École polytechnique de Louvain

Optimal State Estimation from Intermittent Measurements

Author: **Mary Jo El Bared**
Supervisors: **Pierre-Antoine Absil, Gianluca Bianchin**
Reader: **Donatien Hainaut**
Academic year 2024–2025
Master [120] in Mathematical Engineering

Abstract

In this work, we investigate the problem of reconstructing the initial state of a linear time-invariant (LTI) dynamical system from a finite number of noisy measurements. Specifically, we study how to optimally select sampling times to improve estimation accuracy from intermittent measurements. Our approach uses classical estimation theory and the observability Gramian to quantify the informativeness of a certain sampling schedule.

We formulate an optimization problem that minimizes the trace of the inverse of the observability Gramian, equivalently minimizing the trace of the covariance matrix of the initial state least-squares estimator.

To build intuition and interpretations, we derive closed-form expressions for low-dimensional systems. In one-dimensional systems, optimal sampling times entirely depend on system stability: early sampling is optimal for stable systems, while late sampling is preferred for unstable ones. In two-dimensional systems, we analyze stable, unstable, and mixed-stability cases, observing that while the intuition from the 1D case extends, the optimization landscape becomes more complex.

Through numerical simulations, we compare optimized and uniform sampling strategies. The results show better accuracy for the former. They also highlight how the objective balances between maximizing the independence of measurement directions (and thus the conditioning of the observability Gramian) and minimizing the variance in estimation.

This work offers theoretical tools for formulating the optimal state estimation problem from intermittent measurements. It analyzes the properties of the resulting solutions by linking the objective function to the dynamics and structure of the system. Through comparison with baseline sampling strategies, it shows the effectiveness of the optimal solutions proposed from this framework.

Acknowledgements

I would first like to express my gratitude to my supervisors Pr. Pierre-Antoine Absil and Pr. Gianluca Bianchin, for their guidance, their availability, and their support. Throughout this process I always felt encouraged, challenged, and supported.

I would also like to thank my reader Pr. Donatien Hainaut for kindly accepting to be part of my jury and to read my work.

Thank you to my support system throughout this entire journey: my family, who always stood by me even if they did not always understand what I was talking about; to my two best friends, Eléonore and Yasmine, for their constant support and all the laughter that made everything lighter. To Lucas, Arnaud, and Madeliene; thank you for the sleepless nights, the shared thoughts, and the unforgettable memories.

Finally, a special thanks goes to my high school teacher, Madame Van Lierde, who not only taught me math but also passed on her love for it, encouraging me to pursue this path.

AI Usage

In this master's thesis, Github Copilot was used for code writing, and Grammarly was used for grammar correction. ChatGPT was also use to improve sentence formulation and the design of figures. All content was reviewed critically.

Contents

Chapter 1

Introduction

Dynamical systems theory provides a fundamental mathematical framework for describing how the state of a system evolves over time under the influence of internal dynamics and external outputs. Such systems naturally occur in diverse domains, including but not limited to physics, biology, epidemiology, robotics, and economics. In many of these real-world applications, the full internal of a system cannot be directly measured; instead, only partial or noisy observations are available. This motivates the problem of *state estimation* from a limited number of measurements.

For example, in satellite tracking, ground stations collect intermittent measurements of a satellite's position to reconstruct and predict its trajectory. However, continuous measurements are infeasible due to physical, energy, or communication constraints. For such a challenge, Lee et al. [4] focus on a Kalman filter-based approach for estimating the initial state of satellites in formation flying. Another area of interest is epidemiology, where sparse population-level data must be used to estimate the number of infected individuals or predict the spread of a disease over time. In these cases, measurement opportunities can be limited, motivating research interested in the design of observers [10] and controllers that rely on them [12] considering this framework.

Foundational work on observability and state estimation was done by R.E Kalman [1]. A system is said to be observable if the current state can be uniquely determined from a finite sequence of output measurements. However, classical observability is a binary property: a system is either observable or not. In practice, especially in the presence of noise and measurement constraints, it is interesting to look at the quality of the state estimate and to consider quantitative measures of observability.

In practice, however, several constraints hinder the collection of continuous data and precise state estimation. They may be due to energy limitations, logistical issues in measurements, limited communication bandwidth, budget-related constraints, or physical constraints on the frequency of measurements. Observations are typically possible at a limited number of time points, and these time points should be chosen to make the best estimation. This leads to the key question that will drive this master thesis:

> *How can we optimally select the measurement times $t_1, dots, t_n$ in order to maximize the ability to estimate the initial state x_0 of a dynamical system*

Throughout this work, we consider this problem in the context of linear time-invariant (LTI) systems, where the system dynamics and the output equation are known and deterministic. Our goal is to determine the optimal sampling schedule; the time points at which to take measurements, that minimizes the uncertainty in the reconstructed initial state.

To address this question, in Chapter 2, we start by reviewing the existing literature and placing our work within this context. Next, we introduce the least-squares estimator for the initial state and show how it is linked to the observability Gramian. Then we formalize an optimization problem to select the optimal sampling schedule to accurately estimate the initial state of the system.

In Chapter 3, we analyze the properties of the objective function, including its gradient and optimization landscape for low-dimensional systems (1D and 2D). From these closed-forms, we show how the properties of the optimal sampling times relate to the system's stability.

Finally in Chapter 4, we extend the analysis to numerical simulations of multiple two-dimensional systems in the state space. These simulations offer an interpretation of the objective function and its role in ensuring observability and accurate state estimation.

This work combines theoretical analysis to formulate and analyze the optimization problem with numerical simulations to validate the solutions and provide insight from the perspective of dynamical systems.

Chapter 2

Theoretical Background and Problem Formulation

In this chapter, we lay the theoretical foundations necessary to formulate the main problem addressed in this work: selecting optimal sampling times to achieve the most accurate estimation of the initial state of a deterministic linear dynamical system. We begin with a review of the relevant literature on observability, state estimation, intermittent measurements, and sensor scheduling, situating our contribution within this context.

We then introduce the mathematical tools used throughout this work: we discuss linear systems, the concept of observability, the least squares observer, and the observability Gramian. We offer a statistical interpretation of the least-squares estimator's performance and highlight the connection between its covariance matrix and the observability Gramian of the system. The central idea of this chapter is that the observability Gramian provides a measure of how "well" the state of a system can be reconstructed from its outputs over time and that the conditioning of the observability Gramian is directly related to the conditioning of the covariance matrix of the least-squares estimator.

These theoretical elements lead to the formulation of the optimization problem that will drive this work: *selecting a sequence of sampling times that minimizes the trace of the inverse of the observability Gramian, equivalently minimizing the trace of the covariance matrix of the initial state estimation error.* This problem defines the principal objective of this master's thesis and will be studied throughout the chapters.

2.1 State of the art

This section reviews existing work related to state estimation and observability in dynamical systems with a focus on intermittent measurements and the scheduling problem. We also position the contributions of this work within the context of the literature.

2.1.1 State Estimation

State estimation is a popular problem in control theory. Given a dynamical system, it seeks to reconstruct the “best” estimate of the state of a dynamical system from different measurements of its outputs [9].

A foundational result of this problem is the "Kalman filter" proposed by R.E Kalman in 1960 [1], who introduced a recursive method to estimate the state of linear and Gaussian dynamical systems. However, that might not always be the case in real-world applications.

Indeed, with practical applications, measurements are often intermittent, sparse, or irregular because of communication constraints, sensor failures, or limited energy budgets, among other reasons. For example, in epidemiology, the evaluation of disease spread through testing can be limited by policy, logistics, or budget constraints and can provide partial information, resulting in discontinuous measurements in time. For satellites, on the other hand, their position could be estimated by receiving information at certain ground stations and with gaps in the measurement times that can also be related to natural constraints, budgetary constraints, or even location constraints.

Similar examples push the study of state estimation from intermittent or sparse measurements, where the timings or choice of sensors and placement play a crucial role in ensuring the observability of a system and the reconstruction of the state.

2.1.2 Observability and Intermittent Measurements

Observability refers to the ability to determine the temporal trajectories of the states of a dynamical system based on information on inputs and outputs [1]. If a system is observable, the ability to design a stable state observer is considered a sufficient and necessary condition [15].

The standard method for verifying observability in linear systems uses the Observability Matrix, which provides a binary criterion: the system is observable or not [3]. This matrix is rigorously defined later in the chapter in Equation 2.8. However, it does not provide nuance on how "well" a system can be observed and consequently estimated; it confirms whether complete state reconstruction is theoretically possible. This binary evaluation led to the need for a physical interpretation and a nuanced evaluation of the properties of system observability with the Observability Gramian [2].

The Observability Gramian, introduced and formally detailed in Section 2.6 in Equation 2.17, can be evaluated for an infinite horizon or over a certain time interval and for the cases of continuous or intermittent measurements. Recent literature has tried to push the quantification of observability properties beyond a rank measure of the observability matrix and has tried to study their degradation or preservation under different conditions, particularly in the context of optimal experiment design or state observers design.

A first example can be seen in Staiger et al. [13] where they discuss the problem of optimal sensor placement for experiment design. They do so using the observability Gramian as a tool and, more specifically, its geometric ellipsoid interpretation, also detailed in Section 2.6, to find how different sensor placements impact the observability of specific states. They then applied this method to optimize sensor placement in mechanical systems for maximal state estimation accuracy.

As another example for a second direction of the field, Wang et al. [7] show that for LTI systems, considering the initial system to be observable, irregular sampling can still preserve observability, under the condition that the sampling density is higher than a certain threshold. The work extends Shannon's sampling theorem for periodic sampling and state reconstruction to observability for general cases of sampling that can be irregular.

2.1.3 Scheduled and Intermittent Sensing for state estimation

A relevant class of problems is the one in which measurement times are considered design variables or the goal is to choose a subset of sensors at a certain timestep; this was defined in [16] as a Sensor Scheduling problem.

Often, the goal is to ensure observability by designing the sensor schedules so that the Observability Gramian is well-conditioned. Multiple scalar measures of the Observability Gramian can be used; Baggio et al. [11] developed such measures for the controllability Gramian and their physical interpretation. These measures can also be applied to the Observability Gramian by the Controllability-Observability duality. Examples of such scalar measures of the Gramian are: minimum eigenvalue, the determinant, or the trace.

However, due to combinatorial growth of the problem with respect to system dimensions and complexity, optimizing the selection and measurement schedules while they are subject to cardinality constraints is generally NP-hard as was proven by Zhang et al. [8]. Indeed, their work shows that choosing a fixed subset of sensors and design times for Kalman filtering is NP-hard even for stable systems. And they show that a greedy algorithm has good performance in solving such problems in practice.

The field explores problems for finite- and infinite-horizon cases. In the case of finite-horizon systems and observability, only a fixed number of measurements is allowed over a time interval, and the problem becomes a combinatorial optimization one. An example is studied by Vitus et al. [6] for linear dynamical systems and the optimization problem of choosing one sensor out of several at each time step to minimize the weighted trace of error covariance for each of these time steps. They then derive a search algorithm for optimal and suboptimal schedules.

As a second example, one that considers measuring times as the design variables of the optimization problem, Aspeel et al. [16] look at stochastic linear time systems, and the optimization problem is minimizing the variance of the Kalman filter over a finite time horizon with regard to a fixed total number of measurements by formulating it as a discrete combinatorial problem and solving it with different algorithms, with a special focus on the performance of the genetic algorithm.

In the case where the time horizon grows infinitely large, the interest shifts to long-term schedules. Zhang et al. [5] study this for linear Gaussian systems and show that, under mild conditions, periodic schedules with a finite period can approximate the optimal estimation cost arbitrarily well. And, for such sampling, the estimation error and stability are independent of the covariance of the initial state.

2.1.4 Contributions of this Work

This works finds itself at the intersection of observability theory and resource-aware sensing. It addresses a specific problem within the field of resource constrained estimation: **how to optimally select a limited number of sampling times to accurately estimate the initial state of a deterministic linear system subject to measurement noise**. In multiple real-world scenarios, as mentioned previously, it is not feasible to measure continuously in time, due to multiple different physical or economical constraints. Understanding how to optimally place a small number of measurement times while preserving estimation quality is therefore a question that has both theoretical and practical importance.

We contribute to the growing literature by building on the theory of estimation and observability to formulate an optimization problem that selects the optimal sampling times for accurate state estimation. Then we analyze the objective function to offer intuition and interpretability on the solutions. First, we analyze the closed form of the objective function for low-dimensional systems to offer links between it and system stability. Second, we extend this analysis by simulating the state estimation for two-dimensional systems and interpreting the solution spaces geometrically and linking the solution to system stability and system interpretation. This analysis reveals connections between the objective function, sampling strategies, system stability and dynamics, and links statistical estimation concepts with dynamical system perspectives.

2.2 Discrete-time systems and the State Estimation Problem

We consider the general form of a discrete linear time-invariant (LTI) system with the following system dynamics:

$$\begin{aligned} x[k+1] &= Ax[k] + Bu[k] \\ y[k] &= Cx[k] + Du[k] + v[k], \end{aligned} \tag{2.1}$$

where:

- $x[k] \in \mathbb{R}^n$ is the state vector at time k.
- $u[k] \in \mathbb{R}^m$ is the system input at time k.
- $y[k] \in \mathbb{R}^p$ is the output measured at time k.
- $v[k] \in \mathbb{R}^p$ is the sensor error (ie. noise) at time k.
- $A \in \mathbb{R}^{n\times n}$ is the state transition matrix.
- $B \in \mathbb{R}^{n\times m}$ is the input matrix.
- $C \in \mathbb{R}^{p\times n}$ is the output matrix.
- $D \in \mathbb{R}^{p\times m}$ is the direct transmission (or feedthrough) matrix.

In many practical applications, only the outputs $y[k]$ can be observed and the full state $x[k]$ is not directly measurable.

Considering that the inputs u and outputs y are observed over a certain time interval that we qualify by $[0, t-1]$, we can formally state the estimation problem as:

Given the system input information $u[0], \cdots, u[t-1]$ and the measurements $y[0], \cdots, y[t-1]$, we want to estimate the state $x[s]$ at a given time s.

A method that produces an estimate of the state, denoted $\hat{x}[s]$, is called an *observer* or a *state estimator*.

In this work, the objective is to find *an estimator that provides the most accurate estimate of the initial state* $x[0]$, denoted by $\hat{x}[0]$, using the collected input and output data.

2.3 State estimation for LTI discrete-time systems in the noiseless case

As a first step, we consider the design of a state observer for a deterministic LTI system, assuming there is no measurement noise.

Such a system is described by the following state-space equations:

$$\begin{aligned} x[k+1] &= Ax[k] + Bu[k] \\ y[k] &= Cx[k] + Du[k]. \end{aligned} \tag{2.2}$$

Assume inputs u and measurements y are observed over time interval $[0, t-1]$: the objective is to recover the initial state $x[0]$. We can write the relationship between the initial state, the inputs, and the measurements as:

$$\begin{pmatrix} y[0] \\ \vdots \\ y[t-1] \end{pmatrix} = \mathcal{O}_{\mathbf{t}} x[0] + \mathcal{T}_t \begin{pmatrix} u[0] \\ \vdots \\ u[t-1] \end{pmatrix}, \tag{2.3}$$

where $\mathcal{O}_\mathbf{t}$ is the **t-step observability matrix** defined by

$$\mathcal{O}_\mathbf{t} = \begin{pmatrix} C \\ CA \\ \vdots \\ CA^{t-1} \end{pmatrix}, \tag{2.4}$$

and $\mathcal{T}_t$ captures the effect of past inputs on the outputs with

$$\mathcal{T}_t = \begin{pmatrix} D & 0 & \cdots & & \\ CB & D & 0 & \cdots & \\ \vdots & & & & \\ CA^{t-2}B & CA^{t-3}B & \cdots & CB & D \end{pmatrix}. \tag{2.5}$$

By rearranging the terms in 2.3, we can write

$$\mathcal{O}_\mathbf{t} x[0] = \begin{pmatrix} y[0] \\ \vdots \\ y[t-1] \end{pmatrix} - \mathcal{T}_t \begin{pmatrix} u[0] \\ \vdots \\ u[t-1] \end{pmatrix}, \tag{2.6}$$

where the right-hand side is known.

