

iRiSC: Iterative Risk Sensitive Control for Nonlinear Systems with Imperfect Observations

Bilal Hammoud^{1,2}, Armand Jordana¹ and Ludovic Righetti^{1,2}

Abstract—This work addresses the problem of risk-sensitive control for nonlinear systems with imperfect state observations, extending results for the linear case. In particular, we derive an algorithm that can compute local solutions with computational complexity similar to the iterative linear quadratic regulator algorithm. The proposed algorithm introduces feasibility gaps to allow the initialization with non-feasible trajectories. Moreover, an approximation for the expectation of the general nonlinear cost is proposed to enable an iterative line search solution to the planning problem. The optimal estimator is also derived along with the controls minimizing the general stochastic nonlinear cost. Finally extensive simulations are carried out to show the increased robustness the proposed framework provides when compared to the risk neutral iLQG counter part. To the authors’ best knowledge, this is the first algorithm that computes risk aware optimal controls that are a function of both the process noise and measurement uncertainty.

I. INTRODUCTION

Recently there has been an increased interest in computing optimal decisions that reason not only about the mean of a certain outcome or cost [1], [2] but also about the higher order statistics of the problem. A particular class of such algorithms is known as risk sensitive optimal control, which takes into account the variability of a certain cost and is concerned with the infrequent occurrences of undesired events besides the frequent most common outcomes. Risk sensitive problems have been widely used in economics [3], [4] and have recently gained popularity in robotics [5], [6] and reinforcement learning [7].

In linear quadratic Gaussian (LQG) control [8], the expectation of a quadratic cost function $\mathbb{E}[\mathcal{L}]$ is minimized over the decision variables. The main result in LQG theory, is that the optimal control is the same as that of the deterministic linear quadratic regulator (LQR) case. This is not the case anymore in the linear risk-sensitive control case with noisy measurements.

The risk-sensitive optimal control problem of interest in this paper was introduced by Jacobson [9] which instead minimizes the cost

$$\mathcal{J} = -\sigma^{-1} \ln \mathbb{E} [e^{-\sigma \mathcal{L}}] \quad (1)$$

where σ is a scalar whose role will become clear later. Under this particular transformation, and in the limit as $\sigma \rightarrow 0$ this

transformed cost can be approximated as

$$\mathcal{J} \approx \mathbb{E}[\mathcal{L}] - \frac{\sigma}{2} \text{Var}[\mathcal{L}] + \dots \quad (2)$$

where $\dots$ contains terms that are functions of higher order moments of $\mathcal{L}$. For $\sigma < 0$ (risk-averse case) the variability of the original cost $\mathcal{L}$ contributes positively to the newly transformed cost, thus an optimizer will seek a solution reducing the variability of the original cost and minimizing the possibility of infrequent events occurring. The opposite is true for the case of $\sigma > 0$ (risk-seeking case). Other forms of risk sensitivity have been studied [10], [11] and will not be considered in this paper. The solution for the stochastic optimal control problem under the exponential transformation given linear process model with additive Gaussian noise and a quadratic cost which we shall denote LQEG was derived by Jacobson [9].

The case with noisy measurements (and no direct access to the state) was studied in Speyer [12], who introduced a solution that grows with the history of observations for linear process and measurement models. The problem involving linear process and measurement models was eventually solved satisfactorily by Whittle [13] for the quadratic cost case. Variations of this problem were later studied either for the case of partially observable systems [14], [15] or the nonlinear case with full state knowledge [16]. Ponton [17] proposed an iterative algorithm to solve the case with measurement noise that requires augmenting the true state dynamics with the dynamics of an Extended Kalman Filter. However the derivation was heuristic with an incorrect line-search step which hinder the convergence along with an incorrect update of the feedforward control.

This work aims to provide the first iterative algorithm that computes a locally optimal solution to the nonlinear risk-sensitive optimal control problem with measurement uncertainty. We extend the results of Whittle [13] to the general nonlinear case of process and measurement dynamics with additive Gaussian noise. We propose a two-stage algorithm, where the first stage computes a nominal trajectory given the observations available at the starting point using an efficient line-search method. The second stage computes the optimal estimator along with the control optimizing the risk sensitive cost function as measurements are made during execution. Finally, extensive simulations of a nonlinear dynamical system involving stiff interactions with the environment demonstrate the capability of such controllers.

Notation: The gradient of a function f with respect to a vector v is denoted as f^v , similarly for second order

*This work was supported by New York University, the European Union’s Horizon 2020 research and innovation program (grant agreement 780684) and the National Science Foundation (grants 1825993, 1932187, 1925079 and 2026479).

¹Tandon School of Engineering, New York University, USA
bah436, aj2988, lr114@nyu.edu

²Max Planck Institute for Intelligent Systems, Tübingen, Germany

derivatives w.r.t vectors u, v will be denoted as f^{uv} . The determinant of a matrix M will be denoted by $|M|$. All functions are assumed to be C^2 unless otherwise stated. If $(v_i)_{i \in \mathbb{N}}$ is a sequence of vectors, then $v_{k:t}$ denotes the batch vector of all v_i for $k \leq i \leq t$.

II. BACKGROUND

This section gives an overview of the Linear Quadratic Exponential Gaussian problem (LQEG), where both process and measurement models are linear and the cost is an exponential of a quadratic in the state and control variables. We go over the main results from Jacobson [9] and Whittle [13] which we will build upon for our extension to the general nonlinear problem. Consider the following discrete time process and observation models,

$$x_{t+1} = f_t^x x_t + f_t^u u_t + \omega_{t+1} \quad (3a)$$

$$y_{t+1} = h_t^x x_t + \gamma_{t+1} \quad (3b)$$

where x_t , y_t and u_t are the state, observation and control at time t respectively. Given a belief $\hat{x}_0$, the initial state is assumed to follow a Gaussian distribution $x_0 \sim \mathcal{N}(\hat{x}_0, \chi_0)$. The process and measurement disturbances are considered to be Gaussian and denoted by $\omega_t \sim \mathcal{N}(0, \Omega_t)$ and $\gamma_t \sim \mathcal{N}(0, \Gamma_t)$. Then, given control inputs (or a policy), the joint probability density of the entire trajectory can be written

$$p(x_0, w_{1:T}, \gamma_{1:T}) = \frac{1}{\kappa} \exp(-\mathcal{D}) \quad (4)$$

$$\text{where } \mathcal{D} = d_0 + \sum_{t=1}^T d_t$$

$$d_0 = \frac{1}{2} (x_0 - \hat{x}_0)^T \chi_0^{-1} (x_0 - \hat{x}_0)$$

$$d_t = \frac{1}{2} \omega_t^T \Omega_t^{-1} \omega_t + \frac{1}{2} \gamma_t^T \Gamma_t^{-1} \gamma_t$$

$$\kappa = |2\pi\chi_0|^{\frac{1}{2}} \prod_{t=1}^T |2\pi\Omega_t|^{\frac{1}{2}} |2\pi\Gamma_t|^{\frac{1}{2}}$$

