

Minimal regret state estimation of time-varying systems

Conference Paper**Author(s):**

Brouillon, Jean-Sébastien; [Dörfler, Florian](#) ; Ferrari-Trecate, Giancarlo

Publication date:

2023-11-22

Permanent link:

<https://doi.org/10.3929/ethz-b-000660769>

Rights / license:

[Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International](#)

Originally published in:

IFAC-PapersOnLine 56(2), <https://doi.org/10.1016/j.ifacol.2023.10.1345>

Funding acknowledgement:

180545 - NFS Dependable Ubiquitous Automation (SNF)

Minimal regret state estimation of time-varying systems

Jean-Sébastien Brouillon* Florian Dörfler**
Giancarlo Ferrari Trecate*

* *Institute of Mechanical Engineering, École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland (e-mail: jean-sebastien.brouillon@epfl.ch, giancarlo.ferraritrecate@epfl.ch)*

** *Automatic Control Laboratory, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland (dorfler@control.ee.ethz.ch)*

Abstract: Kalman and \mathcal{H}_∞ filters, the most popular paradigms for linear state estimation, are designed for very specific noise and disturbance patterns, which may not appear in practice. State observers based on the minimization of regret measures are a promising alternative, as they aim to adapt to recognizable patterns in the estimation error. In this paper, we show that the regret minimization problem for finite horizon estimation can be cast into a simple convex optimization problem. For this purpose, we first rewrite linear time-varying system dynamics using a novel system level synthesis parametrization for state estimation, capable of handling both disturbance and measurement noise. We then provide a tractable formulation for the minimization of regret based on semi-definite programming. Both contributions make the minimal regret observer design easily implementable in practice. Finally, numerical experiments show that the computed observer can significantly outperform both \mathcal{H}_2 and \mathcal{H}_∞ filters.

Copyright © 2023 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: minimal regret, observer design, state estimation

1. INTRODUCTION

A key contribution on the problem of estimating the state and predicting future outputs of a linear system was provided in the seminal paper by Kalman (1960). This work provided the basis for \mathcal{H}_2 (i.e., least squares error) state estimation, which had a tremendous success in aerospace, and in many other applications later. However, for several industrial processes, the performance of \mathcal{H}_2 filtering can be questionable due to the poor modeling of the systems involved (Simon, 2006). In several industrial plants, the process noise (also called disturbance) often includes different forms of uncertainty, therefore rendering the Gaussian assumption flawed. This has motivated the industry to also turn towards estimators optimized for worst case scenarios (\mathcal{H}_∞), as described in the pioneering work of Zames (1981). Another approach introduced in Luenberger (1971) consists in stabilizing the error dynamics, such that errors in the initial state estimate are eventually rejected by the estimator. This approach is much simpler to implement but lacks optimal performance or robustness. Today, the \mathcal{H}_2 , \mathcal{H}_∞ , and Luenberger observers are the dominating paradigms in industrial applications (Simon, 2006).

The \mathcal{H}_∞ filter has proven to work well in practice, and it addresses the reliance of \mathcal{H}_2 filters on Gaussian noise and known system dynamics. However, doing so, it loses the optimality properties of \mathcal{H}_2 methods. The tradeoff between optimal performance and robustness can be optimized us-

ing mixed $\mathcal{H}_2/\mathcal{H}_\infty$ methods (Bernstein and Haddad, 1989; Yaz et al., 1992). However, the relative weights of both components are not tuned according to the disturbance patterns affecting the system, but need to be set *a-priori*. In order to optimize the benefits and drawbacks of each method automatically, one can use a measure of regret compared to an observer that would know the future in advance.

Regret is the difference of performance between a practical observer and its ideal counterpart, which would anticipate all possible uncertainties a system will encounter. This ideal noncausal observer (often called clairvoyant, as it sees the future) is not implementable, but can give useful information on how to adapt to various disturbance profiles. The notion of regret has first been developed by economists (Loomes and Sugden, 1982), for studying human interactions that are highly subject to changes and modelling errors. This motivation is also relevant for industrial processes, hence fostering the interest for the notion of regret measures in engineering. In the estimation setting, because the clairvoyant observer adapts to specific error patterns (e.g. constant, drifting, ...), it can compensate for errors and simplifications made during modeling. Hence, minimizing the worst case difference with this observer allows reducing the conservativeness of \mathcal{H}_∞ filters, while keeping the greater robustness than \mathcal{H}_2 .

In recent years, researchers in the field of estimation have attempted to design new observers with reduced regret. The idea of finding a versatile estimator that adapts to both Gaussian and worst-case noise through regret appears in Vanli et al. (2015). Nevertheless, this

¹ This research is supported by the Swiss National Science Foundation under the NCCR Automation (grant agreement 51NF40_180545).

and most other works aim to improve regret bounds for given estimation methods. Among the works using regret measures for observer design, Ouhamma et al. (2021) provide an improvement to recursive ridge regression, and Gharbi et al. (2021) present an Moving Horizon Estimation (MHE) method, which is suboptimal in terms of the cost function, but has a lower regret bound than the optimal estimator. The work of Agrawal et al. (2021) makes a step towards regret-optimal estimation, as it gives an iterative algorithm that improves the regret bound at each iteration. However, the solution presented by the authors is iterative and only optimizes regret asymptotically. So far, only Sabag and Hassibi (2021) and Goel and Hassibi (2021) have directly computed a state observer that minimizes regret. However, their method is only γ -suboptimal and requires specific conditions for a solution to exist. Didier et al. (2022) and Martin et al. (2022) address this limitation in the context of regret-based controller design by formulating the minimization of regret in a convex way, and provide a simple way to introduce constraints. However, their approach cannot be applied in a straightforward way to state estimation problems, due to the presence of measurement noise.