Under deterministic dynamics and in the absence of measurement noise, exact state estimation is possible *if and only if* the system is observable. This is equivalent to the observability matrix $\mathcal{O}_\mathbf{t}$ having full column rank equal to n. In that case, the initial state $x[0]$ can be uniquely determined by solving the linear system :

$$x[0] = \mathcal{O}_\mathbf{t}^\dagger \left(\begin{pmatrix} y[0] \\ \vdots \\ y[t-1] \end{pmatrix} - \mathcal{T}_t \begin{pmatrix} u[0] \\ \vdots \\ u[t-1] \end{pmatrix} \right), \tag{2.7}$$

where $\mathcal{O}_\mathbf{t}^\dagger$ denotes the left Moore-Penrose pseudo-inverse of $\mathcal{O}_\mathbf{t}$:

$$\mathcal{O}_\mathbf{t}^\dagger = (\mathcal{O}_\mathbf{t}^\top \mathcal{O}_\mathbf{t}^\top)^{-1} \mathcal{O}_\mathbf{t}^\top.$$

Observability Criterion. The concept of observability formalizes the ability to reconstruct the state from measurements. The n-dimensional system determined by matrices A and C is said to be observable if, for any initial state $x[0]$, knowledge of the input sequence $u[0], \cdots, u[n-1]$ and the measurement sequence $y[0], \cdots, y[n-1]$

uniquely determines $x[0]$. The system is observable if and only if the observability matrix:

$$\mathcal{O} = \mathcal{O}_n = \begin{pmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{pmatrix} \tag{2.8}$$

is of full rank, with $\text{rank}(\mathcal{O}) = n$, or equivalently

$$\ker(\mathcal{O}) = \{0\}.$$

If the system is observable, then $x[0]$ can be reconstructed from the measurements *over n time steps*. This matrix provides a binary notion of observability: it determines whether or not the state can be uniquely reconstructed without quantifying the quality of the reconstruction.

The Cayley-Hamilton Theorem. We consider the case where our system is n-dimensional and take t measurements over the interval $[0, t-1]$. The *Cayley-Hamilton theorem* states that any square matrix A satisfies its own characteristic equation [19]:

$$p_A(A) = 0,$$

where $p_A(\lambda)$ is the characteristic polynomial of A.

As a consequence, any power of A^k for $k \geq n$ can be expressed as a linear combination of the lower powers $I, A, \cdots, A^{n-1}$. Thus, for $t \geq n$, the t-step observability matrix satisfies:

$$\ker(\mathcal{O}_t) = \ker(\mathcal{O}),$$

and measuring the outputs beyond n steps does not reveal additional information about the initial state $x[0]$.

Now knowing that to establish the observer, we need at least n measurements and supposing that the t-step Observability matrix $\mathcal{O}_t$ is of full column rank n; let F be any left inverse of $\mathcal{O}_t$. We can estimate $x[0]$ using the observer:

$$x[0] = \hat{x}[0] = F\left(\begin{pmatrix} y[0] \\ \vdots \\ y[t-1] \end{pmatrix} - \mathcal{T}_t \begin{pmatrix} u(0) \\ \vdots \\ u[t-1] \end{pmatrix}\right), \tag{2.9}$$

and in a more general case if we want to estimate a state $t-1$ instances before:

$$x[\tau - t + 1] = F\left(\begin{pmatrix} y[\tau - t + 1] \\ \vdots \\ y[\tau] \end{pmatrix} - \mathcal{T}_t \begin{pmatrix} u[\tau - t + 1] \\ \vdots \\ u[\tau] \end{pmatrix}\right). \tag{2.10}$$

This approach provides an exact reconstruction of the state in finite time, considering the idealized noiseless setting.

2.4 State estimation for LTI discrete-time systems with noise on the measurements

In this section, we extend the previous analysis by considering a slightly more realistic scenario in which the measurements are corrupted by Gaussian sensor noise. While the system dynamics remain deterministic, the noisy outputs prevent exact state reconstruction, as was the case in the previous section.

We return to the discrete LTI system dynamics defined in 2.2, now including measurement noise. Specifically, we assume that the output noise is Gaussian, independent and identically distributed (i.i.d) over time, with zero mean and variance $\sigma^2 I_p$ where p is the number of rows of the output vector $y[t]$:

$$v[t] \sim \mathcal{N}(0, \sigma^2 I_p). \tag{2.11}$$

The relationship between the initial state, the input sequence, and the measurements sequence over the time interval $[0, \cdots, t-1]$, which was previously derived for the noiseless case in Equation 2.3, is now extended to the case with measurement noise and can be rewritten as:

$$\begin{pmatrix} y[0] \\ \vdots \\ y[t-1] \end{pmatrix} = \mathcal{O}_t x[0] + \mathcal{T}_t \begin{pmatrix} u[0] \\ \vdots \\ u[t-1] \end{pmatrix} + \begin{pmatrix} v[0] \\ \vdots \\ v[t-1] \end{pmatrix}, \tag{2.12}$$

where $\mathcal{O}_t$ and $\mathcal{T}_t$ are the respectively previously defined observability matrix 2.4 and input matrix 2.5.

We can observe that the measurements y consist of a *deterministic* part depending on $x[0]$ and input u as well as an *additive stochastic part* depending on measurement noise v.

The aim is to reconstruct the initial state $x[0]$ from the noisy measurement data. However, in the presence of noise, an exact reconstruction is no longer possible. The objective is then to find an estimate $\hat{x}[0]$ that approximates $x[0]$ as accurately as possible. An approach is to find the estimate that minimizes the error between the observed measurements y and the predicted outputs based on the estimated initial state $\tilde{y}$. This leads to the **least-squares formulation**:

$$\hat{x}[0] = \arg\min_{z\in\mathbb{R}^n} \| \begin{pmatrix} y[0] \\ \vdots \\ y[t-1] \end{pmatrix} - \mathcal{O}_{\mathbf{t}} z - \mathcal{T}_t \begin{pmatrix} u[0] \\ \vdots \\ u[t-1] \end{pmatrix} \|^2. \tag{2.13}$$

To simplify the notation, we define:

$$Y = \begin{pmatrix} y[0] \\ \vdots \\ y[t-1] \end{pmatrix}, \quad U = \begin{pmatrix} u[0] \\ \vdots \\ u[t-1] \end{pmatrix}, \quad V = \begin{pmatrix} v[0] \\ \vdots \\ v[t-1] \end{pmatrix}.$$

Under the assumption that the t-step observability matrix $\mathcal{O}_t$ has full column rank n, the least squares problem admits a *unique closed-form solution*. The optimal initial state estimate is given by:

$$\hat{x}[0] = \mathcal{O}_{\mathbf{t}}^{\dagger}(Y - \mathcal{T}_t U), \tag{2.14}$$

where $\mathcal{O}_t^{\dagger} = (\mathcal{O}_t^{\top}\mathcal{O}_t)^{-1}\mathcal{O}_t^{\top}$ is the Moore-Penrose inverse of $\mathcal{O}_t$.

Then, by substituting with the measurement equation 2.12 in vector form as

$$Y = \mathcal{O}_t x[0] + \mathcal{T}_t U + V,$$

we obtain the expression of the initial state least-squares estimate:

$$\hat{x}(0) = x[0] + \mathcal{O}_t^{\dagger} V. \tag{2.15}$$

Equation 2.15 shows that the least-squares estimate $\hat{x}[0]$ consists of the true initial state $x[0]$ added to an estimation error term $\mathcal{O}_t^{\dagger} V$ due to measurement noise. In the noiseless case, this estimate exactly becomes the true initial state. However, with noise, the estimate becomes a random variable whose distribution depends on the statistical properties of the noise.

The previous conclusion leads us to analyze the distribution of the estimation error under the Gaussian probabilistic noise model defined in 2.11. We obtain the vector V by stacking the noise vectors over the interval $[0, t-1]$ with

$$V = \begin{pmatrix} v[0] \\ \vdots \\ v[t-1] \end{pmatrix} \in \mathbb{R}^{tp},$$

which follows a multivariate Gaussian distribution

$$V \sim \mathcal{N}(0, \sigma^2 I_{tp}).$$

The estimation error vector associated with the least-squares estimate $\hat{x}[0]$ is defined as

$$z := \hat{x}[0] - x[0] = \mathcal{O}_t^{\dagger} V.$$

Since V is Gaussian and $\mathcal{O}_t^{\dagger}$ is a linear operator, z is also a Gaussian random vector and its distribution is characterized by the same distribution as the initial state estimate $\hat{x}[0]$:

$$z \sim \mathcal{N}(0, \Sigma_z),$$

where the covariance matrix Σ_z is defined by

$$\Sigma_z = \mathbb{E}[zz^\top] = \mathbb{E}[\mathcal{O}_t^{\dagger} V V^\top (\mathcal{O}_t^{\dagger})^\top] = \mathcal{O}_t^{\dagger} \mathbb{E}[V V^\top](\mathcal{O}_t^{\dagger})^\top.$$

Using the fact that $\mathbb{E}[VV^\top] = \sigma^2 I_{tp}$, we obtain the expression for the covariance matrix of the estimation error:

$$\Sigma_z = \sigma^2 \mathcal{O}_t^{\dagger} (\mathcal{O}_t^{\dagger})^\top.$$

To express this covariance in terms of the observability matrix, we use the definition of the Moore-Penrose pseudoinverse and substitute it in the covariance expression:

$$\Sigma_z = \sigma^2 \mathcal{O}_t^{\dagger} (\mathcal{O}_t^{\dagger})^\top = \sigma^2 (\mathcal{O}_t^\top \mathcal{O}_t)^{-1} \mathcal{O}_t^\top ((\mathcal{O}_t^\top \mathcal{O}_t)^{-1} \mathcal{O}_t^\top)^\top = \sigma^2 (\mathcal{O}_t^T \mathcal{O}_t)^{-1}. \quad (2.16)$$

This result shows that the estimation uncertainty depends critically on the structure of the observability matrix. If the eigenvalues of $(\mathcal{O}_t^T \mathcal{O}_t)^{-1}$ are small, the covariance of the estimation error will be small, leading to more accurate state estimates. Conversely, if the covariance matrix eigenvalues are large, the error covariance will be large, and the estimate will be sensitive to noise.

This is a key observation of this work that motivates optimization-based design aiming to select sampling sequences that improve the conditioning of the observability matrix and, by link, the covariance matrix to improve the robustness of the state observer.

In summary, in this section we introduced the least squares formulation 2.13 as a method to estimate the initial state $x[0]$ from noisy output measurements. In particular, when sensor noise is modeled as i.i.d. Gaussian noise, we derived an explicit expression for the estimator 2.14 that minimizes the MSE between the predicted and observed measurements 2.13. We also computed the covariance matrix of the estimation error vector z, which quantifies the uncertainty around the estimate.

2.5 Statistical Overview of the Least-Squares Observer

In this section, we take a closer look at the estimator 2.14 previously obtained from the least-squares formulation and at its statistical properties. By keeping our assumption on the measurement noise being zero-mean and Gaussian, we show that the least-squares estimator is also the maximum likelihood estimator (MLE) of the initial state. Under these assumptions, it can also be shown that the least-squares estimator is the minimum-variance unbiased estimator (MVUE) of the initial state $x[0]$. These statistical properties justify the use of this estimator and will offer information on the quality of this reconstructed initial state.

This section relies on key results from the course "LINMA1731 - Stochastic processes: Estimation and prediction" [17] and in particular on Theorem 3.6, which characterizes the optimality of the MLE for linear Gaussian models. The theorem we are referring to states:

> Assuming the linear model:
>
> $$m = H\theta + \omega,$$
>
> where the PDF of ω is $\mathcal{N}(0, \Omega)$, then the MLE of θ is
>
> $$\hat{\theta} = (H^T\Omega^{-1}H)^{-1}H^T\Omega^{-1}m,$$
>
> where $\hat{\theta}$ is the minimum variance unbiased estimator (MVU) and is the efficient estimator (it attains the Cramér-Rao bound).
>
> Its probability density function is:
>
> $$\hat{\theta} \sim \mathcal{N}\left(\theta, (H^T\Omega^{-1}H)^{-1}\right)$$

We can make a parallel between this result and the least squares estimation problem 2.15 by the identification :

- $\theta = x[0]$; the initial state.
- $H = \mathcal{O}_t$; the t-step observability matrix.
- $e = Y - \mathcal{T}_t U$; the measurements vector after subtracting the known input information.

- $\omega = V \sim \mathcal{N}(0, \sigma^2 I_{tp})$; the stacked noise vector.

According to further developed results in the course [17], the theoretical least-squares estimator with Gaussian noise is analogous to the MLE and allows the recovery of the MVUE.

This means that the estimator $\hat{x}[0]$ obtained by solving 2.13 is not only a solution to the least-squares problem but also *coincides with the maximum likelihood estimator and the minimum variance unbiased estimator*. It is unbiased and efficient; it achieves the lowest possible estimation variance among all unbiased estimators.

This lower bound on the estimation variance is given by the *Cramér-Rao bound* and is written as:

$$(H^\top \Omega^{-1} H)^{-1} = \sigma^2 (\mathcal{O}_t^\top \mathcal{O}_t)^{-1}.$$

This lower bound brings us back to the conclusion of the previous section: the quality of the estimation depends directly on the conditioning of the matrix $(\mathcal{O}_t^\top \mathcal{O}_t)^{-1}$. We can conclude that this matrix plays an important role in the problem of initial state estimation, and it will be introduced and defined in the next section as the *observability Gramian*. We will study its properties and interpretations and its link with the least-squares observer. It will be our main tool to formulate the optimization problem to select the optimal sampling times for initial state estimation.

2.6 The Observability Gramian

In this section, building on the previous discussion on the least-squares observer, we introduce the **observability Gramian** [18] and its central role in the initial state estimation problem for linear dynamical systems. As we established in the previous section, the covariance matrix of the estimation error of the least-squares estimator is given by

$$\Sigma_z = \sigma^2 (\mathcal{O}_t^\top \mathcal{O}_t)^{-1}.$$

We will come to define the matrix product $\mathcal{O}_t^\top \mathcal{O}_t$ as the observability Gramian.

This link highlights the fact that a *well-conditioned Gramian leads to a more accurate and robust estimate of the initial state* $x[0]$. The role of this section is then to define the observability Gramian, to explore its properties and interpretations, and to use them to formulate the optimization problem that drives this work:

selecting the optimal sampling times to have the most accurate estimate of the initial state of the system.

The t**-step observability Gramian** defined over the interval $[0, t-1]$ is defined as

$$W_O(t) = \sum_{k=0}^{t-1} (A^k)^\top C^\top C A^k = \mathcal{O}_t^\top \mathcal{O}_t. \tag{2.17}$$

The matrix $W_O(t)$ captures how much information about the initial state $x[0]$ is contained in the system's outputs collected over the sampling sequence.
In the case where the system is detectable, meaning that all unstable modes of A are observable, then the infinite horizon observability Gramian is defined as :

$$W_O = \lim_{t\to\infty} \sum_{k=0}^{\top} (A^k)^\top C^\top C A^k, \tag{2.18}$$

and the infinite sum converges and satisfies the discrete-time Lyapunov equation [3]

$$W_O - A^\top W_O A = C^\top C. \tag{2.19}$$

An interesting observation to make is that the *t-step observability Gramian is directly linked to the statistical properties of the least-squares observer.* We have shown that the least-squares estimator $\hat{x}[0]$ is unbiased and achieves minimum variance among all unbiased estimators, and the covariance of the estimation error 2.16 can be written in terms of the observability Gramian:

$$\Sigma_z = \sigma^2 (\mathcal{O}_t^T \mathcal{O}_t)^{-1} = \sigma^2 W_O(t)^{-1}. \tag{2.20}$$

This relationship shows that the observability Gramian directly influences the covariance of the estimation.

We can also look at the geometric interpretation of the observability Gramian; this interpretation offers information about how the initial state affects the output over time. More specifically, the Gramian defines an ellipsoid in the state space known as the **Observability Ellipsoid** [14] which characterizes how well different state directions are observed through the output measurements. It is a *positive semi-definite* matrix.

Two equivalent characterizations of this ellipsoid are:

$$E_1 = \{x \in \mathbb{R}^n : x^\top W_O(t)^{-1} x \le 1\} \quad \text{and} \quad E_2 = \{W_O(t)^{\frac{1}{2}} y : y \in \mathbb{R}^n, \|y\|_2 \le 1\}.$$

By considering these ellipsoids, we can assign a geometrical interpretation to the eigenvectors and eigenvalues of the observability Gramian. As stated in [14], it gives information on how much output energy is obtained when the system starts from a particular state.

- The **eigenvectors** of $W_O(t)$ determine the principal direction or axes of the ellipsoid, corresponding to orthogonal directions of observability in the state space.
- The square root of the eigenvalues of $W_O(t)$; its **singular values**, determine the lengths of these axes and can serve as a measure of the quality of the observation across the observability directions.