We denote the observations available at time t by W_t which takes the form

$$W_t = (\hat{x}_0, y_{1:t}, u_{0:t-1}). \quad (5)$$

It is important to note that the observation history also includes $\hat{x}_0$ the belief of the initial state. Then, the risk-sensitive optimal control problem of the imperfectly observable system at time t becomes a problem of computing the optimal policy $\pi(x)$, where $\pi(x)$ is a control law assigning values to u_t as a function of x_t such that the trajectory followed minimizes the following cost function

$$\mathcal{J}(\pi, W_t) = -\sigma^{-1} \ln(\mathbb{E}_\pi[\exp(-\sigma \mathcal{L}) \mid W_t]) \quad (6)$$

$$\text{where } \mathcal{L} = \frac{1}{2} x_T^T l_T^{xx} x_T + \frac{1}{2} \sum_{t=0}^{T-1} x_t^T l_t^{xx} x_t + u_t^T l_t^{uu} u_t$$

where the expectation is taken over every w_t , γ_t and x_0 . Taking the expectation conditioned on W_t determines $u_{0:t-1}$, however, control inputs, $u_{t:T}$, are taken according to the policy π . We consider policies that only depend on previous

observations. Before time t , observations are determined by W_t but after time t , the observations are random variable therefore future control inputs are also random variables. The optimal policy is then a function of the current available observations

$$\pi^*(W_t) = \underset{\pi}{\operatorname{argmin}} \mathcal{J}(\pi, W_t) \quad (7)$$

A. Optimal Control with Imperfect Observations

The expectation at time $t = 0$ of the total cost of the control problem can be expressed as

$$\begin{aligned} \mathbb{E}_\pi[e^{-\sigma \mathcal{L}} \mid W_0] &= \int e^{-\sigma \mathcal{L}} p(x_0, w_{1:T}, \gamma_{1:T}) dx_0 dw_{1:T} d\gamma_{1:T} \\ &= \frac{1}{\kappa} \int e^{-\sigma(\mathcal{L} + \sigma^{-1} \mathcal{D})} dx_0 dw_{1:T} d\gamma_{1:T}. \end{aligned} \quad (8)$$

Defining the total stress as $\mathcal{S} = \mathcal{L} + \sigma^{-1} \mathcal{D}$, Whittle then uses the quadratic lemma introduced by Jacobson [9] and summarized in Appendix A in order to compute the expectation in (8). The main idea is that an integration over an exponential of a quadratic of a variable x can be replaced by an optimization over the exponential of a quadratic of this variable. Then using induction, Whittle [13] proves that

$$\mathbb{E}_{\pi^*}[e^{(-\sigma \mathcal{L})} \mid W_t] = \frac{\alpha_t}{p(W_t)} e^{(-\sigma \mathcal{S}_t(W_t))} \quad (9)$$

where α_t is provided in appendix B and

$$\mathcal{S}_t(W_t) = \min_{u_t} \text{ext}_{y_{t+1}} \mathcal{S}_{t+1}(W_{t+1}) \quad (10a)$$

$$\mathcal{S}_T(W_T) = \text{ext}_{x_0, \dots, x_T} \mathcal{S} \quad (10b)$$

Here ext denotes an extremization, i.e. computation of either a maximum or a minimum. We say that a variable extremizes $\mathcal{S}$ if it minimises $\mathcal{S}$ when $\sigma < 0$ and maximizes $\mathcal{S}$ when $\sigma > 0$. This is due to the fact that the total stress could possess either a maximum or a minimum in the random variables depending on the sign of σ . Obtaining the optimal policy then boils down to the following minimization

$$\min_{u_t, \dots, u_{T-1}} \text{ext}_{x_0, \dots, x_T} \text{ext}_{y_{t+1}, \dots, y_T} \mathcal{S} \quad (11)$$

The optimal u_t is then the optimal control at time t given the current observations. Whittle's certainty equivalence principle says that the order of the successive extremization and minimization can be interchanged. If σ is negative, then the minimization problem is well defined only if $\mathcal{S}$ is negative definite in the extremizing variables. This condition ensures that the optimal cost is finite. More precisely, there exist a threshold value $\bar{\sigma} < 0$ depending on the problem parameters such that for all $\sigma > \bar{\sigma}$, the problem 7 is well defined. The implications of this condition on σ will become obvious later.

We provide a brief proof of these results in Appendix B. Unlike the iterative linear quadratic Gaussian iLQG [2], or the fully observable case of LQEG [9], the optimal control cannot be obtained with a single backward recursion in a dynamic programming fashion. This is due to the fact that the terminal condition (10b) includes an optimization over the entire state trajectory.

