+1}^{(1)}(x_k)^{\mathrm{T}} M_{k+1}^{-1} g_{k+1}^{(1)}(x_k) \end{array} \right\} \\ M_k &= \left\{ \begin{array}{ll} V_k & \text{if} \quad k = q \\ Q_k & \text{otherwise} \end{array} \right. \end{aligned} \tag{7}$$

Then, the matrix $C^K \in \mathbf{R}^{nM\times nM}$ given by

$$C^K = \begin{pmatrix} C_q & A_{q+1}^{\mathrm{T}} & 0 & \\ A_{q+1} & C_{q+1} & A_{q+2}^{\mathrm{T}} & 0 \\ 0 & \ddots & \ddots & \ddots \\ & 0 & A_K & C_K \end{pmatrix} \tag{8}$$

is just the Hessian of the objective $\tilde{S}^K$ in (6) with respect to y^K, see also [3] for details. Now, define the vector $a_k \in \mathbf{R}^n$ by

$$\begin{aligned} a_k \;=\; & M_k^{-1}[x_k - g_k(x_{k-1})] - h_k^{(1)}(x_k)^{\mathrm{T}} R_k^{-1} [z_k - h_k(x_k)] \\ & -g_{k+1}^{(1)}(x_k)^{\mathrm{T}} M_{k+1}^{-1} [x_{k+1} - g_{k+1}(x_k)] \end{aligned}$$

Let $a^K \in \mathbf{R}^{nM}$ be the column vector representing $\{a_q, \ldots, a_K\}$. Then, notice that a^K represents the gradient of $\tilde{S}^K(x^K, y^K)$ with respect to y^K at $y_k = x_k$, $k = q, \ldots, K$. Now, let

$$B^K = \begin{pmatrix} f_q^{(1)}(x_q) & 0 & \\ 0 & \ddots & 0 \\ & 0 & f_K^{(1)}(x_K) \end{pmatrix} \tag{9}$$

and notice that the affine approximation to the constraints (3) for the K-th moving horizon problem is given by $b^K + B^K y^K \leq 0$. Thus, the QP subproblem (6) becomes

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}(y^K)^{\mathrm{T}} C^K y^K + (d^K)^{\mathrm{T}} y^K \quad \text{w.r.t. } y^K \in \mathbf{R}^{nM} \\ \text{subject to} & b^K + B^K y^K \leq 0 \end{array} \tag{10}$$

where $d^K = a^K - C^K x^K$. The QP subproblem (10) can now be solved using the interior point approach presented in [3] with $O(Mn^3)$ operations. Here, we just recall that the interior point approaches apply a damped Newton's method to a relaxation of the Karush-Kuhn-Tucker (KKT) conditions. In our case, the relaxed subproblem (that contains a log barrier) is given by:

$$\begin{array}{c} \text{minimize } (1/2)(y^K)^{\mathrm{T}} C^K y^K + (d^K)^{\mathrm{T}} y^K \\ -\mu \sum_{i=1}^{\ell M} \log(s_i) \text{ w.r.t} \\ (y^K, s) \in \mathbf{R}^{nM} \times \mathbf{R}_+^{\ell M} \quad \text{s.t.} \quad s + b^K + B^K y^K = 0 \end{array} \tag{11}$$

where μ is the relaxation parameter and s is the vector containing the slack variables.

B. The Nonlinear Algorithm for solving the moving horizon problems

The termination criteria for the K-th nonlinear moving horizon problem where a sequence of M states have to be estimated rely upon the KKT conditions for problem (4). Given the current value for K, we use p to denote the iteration counter of our optimization scheme. Let $x^K(p) \in \mathbf{R}^{nM}$ where $x^K(p)$ contains the blocks $x_k^K(p) \in \mathbf{R}^n$, $k = q, \ldots, K$ while $u^K(p) \in \mathbf{R}_+^{\ell M}$ is the Lagrange multiplier vector containing the blocks $u_k^K(p) \in \mathbf{R}^{\ell}$, $k = q, \ldots, K$. Then, for each K, the algorithm terminates at a primal vector $x^K(p)$ and Lagrange multiplier vector $u^K(p) \in \mathbf{R}_+^{\ell M}$ such that for $k = q, \ldots, K$