As a consequence, directions in the state space that are associated with large eigenvalues correspond to state components that can be observed "well" through the output measurements, while directions with small eigenvalues are poorly observed. If $W_O(t)$ is singular, then its null space corresponds to unobservable directions, those from which no output energy can be extracted.

This geometric interpretation of the observability ellipsoid can be directly linked to the least-squares estimator. In particular, the equality 2.20 highlights that the eigenvalues of the observability Gramian influence the eigenvalues of the covariance matrix of the initial state estimator. Small eigenvalues of $W_O(t)$, which correspond to poorly observable directions in the state space, lead to large eigenvalues in error covariance Σ_z and consequently to large estimation variance along these directions.

This link shows the importance of having a well-conditioned observability matrix and, by extension, a well-conditioned observability Gramian. Ensuring that $W_O(t)$ does not have small eigenvalues and is far from singular leads to smaller eigenvalues in its inverse and consequently to a more accurate and robust initial state estimation.

We recall that the main objective of this chapter is to formulate an optimization problem to select the optimal sampling times to achieve the most accurate estimate of the initial state of a system. These observations on the observability Gramian $W_O(t)$ and the covariance Σ_z motivate the design of an objective function that optimizes the spectral property of the estimation error covariance matrix. Equivalently, this implies improving the spectral properties of the inverse of the observability Gramian, which also optimizes the system's observability through the optimal sampling strategy.

Based on the literature review, multiple scalar metrics have been proposed to quantify the quality of controllability or observability Gramians in optimal design problems. Results such as those in Baggio et al. [11] discuss these scalar metrics for controllability; we can transpose the metrics for the observability Gramian due to the controllability-observability duality. Metrics among the ones discussed in Baggio et al. [11] are:

$$\frac{1}{\lambda_{\min}}, \quad \frac{1}{\lambda_{\max}}, \quad \mathrm{tr}(W_O(t)^{-1}), \text{ and } \det(W_O(t)^{-1}),$$

where $\lambda_{\min}$ and $\lambda_{\max}$ are the smallest and largest eigenvalues of the observability Gramian $W_O(t)$. These metrics capture different aspects of the system's observability, but they have different interpretations. For example, $\lambda_{\min}^{-1}$ and $\lambda_{\max}^{-1}$ capture the worst-case and best-case observability or estimation variance, and $\det(W_O(t)^{-1}$ measures the overall volume of the covariance ellipsoid. The trace of the inverse of the observability Gramian $\mathrm{tr}(W_O(t)^{-1})$ measures the average estimation variance across all state directions and is directly related to the covariance matrix of the least-squares estimator. This measure offers interpretability in the context of our problem of interest. We therefore adopt it as the objective function for the optimal selection of sampling times to accurately estimate the initial state:

$$\mathrm{tr}(W_O(t)^{-1}).$$

The objective selects the sequence of measurement times at which $\mathrm{tr}(W_O(t)^{-1})$ is minimized, to try to reduce the uncertainty of the least-squares estimate of the initial state in the presence of measurement noise.

2.7 Problem Formulation: Optimal Sampling for State Estimation

This section uses previous developments to formulate an optimization problem that selects sampling times that lead to the most accurate estimation of the initial state of the system, considering that we have prior knowledge on the system dynamics and a fixed number of measurements.

In earlier sections, we defined the t-step observability Gramian $W_O(t)$ over a uniformly spaced time interval $[0, t-1]$ 2.17. However, in many practical scenarios, measurements can be limited or sparse and cannot be made at every time step. We thus consider the problem in the setting where the outputs are measured at a selected subset of time instants.

Let us define the sequence of sampling times, that is a *sampling schedule*;

$$\mathcal{S} := \{\tau_0, \tau_1, \cdots, \tau_m\}, \quad \tau_i \in \mathbb{N}, \tau_i < \tau_{\max},$$

where each $\tau_{\max}$ denotes the final time horizon and $m+1$ is the total number of measurements.

For this sampling schedule, the observability Gramian is computed as:

$$W_O(\mathcal{S}) = \sum_{\tau \in \mathcal{S}} (A^\tau)^\top C^\top C A^\tau, \tag{2.21}$$

and the previous definition of 2.17 is now generalized to any sequence of sampling times.

It is now possible to formulate the desired optimization problem. As established previously, the least-squares estimator in equation $\hat{x}[0]$ 2.14 is obtained by minimizing the MSE between the observed outputs y and the predicted outputs based on the estimated initial state $\tilde{y}$, as formulated in Equation 2.13. *The covariance matrix of this estimator is proportional to the inverse of the observability Gramian.* Thus, the accuracy of the estimation is closely linked to the conditioning of the observability Gramian.

This conditioning is determined by the eigenvalues of the matrix. A meaningful scalar metric that tries to capture the overall uncertainty in the estimation is the trace of the estimation error covariance matrix. By equality 2.20, minimizing the trace of the inverse observability Gramian provides an objective function for minimizing estimation uncertainty on average across directions:

$$\min_{\mathcal{S} \subset \{0, \cdots, \tau_{\max}\}, |\mathcal{S}| = m+1} \operatorname{tr}(W_O(\mathcal{S})^{-1}). \tag{2.22}$$

It is interesting to note that similar development can be conducted for continuous-time systems. And, in this case, the observability Gramian is defined analogously as

$$W_O^{\text{cont}}(\mathcal{S}) = \sum_{\tau \in \mathcal{S}} e^{A^\top \tau} C^\top C e^{A\tau}. \tag{2.23}$$

and the objective function can be written:

$$\min_{\mathcal{S} \subset [0, \tau_{\max}], |\mathcal{S}| = m+1} \operatorname{tr}((W_O^{\text{cont}}(\mathcal{S}))^{-1}). \tag{2.24}$$

This objective function also has an advantage of the implicit assumptions derived from it. In particular:

- Full Rank of Observability matrix: it is implied that $W_O(\mathcal{S})$ is invertible equivalently is of full rank. This ensures that the initial state is uniquely identifiable from the measurements. $\mathcal{S}$ must then be chosen to guarantee observability.

- Trade-off of eigenvalues: The sampling times are selected to balance the different directions of uncertainty. By minimizing the trace of the estimation error covariance, the total uncertainty across all state directions is minimized to try to encourage general observability.

2.8 Conclusion

In this chapter, we discussed the theoretical foundation for the problem of state estimation for LTI dynamical systems and formulated an optimization problem to select optimal sampling times for accurate initial state estimation. We began with a review of the state of the art and placed this work within the context of the field.

We then formalized the state estimation problem for both the noiseless and noisy measurement cases. After introducing theoretical concepts such as observability and the observability Gramian, we focused on the least-squares estimator to recover the initial state from a finite set of output measurements. In the noisy measurements case, we show that this estimator remains unbiased and is the minimum variance unbiased estimator.

Then, we made the link between this statistical estimator and observability theory by showing that the estimation error covariance matrix can be expressed in terms of the inverse of the t-step observability Gramian 2.20. This relation shows that poorly observable directions that are associated with small eigenvalues of $W_O(t)$ lead to high estimation variance, and well-observed directions associated with large eigenvalues lead to low estimation variance. This pushes the idea of optimizing the structure of $W_O(t)^{-1}$ and its spectral properties that link observability and state estimation quality.

Building on these conclusions, we propose to optimize a scalar measure of this covariance matrix: **the trace of the inverse of the observability Gramian** becomes the principal objective function used to select the optimal sampling times for accurate initial state reconstruction. This metric approximates the average estimation error variance across the directions of the covariance ellipsoid and provides an exploitable formulation that we will try to solve in later chapters.

Chapter 3

Analysis of the Objective Function in Particular Dynamical Systems

In the previous chapter, we laid out the theoretical foundations for the state estimation problem for LTI dynamical systems and derived an optimization problem for selecting sampling times to achieve the most accurate initial state estimator. In this chapter, we focus on analyzing the behavior of the objective function in different cases of simple, low-dimensional dynamical systems. The goal of these cases is to build intuition about the optimization landscape and on the types of solutions that arise from the defined optimization problem for different system properties.

We specifically focus on one-dimensional and two-dimensional systems, where explicit computations are possible and can provide insights. By studying these cases, we try to see how the objective function responds to different cases of stability of systems and we study the optimal sampling schedules derived from it. We also study the gradient field to better visualize and understand the geometry of the problem.

We analyze both discrete-time and continuous-time systems. In addition to comparing both of the systems, discrete-time systems' sampling times are restricted to integer values, which introduces combinatorial aspects to the optimization problem, and the continuous-time formulation allows the choice of sampling times to be unconstrained and for solutions that might be easier to understand.

We begin this chapter with the one-dimensional case, where the objective function exhibits monotonic behavior and admits a closed-form solution. We then proceed to analyze simple two-dimensional systems, where some cases offer greater complexity and difficulty in interpretation. Even in these simple settings, we observe rich behaviors and the optimal sampling sequences are not trivial. Nevertheless, these examples provide valuable information on how system dynamics influences the optimization landscape and optimal sampling schedules.

3.1 Closed-form solution for One-Dimensional State Space

3.1.1 Discrete-time Case

We consider a simple case of the problem: the discrete one-dimensional LTI system defined by

$$\begin{aligned} x[t+1] &= ax[t] \\ y[t] &= cx[t], \end{aligned} \tag{3.1}$$

where $x[t], y[t], a, c \in \mathbb{R}$.

In this setting, the state dimension is $n = 1$ and the observability matrix reduces to a scalar value given simply by

$$O_1 = c,$$

this matrix is trivially of full rank if and only if $c \neq 0$. Therefore, a single observation for $c \neq 0$ is sufficient to guarantee the observability of the system. This is consistent with the Cayley-Hamilton theorem 2.3.

Without loss of generality, we set $c = 1$, since scaling by c does not affect the optimization strategy. With this assumption, the objective function 2.22 can be written as:

$$\min_{\tau \in \{0,1,\cdots,\tau_{\max}\}} W_O(\tau)^{-1} \iff \min_{\tau \in \{0,1,\cdots,\tau_{\max}\}} \frac{1}{a^{2\tau}} \iff \max_{\tau \in \{0,1,\cdots,\tau_{\max}\}} a^{2\tau}. \tag{3.2}$$

This leads to a simple scalar optimization problem over the discrete time index τ, constrained on the finite time horizon $\tau \in \{0, 1, \cdots, \tau_{\max}\}$.

The monotonic behavior of the objective function 3.2 is entirely dependent on the value of the system parameter a, which also determines the *stability* of the system:

- **Case 1: Stable system ($|\mathbf{a}| < \mathbf{1}$)**
  The objective function is strictly increasing in τ. Therefore, the earliest

possible observation provides the most information about the initial state. The optimal sampling time is:

$$\tau^* = 0.$$

- **Case 2: Unstable system ($|\mathbf{a}| > \mathbf{1}$)**
  The objective function is strictly decreasing in τ. Therefore, the latest possible observation offers the most information about the initial state. The optimal sampling time is:
  $$\tau^* = \tau_{\max}.$$

- **Case 3: Marginally stable system ($|\mathbf{a}| = \mathbf{1}$)**
  The objective function is constant across the time horizon. Therefore, any choice of sampling time is equally optimal:
  $$\tau^* = \tau, \quad \forall \tau \in \{0, \cdots, \tau_{\max}\}.$$

This analysis reveals that *the optimal sampling time is entirely determined by the stability of the system.* In fact, it reflects the amplification or decay of state information over time. In stable systems, the signal converges and the initial state becomes harder to recover as time goes on. In unstable systems, the signal grows and becomes increasingly easier to detect.

3.1.2 Continuous-time Case

We now consider the continuous-time version of the previous problem: the one-dimensional continuous-time LTI system defined by:

$$\begin{aligned} \dot{x}(t) &= ax(t) \\ y(t) &= cx(t), \end{aligned} \tag{3.3}$$

where $x(t), y(t), a, c \in \mathbb{R}$.

As in the discrete case, the system is scalar and the observability matrix reduces to the same value:

$$\mathcal{O}_1 = c.$$

In this case, the observability Gramian for a sampling time τ is given by 2.23

$$W_O^{\text{cont}}(\tau) = e^{2a\tau},$$

and the optimization problem can be written:

$$\min_{\tau \in [0,\tau_{\max}]} W_O(\tau)^{-1} \quad \Longleftrightarrow \quad \min_{\tau \in [0,\tau_{\max}]} \frac{1}{e^{2a\tau}} \quad \Longleftrightarrow \quad \min_{\tau \in [0,\tau_{\max}]} e^{2a\tau}. \tag{3.4}$$

The behavior of the objective function 3.4 is completely determined by the stability of the system, in the same way as it was characterized in the discrete time case:

- **Stable system (a < 0)**
  The objective function is strictly increasing in τ, and the optimal sampling time is the earliest possible in the time horizon:
  $$\tau^* = 0.$$

- **Unstable system (a > 0)**
  The objective function is strictly decreasing in τ, and the optimal sampling time is the latest available in the time horizon:
  $$\tau^* = \tau_{\max}.$$

- **Marginally stable system (a = 0)**
  The function objective function is constant, and all sampling times are equally optimal:
  $$\tau^* = \tau, \quad \forall \tau \in [0, \tau_{\max}].$$

In the same way as for the discrete-time case, this analysis highlights how the optimal sampling strategy is entirely decided by the system's stability.

It is interesting to note that for both cases the observability Gramian is always of full rank as long as $c \neq 0$, regardless of τ. Therefore, optimizing the trace of its inverse serves to guarantee observability as well as ensure robustness in state reconstruction without additional constraints. The general one-dimensional case provides a complete closed-form characterization of the optimization problem. Although very simple, it shows that the stability of the dynamics have a complete influence in determining an optimal measurement strategy. It is interesting to look at this analysis for higher-dimensional systems and see the effects of their stability properties on the objective function.

3.2 Closed-form solution for Discrete-Time Two-Dimensional State Space

Now, we extend the previous analysis to the case where the state space is two-dimensional; $n = 2$.

We consider the discrete LTI system 2.2 for the general two-dimensional case:

$$\begin{aligned} x[t+1] &= Ax[t], \\ y[t] &= Cx[t] \end{aligned} \tag{3.5}$$

where $x[t] \in \mathbb{R}^2$, $A \in \mathbb{R}^{2\times 2}$ and $y[t] \in \mathbb{R}^p$ and $C \in \mathbb{R}^{p\times 2}$ with p the number of measurement outputs.

Two measurements are sufficient to ensure full column rank $n = 2$ of the observability matrix if the system is well conditioned. The observability matrix corresponding to the sampling schedule $\mathcal{S} := \{\tau_1, \tau_2\}$ is:

$$\mathcal{O}_{\mathcal{S}} = \begin{pmatrix} CA^{\tau_1} \\ CA^{\tau_2} \end{pmatrix}. \tag{3.6}$$

The corresponding discrete observability Gramian is given by:

$$W_O(\mathcal{S}) = (A^\top)^{\tau_1} C^\top C A^{\tau_1} + (A^\top)^{\tau_2} C^\top C A^{\tau_2}.$$

For the rest of the section, we evaluate a simplified case where the state matrix A is diagonal:

$$A = \begin{pmatrix} a_{11} & 0 \\ 0 & a_{22} \end{pmatrix}, \tag{3.7}$$

we also examine two different cases for the output matrix C:

1. $C = I_2$: this corresponds to full state observation, where both state variables are measured independently.
2. $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$: this corresponds to a partial observation of the individual state variables, where the output is a scalar equal to the sum of the two state variables. Individual state variables cannot be directly measured through the output since it is a linear combination of the states.

3.2.1 Case 1: $\mathbf{C = I_2}$

First, we consider the case where the output matrix is the identity matrix, $C = I_2$. Here, the full state is directly measured, simplifying the observability conditions.

Since $C = I_2$, the observability matrix after a single observation is already of full column rank with $\mathcal{O}_1 = I_2$, and thus the system is observable for any choice of sampling times. This removes the problem of selecting times that ensure observability, and shifts the focus to optimizing the quality of state estimation.

We now develop a closed form for the objective function for a sampling schedule $\mathcal{S} := \{\tau_1, \tau_2\}$. The observability Gramian simplifies to a diagonal matrix:

$$W_O(\mathcal{S}) = \begin{pmatrix} a_{11}^{2\tau_1} + a_{11}^{2\tau_2} & 0 \\ 0 & a_{22}^{2\tau_1} + a_{22}^{2\tau_2} \end{pmatrix},$$

and its inverse is

$$W_O^{-1}(\mathcal{S}) = \begin{pmatrix} \frac{1}{a_{11}^{2\tau_1} + a_{11}^{2\tau_2}} & 0 \\ 0 & \frac{1}{a_{22}^{2\tau_1} + a_{22}^{2\tau_2}} \end{pmatrix}. \tag{3.8}$$

The closed form of the objective function defined as $f(\tau_1, \tau_2)$ in this setting is:

$$f(\tau_1, \tau_2) = \min_{\tau_1, \tau_2 \in \mathbb{N}^+} \frac{1}{a_{11}^{2\tau_1} + a_{11}^{2\tau_2}} + \frac{1}{a_{22}^{2\tau_1} + a_{22}^{2\tau_2}}. \tag{3.9}$$

It is interesting to note that in this case, there are no constraints on the sampling times and that the objective function is always defined regardless of the specific values of τ_1 and τ_2 as long as $a_{11} \neq 0$ and $a_{22} \neq 0$.