In this paper, we provide a new convex formulation of the regret minimization problem with both process and measurement noise, which allows computing the regret-optimal observer using standard optimization tools. This formulation relies on a new parametrization adapting System Level Synthesis (SLS) (Anderson et al., 2019) for the state estimation problem. We test this observer on several systems, and show that it can improve the state estimation error compared to Kalman/ \mathcal{H}_2 and \mathcal{H}_∞ filters for various deterministic and stochastic disturbance profiles. Similar formulations for controller design show that it is possible to introduce constraints, although this aspect will not be discussed in this paper.

The paper proceeds as follows. Section 2 states the problem and provides the necessary tools and assumptions. Section 3 shows how to compute the \mathcal{H}_2 , \mathcal{H}_∞ , clairvoyant, and regret-optimal observers in the SLS framework. Section 4 discusses the numerical results, and Section 5 concludes the paper.

1.1 Preliminaries and notations

Time indices are denoted by subscripts (e.g., x_t), and boldface letters denote the stacked vectors at all times, e.g. $\mathbf{x} = [x_0^\top, \dots, x_T^\top]^\top$ for a window $[0, T]$. Similarly, calligraphic letters denote linear operators applying to such stacked vectors (e.g. \mathcal{A} such that $\mathbf{y} = \mathcal{A}\mathbf{x}$). The ℓ_2 -norm of a vector is denoted by $\|\cdot\|_2$, which also denotes the spectral norm of a matrix (its largest eigenvalue). The Frobenius norm of a matrix is denoted by $\|\cdot\|_F$. Blackboard bold letters denote sets (e.g. \mathbb{R} for all real numbers), except \mathbb{E} that denotes the expectation of a random variable.

2. SYSTEM MODEL AND PROBLEM STATEMENT

2.1 State space equations

We consider a linear time-varying system given by the equations

$$x_{t+1} = A_t x_t + B_t u_t + w_t, \quad (1a)$$

$$y_t = C_t x_t + v_t, \quad (1b)$$

where $x_t \in \mathbb{R}^n$ is the state, $u_t \in \mathbb{R}^p$ the input, and $y_t \in \mathbb{R}^m$ the output for each time instant $t = 0, \dots, T$. The system dynamics are characterized by the time-varying parameters $A_t \in \mathbb{R}^{n \times n}$, $B_t \in \mathbb{R}^{n \times p}$, and $C_t \in \mathbb{R}^{m \times n}$, and subject to disturbance w_t and measurement noise v_t .

The only restriction that we make on w_t and v_t are the following.

Assumption 1 (Disturbance boundedness). *There exists non-empty full-dimensional² ellipsoids $\mathbf{v} = \{v : \|H_v v\|_2 \leq 1\}$ and $\mathbf{w} = \{w : \|H_w w\|_2 \leq 1\}$ such that $v_t \in \mathbf{v}$ and $w_t \in \mathbf{w}$ for all $t = 0, \dots, T$, respectively.*

Assumption 1 is common in the field of \mathcal{H}_∞ filtering for linear systems. However, it can lead to poor performance as \mathcal{H}_∞ state observers and filters are designed to mitigate worst case disturbances, even if they are not encountered during the system operation. This paper aims to provide a more versatile approach that adapts to the encountered disturbance patterns for better performance.

2.2 Estimation problem

We are interested in computing the state predictions \hat{x}_{t+1} of x_{t+1} for $t = 0, \dots, T$, when only knowing y_t and u_t but not $v_t \in \mathbf{v}$ and $w_t \in \mathbf{w}$, and minimizing the total cost

$$\text{cost}(\hat{x}_1 - x_1, \dots, \hat{x}_{T+1} - x_{T+1}) = \sum_{t=1}^{T+1} \ell_t(\hat{x}_t - x_t), \quad (2)$$

where ℓ_t is a given loss function (e.g. least squares). To do so, we design a dynamic Luenberger-type observer defined by the gains $L_{\tau|t} \in \mathbb{R}^{n \times m}$ for $\tau = 0, \dots, t$, and given by

$$\hat{x}_{t+1} = A_t \hat{x}_t + B_t u_t - \sum_{\tau=0}^t L_{\tau|t} (C_\tau \hat{x}_\tau - y_\tau). \quad (3)$$

The gains are chosen to stabilize the prediction error $e_t = \hat{x}_t - x_t$ following the dynamics

$$e_{t+1} = A_t e_t - w_t - \sum_{\tau=0}^t L_{\tau|t} (C_\tau e_\tau - v_\tau), \quad (4)$$

while minimizing the cost (2).