$$\begin{array}{l} f_k(x_k^K(p)) \leq \varepsilon \;, \; \|u_k^K(p) \cdot f_k(x_k^K(p))\|_\infty \leq \varepsilon \text{ and} \\ \|(u_k^K(p))^{\mathrm{T}} f_k^{(1)}(x_k^K(p)) + \partial_{x_k^K(p)} S^K(x^K(p))\|_\infty \leq \varepsilon \end{array} \tag{12}$$

where ε is a termination tolerance. Given a vector $w \in \mathbf{R}^m$, $\max(0, w) \in \mathbf{R}^m$ denotes the vector with i-th component equal to $\max(0, w_i)$. Given a $x^K \in \mathbf{R}^{nM}$, the $\boldsymbol{\ell}_1$ distance from the constraint function values $\{f_k(x_k^K)\}_{k=q}^K$ to the constraint set is

$$\phi(x^K) = \sum_{k=q}^{K} \sum_{i=1}^{\ell} \max([f_k(x_k^K)]_i, 0)$$

with its approximation given by

$$\tilde{\phi}(x^K; y^K) = \sum_{k=q}^{K} \sum_{i=1}^{\ell} \max([\tilde{f}_k(x_k^K; y_k^K)]_i, 0)$$

We are now in a position to describe the algorithm that solves in an online-manner the sequence of nonlinear moving horizons problems (4). Needless to say, the unconstrained version becomes just a special case of the scheme described below.

Algorithm 2: Moving Horizon Version of the Inequality Constrained Nonlinear Smoother

1) Initialization: Set $K = 0$
2) Set $K = K + 1$, $q = \max\{1, K - M + 1\}$, x_{q-1} to its estimate and V_q to a numerical approximation of its posterior autocovariance[1]. In addition, set the iteration counter $p = 0$ and the initial penalty parameter $\alpha_0 = 0$.
3) Affine approximation: Substitute $x^K(p)$ for x^K in equations (8), (9) and let $a(p)$, $b(p)$, $B(p)$, $C(p)$, and $d(p)$ be the corresponding values for a^K, b^K, B^K, C^K, and d^K in QP (10).
4) Solve this QP using Algorithm 4 described in [3] with inputs $\delta = \varepsilon \times 10^{-2}$ and let $y(p)$ and $u(p)$ be the resulting solution.
5) If the convergence criteria (12) are satisfied, return $x(p)$, $u(p)$ as the solution for the K-th moving horizon problem and go to Step 2.
6) If $\alpha_p > 0$, set $\hat{\alpha}_p = \alpha_p$; otherwise, $\hat{\alpha}_p = \|u(p)\|_\infty$. Define the value
$$\begin{array}{rcl} \zeta_p & = & (y(p) - x(p))^{\mathrm{T}} C(p)(y(p) - x(p)) \\ & + & (a(p))^{\mathrm{T}}(y(p) - x(p)) \end{array}$$
If $\zeta_p \leq \hat{\alpha}_p \phi(x(p))$, set $\alpha_{p+1} = \hat{\alpha}_p$; otherwise, $\alpha_{p+1} = \max[\zeta_p / \phi(x(p)) \;,\; 2\hat{\alpha}_p]$.
7) Compute the line search step size λ_p as follows:
$$\begin{array}{l} \eta_p = (a(p))^{\mathrm{T}}(y(p) - x(p)) \\ \qquad + \alpha_{p+1}[\tilde{\phi}(x(p); y(p)) - \phi(x(p))] \\ H_p(\lambda) = S^K[x(p) + \lambda(y(p) - x(p))] \\ \qquad + \alpha_{p+1}\phi[x(p) + \lambda(y(p) - x(p))] \\ \lambda_p = \max\{\; 2^{-r} \mid r \in \mathbf{Z}_+ \text{ and} \\ \qquad H_p(2^{-r}) - H_p(0) \leq 2^{-r}\eta_p/10 \;\} \end{array}$$
8) Set $x^{p+1} = x(p) + \lambda_p(y(p) - x(p))$, then set $p = p+1$ and go to step 3.

The following convergence results is taken from [3], see also [5].

Theorem 3: Suppose $\varepsilon = 0$, for each K all the quadratic subproblems in step 3 have feasible solutions, the corresponding sequence $\{y(p)\}$ is bounded, and every cluster point of $\{x(p)\}$ satisfies the Mangasarian-Fromowitz Constraint Qualification. Then the sequence $\{x(p)\}$ is bounded and each of its cluster points is a KKT point for problem (4), i.e. satisfies convergence criteria (12) for some vector of Lagrange multipliers.