To better understand the behavior of the objective function and the nature of the solutions obtained, we study this objective function and its gradient with respect to the sampling times τ_1 and τ_2.

We compute the gradient $\nabla f(\tau_1, \tau_2)$. Using the identity $a^b = \exp(b \ln(a))$, the partial derivatives are computed and the gradient is:

$$\nabla f(\tau_1, \tau_2) = -\begin{pmatrix} \frac{2a_{11}^{2\tau_1} \ln(a_{11})}{\left(a_{11}^{2\tau_1} + a_{11}^{2\tau_2}\right)^2} + \frac{2a_{22}^{2\tau_1} \ln(a_{22})}{\left(a_{22}^{2\tau_1} + a_{22}^{2\tau_2}\right)^2} \\ \frac{2a_{11}^{2\tau_2} \ln(a_{11})}{\left(a_{11}^{2\tau_1} + a_{11}^{2\tau_2}\right)^2} + \frac{2a_{22}^{2\tau_2} \ln(a_{22})}{\left(a_{22}^{2\tau_1} + a_{22}^{2\tau_2}\right)^2} \end{pmatrix}. \tag{3.10}$$

By analyzing the sign of each term in the gradient, we can characterize the monotonicity of f and deduce an optimal sampling strategy based on the eigenvalues a_{11} and a_{22} of the system. We can see how the sampling strategy is influenced by the stability of the system. The analysis is performed on *positive* values of the entries of matrix A to ensure that the gradient 3.10 is well defined, since each term contains logarithms of the form $\ln(a_{ii})$.

Case 1: Stable system ($0 < a_{11}, a_{22} < 1$): Here, $\ln(a_{11}) < 0$ and $\ln(a_{22}) < 0$ and, as a consequence, each gradient entry is positive. The objective function is strictly increasing in both τ_1 and τ_2. Since the observability Gramian and the objective function are both well defined for any value of the sampling times, to minimize f the earliest possible sampling times are optimal:

$$\tau_1^* = \tau_2^* = 0.$$

Case 2: Unstable system ($a_{11}, a_{22} > 1$): Now, $\ln(a_{11}) > 0$ and $\ln(a_{22}) > 0$, so each gradient entry is negative. The objective function is strictly decreasing in both τ_1 and τ_2. The optimal sampling times are at the latest time allowed within the time horizon:

$$\tau_1^* = \tau_2^* = \tau_{\max}$$

Case 3: Mixed stability ($a_{11} < 1$, $a_{22} > 1$): In this case, $\ln(a_{11}) < 0$ and $\ln(a_{22}) > 0$, so the two terms in each gradient entry have opposite signs. The sign of each gradient component for the sampling times depends on the magnitude of the two terms. For instance, the first entry of the gradient is:

$$\frac{\partial f(\tau_1, \tau_2)}{\partial \tau_1} = -\left(\frac{2a_{11}^{2\tau_1} \ln(a_{11})}{(a_{11}^{2\tau_1} + a_{11}^{2\tau_2})^2} + \frac{2a_{22}^{2\tau_1} \ln(a_{22})}{(a_{22}^{2\tau_1} + a_{22}^{2\tau_2})^2} \right), \tag{3.11}$$

the sign of this entry depends on the relative magnitudes of its two terms. If the term associated with a_{11} is larger in absolute value than the second term associated with a_{22}, the entry becomes positive, otherwise it is negative.

More precisely, the first gradient component $\frac{\partial f(\tau_1,\tau_2)}{\partial \tau_1}$ is positive if:

$$\left| \frac{2a_{11}^{2\tau_1} \ln(a_{11})}{(a_{11}^{2\tau_1} + a_{11}^{2\tau_2})^2} \right| > \left| \frac{2a_{22}^{2\tau_1} \ln(a_{22})}{(a_{22}^{2\tau_1} + a_{22}^{2\tau_2})^2} \right|, \tag{3.12}$$

and negative otherwise.
And the second gradient component $\frac{\partial f(\tau_1,\tau_2)}{\partial \tau_2}$ is positive if:

$$\left| \frac{2a_{11}^{2\tau_2} \ln(a_{11})}{(a_{11}^{2\tau_1} + a_{11}^{2\tau_2})^2} \right| > \left| \frac{2a_{22}^{2\tau_2} \ln(a_{22})}{(a_{22}^{2\tau_1} + a_{22}^{2\tau_2})^2} \right|, \tag{3.13}$$

and negative otherwise.

In this case, the objective function is not simply monotonic in either variable, and the optimum is attained at asymmetric sampling times. A closed-form solution is not trivial, and numerical optimization is generally required to determine the optimal values.

To deepen the insights offered by the analytical analysis of the objective function and the nature of optimal solutions, we visualize the optimization landscape in the space of the sampling times. For each of the three cases discussed above, we generate plots showing the level curves of the objective function, the corresponding gradient vectors, and the optimal sampling times obtained by solving the optimization problem. The goal of these visualizations is to provide geometric intuition on the optimization landscape and to have system-specific intuition to interpret the nature of the solutions obtained beyond the pure analytical way. The details of the construction of these plots are presented in Chapter 4 in Section 4.1.

In the fully stable case in Figure 3.1a or the fully unstable case in Figure 3.1b, the behavior of the gradient is relatively straightforward. The direction of steepest descent is clearly defined, and optimal sampling times respect the expected analytical solutions. Indeed, we can see that in these cases the results join those of the one-dimensional case: in stable systems, the system's memory of its initial state decays exponentially and later outputs carry less information about this state, and thus the earliest sampling times are preferable, in unstable systems the output is exponentially amplified over time, and the initial state is recovered more precisely with the latest possible observation.

In contrast, for the mixed stability case in Figure 3.1c, we observe a more complex behavior. The gradient field is no longer monotonic and this can be seen in regions that have opposing gradient directions, and the local minima found by the optimization function do not correspond to zero gradient points. These local minima found can also be due to the constraint of having integer sampling times for the discrete-case. In Section 3.3,the continuous case is then explored to discuss solutions that are free from the integer constraint.

We can see that even for a relatively simple case, the solutions quickly become less intuitive to interpret. To push the understanding further, the next chapter presents numerical experiments to help interpret the results through system theory.

3.2.2 Case 2: $\mathbf{C} = \begin{pmatrix} \mathbf{1} & \mathbf{1} \end{pmatrix}$

We now examine the second case where the system has a single scalar output. Specifically, we consider the case where the output matrix is given by

$$C = \begin{pmatrix} 1 & 1 \end{pmatrix},$$

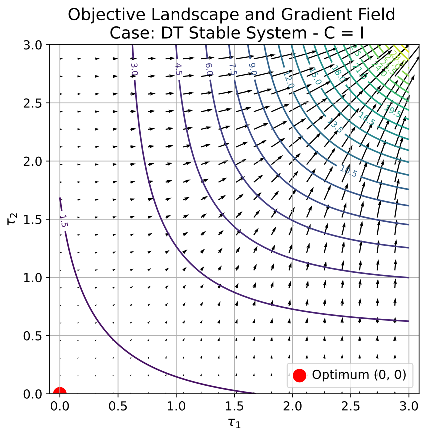

(a) Stable System

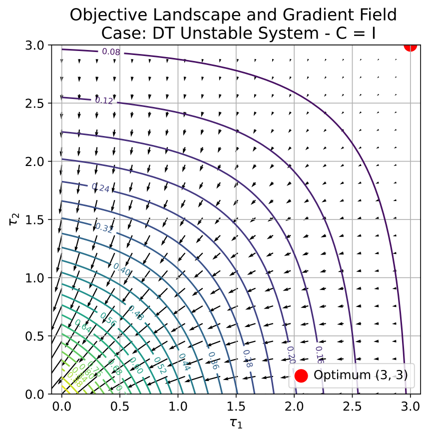

(b) Unstable System

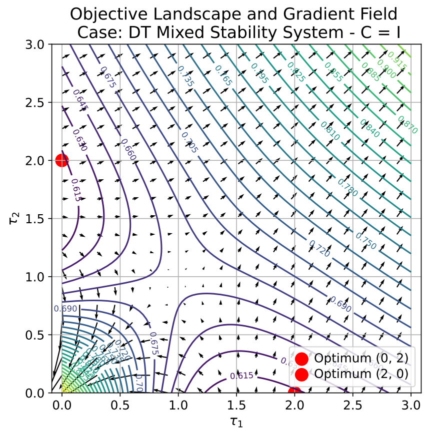

(c) Mixed Stability System

Figure 3.1: Gradient field visualization for Discrete-Time System with Output Matrix $C = I_2$.

then, the scalar output is a linear combination of the full state.

The observability Gramian evaluated at a sampling sequence $\mathcal{S} := \{\tau_1, \tau_2\}$ then takes the form:

$$W_O(\mathcal{S}) = \begin{pmatrix} a_{11}^{2\tau_1} + a_{11}^{2\tau_2} & (a_{11}a_{22})^{\tau_1} + (a_{11}a_{22})^{\tau_2} \\ (a_{11}a_{22})^{\tau_1} + (a_{11}a_{22})^{\tau_2} & a_{22}^{2\tau_1} + a_{22}^{2\tau_2} \end{pmatrix},$$

we compute its inverse as

$$W_O(\mathcal{S})^{-1} = \frac{1}{\det W_O(\mathcal{S})} \mathrm{Adj}(W_O(\mathcal{S})),$$

as a result, the closed form of the objective function becomes:

$$f(\tau_1, \tau_2) = \min_{\tau_1, \tau_2 \in \mathbb{N}^+} \frac{a_{11}^{2\tau_1} + a_{11}^{2\tau_2} + a_{22}^{2\tau_1} + a_{22}^{2\tau_2}}{(a_{11}^{\tau_1} a_{22}^{\tau_2} - a_{11}^{\tau_2} a_{22}^{\tau_1})^2}. \tag{3.14}$$

An important property of this expression is that the **denominator is zero when both sampling times are equal**:

$$\tau_1 = \tau_2 \implies (a_{11}^{\tau_1} a_{22}^{\tau_2} - a_{11}^{\tau_2} a_{22}^{\tau_1})^2 = 0,$$

equivalently, two equal sampling times also make the system unobservable:

$$\tau_1 = \tau_2 \implies \det(W_{\mathcal{O}}(\mathcal{S})) = 0 \iff W_{\mathcal{O}}(\mathcal{S}) \text{ is singular and non-invertible.}$$

This singularity reflects a fundamental limitation: when the two sampling times are identical, the observability matrix loses column rank, and the system becomes unobservable. This aligns with the Cayley-Hamilton theorem stated in 2.3. This is directly implied in the objective function, where the denominator vanishes when $\tau_1 = \tau_2$, leading to a divergence in cost. The structure of the objective function penalizes redundant measurements by assigning them an infinite cost, without the need to explicitly write the constraint, which becomes an advantage in the computational complexity of the problem and in creating useful choices of sampling times.

We perform a similar analysis as for the previous case by looking at the gradient of our objective function. By computing the partial derivatives of the objective function with respect to the optimization variables τ_1 and τ_2 we obtain a closed form for each entry of its gradient:

$$\frac{\partial f(\tau_1, \tau_2)}{\partial \tau_1} = \frac{2a_{11}^{\tau_1} \ln(a_{11}) \Big(a_{22}^{\tau_1} a_{11}^{\tau_1+\tau_2} + a_{11}^{2\tau_2} a_{22}^{\tau_2} + a_{22}^{2\tau_1+\tau_2} + a_{22}^{3\tau_2} \Big)}{(a_{11}^{\tau_2} a_{22}^{\tau_1} - a_{11}^{\tau_1} a_{22}^{\tau_2})^3} - \frac{2a_{22}^{\tau_1} \ln(a_{22}) \Big(a_{11}^{\tau_1} a_{22}^{\tau_1+\tau_2} + a_{11}^{\tau_2} a_{22}^{2\tau_2} + a_{11}^{2\tau_1+\tau_2} + a_{11}^{3\tau_2} \Big)}{(a_{11}^{\tau_2} a_{22}^{\tau_1} - a_{11}^{\tau_1} a_{22}^{\tau_2})^3} , \tag{3.15}$$

$$\frac{\partial f(\tau_1, \tau_2)}{\partial \tau_2} = \frac{2a_{22}^{\tau_2} \ln(a_{22}) \Big(a_{11}^{\tau_2} a_{22}^{\tau_2+\tau_1} + a_{11}^{\tau_1} a_{22}^{2\tau_1} + a_{11}^{2\tau_2+\tau_1} + a_{11}^{3\tau_1} \Big)}{(a_{11}^{\tau_2} a_{22}^{\tau_1} - a_{11}^{\tau_1} a_{22}^{\tau_2})^3} - \frac{2a_{11}^{\tau_2} \ln(a_{11}) \Big(a_{22}^{\tau_2} a_{11}^{\tau_2+\tau_1} + a_{11}^{2\tau_1} a_{22}^{\tau_1} + a_{22}^{2\tau_2+\tau_1} + a_{22}^{3\tau_1} \Big)}{(a_{11}^{\tau_2} a_{22}^{\tau_1} - a_{11}^{\tau_1} a_{22}^{\tau_2})^3} . \tag{3.16}$$

Contrary to the previous case, this closed form expression of the gradient is more difficult to analyze directly. Even in the fully stable and fully unstable cases, it is not straightforward to derive an analytical solution for the optimal sampling times. Thus, we turn to visual analysis of the gradient field and contour lines of the objective function, relating them with the numerically computed solutions for the three cases of a fully stable system, a fully unstable one, and one with mixed stability (one stable eigenvalue and an unstable one).

In all three cases of Figure 3.2, we can see that the objective function is undefined along the diagonal $\tau_1 = \tau_2$. Near this diagonal, the function exponentially increases as can be seen in the contour lines. This behavior is consistent with the analytical analysis since the denominator vanishes for $\tau_1 = \tau_2$. Due to the symmetry on each part of this diagonal, which can also be seen in the analytical expression, it is sufficient to consider the region $\tau_1 < \tau_2$ for the analysis.

Since the magnitude of the gradient becomes extremely large around $\tau_1 = \tau_2$, making directional information in the rest of the graph not visible, we look at *normalized vectors* in the plots of Figure 3.2.

In the mixed stability case in Figure 3.2c, the gradient direction appears to be monotonic on each side of the diagonal, and the optimal solution is found at the boundary: one sampling time is at the beginning of the interval and the second is at the end. Thus, each measurement, respectively, captures the decaying stable mode and the growing unstable mode of the system.

However, the fully stable (Figure 3.2a) and fully unstable cases (Figure 3.2b) do not necessarily exhibit a straightforward interpretation of the landscape. The optimal sampling times are located in the early part of the interval for the fully stable system and in the later part of the interval for the fully unstable system; which is coherent regarding the previous link between the optimal sampling times and the stability of systems. However, their precise positions vary depending on the specific eigenvalues of the system and are less clear to interpret. This also brings back the constraint of integer times of the discrete framework and motivates the exploration of the continuous-time case in the next section. Nevertheless, the overall qualitative behavior of these cases, seen for the previously explored examples, is still present: the objective function favors early sampling times for stable systems and late sampling times for unstable ones.

In summary, this analysis once again highlights the influence of system stability on the optimal sampling schedules for state reconstruction. Although qualitative trends for fully stable and unstable systems persist, the precise location of the sampling times is sensitive to the eigenvalues of the system and is not always easy to interpret analytically,and this difficulty in interpretation is reinforced by the integer constraint on the sampling times. This motivates the exploration of the solutions in the continuous-time framework explored in the next section.

3.3 Closed-form solution for Continuous-Time Two-Dimensional State Space

We perform an analysis analogous to the previous section, focusing on two-dimensional systems in a continuous-time framework. The main property of this framework is that sampling times are no longer limited to integer values anymore. It allows us to compare the objective functions and their behavior for different cases of system stability.