B. Separation principle

As Whittle concluded, at time t the problem includes minimizing over future controls, and extremizing over future observations and the entire state trajectory. Whittle [13] then suggested breaking the problem into two sub problems while keeping the current state x_t as a free variable, thus introducing the past stress $\mathcal{P}(x_t, W_t)$ and future stress $\mathcal{F}(x_t)$ at each time t^1 , which are defined by

$$\mathcal{P}(x_t, W_t) = \text{ext}_{x_{0:t-1}} \sum_{i=0}^{t-1} l_i + \sigma^{-1} \sum_{i=0}^t d_i, \quad (12)$$

$$\mathcal{F}(x_t) = \min_{u_{t:T-1}} \text{ext}_{\substack{x_{t+1:T} \\ y_{t+1:T}}} \sum_{i=t}^T l_i + \sigma^{-1} \sum_{i=t+1}^T d_i. \quad (13)$$

The past and future stress problems can then be computed sequentially separately as

$$\mathcal{P}(x_t, W_t) = \text{ext}_{x_{t-1}} l_{t-1} + \sigma^{-1} d_t + \mathcal{P}(x_{t-1}, W_{t-1}), \quad (14)$$

$$\mathcal{F}(x_t) = \min_{u_t} \text{ext}_{x_{t+1}} l_t + \sigma^{-1} d_{t+1} + \mathcal{F}(x_{t+1}), \quad (15)$$

where Whittle shows that optimizing over future observations $y_{t+1:T}$ yields a prediction of the form $y_{t+1} = h_t^x x_t$. In other words, for all $s > t + 1$, $\gamma_s = 0$. The recursions have the following boundary conditions:

$$\sigma \mathcal{P}(x_0, W_0) = \frac{1}{2} (x_0 - \hat{x}_0)^T \chi_0^{-1} (x_0 - \hat{x}_0), \quad (16)$$

$$\mathcal{F}(x_T) = \frac{1}{2} x_T^T l_T^{xx} x_T. \quad (17)$$

As the current state x_t has been kept as a free variable in optimization, the optimal control can be written as a function of x_t : $u_t = \pi_{x_t}(W_t)$. x_t is then determined by the last extremization

$$\tilde{x}_t = \underset{x_t}{\text{argmin}} \mathcal{P}(x_t, W_t) + \mathcal{F}(x_t) \quad (18)$$

The solution $\tilde{x}_t$ of this extremization then gives us the optimal control: $u_t^* = \pi_{\tilde{x}_t}(W_t) = \pi^*(W_t)$. This means that the optimal control is not obtained by only solving a backward recursion, but instead by optimizing an entire trajectory of states, while the policy as a function of these optimal states is obtained during the backward recursion. We will appeal next to these results to deriving an iterative algorithm for the nonlinear case.

III. ALGORITHM OVERVIEW

To solve the nonlinear problem, we propose an algorithm consisting of two stages. The first stage (discussed in detail in Sec. IV) optimizes a nominal trajectory $\mathcal{X}, \mathcal{U}$ and a policy $\pi(\delta \tilde{x}_t)$ where $\delta \tilde{x}_t$ is the deviation of the state minimizing the total stress $\mathcal{S}$ from the nominal trajectory. The second stage (discussed in Sec. V) is executed at run-time as measurements are made available and includes computing the

minimum stress deviations $\delta \tilde{x}_t$ along with the corresponding optimal policy.

The second stage happens at run-time. As each observations y_t becomes available, the past stress is propagated one time step forward as described in Algorithm 2 to obtain $\hat{x}_t$ and P_t , then the minimum stress estimate $\tilde{x}_t$ at time t is computed, and the optimal control is obtained and executed.

IV. NOMINAL TRAJECTORY OPTIMIZATION

Consider the general nonlinear process and measurement models with independent identically distributed (i.i.d) additive Gaussian noise ω_t and γ_t

$$x_{t+1} = f_t(x_t, u_t) + \omega_{t+1} \quad (19a)$$

$$y_{t+1} = h_t(x_t) + \gamma_{t+1} \quad (19b)$$

where f_t and h_t are differentiable. Along with the general nonlinear cost function of the form

$$\mathcal{L} = l_T(x_T) + \sum_{t=0}^{T-1} l_t(x_t, u_t) \quad (20)$$

where $l_t(x_t, u_t)$ is any twice differentiable nonlinear function of the states and controls. With this cost, the minimization (7) generally leads to an intractable optimal control problem that might possess multiple local minima [18]. A common approach to find one such local minimum is to follow an iterative approach similar to that of iterative linear quadratic regulator [1] and perform a line search on the cost function [19] until some convergence criteria is achieved. So, we seek some iterative updates to the state and control trajectories in the form

$$\begin{bmatrix} \mathcal{X} \\ \mathcal{U} \end{bmatrix}^{i+1} = \begin{bmatrix} \mathcal{X} \\ \mathcal{U} \end{bmatrix}^i + \alpha \begin{bmatrix} \delta \mathcal{X} \\ \delta \mathcal{U} \end{bmatrix}^i \quad (21)$$

$\mathcal{X}$ and $\mathcal{U}$ denote the nominal state and control sequences $x_{0:T}$ and $u_{0:T-1}$ at the indicated iterations respectively. The δ is used to denote a small change around these trajectories and α is the step length corresponding to the search direction $\delta \mathcal{X}, \delta \mathcal{U}$ and determined via line search.

A. Step Direction

In order to generate a descent direction for the line search a linear approximation of the dynamics and observations along the nominal trajectory is computed and a quadratic approximation of the cost is considered. These approximations are computed along the nominal trajectory at iteration i denoted by $\mathcal{X}^i, \mathcal{U}^i$ and are given by

$$x_{t+1} \approx f_t(x_t^i, u_t^i) + f_t^x \delta x_t + f_t^u \delta u_t + \omega_{t+1} \quad (22a)$$

$$y_{t+1} \approx h_t(x_t^i) + h_t^x \delta x_t + \gamma_{t+1} \quad (22b)$$

and

$$l_t(x_t^i + \delta x_t, u_t^i + \delta u_t) \approx l_t(x_t^i, u_t^i) + \delta l_t \quad (23)$$

where $\delta x_t = x_t - x_t^i$, $\delta u_t = u_t - u_t^i$ and $\delta y_t = y_t - h_t(x_t^i)$. For the remainder of the paper we will omit the dependence

¹Note that $\mathcal{P}$ and $\mathcal{F}$ should be index by t , but in order to simplify the notations, we omit this time index as the latter can be inferred thanks to x_t .