IV. Experimental results

In this section we offer an extensive set of simulation results aimed to outline the properties of the proposed algorithm (MH – moving horizon), and we also contrast it with the iterative Kalman filter. The chosen benchmark problem is a classical localization problem in a given map with landmarks located at known positions. All the code and data needed to replicate the results presented in this section are available for download at `http://robotics.ucmerced.edu`.

[1] In the numerical implementation, this autocovariance has been obtained by the last affine approximation providing the solution of the nonlinear problem (4) without constraints.

A. System model

We consider a differential drive robot moving on a flat terrain populated with m landmarks placed at known locations. Let the location of the i-th landmark be (p_x^i, p_y^i). As usual, the pose of the robot is indicated as $(x, y, \vartheta) \in \mathbf{R}^3$, and we hypothesize two inputs control the system, namely the translational speed u_v and the rotational speed u_w. Consistently with the formerly depicted framework, we model the system with discrete time equations describing how the state evolves over time. When u_w is different from 0, the following relationships[2] hold (see [23], chapter 5):

$$
\begin{aligned}
x_k &= x_{k-1} - \frac{u_v}{u_w}\sin\theta_{k-1} + \frac{u_v}{u_w}\sin(\theta_{k-1} + u_w) \\
y_k &= y_{k-1} + \frac{u_v}{u_w}\cos\theta_{k-1} - \frac{u_v}{u_w}\cos(\theta_{k-1} + u_w) \\
\theta_k &= \theta_{k-1} + u_w
\end{aligned}
$$

When u_w is equal to 0 the robot simply moves forward, so these relationships simplify in a straightforward way with θ_k remaining equal to θ_{k-1}, and x_k, y_k changing according to the heading. Following the hypotheses presented while introducing the problem, we assume that state evolution is affected by Gaussian noise $w_k \sim \mathbf{N}(0, Q_k)$. Q_k is a known 3×3 diagonal covariance matrix whose values on the diagonal are not all necessarily equal. It is furthermore assumed that Q_k is the same for each k, though this is not necessary. The robot is supposed to be equipped with a single sensor returning the distance from the landmarks, provided that they are closer than a known constant threshold T. That is, at step k the sensor returns a vector $z_k \in \mathbf{R}^m$ where the i-th entry is either

$$
d_k^i = \sqrt{(x_k - p_x^i)^2 + (y_k - p_y^i)^2}
$$

if $d_k^i < T$, or 0 otherwise. In all examples presented in this section, $T = 7m$. Entries larger than 0, i.e. entries corresponding to actual readings, are corrupted by Gaussian noise $v_k \sim \mathbf{N}(0, R_k)$. R_k is assumed to be an $m \times m$ diagonal matrix with identical values on the main diagonal, although it is possible to consider situations where the values are not the same. This extension will not be considered on this paper, but the price to pay is just a slightly more articulated implementation. The reader should note that the sensor does not return any information about the heading of the robot, whereas the sensor values are exclusively dependent from the x, y components. This aspect will be important to consider while evaluating the impact of the proposed technique, and also when comparing it with the IKF. Also, we implicitly assumed that the landmark correspondence problem does not occur, i.e. at every time step the sensor knows the identity of the landmarks being seen. This hypothesis is realistic if the landmarks are properly designed (see e.g. [8]), and when this is not the case the problem can be addressed using a maximum likelihood approach, as evidenced in [23].

[2] In this section x_k indicates the x component of the pose at time k, while in Section II x_k indicated the whole state. This slight abuse of notation is accepted to give an immediate physical meaning to the individual components of the state.

In this section we will consider the three classical problems contemplated in localization literature, i.e.

- tracking: the robot starts from a known location;
- global localization: the robot starts from an unknown location;
- kidnapped robot: the localization algorithm starts with a strong confidence about the position of the robot, but such position is wrong.

Global localization and the kidnapped robot are considered even if they are problematic for the IKF estimator in order to outline that the proposed method can handle those as well. Moreover, we considered different scenarios with varying degrees of noise.

B. Sensitivity to horizon length

The first experiment aims to verify the impact of the moving horizon length on the localization error. Figure 1 displays the root mean square error (RMSE) error obtained in 100 runs with the moving horizon length varying from 2 to 10. The three subplots display the errors for the x, y, and ϑ components of the state. The specific problem was the kidnapped robot problem.