Thus, we consider the continuous-time linear system, for the previously defined diagonal system matrix A 3.7:

$$\begin{aligned} \dot{x}(t) &= Ax(t), \\ y(t) &= Cx(t). \end{aligned} \tag{3.17}$$

As in the discrete-time case, two measurements are sufficient to ensure full column rank of the observability matrix. For the sampling schedule $\mathcal{S} := \{\tau_1, \tau_2\}$ it is of the form

$$\mathcal{O}_{\mathcal{S}} = \begin{pmatrix} Ce^{A\tau_1} \\ Ce^{A\tau_2 2} \end{pmatrix}, \tag{3.18}$$

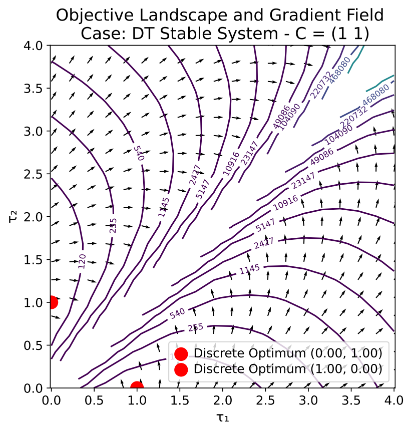

(a) Stable System

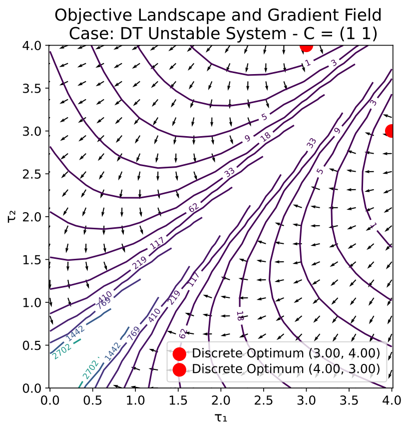

(b) Unstable System

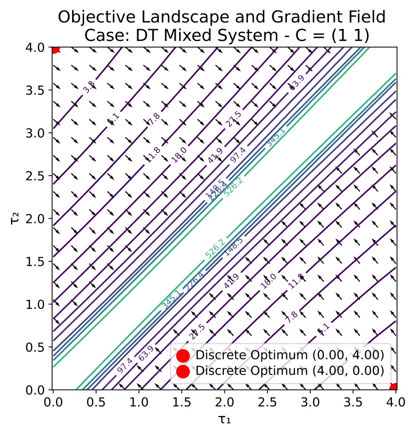

(c) Mixed Stability System

Figure 3.2: Gradient field visualization for Discrete-Time System with Output Matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$.

the corresponding observability Gramian is given by:

$$W_O^{\text{cont}}(\mathcal{S}) = e^{A^\top \tau_1} C^\top C e^{A\tau_1} + e^{A^\top \tau_1} C^\top C e^{A\tau_1}.$$

In this section, we evaluate the same cases as for the discrete time case, where A is the diagonal matrix 3.7 and the output matrix is either $C = I_2$ or $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$.

3.3.1 Case 1: $\mathbf{C = I_2}$

First, we consider the case where the output matrix is the identity matrix, $C = I_2$. As for the discrete case, a single measurement is sufficient to ensure observability of the system.

The closed form of the observability Gramian simplifies to the diagonal matrix

$$W_O(\mathcal{S}) = \begin{pmatrix} e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}} & 0 \\ 0 & e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}} \end{pmatrix},$$

and its inverse is:

$$W_O^{-1}(\mathcal{S}) = \begin{pmatrix} \dfrac{1}{e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}}} & 0 \\ 0 & \dfrac{1}{e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}}} \end{pmatrix}. \tag{3.19}$$

The closed form of the objective function for this setting becomes:

$$f(\tau_1, \tau_2) = \min_{\tau_1, \tau_2} \frac{1}{e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}}} + \frac{1}{e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}}} \tag{3.20}$$

and there are *no constraints on the sampling times*; the objective function is always defined regardless of the specific values of τ_1 and τ_2.

A similar analysis to the one of the discrete-time case is performed to understand and compare the behavior of the objective function for this particular case.

We compute the gradient $\nabla f(\tau_1, \tau_2)$ using the the partial derivatives:

$$\nabla f(\tau_1, \tau_2) = - \begin{bmatrix} \dfrac{2a_{11} e^{2\tau_1 a_{11}}}{(e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}})^2} + \dfrac{2a_{22} e^{2\tau_1 a_{22}}}{(e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}})^2} \\ \\ \dfrac{2a_{11} e^{2\tau_2 a_{11}}}{(e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}})^2} + \dfrac{2a_{22} e^{2\tau_2 a_{22}}}{(e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}})^2} \end{bmatrix} \tag{3.21}$$

We analyze the gradient of the continuous-time objective function and observe results similar to those in the discrete-time case, with a *sampling strategy that depends on system stability.*

Case 1: Stable system ($a_{11}, a_{22} < 0$): Here each gradient entry is positive. The objective function is strictly increasing in both τ_1 and τ_2. The earliest possible sampling times are optimal:

$$\tau_1^* = \tau_2^* = 0$$

Case 2: Unstable system ($a_{11}, a_{22} > 0$): Now, each gradient entry is negative. The objective function is strictly decreasing in both τ_1 and τ_2. The optimal sampling times are at the latest time allowed within the time horizon:

$$\tau_1^* = \tau_2^* = \tau_{\max}$$

Case 3: Mixed stability ($a_{11} < 0$, $a_{22} > 0$): In this case, the sign of each gradient component for the sampling times depends on the magnitude of each of their two terms.

For instance, the first entry of the gradient is:

$$-\left(\frac{2a_{11}e^{2\tau_1 a_{11}}}{(e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}})^2} + \frac{2a_{22}e^{2\tau_1 a_{22}}}{(e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}})^2}\right). \tag{3.22}$$

The sign of this entry depends on the relative magnitudes of the two terms. More precisely, the first gradient component is positive if:

$$\left|\frac{2a_{11}e^{2\tau_1 a_{11}}}{(e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}})^2}\right| > \left|\frac{2a_{22}e^{2\tau_1 a_{22}}}{(e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}})^2}\right|, \tag{3.23}$$

and negative otherwise.

And the second gradient component is positive if:

$$\left|\frac{2a_{11}e^{2\tau_2 a_{11}}}{(e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}})^2}\right| > \left|\frac{2a_{22}e^{2\tau_2 a_{22}}}{(e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}})^2}\right|, \tag{3.24}$$

and negative otherwise.

This leads to an analysis analogous to the discrete-time case where the objective function is not monotonous and with asymmetrical times.

For the three cases detailed above, we perform the same visual analysis of the optimization landscape as for the discrete time. The results are qualitatively similar to those observed in discrete time.

For the fully stable (Figure 3.3a) and fully unstable (Figure 3.3b) cases, the behavior is the same as for the discrete time system, and the analytical solutions found are equal. In stable systems, the earliest sampling time is optimal to avoid the decay of the initial condition, and in the unstable system the latest sampling time allows for better recovery of the initial condition which is exponentially increased.

In the mixed stability case (Figure 3.3c), we can observe a similar optimization landscape as for the discrete-time case with non-monotonicity in the objective function. Although the sampling times are no longer limited to integers, the solution remains difficult to interpret and intuitively explain.

Thus, the continuous formulation removes the integer constraint on the sampling times, but the main properties of the solutions remain dependent on the stability and show similar behavior across both formulations. This makes the conclusions applicable to both discrete and continuous systems.

3.3.2 Case 2: $\mathbf{C} = (\mathbf{1} \quad \mathbf{1})$

We now examine the continuous case for the output matrix:

$$C = \begin{pmatrix} 1 & 1 \end{pmatrix},$$

where the scalar output is a linear combination of the full state.

In this case the observability Gramian evaluated at a sampling sequence $\mathcal{S} := \{\tau_1, \tau_2\}$ takes the form:

$$W_O^{\text{cont}}(\mathcal{S}) = \begin{pmatrix} e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}} & e^{\tau_1(a_{11}+a_{22})} + e^{\tau_2(a_{11}+a_{22})} \\ e^{\tau_1(a_{11}+a_{22})} + e^{\tau_2(a_{11}+a_{22})} & e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}} \end{pmatrix},$$

and the closed form of the objective function then becomes:

$$f(\tau_1, \tau_2) = \min_{\tau_1, \tau_2} \frac{e^{2\tau_1 a_{11}} + e^{2\tau_2 a_{11}} + e^{2\tau_1 a_{22}} + e^{2\tau_2 a_{22}}}{\left(e^{a_{11}\tau_1 + a_{22}\tau_2} - e^{a_{11}\tau_2 + a_{22}\tau_1}\right)^2}. \tag{3.25}$$

This formulation has the same property of implicitly enforcing different sampling times, since two equal sampling times make the system unobservable and the objective function diverges in cost:

$$\tau_1 = \tau_2 \implies \left(e^{a_{11}\tau_1 + a_{22}\tau_2} - e^{a_{11}\tau_2 + a_{22}\tau_1}\right)^2 = 0 \iff W_{\mathcal{O}}(\mathcal{S}) \text{ is singular and non-invertible.}$$

Then a gradient analysis is performed on the objective function to characterize its behavior. By computing the partial derivatives of the objective function with

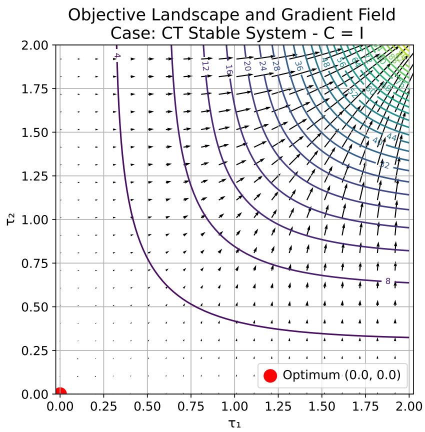

(a) Stable System

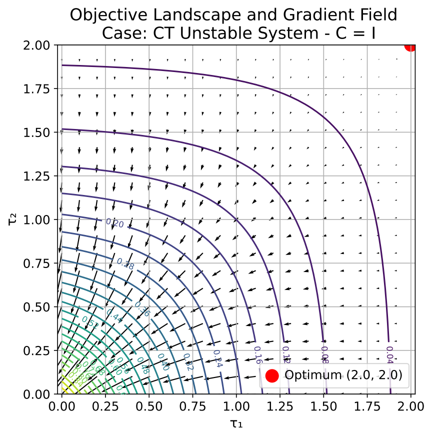

(b) Unstable System

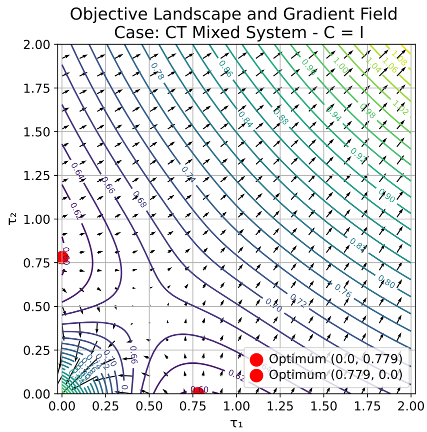

(c) Mixed Stability System

Figure 3.3: Gradient field visualization for Continuous-Time System with Output Matrix $C = I_2$.

respect to the optimization variables τ_1 and τ_2 we obtain a closed form for its gradient:

$$\begin{aligned}\frac{\partial f(\tau_1,\tau_2)}{\partial \tau_1} =& \frac{2a_{11}e^{\tau_1 a_{11}}\left(e^{a_{22}\tau_1+a_{11}(\tau_1+\tau_2)}+e^{\tau_2(2a_{11}+a_{22})}+e^{a_{22}(2\tau_1+\tau_2)}+e^{3a_{22}\tau_2}\right)}{(e^{a_{11}\tau_2+a_{22}\tau_1}-e^{a_{11}\tau_1+a_{22}\tau_2})^3}\\ &-\frac{2a_{22}e^{\tau_1 a_{22}}\left(e^{a_{11}\tau_1+a_{22}(\tau_1+\tau_2)}+e^{\tau_2(2a_{22}+a_{11})}+e^{a_{11}(2\tau_1+\tau_2)}+e^{3a_{11}\tau_2}\right)}{(e^{a_{11}\tau_2+a_{22}\tau_1}-e^{a_{11}\tau_1+a_{22}\tau_2})^3}\end{aligned} \tag{3.26}$$

$$\begin{aligned}\frac{\partial f(\tau_1,\tau_2)}{\partial \tau_2} =& \frac{2a_{22}e^{\tau_2 a_{22}}\left(e^{a_{11}\tau_2+a_{22}(\tau_1+\tau_2)}+e^{\tau_1(2a_{22}+a_{11})}+e^{a_{11}(2\tau_2+\tau_1)}+e^{3a_{11}\tau_1}\right)}{(e^{a_{11}\tau_2+a_{22}\tau_1}-e^{a_{11}\tau_1+a_{22}\tau_2})^3}\\ &-\frac{2a_{11}e^{\tau_2 a_{11}}\left(e^{a_{22}\tau_2+a_{11}(\tau_1+\tau_2)}+e^{\tau_1(2a_{11}+a_{22})}+e^{a_{22}(2\tau_2+\tau_1)}+e^{3a_{22}\tau_1}\right)}{(e^{a_{11}\tau_2+a_{22}\tau_1}-e^{a_{11}\tau_1+a_{22}\tau_2})^3}\end{aligned} \tag{3.27}$$

For this continuous-time case, we also encounter an analytical expression of the gradient that is challenging to analyze. We repeat the same analysis of the discrete-time case where we visualize the gradient field and contour lines of the objective function for the three stability cases.

Similarly to the discrete-time framework, the objective function is undefined along the diagonal $\tau_1 = \tau_2$, and it diverges close to it due to the vanishing denominator. We also find the same symmetry on each part of this diagonal, and then we can restrict our analysis to the natural case of $\tau_1 < \tau_2$.

The qualitative behavior of the objective function and the gradient is consistent with the discrete-time case. For the mixed stability system (Figure 3.4c), the optimal solution is at the earliest and latest sampling times, with the prior capturing the stable mode and the latter capturing the unstable mode.

In the fully stable case (Figure 3.4a) the sampling times are chosen near the beginning of the time interval, and for the fully unstable case (Fig 3.4b) they are chosen near the latest possible sampling times. However, the exact placement of these sampling times remains sensible to the eigenvalues of the system and is not easily explained; this stays consistent with the discrete framework. One hypothesis is that the divergence near the diagonal pushes these solutions to have separate sampling times. Even in cases where both times tend towards the beginning or the end of the interval, the strong gradient repulsion near the diagonal prevents sampling times from being too close. The exact spacing, however, is not easily explained from this analysis.

In summary, this analysis confirms that the qualitative behavior of these two-dimensional systems holds consistently in both continuous- and discrete-time cases. While the general patterns of optimal sampling strategies can be linked to system stability, some solutions remain difficult to interpret analytically or from the explored plots. This motivates the next chapter, which examines the sampling schedules in the state space. By simulating the resulting initial state estimates and trajectories in the state space, we aim to better understand the role of the optimal sampling times.

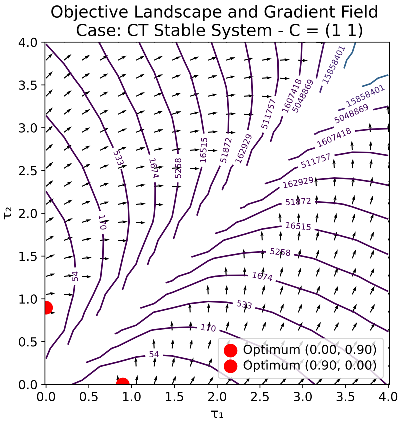

(a) Stable System

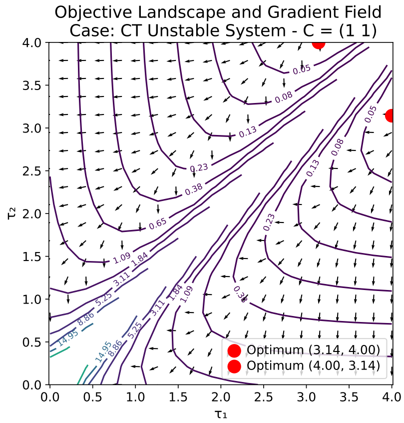

(b) Unstable System

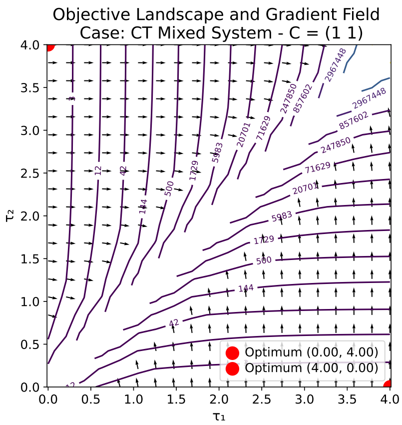

(c) Mixed Stability System

Figure 3.4: Gradient field visualization for Continuous-Time System with Output Matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$.