Remark. *We highlight that the error on the initial state is embedded without loss of generality in the first disturbance, i.e. $w_0 = A_0(\hat{x}_0 - x_0)$ and $e_0 = 0$.*

Substituting (4) in the cost (2) gives the final problem

$$\arg \min_{\mathbf{v}, \mathbf{w}, \{L_{\tau|t} : t \in [0, T], \tau \in [0, t]\}, \mathbf{e}} \text{cost}(\mathbf{e}), \quad (5)$$

$$\text{s.t. } e_{t+1} = A_t e_t - w_t - \sum_{\tau=0}^t L_{\tau|t} (C_\tau e_\tau - v_\tau), \\ e_0 = 0,$$

where

² The ellipsoids \mathbf{v} and \mathbf{w} contain a m - and n -dimensional ball, respectively.

$$\mathbf{v} = \begin{bmatrix} v_0 \\ v_1 \\ \vdots \\ v_T \end{bmatrix}, \mathbf{w} = - \begin{bmatrix} A_0 e_0 \\ w_1 \\ \vdots \\ w_T \end{bmatrix}, \mathbf{e} = \begin{bmatrix} e_1 \\ \vdots \\ e_T \\ e_{T+1} \end{bmatrix}. \quad (6)$$

Remark. While in the sequel we focus on the finite horizon estimation problem (5), we notice that it provides the building block for defining MHE methods that are applicable for $t = 0, \dots, \infty$ (Alessandri et al., 2008; Ferrari-Trecate et al., 2002; Farina et al., 2010).

We also define the set \mathbb{V} such that each element \mathbf{v}_i of $\mathbf{v} \in \mathbb{V}$ is in \mathbf{v} , and similarly, the set \mathbb{W} . From Assumption 1, it follows that \mathbb{V} and \mathbb{W} are non-empty full-dimensional ellipsoids, and can therefore be written as $\mathbb{V} = \{\mathbf{v} \mid \|\mathcal{H}_v \mathbf{v}\|_2 \leq 1\}$ and $\mathbb{W} = \{\mathbf{w} \mid \|\mathcal{H}_w \mathbf{w}\|_2 \leq 1\}$, where \mathcal{H}_v and \mathcal{H}_w are square invertible block-diagonal matrices.

2.3 System level synthesis

In order to formulate a convex problem to compute all $L_{\tau|t}$ minimizing (2), we introduce a parametrization similar to SLS (Anderson et al., 2019), which was proposed for controller design. We will show the necessary modifications for addressing observer design in the this section.

Let \mathcal{Z} be the block-downshift operator, namely a matrix with identity matrices along its first block sub-diagonal and zeros elsewhere, and let $\mathcal{A} = \text{blkdiag}(A_0, \dots, A_T, 0_{n \times n})$ and $\mathcal{C} = \text{blkdiag}(0_{m \times n}, C_0, \dots, C_T)$. For $t = 0, \dots, T$, the model (5) is equivalent to

$$\mathbf{e} = \mathcal{Z}\mathcal{A}\mathbf{e} - \mathcal{L}\mathcal{C}\mathcal{Z}\mathbf{e} + \mathcal{L}\mathbf{v} + \mathbf{w}, \quad (7)$$

where

$$\mathcal{L} = \begin{pmatrix} L_{0|0} & 0 & \cdots & 0 \\ L_{0|1} & L_{1|1} & \ddots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ L_{0|T} & L_{1|T} & \cdots & L_{T|T} \end{pmatrix}. \quad (8)$$

Remark. With this notation, a standard Luenberger observer is modeled as a block-diagonal matrix with the same block $L_{t|t} = L$ repeated in all diagonal elements. A Kalman filter is modeled as a block-diagonal \mathcal{L} , in which each block $L_{t|t}$ is a pre-computed Kalman gain for time t . The block lower-triangular formulation adopted in Sabag and Hassibi (2021) and in this paper is more general and allows one to consider noise patterns with dependencies between time instants.

The SLS model is defined by the matrices

$$\Phi_w = (I - \mathcal{Z}(\mathcal{A} - \mathcal{L}\mathcal{C}))^{-1}, \quad (9a)$$

$$\Phi_v = (I - \mathcal{Z}(\mathcal{A} - \mathcal{L}\mathcal{C}))^{-1}\mathcal{L}, \quad (9b)$$

where, by construction, Φ_w and Φ_v satisfy (7), i.e.,

$$\mathbf{e} = \Phi_w \mathbf{w} + \Phi_v \mathbf{v}. \quad (10)$$

The lower block-triangular matrices Φ_v and Φ_w map the noise components \mathbf{v} and \mathbf{w} to their respective shares of the estimation error \mathbf{e} . The maps Φ_v and Φ_w will be called prediction error maps in the rest of the paper. In contrast, when SLS is used for controller design, the closed-loop state and input maps Φ_x and Φ_u give the state and input trajectories depending solely on \mathbf{w} Anderson et al. (2019).

Theorem 2. Consider the dynamics (7), there exist a unique matrix \mathcal{L} satisfying (9a) and (9b) if and only if the following achievability condition is satisfied

$$\Phi_w(I - \mathcal{Z}\mathcal{A}) + \Phi_v\mathcal{C}\mathcal{Z} = I. \quad (11)$$

Moreover $\mathcal{L} = \Phi_w^{-1}\Phi_v$.