of $f(x, u)$ and $l(x, u)$ on x and u for brevity. The quadratic approximation of the cost is then given by

$$\delta l_t = \frac{1}{2} \begin{bmatrix} 1 \\ \delta x_t \\ \delta u_t \end{bmatrix}^T \begin{bmatrix} 0 & l_t^x & l_t^u \\ l_t^x & l_t^{xx} & l_t^{xu} \\ l_t^u & l_t^{xu} & l_t^{uu} \end{bmatrix} \begin{bmatrix} 1 \\ \delta x_t \\ \delta u_t \end{bmatrix}. \quad (24)$$

Following [20] we introduce feasibility gaps along the nominal trajectory to form

$$\bar{f}_{t+1} = f(x_t^i, u_t^i) - x_{t+1}^i \quad (25)$$

Hence, we optimize $\mathbb{E}_\pi [e^{-\sigma \delta \mathcal{L}}]$ subject to the process and observation deviation dynamics

$$\delta x_{t+1} = f_t^x \delta x_t + f_t^u \delta u_t + \bar{f}_{t+1} + \omega_{t+1} \quad (26a)$$

$$\delta y_{t+1} = h_t^x \delta x_t + \gamma_{t+1} \quad (26b)$$

The optimal search direction at iteration i is obtained by computing the optimal control deviation δu_t^* according to the future stress. We provide an alternative derivation to that of Whittle's in an extended document provided along with the code in ². The document proves that the future stress optimization is independent of the future observations. The optimal control deviations are then given by the following theorem

Theorem 4.1: Assuming that $\mathcal{F}(\delta x_{t+1})$ has the form

$$\mathcal{F}(\delta x_{t+1}) = \frac{1}{2} \delta x_{t+1}^T V_{t+1} \delta x_{t+1} + \delta x_{t+1}^T v_{t+1} + \bar{v}_{t+1} \quad (27)$$

Then the optimal control deviations δu_t^* minimizing $\mathbb{E}_\pi [e^{-\sigma \delta \mathcal{L}} | W_0]$ are given recursively by

$$\delta u_t^* = - \underbrace{Q_t^{uu^{-1}} Q_t^u}_{k_t} - \underbrace{Q_t^{uu^{-1}} Q_t^{ux}}_{K_t} \delta x_t \quad (28)$$

where Q_t^u , Q_t^{uu} and Q_t^{ux} are defined in (31).

Proof:

We start by writing the future stress, we know from [13] that the measurement uncertainty yields a prediction and that the future stress takes the form

$$\begin{aligned} \mathcal{F}(\delta x_t) = \min_{\delta u_t} \text{ext}_{\delta x_{t+1}} \delta l_t + \mathcal{F}(\delta x_{t+1}) \\ + \frac{\sigma^{-1}}{2} (\delta x_{t+1} - z_t)^T \Omega_{t+1}^{-1} (\delta x_{t+1} - z_t) \end{aligned} \quad (29)$$

where $z_t = f_t^x \delta x_t + f_t^u \delta u_t + \bar{f}_{t+1}$. The first step is to optimize with respect to δx_{t+1} and rearrange the partially optimized future stress into a quadratic in the current controls and states δu_t and δx_t respectively

$$\mathcal{F}(\delta x_t) = \min_{\delta u_t} \frac{1}{2} \begin{bmatrix} 1 \\ \delta x_t \\ \delta u_t \end{bmatrix}^T \begin{bmatrix} \bar{q}_t & Q_t^x & Q_t^{xu} \\ Q_t^x & Q_t^{xx} & Q_t^{xu} \\ Q_t^u & Q_t^{xu} & Q_t^{uu} \end{bmatrix} \begin{bmatrix} 1 \\ \delta x_t \\ \delta u_t \end{bmatrix} \quad (30)$$

where

$$Q_t^x = l_t^x + f_t^{xT} M_t \bar{f}_{t+1} + f_t^{xT} N_t \quad (31a)$$

$$Q_t^u = l_t^u + f_t^{uT} N_t + f_t^{uT} M_t \bar{f}_{t+1} \quad (31b)$$

$$Q_t^{uu} = l_t^{uu} + f_t^{uT} M_t f_t^u \quad (31c)$$

$$Q_t^{ux} = l_t^{ux} + f_t^{uT} M_t f_t^x \quad (31d)$$

$$Q_t^{xx} = l_t^{xx} + f_t^{xT} M_t f_t^x \quad (31e)$$

$$M_t = (\sigma \Omega_{t+1} + V_{t+1}^{-1})^{-1} \quad (31f)$$

$$N_t = v_{t+1} - \sigma M_t \Omega_{t+1} v_{t+1} \quad (31g)$$

And where $\bar{q}_t$ is a constant. Being quadratic in δu_t , the remaining optimization is obtained by solving $\nabla_{\delta u_t} \mathcal{F}(\delta x_t) = 0$ which concludes the proof for the control recursions. A more detailed version of the proof that includes straightforward yet lengthy algebraic manipulations is provided in our public repository along with the code implementing all the simulations. ■

Theorem 4.2: The future stress is a quadratic in the state deviations δx_t and takes the form

$$\mathcal{F}(\delta x_t) = \frac{1}{2} \delta x_t^T V_t \delta x_t + \delta x_t^T v_t + \bar{v}_t \quad (32)$$

where

$$V_t = Q_t^{xx} + K_t^T Q_t^{uu} K_t - K_t^T Q_t^{ux} - Q_t^{xu} K_t \quad (33a)$$

$$v_t = Q_t^x + K_t^T Q_t^{uu} k_t - K_t^T Q_t^u - Q_t^{xu} k_t \quad (33b)$$

$$\Delta \bar{v}_t = \bar{q}_t + k_t^T Q_t^u + \frac{1}{2} k_t^T Q_t^{uu} k_t \quad (33c)$$

Proof: The proof follows directly by replacing the optimal controls in $\mathcal{F}(\delta x_t)$ and regrouping, details of the calculations are provided in the accompanying repository. ■