3.4 Effect of System Stability on Optimal Objective Values

In the previous sections, closed-form expressions of the objective function of the optimization problem were derived and analyzed for different stability properties and output matrices C. While the one-dimensional case offered a relatively straightforward analytical interpretation, the two-dimensional case introduced more nuanced behaviors.

In this section, we examine how the stability of the system impacts the *optimal value of the objective function*, which, as introduced in Chapter 2, is defined as the trace of the error covariance matrix of the least-squares initial state estimation problem. We study how this scalar measure of state reconstruction quality varies as a function of system stability.

We focus on the two-dimensional cases explored in previous sections, with A a diagonal system matrix

$$A = \begin{pmatrix} a_{11} & 0 \\ 0 & a_{22} \end{pmatrix}$$

and considering output matrices $C = I_2$ and $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$. For both discrete- and continuous-time settings, we use the closed-form expressions of the objective function derived in Sections 3.2 and 3.3. For each pair (a_{11}, a_{22}) ranging logarithmically from 0.1 to 100, we solve the optimization problem, as detailed in Section 4.2, to determine the optimal sampling times and compute the corresponding minimal value of the objective function.

These optimal values, plotted as a function of the system's eigenvalues; a_{11} and a_{22}, serve as a quantitative measure of the estimation error variance. A higher optimal value represents larger uncertainty in state estimation, while a lower value represents a tighter reconstruction of the initial state.

The resulting plots reveal a consistent behavior across all cases (discrete- and continuous-time, for both $C = I_2$ and $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$). When both eigenvalues are stable, the optimal value of the objective function reaches its maximum. As one of the two eigenvalues increases until it introduces instability, the optimal value decreases sharply. A single unstable eigenvalue significantly reduces the optimal objective value, and when both are unstable, the objective value tends toward zero.

It is interesting to discuss this behavior from the perspective of observability theory. As discussed in Chapter 2, minimizing the trace of the error covariance matrix is equivalent to minimizing the trace of the inverse of the observability Gramian. For unstable systems, the state grows over time, leading to large observable outputs that amplify the effect of the initial state and, consequently, they improve the conditioning of the observability Gramian. This leads to a lower estimation error, hence a *lower optimal objective value.*

In contrast, for stable systems, state trajectories decay and converge towards the origin. This decay reflects the system "forgetting" the information on the initial state, especially in later times where the information the output carries about the initial state is reduced. Even when sampling early, the observability Gramian remains relatively ill-conditioned, resulting in larger estimation variance, and higher values of the objective function.

The intermediate case of mixed stability reveals more complex behavior. As we saw in the gradient plots of Figures 3.10 and 3.3, the optimal sampling schedule places an early sample and a later one. This reflects a balance between capturing the decaying mode before it vanishes and capturing the unstable mode that will dominate the output.

It is interesting to note that in the specific case where the output matrix is the *identity* $C = I_2$ the optimal value of the objective function is bounded between 0 and 1. This can be linked to the state being fully measurable and equally observable in the state directions. Unlike the case $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$ where the information we have on each state variable is partial and where we see a divergence in the objective function.

The results highlight the quantitative link between system dynamics and state estimation performance, and they show that the objective function introduced reflects this relationship which helps to validate its usage.

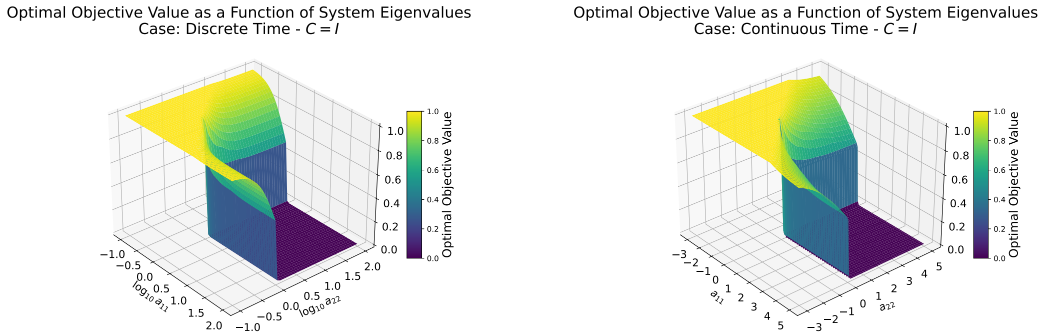

(a) Discrete-Time System (b) Continuous-Time System

Figure 3.5: Optimal objective values for output matrix $C = I$.

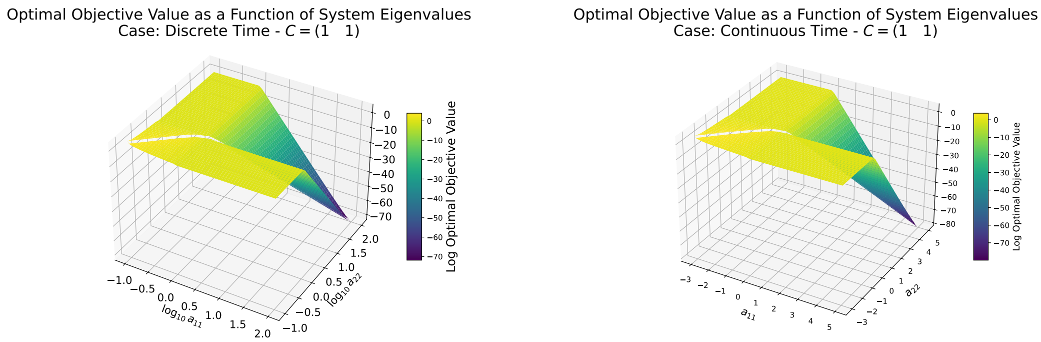

(a) Discrete-Time System (b) Continuous-Time System

Figure 3.6: Optimal objective values for output matrix $C = (1 \quad 1)$.

3.5 Conclusion

In this chapter, we performed a detailed analytical study of the objective function used to evaluate the quality of initial state estimation in LTI dynamical systems, focusing on both discrete- and continuous-time settings for particular cases of dynamical systems. We started by analyzing one-dimensional systems, and we derived closed-form expressions for the objective function and for the optimal sampling times for different system stability cases.

We then extended this analysis to two-dimensional systems with a diagonal state matrix whee the eigenvalues of the system and the choice of output matrix

C introduce more complex behaviors. We explored two output configurations: a fully observable case $C = I_2$ and a partially observable case $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$. After deriving the closed form of the objective functions and their gradients, we visualized the optimization landscape to interpret the optimal sampling schedules and their variation across different system dynamics and output measurements.

Finally, we studied how the stability of the system influences the optimal value of the objective function. We showed that instability promoted lower optimal objective values while stability leads to higher ones; which is directly related to the behavior of stable and unstable systems which, respectively, decay or amplify the value of the initial condition. Showing that the objective function captures this property of dynamical systems. We also showed that in the case $C = I_2$ the objective function is naturally bounded between 0 and 1, reflecting the well-conditioned estimation problem due to the independence of the measurements that capture the full state, unlike in partially observable scenarios where the estimation error can diverge.

These results reinforce the link between the statistical meaning of the objective function and its observability-related interpretation. They show that this objective function captures the different properties of the system dynamics and their link to system observability as well as the quality of the initial state reconstruction variance. However, the solutions are not always easy to interpret in this system-oriented way, and this motivates the next chapter, where we will simulate system dynamics to interpret the optimal solutions directly in the state space.

Chapter 4

Numerical Experiments

This chapter presents a series of numerical experiments exploring the behavior of the solutions of the proposed optimization problem 2.22. The goal of this optimization problem is to minimize the covariance of the initial state estimation error for a linear dynamical system observed through a least-squares observer. Thus, the aim is to find a set of sampling times, of predefined cardinality, that offers the most accurate reconstruction of the initial state.

After developing this problem and looking at the closed form of its objective function for the one-dimensional case and the two-dimensional case for particular diagonal systems, the objective of these numerical experiments is then to study the state reconstruction in the state space: we evaluate the quality of initial state reconstruction obtained using the optimal sampling times identified by solving the optimization problem. These results are compared with a baseline strategy that uses uniformly spaced sampling times. This comparison allows us to interpret the behavior of the solutions in the state space and to interpret them.

The code used to perform all simulations presented in this work is available at: `https://github.com/MJ-bared/Optimal_State_Estimation_From_Intermittent_Measures.git`.

4.1 Gradient Analysis

In Chapter 3, we derived closed-form expressions for the objective function 2.22 and its gradient in the examples of one-dimensional and two-dimensional systems, for discrete and continuous time. Through these expressions, an analytical analysis of the function was possible under different assumptions about the dynamics of the system matrix A and the observation matrix C.

To complement this theoretical analysis, we visualized the gradient field of the objective function with respect to the sampling times τ_1 and τ_2. The main goal of these visualizations is to confirm and illustrate the analytical solutions derived in the simpler cases and to gain qualitative understanding of the more complex cases. These plots were generated for different stability cases; stable, unstable, or mixed stability, and for the two observation matrices explored throughout the chapters $C = I_2$ and $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$. This section explains how these graphs seen in Figure 3.1, 3.2, 3.3 and 3.4 were constructed.

For each case, the gradient was computed on a grid of sampling times within a fixed time horizon. They are then represented as a vector field that illustrates the direction and relative magnitude of the gradient at each point. This allows us to characterize the shape of the objective function to try to make sense of the optimal times found in terms of local or global optima. We also plot the level curves of the objective function and the optimal sampling times obtained through numerical optimization as detailed in the next section 4.2, to make the link between the optimization landscape and the solutions.

4.2 Solving the Sampling Optimization Problem

In this section, we detail the numerical approach used to solve the optimization problem 2.22 formulated in Chapter 2. This problem consists in selecting a fixed number of sampling times within a finite interval $[0, \tau_{\max}]$ to minimize the uncertainty in the reconstructed initial state. More precisely, the aim is to minimize the trace of the inverse of the observability Gramian, or equivalently, the trace of the error covariance matrix, which serves as a metric for the quality of the estimation achieved by the least-squares observer of the initial state.

This approach is used throughout this chapter and in the numerical simulation and for the gradient plots discussed previously, where the computed optima are marked.

We focus on solving for the two-dimensional case where the problem is written as

$$\min_{\tau_1, \tau_2 \in [0, \tau_{\max}]} \quad \operatorname{tr}\left(W_{\mathcal{O}}(\tau_1, \tau_2)^{-1} \right).$$

This section explains how this problem was numerically solved for the different system dynamics and for the two observation scenarios in both discrete-time and continuous-time settings.

4.2.1 Discrete-Time Case

In the discrete-time setting, sampling times are restricted to integer values; $\tau_1, \tau_2 \in \{0, 1, \cdots, \tau_{\max}\}$. The optimization problem then becomes a discrete combinatorial minimization problem for all valid pairs in the set. The objective function can then generally exhibit multiple local minima.

Due to this complexity of the function and the discrete nature of the solutions, classical optimization methods are not best suited or evident to implement. Since we are generally studying two-dimensional systems on reasonable time horizons, we use an *exhaustive grid search strategy* for all possible pairs of sampling times $(\tau_1, \tau_2) \in \{0, \cdots, \tau_{\max}\}^2$. In two dimensions, this results in a computational complexity of $\mathcal{O}(\tau_{\max}^2)$ which is tractable for the moderate sizes of the time horizon used throughout the experiments.

However, this brute-force approach does not scale well to higher-dimensional systems. For n sampling times, the complexity grows as $\mathcal{O}(\tau_{\max}^n)$ making the computations impractical. This limitation motivates the exploration of more scalable solutions.

4.2.2 Continuous-Time Case

In the continuous-time case, the optimization problem becomes a continuous nonlinear minimization over the interval $[0, \tau_{\max}]$. To use the exhaustive grid search approach as for the discrete-time case, we discretize the interval into a grid of m points;

$$\mathcal{G} := \{t_1, t_2, \cdots, t_m\} \subset [0, \tau_{\max}].$$

We then evaluate the objective function on all possible pairs $(\tau_1, \tau_2) \in \mathcal{G} \times \mathcal{G}$, reducing the problem to a discrete one over a discretization of the continuous domain.

This results in an increased complexity compared to the discrete-time problem with a complexity of $\mathcal{O}(m^2)$ for two-dimensional systems and $\mathcal{O}(m^n)$ for the general n-dimensional case. Since continuous-time dynamics often require fine temporal resolution to accurately reflect the continuous behavior of the system, the grid $\mathcal{G}$ must be sufficiently dense, which increases the computational cost considerably.

Nevertheless, for the two-dimensional systems considered throughout this work, this approach remains tractable and has the advantage of being exhaustive and providing the optimal solution within the grid. It serves its purpose of providing the optimal sampling schedule solutions that we are interested in analyzing.

4.3 The Optimal Sampling Strategy and Initial State Estimation

In this section, we implement the least-squares estimator described in Chapter 2 to estimate the initial state of a two-dimensional LTI system from a fixed number of measurements taken following a certain sampling schedule. In this previous chapter, we introduced the least-squares estimator and computed the residual estimation error and its covariance. We then derived the objective function designed to minimize on average the covariance of this estimation error. In the previous section, we discussed the numerical approach to find the optimal sampling schedule that leads to a minimized trace of this covariance matrix.

Building on this previous work, we apply the reconstruction procedure for different dynamical systems and for different sampling schedules, then visualize the estimated initial state in the state space. The goal of this analysis is first to interpret the optimal sampling schedules in the state space to complement the gradient analysis performed previously. The second is to compare the reconstruction performance of the optimized sampling schedule compared to a simpler baseline schedule of uniformly spaced samples. This section details the reconstruction protocol and how the numerical experiments were performed, then we present and discuss the results of the simulations in the next one.

4.3.1 Problem Setup

We consider the reconstruction of the initial state $x_0 \in \mathbb{R}^2$ of a two-dimensional LTI system defined previously in discrete 3.5 or continuous time 3.17. Given a certain set of sampling times $\mathcal{S} = \{\tau_1, \tau_2\}$, the corresponding observability matrices $\mathcal{O}_{\mathcal{S}}$, 3.6 and 3.18, are constructed for each setting. The observed outputs at the selected sampling times are stacked in a vector $y_{\mathcal{S}}$.

The initial state is estimated using least-squares estimator 2.14, by using the Moore-Penrose pseudo-inverse of the observability matrix:

$$\hat{x}_0 = (\mathcal{O}_{\mathcal{S}}^\top \mathcal{O}_{\mathcal{S}})^{-1} \mathcal{O}_{\mathcal{S}}^\top y_{\int}.$$

To evaluate the quality of the reconstruction of the initial state, we compute the Relative Root Mean Squared Error (RRMSE) between the true initial state x_0 and the estimate $\hat{x}_0$. Beyond performance evaluation, we also analyze the geometry of the estimation process in the state space by studying how different output matrices C, sampling schedules, and noise levels affect the shape and spread of the least-squares solution.

4.3.2 Simulation Protocol

We present the simulation protocol used to evaluate the different sampling strategies for initial state estimation and their performance. These experiments are on two-dimensional LTI systems, for discrete- and continuous-time. However, since they have analogous results as was seen in the gradient analysis, we restrict ourselves to the discrete-time case since it is more restrictive with the integer constraints. We evaluate different cases for the system matrix A as well as different output matrices: the full state output matrix $C = I_2$ or the vector $\begin{pmatrix} 1 & 1 \end{pmatrix}$.

Each simulation follows the steps:

1. *System definition*: We choose a pair of matrices A and C, and define a true initial state $x[0] \in \mathbb{R}^2$.

2. *Noisy Measurements Generation*: For a given sampling schedule $\mathcal{S}$ and a noise level σ, we simulate noisy measurements of the outputs:

$$y_{\mathcal{S}} = \begin{pmatrix} y(\tau_1) \\ y(\tau_2) \end{pmatrix}, \qquad \text{with } y(\tau_i) = CA_i^{\tau} x[0] + v_i, \quad v_i \sim \mathcal{N}(0, \sigma^2 I).$$

3. *State estimation*: We apply the least-squares estimator to reconstruct an estimate $\hat{x}_0$ of the initial state

$$(\mathcal{O}_{\mathcal{S}}^{\top}\mathcal{O}_{\mathcal{S}})^{-1}\mathcal{O}_{\mathcal{S}}^{\top} y_{\mathcal{S}}.$$

4. *RRMSE evaluation*: We compute the Relative Root Mean Squared Error (RRMSE) between the estimated and true initial state, defined as:

$$\text{RRMSE} = \left(\mathbb{E}\left[\left\| \frac{x[0] - \hat{x}[0]}{\|x[0]\|} \right\|^2 \right] \right)^{\frac{1}{2}}.$$

This process is repeated for multiple random trials to evaluate the robustness of the estimation.