Proof. From (9), one has that $\Phi_v = \Phi_w\mathcal{L}$, and (7) can be rewritten as

$$\Phi_w(\mathbf{w} + \mathcal{L}\mathbf{v}) = \mathcal{Z}\mathcal{A}\Phi_w(\mathbf{w} + \mathcal{L}\mathbf{v}) - \mathcal{L}\mathcal{C}\mathcal{Z}\Phi_w(\mathbf{w} + \mathcal{L}\mathbf{v}) + (\mathbf{w} + \mathcal{L}\mathbf{v}), \quad (12)$$

which is satisfied for all $\mathbf{w} \in \mathbb{W}$ and $\mathbf{v} \in \mathbb{V}$ if and only if

$$\Phi_w = \mathcal{Z}\mathcal{A}\Phi_w - \mathcal{L}\mathcal{C}\mathcal{Z}\Phi_w + I. \quad (13)$$

In order to prove the theorem, we only need to prove that Φ_w is invertible and that (11) is equivalent to (13) with $\mathcal{L} = \Phi_w^{-1}\Phi_v$.

First, $\Phi_w\mathcal{Z}\mathcal{A}$ is lower block-triangular with a zero block-diagonal because it is the product of (i) a lower block-triangular (ii) a block downshift operator, which zeroes the block-diagonal of $\Phi_w\mathcal{Z}$, and (iii) a block diagonal matrix which does not change the sparsity pattern. Moreover, $\Phi_v\mathcal{C}\mathcal{Z}$ follows the same pattern because Φ_v and \mathcal{C} are also lower block-triangular and block-diagonal, respectively. Hence, the diagonal of Φ_w must be equal to the identity on the right hand side of (11), which means that all its eigenvalues are equal to one if (11) is satisfied. This proves that Φ_w is invertible.

Second, by post-multiplying and pre-multiplying (13) by Φ_w^{-1} and Φ_w , respectively, we obtain

$$\Phi_w = \Phi_w\mathcal{Z}\mathcal{A} - \Phi_w\mathcal{L}\mathcal{C}\mathcal{Z} + I. \quad (14)$$

To obtain (11) from (14), one must replace \mathcal{L} by $\Phi_w^{-1}\Phi_v$ and rearrange the terms, which proves the theorem. \square

Transforming the optimal estimation problem (5) into the SLS framework gives the following convex problem.

$$\arg \min_{\Phi_w, \Phi_v, \mathbf{w}, \mathbf{v}} \text{cost}(\Phi_w \mathbf{w} + \Phi_v \mathbf{v}), \quad (15a)$$

$$\text{s.t. } \Phi_w(I - \mathcal{Z}\mathcal{A}) + \Phi_v\mathcal{C}\mathcal{Z} = I, \quad (15b)$$

$$\Phi_w, \Phi_v \text{ lower block-triangular.} \quad (15c)$$

The constraint (15c) is added for the prediction error maps Φ_v and Φ_w to be causal, i.e. e_t only depends on v_0, \dots, v_t and w_0, \dots, w_t but not on future noise. Hence, the observer $\mathcal{L} = \Phi_w^{-1}\Phi_v$ is also causal. At time t , \mathcal{L} predicts future states x_{t+1}, \dots, x_T only based on past and present measurements y_1, \dots, y_t .

3. MINIMAL REGRET ESTIMATION

We will first explain how to derive two standard causal benchmarks using the SLS, namely the optimal \mathcal{H}_2 and \mathcal{H}_∞ observers.

3.1 Standard benchmarks

\mathcal{H}_2 filtering: \mathcal{H}_2 observers such as the Kalman filter are designed to minimize the mean square error when disturbances and noise are Gaussian. In mathematical terms, this means that the filter gains solve the problem (5) with

$$\text{cost}_2(\mathbf{e}(\mathbf{v}, \mathbf{w}, \mathcal{L})) = \frac{1}{T} \sum_0^T \mathbb{E}_{[v_t, w_t] \sim \mathcal{N}(0, \Sigma_t)} [\|e_t\|_2^2], \quad (16)$$

where $\Sigma_t = \text{blkdiag}(\Sigma_{v,t}, \Sigma_{w,t})$ is the covariance of the noise v_t and the disturbance w_t , assumed Gaussian and uncorrelated in time. Similar to Anderson et al. (2019) and by using (10), one has that (16) can be rewritten in terms of Φ_w and Φ_v as

$$\text{cost}_2(\Phi_v, \Phi_w) = \frac{1}{T} \left\| \left[\Phi_v, \Phi_w \right] \begin{bmatrix} \Sigma_v^{\frac{1}{2}} & 0 \\ 0 & \Sigma_w^{\frac{1}{2}} \end{bmatrix} \right\|_F^2. \quad (17)$$

where $\Sigma_v = \text{blkdiag}(\Sigma_{v,0}, \dots, \Sigma_{v,T})$ and $\Sigma_w = \text{blkdiag}(\Sigma_{w,0}, \dots, \Sigma_{w,T})$. Theorem 2 shows that one can replace the problem (5) by (15), which finally gives.

$$\{\Phi_{v,2}, \Phi_{w,2}\} = \arg \min_{\Phi_v, \Phi_w} \left\| \left[\Phi_v, \Phi_w \right] \begin{bmatrix} \Sigma_v^{\frac{1}{2}} & 0 \\ 0 & \Sigma_w^{\frac{1}{2}} \end{bmatrix} \right\|_F^2, \quad (18)$$

s.t. (15b), (15c).