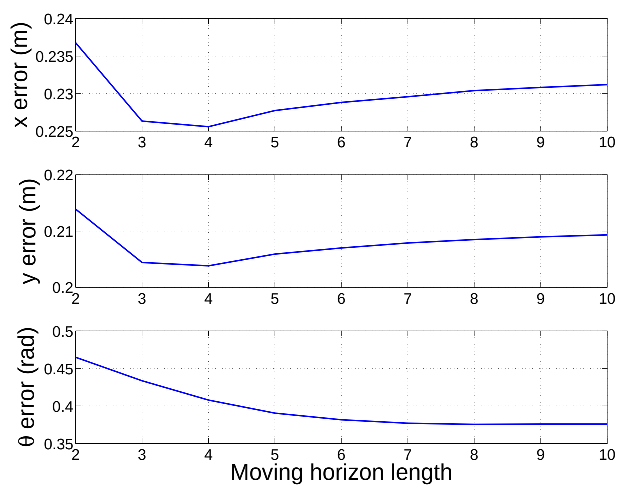

Fig. 1. Error trend for the x, y, ϑ components for different values of the moving horizon length.

Similar trends were obtained when considering the tracking and global localization problems. The chart shows that only modest improvements are obtained by increasing the size of the horizon. Consequently, in order to keep processing times at the minimum, in all subsequent tests presented in this section we used an horizon of length 2.

C. Comparison with the IKF for the unconstrained case

In this subsection we contrast the performance of the proposed estimator with an iterated Kalman filter performing 10 iterations at each estimation step. Both estimators are

		IKF			MH		
		x	y	ϑ	x	y	ϑ
Tracking	Case 1	0.0874	0.0865	0.1331	0.0552	0.0554	0.1394
	Case 2	0.1013	0.1023	0.1393	0.0680	0.0688	0.1443
	Case 3	0.1385	0.1384	0.3998	0.0771	0.0783	0.4012
	Case 4	0.1515	0.1516	0.4060	0.0998	0.1021	0.4062
Global	Case 1	0.2399	0.1171	0.2901	0.0579	0.0573	0.2208
	Case 2	0.2542	0.1517	0.3450	0.0717	0.0725	0.2445
	Case 3	0.2538	0.1544	0.4881	0.0774	0.0786	0.4572
	Case 4	0.2639	0.1667	0.5010	0.1005	0.1029	0.4704
Kidnapped	Case 1	0.3667	0.1654	0.4832	0.2321	0.2090	0.4624
	Case 2	0.4797	0.2094	0.5090	0.2677	0.2471	0.4963
	Case 3	0.2613	0.1563	0.5509	0.1601	0.1503	0.6014
	Case 4	0.2836	0.1728	0.6389	0.1944	0.1817	0.6077

TABLE I

COMPARISON BETWEEN THE PERFORMANCE OF IKF AND THE PROPOSED ESTIMATOR (MH — MOVING HORIZON). EACH ROW DISPLAYS THE RMSE AVERAGED OVER 100 INDEPENDENT TRIALS. ERRORS FOR ϑ ARE EXPRESSED IN RADIANS.

fed with exactly the same data, and a priori information about the matrices Q_k and R_k. For each of the three estimation problems considered, i.e. tracking, global localization and kidnapped robot problem, we consider four different situations characterized by different matrices Q_k and R_k (see the four cases in Table I for every scenario). Table I shows the overall results. Again, we display RMSE obtained averaging 100 runs. For cases where the estimator starts with an erroneous value for x_0, the same wrong value is used for both estimators. In order to get an idea about the relative importance of the plotted numbers the reader is referred to figure 2 for the dimensions about the environment. It is immediate to observe that for the x and y components of the state MH largely outperforms IKF in every instance of the tracking and global localization problems. The situation is less drastic in the the kidnapped case, where IKF prevails in 3 out of 8 cases. However, one should keep in mind that MH has been run with the shortest possible horizon length, i.e. 2. By increasing this parameter a more accurate estimate is expected, whereas for the problems at hand we observed that with 10 iterations the IKF seems to converge to a point where adding more iterations would not help. For the ϑ component of the state the situation is less crisp. That is to say that there is no clear winner. This is somehow expected, since the chosen sensor bears no information about the orientation of the robot. This problem equally affects both estimators. In such scenario IKF mostly relies on prediction, while MH may get stuck in local minima during its search for the minimum. In the next stage of this research we will experiment with a different sensor model returning not only the distance from the landmarks, but also the relative orientation, i.e. a range-bearing sensor. This information would then be used to infer the orientation. Based on the preliminary results seen for the x and y components, we expect MH to again outperform IKF also for this state variable.