In addition to the RRMSE comparisons, we geometrically analyze the estimate in the state space by backpropagation of the measurements through the system. In the case of the identity observation matrix $C = I_2$, each measurement corresponds to a point in the state space. We backpropagate each measurement through the inverse dynamics of the system to obtain a *candidate initial state point*. In the case of the output matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$ each measurement defines a line of potential points in the state space: the set of points $x \in \mathbb{R}^2$ such that $CA^\tau x = y$. By backpropagating each point along this line, we obtain a corresponding set of candidate for the initial state.

This backpropagation visualization is performed to understand the geometry of the measurements and its relation with the system dynamics. It also aims to provide a geometric intuition on the solutions offered by the objective function while comparing it to other sampling strategies.

4.4 Case Studies and Analysis of Results

We now apply the simulation protocol described in the previous section to different system configurations as the ones previously explored. In the following, we consider a two-dimensional diagonal system matrix $A = \mathrm{diag}(a_{11}, a_{22})$, where the entries are also the eigenvalues of the system, allowing us to explore the effect of system stability on the estimation of the initial state. We analyze stable, unstable, and mixed stability systems and the effect of different sampling schedules on initial state reconstruction performance. For each case, we analyze the state reconstruction for the two output matrices $C = I_2$ and $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$.

In addition to the cases we explored through the previous chapters, we also include a non-diagonal case where the system matrix A has imaginary eigenvalues, resulting in a circular state trajectory in the interest of exploring the case of periodic systems.

4.4.1 Case 1: Full-state output matrix $C = I$

In this case, both components of the state are directly observed. The output vector at each sampling time provides full information on the state of the system: These outputs can be plotted as points in the state space.

We visualize the measurements taken following the sampling schedule and their backpropagation through the system. We also mark the reconstructed estimates of the initial state $\hat{x}[0]$ obtained from multiple trials. Over these repeated trials, we can see that the different initial state estimates $\hat{x}[0]$ form Gaussian-like clouds in the state space, reflecting the additive noise in the measurements. We analyze the results for different stability conditions.

Stable Systems (Figure 4.1). For stable systems, the state trajectory converges over time. This leads to a loss of information about the initial condition as time increases (as discussed in the gradient-based analysis 3). When using a uniformly spaced sampling schedule, we observe a higher variance in the estimator: multiple different initial states produce identical outputs, making the reconstruction ambiguous. In contrast, the optimized sampling strategy that tends to select the earliest measurement times preserves information about the initial state and offers a tighter covariance ellipse. We can also see that the RRMSE is generally lower for the optimal sampling schedule.

Unstable Systems (Figure 4.2). In unstable systems, trajectories diverge over time. This divergence magnifies the effect of the initial condition in later measurements, making it easier to estimate the initial state even with noisy observations. As a result, we observe a smaller difference between the performance of the optimized and baseline strategy. The optimized strategy still produces slightly mire compact estimation distributions. We can also see that the RRMSE is generally lower for the optimal sampling schedule.

Mixed-Stability Systems (Figure 4.3). In systems with one stable and one unstable eigenvalue, the behavior becomes asymmetric. In the baseline strategy, the unstable component is well estimated due to the divergence of the initial condition along that direction, but the stable component is poorly identified, leading to large uncertainty along its axis. The optimized strategy takes earlier times and balances the measurement schedule to reduce the overall estimation uncertainty, as defined by the objective function, which can be seen in the cloud of points where the uncertainty along the different directions is more balanced and the estimation error is lower.

Overall, the results show that the optimal sampling strategy improves the estimation precision, especially for stable systems. However, it would be interesting to look at cases where the outputs are not independent for each state variable and to analyze the solutions in the state space.

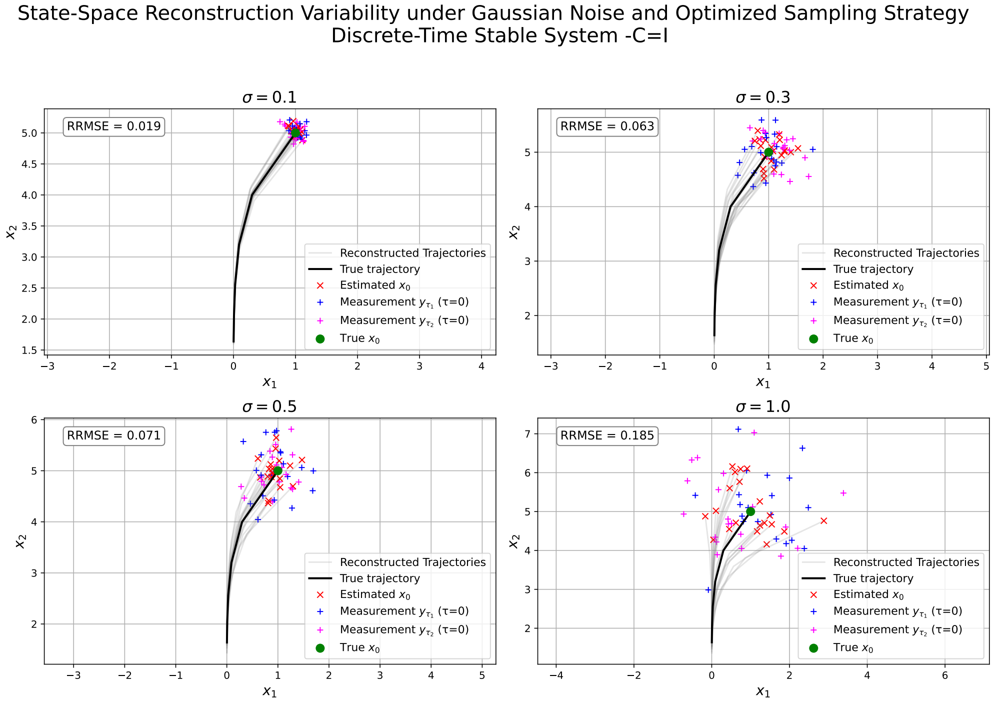

(a) Optimal Sampling Strategy

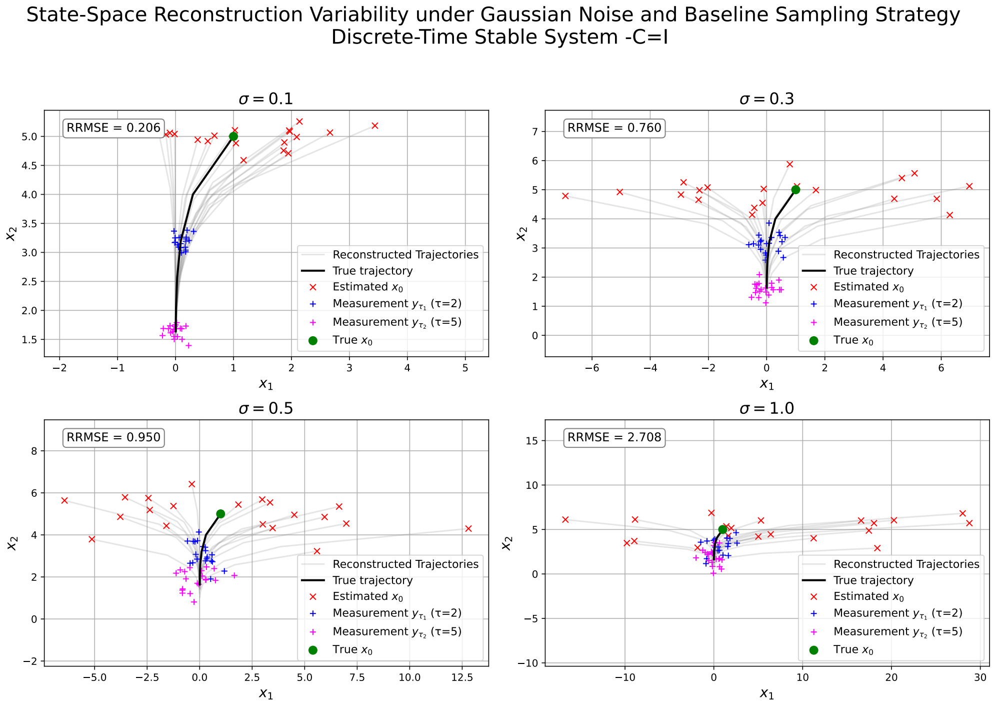

(b) Baseline Sampling Strategy

Figure 4.1: State Space Reconstruction Comparison for Optimal and Baseline Sampling Strategies - Stable System and State output matrix $C = I_2$

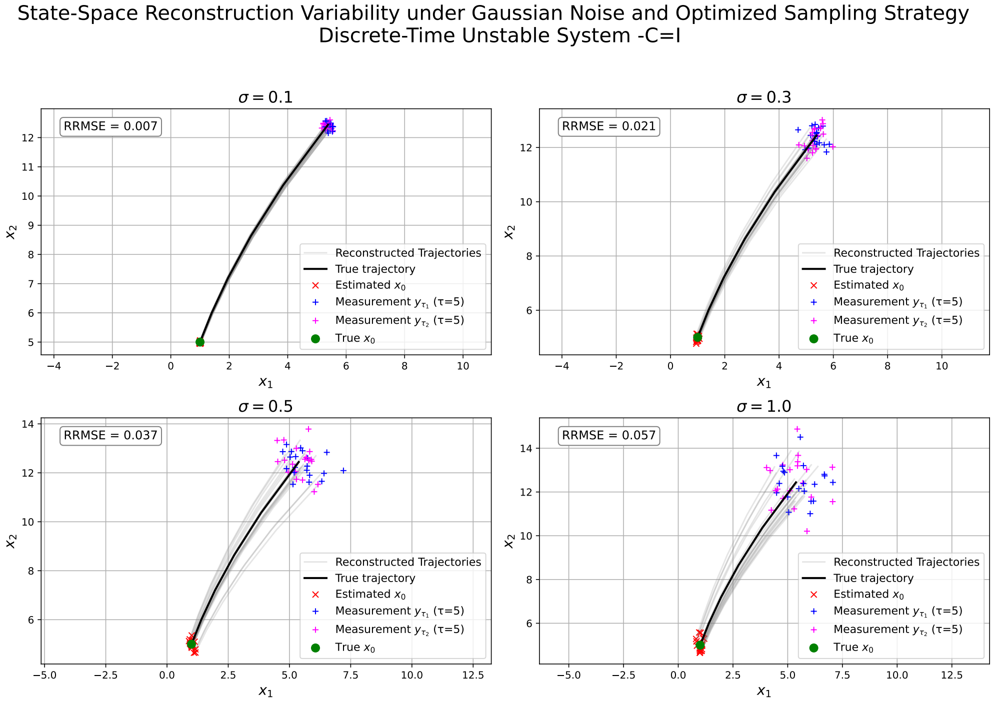

(a) Optimal Sampling Strategy

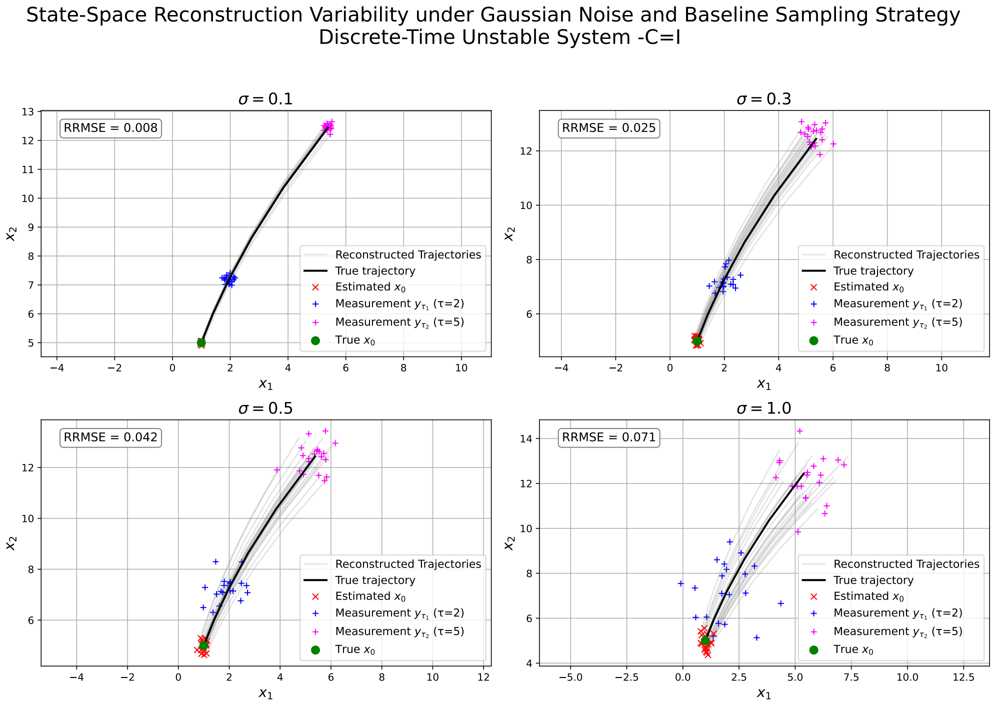

(b) Baseline Sampling Strategy

Figure 4.2: State Space Reconstruction Comparison for Optimal and Baseline Sampling Strategies - Unstable System and State output matrix $C = I_2$

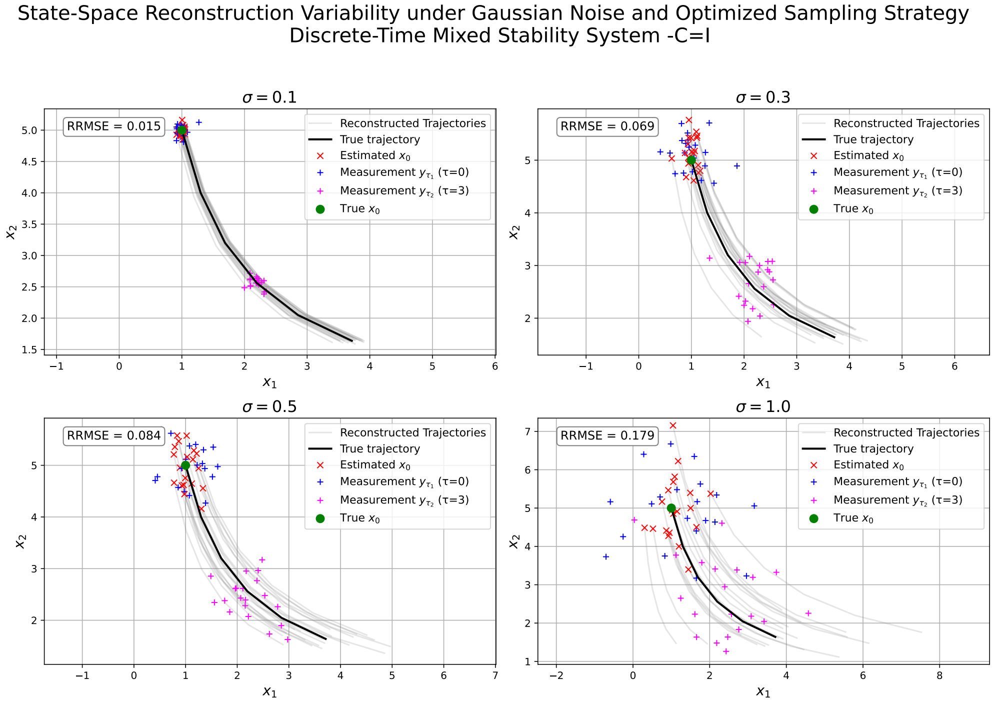

(a) Optimal Sampling Strategy

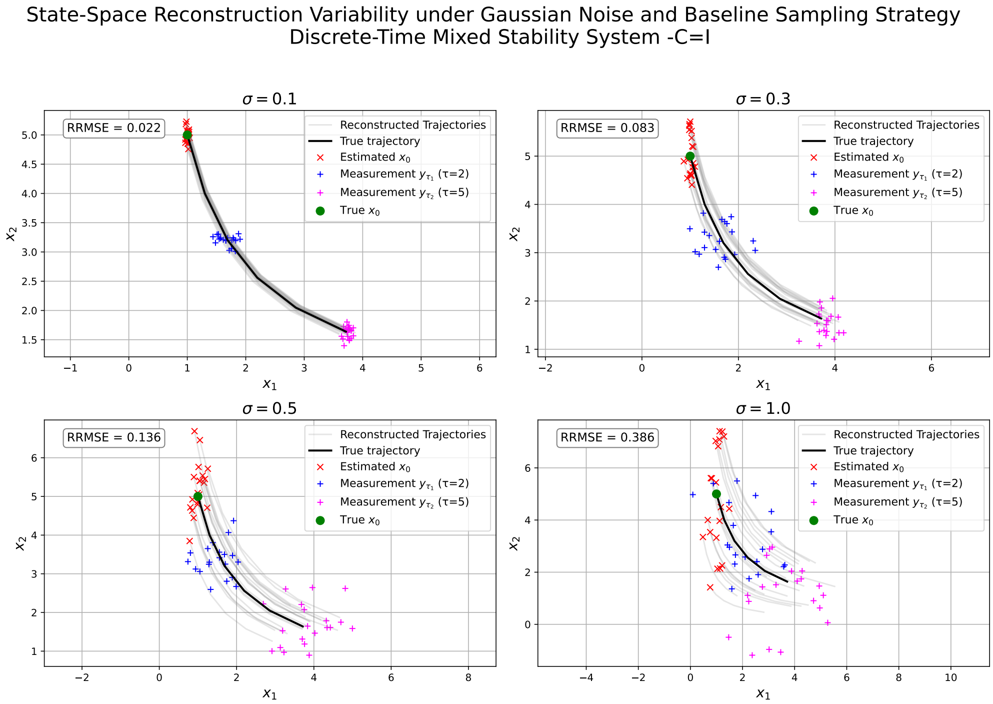

(b) Baseline Sampling Strategy

Figure 4.3: State Space Reconstruction Comparison for Optimal and Baseline Sampling Strategies - Mixed System and State output matrix $C = I_2$

4.4.2 Case 2: Linear combination of states observation matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$

In this case, we consider the output matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$ where the output at each measurement time is given by a scalar linear combination of the state variables. Each measurement offers only partial information about the system's true state.