B. Approximating the Cost Expectation

In order to implement a line search procedure, it is necessary to evaluate the nonlinear risk sensitive cost. However, for a general nonlinear cost, computing the expectation is generally intractable analytically. Therefore, we use a quadratic approximation of the cost and a linearization of the dynamics to obtain an approximation of the risk sensitive cost. The linearization and quadratic approximations are done only over the states as the candidate control trajectory is considered to be fixed. More precisely, given a candidate control sequence $\mathcal{U}$ and a nominal trajectory $\mathcal{X}$, we evaluate:

$$\tilde{J} = -\sigma^{-1} \ln \left(\mathbb{E} \left[e^{-\sigma \tilde{\mathcal{L}}} | u_{0:T-1} \right] \right) \quad (34)$$

subject to

$$\delta x_0 = x_0 - \hat{x}_0 \quad (35a)$$

$$\delta x_{t+1} = f_t^x \delta x_t + \bar{f}_{t+1} + \omega_{t+1} \quad (35b)$$

with:

$$\tilde{\mathcal{L}} = \sum_{t=0}^T \underbrace{l(x_t^n, u_t)}_{\tilde{l}_t} + \underbrace{\frac{1}{2} \delta x_t^T l_t^{xx} \delta x_t + l_t^{xT} \delta x_t}_{\delta l_t} \quad (36)$$

²<https://github.com/hammoudbilal/irisc>

Theorem 4.3:

$$\tilde{\mathcal{J}} = \sigma^{-1} \ln(\alpha) + \bar{v}_0 - \frac{1}{2} v_0^T (V_0 + \sigma^{-1} \chi_0^{-1})^{-1} v_0 \quad (37)$$

with

$$\alpha = |I + \sigma V_0 \chi_0|^{\frac{1}{2}} \Pi_{t=1}^T |I + \sigma V_t \Omega_t|^{\frac{1}{2}} \quad (38)$$

and where $\bar{v}_0, v_0, V_0$ are the terminal values of the recursion:

$$V_t = l_t^{xx} + f_t^{xT} M_t f_t^x \quad (39a)$$

$$v_t = l_t^x + f_t^{xT} M_t \bar{f}_{t+1} + f_t^{xT} N_t \quad (39b)$$

$$\begin{aligned} \bar{v}_t &= \bar{l}_t + \bar{v}_{t+1} + \bar{f}_{t+1}^T N_t + \frac{1}{2} \bar{f}_{t+1}^T M_t \bar{f}_{t+1} \\ &\quad - \frac{1}{2} v_{t+1}^T (V_{t+1} + \sigma^{-1} \Omega_{t+1}^{-1})^{-1} v_{t+1} \end{aligned} \quad (39c)$$

where M_t and N_t have the same form as (31f) and (31g) but evaluated at V_t and v_t given in (39)

Proof: The proof follows a similar logic of theorem 4.1 and is provided fully in the repository. ■

C. Step Length & Partial Update of the Gaps

We choose to update the gaps partially as is done in [21]. This allows for a partial contraction of the gaps unless a full step $\alpha = 1$ is taken thus allowing for a better convergence when initialized from infeasible nominal trajectories, the update equations are then

$$u_t^{i+1} = u_t^i - \alpha k_t - K_t (x_t^{i+1} - x_t^i), \quad (40)$$

$$x_{t+1}^{i+1} = f_t(x_t^{i+1}, u_t^{i+1}) - (1 - \alpha) \bar{f}_{t+1}^i. \quad (41)$$

with the initial condition $x_0^{i+1} = x_0^i$. We perform a simple backtracking line search for a set of values of α denoted $\mathcal{A}$ until no more decrease in the approximated cost is observed. The pseudo-code algorithm is provided in Algorithm 1.

V. ESTIMATION & OPTIMAL CONTROLS

The actual optimal control problem of the imperfectly observable system is not fully solved yet. So far we only have the optimal control if optimal x_t were known. The next step consists of solving for the past stress

A. Past Stress Optimization

The solution to the past stress (12) optimization with the cost approximation (37) and linear dynamics (26) can be obtained using a forward recursion as described in the following theorem:

Theorem 5.1:

$$\sigma \mathcal{P}(\delta x_t, W_t) = \frac{1}{2} (\delta x_t - \delta \hat{x}_t)^T P_t^{-1} (\delta x_t - \delta \hat{x}_t) \quad (42)$$

where P_t and $\delta \hat{x}_t$ are subject to the following recursion

$$\begin{aligned} \delta \hat{x}_{t+1} &= f_t^x \delta \hat{x}_t + f_t^u \delta u_t + \bar{f}_{t+1} + G_t (\delta y_{t+1} - h_t^x \delta \hat{x}_t) \\ &\quad - \sigma f_t^x H_t (l_t^{xx} \delta \hat{x}_t + l_t^{xu} \delta u_t + l_t^x) \end{aligned} \quad (43)$$

Algorithm 1: Planning Phase

Input: nonlinear cost, nonlinear dynamics, nonlinear measurement, sensitivity, $\mathcal{X}^0, \mathcal{U}^0$

```

1 while Not Converged do
2   compute dynamics & cost
   approximations (26), (24)
   // backward pass
3   for  $t = T; t \geq 0; t--$  do
4      $k_t, K_t \leftarrow \text{eqn. (28)}$ 
5      $V_t, v_t, \bar{v}_t \leftarrow \text{eqn. (33)}$ 
6   end
   // forward pass
7   for  $\alpha \in \mathcal{A}$  do
8     for  $t = 0; t \leq T; t++$  do
9        $u_t^{i+1} \leftarrow \text{eqn. (40)}$ 
10       $x_t^{i+1} \leftarrow \text{eqn. (41)}$ 
11    end
12     $\mathcal{J}^{i+1} \leftarrow \text{eqn. (37)}$ 
13    if  $\mathcal{J}^{i+1} \leq \mathcal{J}^i$  then
14       $\mathcal{X}, \mathcal{U} \leftarrow x_{0:T}^{i+1}, u_{0:T-1}^{i+1}$ 
15      break
16    end
17  end
18  if  $|\mathcal{J}^{i+1} - \mathcal{J}^i| \leq 10^{-12}$  then
19     $\text{Converged} \leftarrow \text{True}$ 
20  end
21 end