The optimal \mathcal{H}_2 observer is the Kalman filter, which is such that \mathcal{L} is block-diagonal. This means that, even if $\Phi_{w,2}$ and $\Phi_{v,2}$ can be dense lower block-triangular matrices, $\mathcal{L}_2 = \Phi_{w,2}^{-1} \Phi_{v,2}$ is block diagonal. This is not included as a hard constraint in (18). The block-diagonality of \mathcal{L}_2 appears naturally when w_t and v_t are not correlated between different time instants, because the Kalman filter is proven to be optimal (Kalman, 1960). The problem (18) is a relaxation of general \mathcal{H}_2 observers, which generalizes the Kalman filter and can therefore improve performance if the noise is correlated between different instants. This means that it gives a suitable benchmark for comparison. It will also be used later to formulate regret measures.

\mathcal{H}_∞ filtering: Similar to the \mathcal{H}_2 observer, Simon (2006) describes the \mathcal{H}_∞ observer as a solution to (5) with the cost

$$\begin{aligned} \text{cost}_\infty(e(\mathbf{v}, \mathbf{w}, \mathcal{L})) &= \frac{1}{T} \sum_{t=0}^T \max_{\substack{v_t \in \mathbf{v} \\ w_t \in \mathbf{w}}} \|e_t\|_2^2, \quad (19) \\ &= \frac{1}{T} \max_{\substack{v \|\mathcal{H}_v v\|_2 \leq 1 \\ w \|\mathcal{H}_w w\|_2 \leq 1}} \|e\|_2^2. \end{aligned}$$

By using the change of variable $\mathbf{v}_n = \mathcal{H}_v \mathbf{v}$ and $\mathbf{w}_n = \mathcal{H}_w \mathbf{w}$, one can rewrite the cost in terms of Φ_v and Φ_w as

$$\text{cost}_\infty(\Phi_v, \Phi_w) = \frac{1}{T} \max_{\substack{v_n \|\mathbf{v}_n\|_2 \leq 1 \\ w_n \|\mathbf{w}_n\|_2 \leq 1}} \left\| \left[\Phi_v, \Phi_w \right] \begin{bmatrix} \mathcal{H}_v^{-1} & 0 \\ 0 & \mathcal{H}_w^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{v}_n \\ \mathbf{w}_n \end{bmatrix} \right\|_2^2, \quad (20)$$

where \mathcal{H}_v and \mathcal{H}_w are invertible because of Assumption 1. According to the definition of spectral norms, and as proven by Anderson et al. (2019), (19) is equal to

$$\text{cost}_\infty(\Phi_v, \Phi_w) = \frac{1}{T} \left\| \left[\Phi_v, \Phi_w \right] \begin{bmatrix} \mathcal{H}_v^{-1} & 0 \\ 0 & \mathcal{H}_w^{-1} \end{bmatrix} \right\|_2^2. \quad (21)$$

Hence, the \mathcal{H}_∞ observer is $\mathcal{L}_\infty = \Phi_{w,\infty}^{-1} \Phi_{v,\infty}$ with

$$\{\Phi_{v,\infty}, \Phi_{w,\infty}\} \in \arg \min_{\Phi_v, \Phi_w} \left\| \left[\Phi_v, \Phi_w \right] \begin{bmatrix} \mathcal{H}_v^{-1} & 0 \\ 0 & \mathcal{H}_w^{-1} \end{bmatrix} \right\|_2^2, \quad (22)$$

s.t. (15b), (15c).

Note that the problem (22) may admit multiple solutions because the cost only affects the largest eigenvalue (i.e. the spectral norm) but not the other eigenvalues, which are free to vary.

3.2 Clairvoyant estimator

A non-causal observer does away with the causal constraint (15c), estimating states based on both past and future measurements. This means that the optimal clairvoyant observer is given by a gain matrix $\mathcal{L}_{nc} \in \mathbb{R}^{nT \times mT}$, which is full and such that the resulting estimation error minimizes the cost (2). While not implementable, \mathcal{L}_{nc} can give useful information about how an observer should react to noise. In what follows, we show how to explicitly compute \mathcal{L}_{nc} using the same SLS framework as in Section 3.1.

Lemma 3. For quadratic, positive definite loss functions $\ell_t = e_t^\top Q_t e_t$ with $0 \prec Q_t \in \mathbb{R}^{n \times n}$, the prediction error maps $\Phi_{v,nc}, \Phi_{w,nc}$ of the clairvoyant observer $\mathcal{L}_{nc} = \Phi_{w,nc}^{-1} \Phi_{v,nc}$ are given in the SLS form by

$$\{\Phi_{v,nc}, \Phi_{w,nc}\} = \arg \min_{\Phi_v, \Phi_w} \left\| \mathcal{Q} \left[\Phi_v, \Phi_w \right] \begin{bmatrix} \Sigma_v^{\frac{1}{2}} & 0 \\ 0 & \Sigma_w^{\frac{1}{2}} \end{bmatrix} \right\|_F^2, \quad (23a)$$