D. Pose estimation with constraints

The last experiments aim to show the utility of the proposed estimation algorithm when constraints need to be satisfied while estimating the state. In many scenarios dealing with robots moving on the plane, it may be known upfront that the robot is bound to remain within the convex hull defined by the position of the landmarks. This is for example the case when the robot is moving indoor, and it is known that some landmarks are located on the walls surrounding the working area. Discarding or using this valuable a priori information may make a difference in the estimation process. For example, this information can be exploited in order to filter away pose estimates placing the robot outside the area of interest. Figure 2 illustrates the environment used throughout the tests presented in this section, with the blue segments bounding the working area.

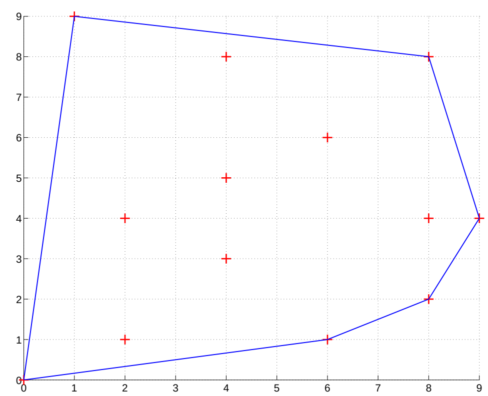

Fig. 2. The experimental environment used in all experiments described in this section. Distances are expressed in meters, red crosses indicate landmarks locations, and blue segments delimit the convex hull.

In the last batch of experiments we generated 100 robot paths that often get close to the boundaries of the area. Next, we used both MH and IKF to track the robot pose. Figure 3 shows a zoomed version of the prototypical results produced by IKF (green) and MH (red). The figure also shows the ground truth (black) and the boundaries (blue). The data has been generated in a situation where significant noise affects the evolution of the state, as it can be guessed from the jagged black path. The figure clearly shows that while

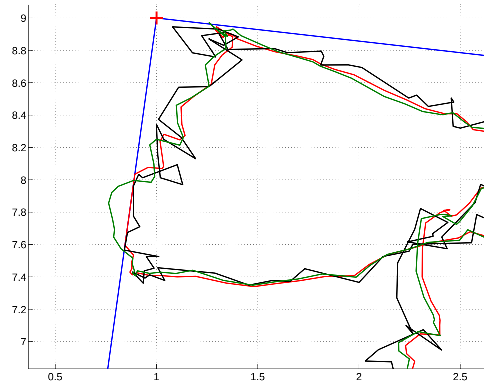

Fig. 3. A detail of a situation where the IKF estimate exits from the boundaries while MH respects the constraints.

IKF estimates a trajectory that exists the boundary of the environment, MH enforces the constraint, to the point that for a certain number of time steps its estimated trajectory coincides exactly with the imposed boundaries (see the part where the red path perfectly overlaps with the blue segment). Throughout the 100 runs, MH produced estimates that never left the boundaries. IKF, instead, being uninformed of the constraints, repeatedly violated them. To be precise, with each run being 400 steps long, on average IKF violated the constraints 15.32 times, with individual runs violating up to 68 times the constraints. To put these numbers into the right context, one should consider that in each run the trajectory stays close to the boundaries less than half of the time (i.e. less than 200 steps). In fact, almost every time the state trajectory approached the boundary, IKF generated an estimate violating them. MH, instead, seamlessly enforced these constraints.

V. Conclusions

When nonlinear dynamic systems are considered, estimates obtained by IKF may be quite distant from the optimal ones. An alternative is the use of particle filters. However, these techniques require delicate tuning of proposal densities in order to improve their convergence rates, and detection of convergence is often uncertain. These design problems can be further complicated when inequality constraints have to be included in the model to improve the state estimate.