```

Output: $\mathcal{X}, \mathcal{U}, K_{0:T-1}, V_{0:T}, v_{0:T}$

$$H_t = \left(P_t^{-1} + h_t^{xT} \Gamma_{t+1}^{-1} h_t^x + \sigma l_t^{xx} \right)^{-1} \quad (44a)$$

$$G_t = f_t^x H_t h_t^{xT} \Gamma_{t+1}^{-1} \quad (44b)$$

$$P_{t+1} = \Omega_{t+1} + f_t^x H_t f_t^{xT} \quad (44c)$$

Proof: The proof is done by induction. Given $\mathcal{P}(\delta x_{t-1}, W_{t-1})$, We first notice that:

$$\begin{aligned} \sigma l_{t-1} + \sigma \mathcal{P}(\delta x_{t-1}, W_{t-1}) &= \\ \frac{1}{2} (\delta x_{t-1} - \delta \tilde{x}_{t-1})^T \tilde{P}_{t-1}^{-1} (\delta x_{t-1} - \delta \tilde{x}_{t-1}) + C^{ste} \end{aligned} \quad (45)$$

where:

$$\tilde{P}_{t-1} = (P_{t-1}^{-1} + \sigma l_{t-1}^{xx})^{-1} \quad (46a)$$

$$\delta \tilde{x}_{t-1} = \tilde{P}_{t-1} (P_{t-1}^{-1} \delta \hat{x}_{t-1} - \sigma l_{t-1}^{xu} \delta u_{t-1} - \sigma l_{t-1}^x) \quad (46b)$$

Hence, (12) can be written as:

$$\begin{aligned} \sigma \mathcal{P}(\delta x_t, W_t) &= C^{ste} \\ &\quad + \underset{\delta x_{t-1}}{\text{ext}} d_t + \frac{1}{2} (\delta x_{t-1} - \delta \tilde{x}_{t-1})^T \tilde{P}_{t-1}^{-1} (\delta x_{t-1} - \delta \tilde{x}_{t-1}) \end{aligned} \quad (47)$$

Carrying out this optimization over δx_{t-1} and rearranging the terms to construct a perfect square concludes the proof. This last step is included in the accompanying repository. ■

B. Minimum Stress State & Optimal Control

Now that both $\mathcal{P}(\delta x_t, W_t)$ and $\mathcal{F}(\delta x_t)$ are obtained, what remains is optimizing the last remaining variable δx_t , i.e. solving

$$\begin{aligned} \delta \tilde{x}_t &= \underset{\delta x_t}{\operatorname{argmin}} \quad \mathcal{P}(\delta x_t, W_t) + \mathcal{F}(\delta x_t) \\ &= (I + \sigma P_t V_t)^{-1} (\delta \hat{x}_t - \sigma P_t v_t) \end{aligned} \quad (48)$$

which results in the minimum stress state deviation $\delta \tilde{x}_t$. With u_t^n being the output of algorithm 1, we obtain the optimal control law

$$u_t^* = u_t^n - K_t (I + \sigma P_t V_t)^{-1} (\delta \hat{x}_t - \sigma P_t v_t) \quad (49)$$

Algorithm 2: Run Time Estimation & Control

Input: $\mathcal{X}, \mathcal{U}, K_{0:T-1}, V_{0:T}, v_{0:T}, \delta y_t, \delta \hat{x}_{t-1}, P_{t-1}$

1 $H_{t-1}, G_{t-1}, P_t \leftarrow \text{eqn. (44)}$

2 $\delta \hat{x}_t \leftarrow \text{eqn. (43)}$

3 $\delta \tilde{x}_t \leftarrow \text{eqn. (48)}$

4 $u_t^* \leftarrow \text{eqn. (49)}$

Output: u_t^*

VI. RESULTS

To demonstrate the capabilities of the algorithm with nonlinear systems with measurement noise, we optimize the motion of a one dimensional pneumatic hopper depicted in 1. The state vector is four dimensional $x^T = [q^h, q^f, v^h, v^f]$ including the hip position relative to the ground, the foot position relative to the hip position, and their velocities respectively. The control vector is one dimensional controlling the acceleration of the foot relative to the hip. The distance from the foot position to the ground can then be defined as $e = q^f - q^h$.




Fig. 1: Simulation snapshots with pneumatic hopper.

The simulation uses a stiff visco-elastic contact model $\lambda(e, \dot{e}) = ke - b\dot{e}$. For optimization we will use the relaxed contact model

$$\lambda(e) = \begin{cases} 0, & \text{if } e < 0. \\ \frac{k}{2\alpha} e^2, & \text{if } 0 \leq e < \alpha. \\ ke - \frac{1}{2}k\alpha & \text{otherwise.} \end{cases} \quad (50)$$

To solve the optimal control problem, we use $\alpha = 0.01m$ and $k = 500N/m$, where the simulation parameters use $k = 10^5 N/m$ and $b = 300 N.s/m$ (we aim to demonstrate that feedback can compensate for the unexpectedly stiff contact in simulation). The cost function used includes three phases, the stance phase, which is along the horizon T except for $t = T/2$, the jumping cost for $t = T/2$ and terminal cost for time $t = T$. We choose to write the original cost as $l_t = \frac{1}{2}(x_t - x_{des})^T l_t^{xx} (x_t - x_{des}) + \frac{1}{2} u_t^T l_t^{uu} u_t$ where $x_{des}^T = [0.5, 0., 0., 0.]$ for the stance and terminal costs and $x_{des}^T = [2., 0., 0., 0.]$ for the jump phase. The cost weights are set as follows, for the stance phase $l^{xx} = \text{diag}(10^1, 1, 10^{-4}, 0.)$, for the jump phase $l^{xx} = \text{diag}(10^1, 10^{-2}, 10^{-1}, 0.)$, for the terminal phase we have $l^{xx} = \text{diag}(10^1, 10^1, 1, 1.)$, and $l^{uu} = 10^{-3}$ for all running costs. The noise is assumed to be uniform with $\Omega_t = \Gamma_t = \text{diag}(10^{-3}, 10^{-3}, 10^{-3}, 10^{-3})$. One thousand simulations are carried out for each of the three scenarios, a risk neutral controller and estimator namely DDP [21] combined with an extended kalman filter. A risk averse controller with $\sigma = -0.5$ and a risk seeking controller with $\sigma = 10$. All the results below are presented in terms of the mean/average and standard deviation of the data obtained. Risk Sensitive behavior achieves superior performance compared to the two other formulations as we will show below.