s.t. (15b),

$$\in \arg \min_{\Phi_v, \Phi_w} \left\| \mathcal{Q} \left[\Phi_v, \Phi_w \right] \begin{bmatrix} \mathcal{H}_v^{-1} & 0 \\ 0 & \mathcal{H}_w^{-1} \end{bmatrix} \right\|_2^2, \quad (23b)$$

s.t. (15b),

where $\mathcal{Q} = \text{blkdiag}(Q_0^{\frac{1}{2}}, \dots, Q_{T+1}^{\frac{1}{2}})$. The corresponding optimal cost is

$$\min_{\mathcal{L}} \text{cost}(e(\mathbf{v}, \mathbf{w}, \mathcal{L})) = \left\| \mathcal{Q} \left[\Phi_{v,nc}, \Phi_{w,nc} \right] \begin{bmatrix} \mathbf{v} \\ \mathbf{w} \end{bmatrix} \right\|_2^2. \quad (24)$$

Proof. Using (11), one can always obtain $\{\Phi_{v,nc}, \Phi_{w,nc}\}$ from \mathcal{L}_{nc} using

$$\Phi_{w,nc} = ((I - \mathcal{Z}\mathcal{A}) - \mathcal{L}_{nc}\mathcal{C}\mathcal{Z})^{-1}, \quad (25a)$$

$$\Phi_{v,nc} = \Phi_{w,nc} \mathcal{L}_{nc}. \quad (25b)$$

The optimal non-causal estimator \mathcal{L}_{nc} is a known result, which can be found in Hassibi et al. (1999). It is unique and provides estimates that minimize both \mathcal{H}_2 and \mathcal{H}_∞ costs. Using (25), this means that the corresponding pair of transfer functions $\{\Phi_{v,nc}, \Phi_{w,nc}\}$ exists, is unique, and solves both problems in (23).

The problem (23a) is strictly convex because, by definition, $\mathcal{Q} \succ 0$, $\Sigma_v \succ 0$, and $\Sigma_w \succ 0$. Hence, its solution is unique. Since a solution solving both problems exists, and the solution to (23a) is unique, we conclude that one only needs to solve (23a). Finally, (24) follows from (10). \square

3.3 Regret optimal observer

In this paper, we aim to minimize the worst-case dynamic regret, defined as

$$\text{regret}(\mathcal{L}, T) = \max_{\substack{v \in \mathbf{V} \\ w \in \mathbf{W}}} (\text{cost}(e(\mathbf{v}, \mathbf{w}, \mathcal{L})) - \underbrace{\min_{\mathcal{K}} \text{cost}(e(\mathbf{v}, \mathbf{w}, \mathcal{K}))}_{=\text{cost}(e(\mathbf{v}, \mathbf{w}, \mathcal{L}_{nc}))}). \quad (26)$$

With only the first term, (26) reduces to the \mathcal{H}_∞ cost. The second term provides a system-dependent baseline, which represents the best performance achievable by the system. From Lemma 3, this baseline is provided by the clairvoyant estimator. For any noise pattern, if the clairvoyant observer gives a perfect estimate, the difference between the

two costs in (26) is equal to the cost of \mathcal{L} only. However, this difference decreases if the clairvoyant observer also generates some error. This allows the regret measure to reduce the importance of situations in which some error is unavoidable, leading to a less conservative estimator than \mathcal{H}_∞ . Moreover, since the clairvoyant estimator is \mathcal{H}_2 optimal (see (23a)), minimizing the regret amounts to minimizing the suboptimality that a robust observer \mathcal{L} induces.

Directly computing the minimal regret observer $\mathcal{L}_r = \arg \min_{\mathcal{L}} \text{regret}(\mathcal{L}, T)$ would be very complex due to the nested optimization problems, as well as the non-convexity of (26) in terms of \mathcal{L} . In the sequel, we compute \mathcal{L}_r from Φ_v and Φ_w , by adopting the approach used in Martin et al. (2022) for designing regret-based controllers.

Theorem 4. For quadratic costs and if either \mathcal{L}_{nc} or $\Phi_{v,nc}$ and $\Phi_{w,nc}$ are known, the minimal regret observer is given by $\mathcal{L}_r = \Phi_{w,r}^{-1} \Phi_{v,r}$, where

$$\{\Phi_{v,r}, \Phi_{w,r}\} = \arg \min_{\Phi_v, \Phi_w} \text{regret}(\Phi_w^{-1} \Phi_v, T), \quad (27)$$

$$\begin{aligned} \text{s.t. } & \Phi_w^{-1} \Phi_v \text{ lower block-triangular,} \\ & = \arg \min_{\Phi_v, \Phi_w} \min_{\lambda} \lambda, \end{aligned} \quad (28)$$

s.t. (15b), (15c), $\lambda > 0$,

$$\begin{bmatrix} I & Q^{\frac{1}{2}} [\Phi_v H_v^{-1}, \Phi_w H_w^{-1}] \\ [\Phi_v H_v^{-1}, \Phi_w H_w^{-1}]^\top Q^{\frac{1}{2}} & \lambda I + \mathcal{J}_{nc} \end{bmatrix} \succeq 0,$$

$$\text{where } \mathcal{J}_{nc} = \begin{bmatrix} H_v^{-\top} \Phi_{v,nc}^\top \\ H_w^{-\top} \Phi_{w,nc}^\top \end{bmatrix} Q \begin{bmatrix} H_v^{-\top} \Phi_{v,nc}^\top \\ H_w^{-\top} \Phi_{w,nc}^\top \end{bmatrix}^\top.$$