In this paper, we have shown that moving horizon approaches represent an interesting alternative. In particular, we have cast in this scenario a recently proposed algorithm able also to efficiently handle inequality constraints on the state. The key feature of this approach is that it exploits the same decomposition used for unconstrained Kalman-Bucy smoothers. Thus, the required operations scale linearly with the horizon length. Results show that when the horizon length is set just to 2 or 3, the quality of the estimates improves significantly and with few additional computational effort in comparison with IKF. In particular, the moving horizon makes the filter much less sensitive to unknown initial conditions and local minima.

References

[1] B.M. Bell. The iterated Kalman smoother as a Gauss-Newton method *SIAM Journal on Optimization*, 4(3): 626-636, 1994.

[2] B.M. Bell. The marginal likelihood for parameters in a discrete Gauss-Markov process. *IEEE Trans. on Signal Processing*, 48(3):870–873, 2000.

[3] B.M. Bell, J. Burke and G. Pillonetto An inequality constrained nonlinear Kalman-Bucy smoother by interior point likelihood maximization *Automatica*, 45: 25-33, 2009.

[4] Y. Boers, H. Driessen, Particle filter track-before-detect application using inequality constraints, *IEEE Transactions on Aerospace and Electronic Systems*, 41(4): 1481-1487, 2005.

[5] J.V. Burke and S-P. Han, A robust sequential quadratic programming method, *Mathematical Programming*, 43(1-3): 277-303, 1989.

[6] M.L. Darby and M. Nikolaou, A parametric programming approach to moving-horizon state estimation, *Automatica*, 43(5): 885-891, 2007.

[7] A. El-Keyi, T. Kirubarajan, A.B. Gershman, Robust adaptive beamforming based on the Kalman filter, *IEEE Transactions on Signal Processing*, 53(8): 3032-3041, 2005.

[8] G. Erinc, G. Pillonetto, S. Carpin. Online estimation of variance parameters: experimental results with application to localization. *Proceedings of the 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1890-1895

[9] A. H. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, New York, 1970.

[10] S. Ko and R. Bitmead, State estimation for linear systems with state equality constraints, *Automatica*, 43(8): 1363-1368, 2007.

[11] C. Kwok, D. Fox, and M. Meila. Real-time particle filters. *Proceedings of the IEEE*, 92(3):469–484, 2004.

[12] J. Leonard and H. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transaction on Robotics and Automation*, 7(3):376–382, 1991.

[13] P.S. Maybeck. *Stochastic models, estimation and control*. Academic Press, 1979.

[14] A.I. Mourikis and S.I. Roumeliotis. Optimal sensor scheduling for resource-constrained localization of mobile robot formations. *IEEE Transactions on Robotics*, 22(5):917–931, 2006.

[15] A.I. Mourikis and S.I. Roumeliotis. Performance analysis of multirobot cooperative localization. *IEEE Transactions on robotics*, 22(4):666–681, 2006.

[16] C.V. Rao, J.B. Rawlings, J.H, Lee, Constrained linear state estimation - A moving horizon approach, *Automatica*, 37(10): 1619-1628, 2001.

[17] C.V. Rao, J.B. Rawlings, D.Q. Mayne, Constrained state estimation for nonlinear discrete-time systems: stability and moving horizon approximations, *IEEE Transactions on Automatic Control*, 48(2): 246-258, 2003.

[18] H.E. Rauch, F. Tung, C. T. Striebel, Maximum likelihood estimates of linear dynamic systems, *J. Amer. Inst. Aeronautics and Astronautics*, 3(8): 1445-1450, 1965.

[19] J.B. Rawlings and B.R. Bakshi, Particle filtering and moving horizon estimation, *Computers and Chemical Engineering*, 30: 1529-1541, 2006.

[20] D. Simon and D.L. Simon, Kalman filtering with inequality constraints for turbofan engine health estimation, *IEEE Proc-Control Theory and Applications*, 153(3): 371-378, 2006.

[21] R. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. In *Proceedings of IJCAI*, pages 1080–1087, 1995.

[22] R.F. Stengel. *Optimal control and estimation*. Dover, 1994.

[23] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2006.

[24] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 31:29–53, 1998. also appeared in Autonomous Robots 5(2):253–271 (joint issue).

[25] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust MonteCarlo localization for mobile robots. *Artificial Intelligence*, 128(1-2):99–141, 2001.

[26] L.S. Wang, Y.T. Chiang, F.R. Chang, Filtering method for nonlinear systems with constraints, *IEE Proceedings: Control Theory and Applications*, 149(6): 525-531, 2002.