Fig. 2: Planned Trajectories, Solid lines represent the hip trajectory and dashed lines represent the foot trajectory.

A. Optimized Trajectories

One advantage of the proposed procedure is that it is capable of generating different trajectories depending on the noise and sensitivity parameters, this is evident from the resulting plans. As the results in Fig. 2 show, both DDP and risk seeking scenarios generate a plan that reaches exactly the desired height Fig. 2a and Fig. 2c, whereas the risk averse scenario generates a higher maximum height Fig. 2b. The penetration into the ground in the plan is due to using a relaxed contact model that does not generate stiff

contact behaviors. The goal is to have the control feedback compensate for the stiff behavior in the simulations.

B. Simulation Statistics

In order to compare performance, the statistics of the accumulated original cost $\mathcal{L}$ is computed for each of the three scenarios studied and summarized in Table I. Risk averse design achieves a lower cost average, and a significantly tighter standard deviation. Whereas risk seeking design generates the highest cost average and ddp generates the highest standard deviation (Fig. 3).

	DDP	$\sigma = -0.5$	$\sigma = 10$
Total average	5.79×10^{-2}	5.13×10^{-2}	13.92×10^{-2}
Maximum Std	13.17×10^{-2}	3.53×10^{-2}	12.92×10^{-2}

TABLE I: Cost Statistics




Fig. 3: Cost Statistics for DDP, Risk Averse and Risk Seeking control designs after 1000 simulations.

The tracking error statistics is computed with respect to the optimized plans for each of the scenarios, and presented for each of the states in Table II. We see that the risk-averse design maintains the lowest average position errors and lowest standard deviations for all four states. Although the planned trajectory aims for a higher height, the actual average hip height reaches the desired hip height as shown in Fig. 4.

State	Max. Average			Max. STD		
	DDP	$\sigma = -0.5$	$\sigma = 10$	DDP	$\sigma = -0.5$	$\sigma = 10$
q^h	0.354	0.334	0.410	0.248	0.230	0.244
q^f	0.351	0.306	0.370	0.133	0.122	0.150
v^h	3.83	3.337	3.135	1.815	1.652	1.585
v^f	2.35	2.182	3.959	1.256	0.743	1.704

TABLE II: Tracking Error Statistics




Fig. 4: Hip and Foot Simulated Position Statistics

On the other hand, DDP comes short from the desired height, and experiences a foot bounce off the ground more often as shown in Fig. 4b. The risk seeking design achieves




Fig. 5: Average contact force

a higher height than desired, much later than planned and also exhibits a foot bouncing back off the ground (Fig. 4).

Another interesting aspect is the maximum average forces generated. Even though the forces are not explicit states that are penalized in the cost, the force plays a crucial role in taking off, reaching the desired height and landing safely. DDP generates the lowest takeoff force, which explains the lower average peak height (Table III). For the landing phase DDP generates the highest impact force (Fig. 5) while risk seeking generates the lowest, but then diverges. Risk averse generates the most stable contact force profiles, a highly sought behavior for systems establishing and breaking contact with the environment instead of avoiding it.

	Takeoff f	Landing f	Hip Height	Foot Height
DDP	118.099	106.902	1.888	0.698
$\sigma = -0.5$	106.146	100.527	2.007	0.801
$\sigma = 10$	103.712	330.583	2.380	0.816

TABLE III: Maximum average force at takeoff and landing, hip height and foot height

VII. CONCLUSION

In this paper we present, to our knowledge, the first computationally efficient iterative algorithm to compute risk-sensitive optimal trajectories and control policies that handle both process and measurement noise. The plan is obtained before any observations are collected, as observations become available, an estimator that takes into account the cost accumulated generates state estimates that are then used to update both the control signal and the feedback controls. We tested the proposed algorithm on a one dimensional hopper in intermittent contact with the environment, subject to both model disturbances and measurement noise. In the risk averse scenario the proposed algorithm proves to be superior to its deterministic counter part namely DDP.

REFERENCES

- [1] A. Sideris and J. E. Bobrow, "An efficient sequential linear quadratic algorithm for solving nonlinear optimal control problems," in *Proceedings of the 2005, American Control Conference, 2005*. IEEE, 2005, pp. 2275–2280.
- [2] E. Todorov and W. Li, "A generalized iterative lqg method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *Proceedings of the 2005, American Control Conference, 2005*. IEEE, 2005, pp. 300–306.
- [3] K. J. Arrow, "The theory of risk aversion," *Essays in the theory of risk-bearing*, pp. 90–120, 1971.
- [4] J. W. Pratt, "Risk aversion in the small and in the large," in *Uncertainty in economics*. Elsevier, 1978, pp. 59–79.