Proof. According to Lemma 3, $\Phi_{v,nc}$ and $\Phi_{w,nc}$ give the second term of (26), and can be computed from \mathcal{L}_{nc} . Hence (26) can be rewritten as

$$\text{regret}(\Phi_w^{-1} \Phi_v, T) = \quad (29)$$

$$\max_{\substack{v \in \mathbb{V} \\ w \in \mathbb{W}}} \begin{bmatrix} v \\ w \end{bmatrix}^\top \left(\begin{bmatrix} \Phi_v^\top \\ \Phi_w^\top \end{bmatrix} Q \begin{bmatrix} \Phi_v^\top \\ \Phi_w^\top \end{bmatrix}^\top - \begin{bmatrix} \Phi_{v,nc}^\top \\ \Phi_{w,nc}^\top \end{bmatrix} Q \begin{bmatrix} \Phi_{v,nc}^\top \\ \Phi_{w,nc}^\top \end{bmatrix}^\top \right) \begin{bmatrix} v \\ w \end{bmatrix}.$$

With the same change of variable as in (20), (29) is equivalent to

$$\text{regret}(\Phi_w^{-1} \Phi_v, T) = \quad (30)$$

$$\max_{\substack{\|v_n\| \leq 1 \\ \|w_n\| \leq 1}} \begin{bmatrix} v_n \\ w_n \end{bmatrix}^\top \underbrace{\left(\begin{bmatrix} H_v^{-\top} \Phi_v^\top \\ H_w^{-\top} \Phi_w^\top \end{bmatrix} Q \begin{bmatrix} H_v^{-\top} \Phi_v^\top \\ H_w^{-\top} \Phi_w^\top \end{bmatrix}^\top - \mathcal{J}_{nc} \right)}_{\mathcal{M}} \begin{bmatrix} v_n \\ w_n \end{bmatrix}.$$

Finding the solution of (27) amounts to minimizing the largest eigenvalue of the matrix \mathcal{M} . From Boyd et al. (1994), such an eigenvalue problem can be solved by writing \mathcal{M} as a Schur complement. The solution is given by the semi-definite program (28), which proves the theorem. \square

Boyd et al. (1994) also show that the problem (28) is convex and can be solved with any convex optimization methods, which makes it very convenient in practice. Note that the γ -suboptimal estimators presented in Sabag and Hassibi (2021) are cheaper to compute, but do not provide an exact solution. For time-invariant systems, the problems (23a) and (28) can be solved once offline, yielding the observer $\mathcal{L}_r = \Phi_{w,r}^{-1} \Phi_{v,r}$, which can be applied using

(3). For time-varying systems, both problems must be solved every time the parameters change.

4. NUMERICAL EXPERIMENTS

To show the benefits of minimal regret estimation, we will simulate the trajectories of the benchmark system "NN4" of order 4 from Complib (Leibfritz, 2004) under nine different disturbance patterns. The script is available in (Brouillon, 2022). In order to assess the performance in practice, we also provide the average performance for the first three fifth order systems in Complib "AC1", "AC2", and "AC3", which model aircrafts. For convenience, we consider only LTI systems, even though our framework also encompasses LTV systems, see (1). All systems are discretized with $T_s = 5ms$ and the prediction horizon is $T = 10$. First, the optimal error maps Φ_v and Φ_w are computed for \mathcal{H}_2 , \mathcal{H}_∞ and minimal regret (denoted by \mathcal{R}) methods. Then, the performance is evaluated using (10) and the cost $\|e\|_2^2$ for each disturbance pattern. Each element of the noises v and w follow independent uniform or Gaussian distributions in three cases, and deterministic sequences (e.g. a sine wave $w_t = [\sin(t), \dots, \sin(t)]^\top$) in the other six cases.

v, w	\mathcal{H}_2	\mathcal{H}_∞	\mathcal{R}	\mathcal{H}_2	\mathcal{H}_∞	\mathcal{R}
$\mathcal{N}(0, 1)$	1	0.83%	11%	1	1.44%	9.70%
$\mathcal{U}_{[0.5, 1]}$	7.60%	8.34%	1	6.33%	4.37%	1
$\mathcal{U}_{[0, 1]}$	6.52%	7.23%	1	5.47%	3.71%	1
1	7.83%	8.55%	1	6.44%	4.46%	1
sin	7.15%	7.64%	1	5.81%	4.09%	1
sawtooth	7.72%	8.51%	1	6.34%	4.35%	1
step	6.90%	8.35%	1	6.68%	4.27%	1
stairs	5.03%	4.37%	1	4.02%	2.71%	1
worst	0.24%	1	0.24%	0.26%	1	0.41%

(a) NN4

(b) average of AC1-3

Table 1. Relative difference in average estimation costs with $Q = I$, for nine difference patterns. The performance for the stochastic patterns (i.e. the first three) is averaged over 1000 realizations. The best estimator is marked by **1** and sets the baseline for the relative percentage increase.