- [5] B. Hammoud, M. Khadiv, and L. Righetti, "Impedance optimization for uncertain contact interactions through risk sensitive optimal control," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4766–4773, 2021.
- [6] A. Majumdar and M. Pavone, "How should a robot assess risk? towards an axiomatic theory of risk in robotics," in *Robotics Research*. Springer, 2020, pp. 75–84.
- [7] S. Bechtle, Y. Lin, A. Rai, L. Righetti, and F. Meier, "Curious ilqr: Resolving uncertainty in model-based rl," in *Proceedings of the Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, vol. 100, Nov 2019, p. 162–171.
- [8] A. E. Bryson and Y.-C. Ho, *Applied optimal control: optimization, estimation, and control*. Routledge, 2018.
- [9] D. Jacobson, "Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games," *IEEE Transactions on Automatic control*, vol. 18, no. 2, pp. 124–131, 1973.
- [10] S. Singh, Y. Chow, A. Majumdar, and M. Pavone, "A framework for time-consistent, risk-sensitive model predictive control: Theory and algorithms," *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2905–2912, 2018.
- [11] A. Tsiamis, D. S. Kalogerias, L. F. Chamon, A. Ribeiro, and G. J. Pappas, "Risk-constrained linear-quadratic regulators," in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 3040–3047.
- [12] J. Speyer, J. Deyst, and D. Jacobson, "Optimization of stochastic linear systems with additive measurement and process noise using exponential performance criteria," *IEEE Transactions on Automatic Control*, vol. 19, no. 4, pp. 358–366, 1974.
- [13] P. Whittle, "Risk-sensitive linear/quadratic/Gaussian control," *Advances in Applied Probability*, pp. 764–777, 1981.
- [14] A. Bensoussan and J. H. van Schuppen, "Optimal control of partially observable stochastic systems with an exponential-of-integral performance index," *SIAM Journal on Control and Optimization*, vol. 23, no. 4, pp. 599–613, 1985.
- [15] M. R. James, J. S. Baras, and R. J. Elliott, "Risk-sensitive control and dynamic games for partially observed discrete-time nonlinear systems," *IEEE transactions on automatic control*, vol. 39, no. 4, pp. 780–792, 1994.
- [16] V. Roulet, M. Fazel, S. Srinivasa, and Z. Harchaoui, "On the convergence of the iterative linear exponential quadratic gaussian algorithm to stationary points," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 132–137.
- [17] B. Pontón, S. Schaal, and L. Righetti, "On the effects of measurement uncertainty in optimal control of contact interactions," in *Algorithmic Foundations of Robotics XII*. Springer, 2020, pp. 784–799.
- [18] M. P. Chapman and K. M. Smith, "Classical risk-averse control for a finite-horizon borel model," *IEEE Control Systems Letters*, 2021.
- [19] J. Nocedal and S. Wright, *Numerical optimization*. Springer Science & Business Media, 2006.
- [20] M. Gifithaler, M. Neunert, M. Stäuble, J. Buchli, and M. Diehl, "A family of iterative gauss-newton shooting methods for nonlinear optimal control," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–9.
- [21] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocodyl: An efficient and versatile framework for multi-contact optimal control," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2536–2542.
- [22] P. Whittle, *Optimization over time*. John Wiley & Sons, Inc., 1982.

APPENDIX

A. Integration of Exponential of a Quadratic form.

Lemma 1.1: For a quadratic form in x and y

$$Q(x, y) = \frac{1}{2} \begin{bmatrix} 1 \\ x \\ y \end{bmatrix}^T \begin{bmatrix} \bar{q} & q_x^T & q_y^T \\ q_x & Q_{xx} & Q_{xy} \\ q_y & Q_{yx} & Q_{yy} \end{bmatrix} \begin{bmatrix} 1 \\ x \\ y \end{bmatrix} \quad (51)$$

We have:

$$\int \exp(-Q(x, y)) dx = |2\pi Q_{xx}^{-1}|^{\frac{1}{2}} \exp(-Q(\hat{x}, y)) \quad (52)$$

$$\hat{x} = \underset{x}{\operatorname{argmin}} Q(x, y) = -Q_{xx}^{-1}(Q_{xy}y + q_x) \quad (53)$$

Proof: We notice that:

$$Q(x, y) - Q(\hat{x}, y) = \frac{1}{2}(x - \hat{x})^T Q_{xx}(x - \hat{x}) \quad (54)$$

Then, the following equality concludes the proof:

$$\int \exp\left(-\frac{1}{2}(x - \hat{x})^T Q_{xx}(x - \hat{x})\right) dx = \sqrt{|2\pi Q_{xx}^{-1}|} \quad (55)$$

B. main theorem

Theorem 1.2:

$$\mathbb{E}_{\pi^*} [\exp(-\sigma \mathcal{L}) | W_t] = \frac{\alpha_t}{p(W_t)} \exp(-\sigma \mathcal{S}_t(W_t)) \quad (56)$$

Proof:

According to [22], we have

$$\begin{aligned} \mathbb{E}_{\pi^*} [\exp(-\sigma \mathcal{L}) | W_t] \\ = \operatorname{ext}_{u_t} \int \mathbb{E}_{\pi^*} [\exp(-\sigma \mathcal{L}) | W_{t+1}] p(y_{t+1} | W_t) dy_{t+1} \end{aligned} \quad (57)$$

The proof is then a backward induction based on the recursion equality (57). We first notice that:

$$\begin{aligned} \mathbb{E}_{\pi^*} [\exp(-\sigma \mathcal{L}) | W_T] \\ = \int e^{-\sigma \mathcal{L}} p(x_{0:T} | W_T) dx_{0:T} \\ = \frac{1}{p(W_T)} \int e^{-\sigma \mathcal{L}} p(x_{0:T}, y_{0:T}) dx_{0:T} \\ = \alpha_T \exp(-\sigma \min_{x_0, \dots, x_T} \mathcal{S}) \end{aligned} \quad (58)$$

The last equality comes from the lemma therefore, we have:

$$\alpha_T = \frac{|2\pi \sigma \frac{\partial^2 \mathcal{S}}{\partial^2 x_{0:T}}|^{\frac{1}{2}}}{p(W_T) |2\pi \chi_0|^{\frac{1}{2}} \prod_{t=1}^T |2\pi \Omega_t|^{\frac{1}{2}}} \quad (59)$$

Now, assuming the Equation 56 for $t + 1$, let's show the property at time t : Thanks to (57), we have

$$\begin{aligned} \mathbb{E}_{\pi^*} [\exp(-\sigma \mathcal{L}) | W_t] \\ = \operatorname{ext}_{u_t} \int \frac{\alpha_{t+1}}{p(W_{t+1})} \exp[-\sigma(\mathcal{S}_{t+1}(W_{t+1}))] p(y_{t+1} | W_t) dy_{t+1} \\ = \frac{\alpha_{t+1}}{p(W_t)} \operatorname{ext}_{u_t} \int \exp[-\sigma \mathcal{S}_{t+1}(W_{t+1})] dy_{t+1} \\ = \frac{\alpha_t}{p(W_t)} \operatorname{ext}_{u_t} \operatorname{ext}_{y_{t+1}} \exp[-\sigma \mathcal{S}_t(W_t)] \end{aligned} \quad (60)$$

The last equality comes from the lemma therefore, we have:

$$\alpha_t = \alpha_{t+1} \left| 2\pi \sigma \frac{\partial^2 \mathcal{S}_{t+1}(W_{t+1})}{\partial^2 y_{t+1}} \right|^{\frac{1}{2}} \quad (61)$$