Table 1 shows that, as expected, \mathcal{H}_2 and \mathcal{H}_∞ observers give the best performance for Gaussian and worst-case disturbances, respectively. However, observe that, unlike in the work of Sabag and Hassibi (2021), the error of the \mathcal{R} observer is not between that of \mathcal{H}_2 and \mathcal{H}_∞ . For all other disturbance patterns, the \mathcal{R} observer achieves a significant improvement compared to the other two.

Note that the difference between the costs of \mathcal{H}_2 and \mathcal{H}_∞ observers is relatively small compared to the difference between both costs and the one of the \mathcal{R} observer. It shows that, for some systems, this new observer can provide a radically different alternative to the two state-of-the-art state estimation methods.

5. CONCLUSIONS

By using a new SLS parametrization, we provided a tractable formulation of the regret minimization problem, allowing one to compute the minimal regret observer

easily. In multiple experiments, we showed that such an observer can be effective in reducing state estimation error when the disturbance and measurement noises are not Gaussian or adversarial.

Even though linear systems are well studied, the effect of noise and disturbance distributions on estimation error is not yet fully understood. Namely, it is not trivial to find the boundaries between the sets of disturbance profiles where \mathcal{H}_2 , \mathcal{H}_∞ , and minimal-regret observers are optimal. Finally, the observers proposed in the paper can provide the foundations for designing minimal-regret MHE methods for nonlinear and hybrid systems.

REFERENCES

- Agrawal, P., Chen, J., and Jiang, N. (2021). Improved worst-case regret bounds for randomized least-squares value iteration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 6566–6573.
- Alessandri, A., Baglietto, M., and Battistelli, G. (2008). Moving-horizon state estimation for nonlinear discrete-time systems: New stability results and approximation schemes. *Automatica*, 44(7), 1753–1765.
- Anderson, J., Doyle, J.C., Low, S.H., and Matni, N. (2019). System level synthesis. *Annual Reviews in Control*, 47, 364–393.
- Bernstein, D.S. and Haddad, W.M. (1989). Steady-state kalman filtering with an H_∞ error bound. *Systems and Control Letters*, 12(1), 9–16.
- Boyd, S., El Ghaoui, L., Feron, E., and Balakrishnan, V. (1994). *Linear matrix inequalities in system and control theory*. SIAM.
- Brouillon, J.S. (2022). Matlab script generating the simulation results. <https://zenodo.org/record/7819342>, doi:10.5281/zenodo.5725215.
- Didier, A., Sieber, J., and Zeilinger, M.N. (2022). A system level approach to regret optimal control. *IEEE Control Systems Letters*.
- Farina, M., Ferrari-Trecate, G., and Scattolini, R. (2010). Distributed moving horizon estimation for linear constrained systems. *IEEE Transactions on Automatic Control*, 55(11), 2462–2475.
- Ferrari-Trecate, G., Mignone, D., and Morari, M. (2002). Moving horizon estimation for hybrid systems. *IEEE transactions on automatic control*, 47(10), 1663–1676.
- Gharbi, M., Gharesifard, B., and Ebenbauer, C. (2021). Anytime proximity moving horizon estimation: Stability and regret for nonlinear systems. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 728–735. IEEE.
- Goel, G. and Hassibi, B. (2021). Regret-optimal estimation and control. *arXiv preprint arXiv:2106.12097*.
- Hassibi, B., Sayed, A.H., and Kailath, T. (1999). *Indefinite-Quadratic estimation and control: a unified approach to \mathcal{H}_2 and \mathcal{H}_∞ theories*. SIAM.
- Kalman, R.E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82(Series D), 35–45.
- Leibfritz, F. (2004). COMPLeIB, CONstraint Matrix-optimization Problem LIBrary—a collection of test examples for nonlinear semidefinite programs, control system design and related problems. *Dept. Math., Univ. Trier, Trier, Germany, Tech. Rep.*, 2004.
- Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The economic journal*, 92(368), 805–824.
- Luenberger, D. (1971). An introduction to observers. *IEEE Transactions on automatic control*, 16(6), 596–602.
- Martin, A., Frieri, L., Dörfler, F., Lygeros, J., and Ferrari-Trecate, G. (2022). Safe control with minimal regret. In *Learning for Dynamics and Control Conference*, 726–738. PMLR.
- Ouhamma, R., Maillard, O.A., and Perchet, V. (2021). Stochastic online linear regression: the forward algorithm to replace ridge. *Advances in Neural Information Processing Systems*, 34, 24430–24441.
- Sabag, O. and Hassibi, B. (2021). Regret-optimal filtering. In *International Conference on Artificial Intelligence and Statistics*, 2629–2637. PMLR.
- Simon, D. (2006). *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. John Wiley & Sons.
- Vanli, N.D., Donmez, M.A., and Kozat, S.S. (2015). Robust least squares methods under bounded data uncertainties. *Digital Signal Processing*, 36, 82–92.
- Yaz, E., Grigoriadis, K., and Skelton, R. (1992). Discrete-time H_2/H_∞ norm-bounding pole placement control and estimation. In *[1992] Proceedings of the 31st IEEE Conference on Decision and Control*, 487–492. IEEE.
- Zames, G. (1981). Feedback and optimal sensitivity: Model reference transformations, multiplicative seminorms, and approximate inverses. *IEEE Transactions on automatic control*, 26(2), 301–320.