Interpretation of Scalar Measurements as Lines in the State Space. We first start by interpreting the representation of an output $y[t]$ in the state space. A measurement taken at time τ_i has the expression

$$y[\tau_i] = Cx[\tau_i] = x_1[\tau_i] + x_2[\tau_i]. \tag{4.1}$$

Hence, each scalar output 4.1 defines a line in the two-dimensional state space:

$$\mathcal{L}_i := \{x[\tau_i] \in \mathbb{R}^2 | x_1[\tau_i] + x_2[\tau_i] = y[\tau_i]\}, \tag{4.2}$$

which is an affine line orthogonal to the vector $C^T = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$. Even though we do not have access to the full state, we know that the *system's true state at measurement time τ_i must lie on this line.*

To understand the information that this measurement line gives us on the initial state of the system $x[0]$, we backpropagate the points of this line through the inverse system. Thus, we obtain the set of possible initial states consistent with measurement $y[\tau_i]$:

$$\mathcal{L}_i^0 := \{x[0] \in \mathbb{R}^2 | CA^{\tau_i} x[0] = y[\tau_i]\}. \tag{4.3}$$

This set is also a line that will be defined for each measurement time τ_i and represents a constraint on the position of the initial state in the state-space. Considering that in the case of our system we take two measurements at distinct times τ_1 and τ_2, we obtain two lines $\mathcal{L}_1^0$, $\mathcal{L}_2^0$, whose intersection ideally gives us a good estimate of the true initial state $x[0]$.

Geometry of Estimation and Sampling Times Each line $\mathcal{L}_i^0$ corresponding to a sampling time τ_i has a direction orthogonal to the vector

$$d_i = CA^{\tau_i}.$$

By backpropagating two measurement lines, their intersection gives an estimate of $x[0]$. The robustness of this estimate depends on the angle between the two lines $\mathcal{L}_1^0$ and $\mathcal{L}_2^0$: their angle is the same as the angle between their respective orthogonal vectors d_1 and d_2, if the lines intersect at a small angle, small noise in the measurements can cause large errors in the estimate. If the vectors are close to orthogonal, the estimate becomes more robust.

This can be directly linked to the observability matrix:

$$\mathcal{O} = \begin{pmatrix} CA^{\tau_1} \\ CA^{\tau_2} \end{pmatrix},$$

the quality of the state estimation depends on the conditioning of this matrix. If its rows are nearly linearly dependent, then it is nearly rank-deficient, and its inverse becomes ill-conditioned. But if the rows are close to orthogonal, then it is well conditioned and the estimate is more stable. We can conclude that selecting sampling times that maximize the angular separation between $\mathcal{L}_1^0$ and $\mathcal{L}_2^0$ improves the information obtained from the measurements.

Discussion of simulation results. We now interpret the results of the simulations conducted for stable, unstable, and mixed-stability systems using the protocol described in Subsection 4.3.2. We tested different levels of measurement noise and visualized in each graph: the true initial state $x[0]$, the estimated initial states $\hat{x}[0]$ of each trial and the corresponding state trajectories, the averaged backpropagated measurement lines across the trials $\mathcal{L}_1^0$ and $\mathcal{L}_2^0$, and the relative root mean squared error (RRMSE) of the estimated initial state across trials.

The goal is to relate the numerical results from the optimization problem to the geometry of the solutions and to understand how optimal sampling times improve the observability and reduce the estimation error compared to a uniformly spaced sampling schedule.

Mixed Stability Case (Figure 4.6): This case highlights the geometrical interpretation of the objective function 3.14. The system has both a stable and an unstable mode, and the optimal sampling schedule selects one early and one later sampling time. This results in one measurement line capturing the decaying mode before it fades, and the other capturing the growing mode after it has diverged. Geometrically, this leads to backpropagated measurement lines that are *closer to orthogonal* than for the baseline sampling schedule. Each line reflects information from a different direction in the state space. We also observe a generally better RRMSE for the different noise levels compared to the baseline sampling strategy.

Fully Stable (Figure 4.4) and Fully Unstable (Figure 4.5) Cases: In both fully stable and fully unstable systems, the angle between the backpropagated lines under the optimized sampling schedule is often smaller than for the baseline. However, it still results in better reconstruction accuracy and with a variance of estimates that is less prominent on only one axis.

This observation motivates us to look at the backpropagated measurement lines of each trial: in Figure 4.7 we consider the stable case since it is the most pronounced one, however results are consistent across the different cases. We can see that for the optimal sampling schedule the measurement lines present a smaller variance than for the baseline. This shows that the optimization problem does not only seek to maximize the angle between measurements, but searches for a tradeoff between:

- **Measurement independence:** maximizing angular separation between the measurement directions to improve conditioning of the observability Gramian and ensure observability;
- **Measurement robustness:** selecting sampling times that reduce the effect of noise by minimizing the variance of the reconstructed state estimate.

Indeed, one might expect that the best reconstruction is always achieved when the measurement lines are as orthogonal as possible. However, the results show that this is not always the case. This tradeoff is interesting to link to the equivalence between the Covariance of the estimation and the observabilty Gramian defined in 2.20 considering that this is the relationship that defines our objective function.

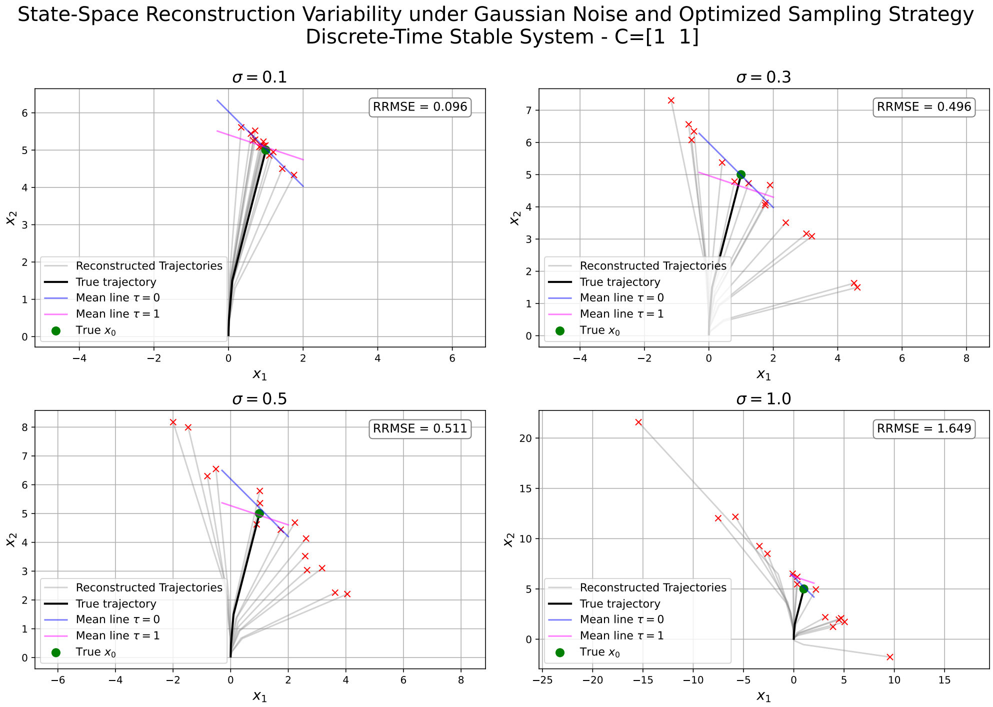

(a) Optimal Sampling Strategy

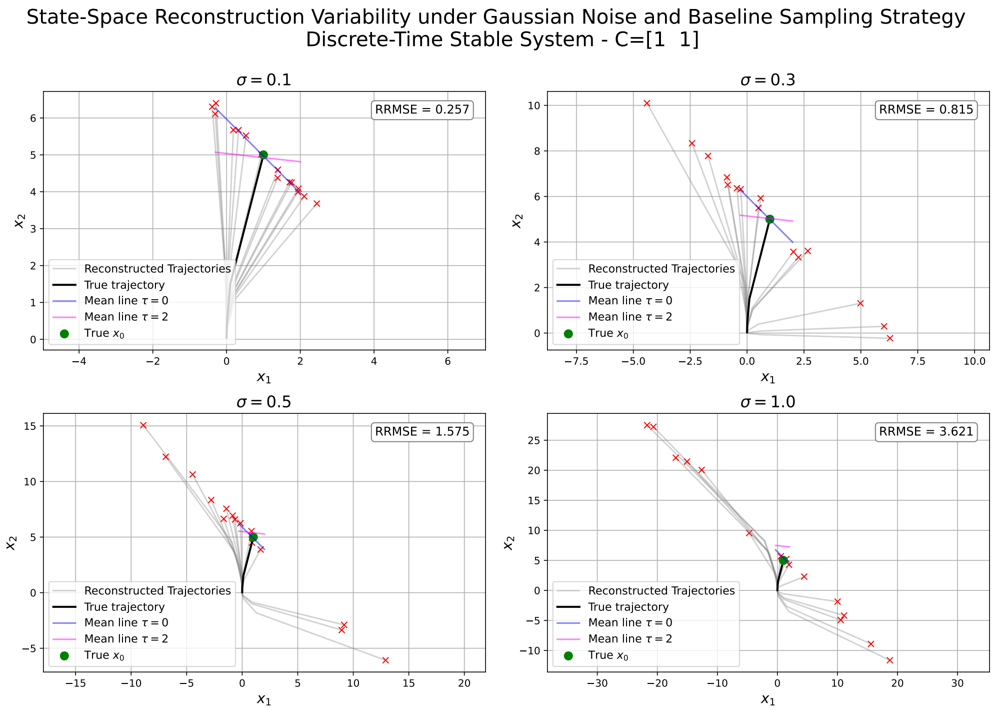

(b) Baseline Sampling Strategy

Figure 4.4: State Space Reconstruction Comparison for Optimal and Baseline Sampling Strategies - Stable System and State output matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$

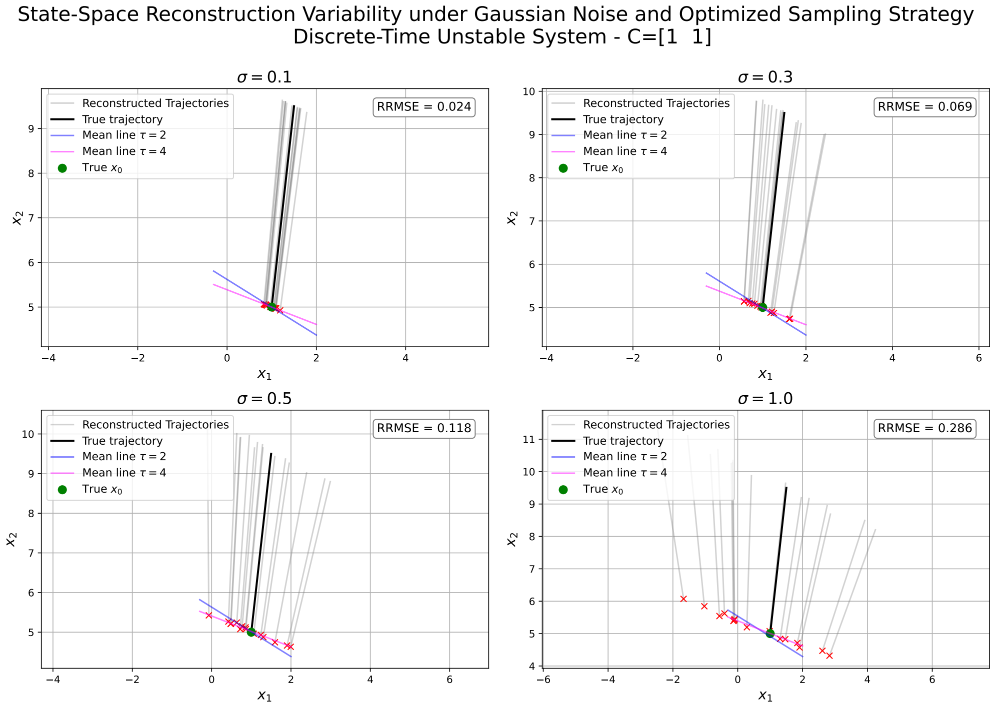

(a) Optimal Sampling Strategy

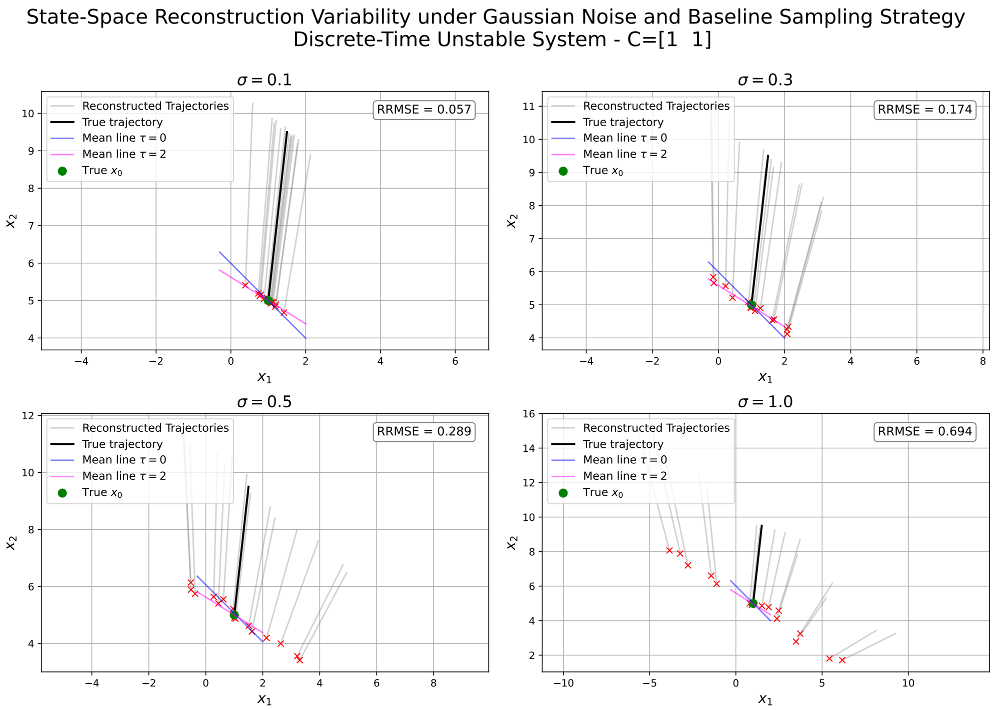

(b) Baseline Sampling Strategy

Figure 4.5: State Space Reconstruction Comparison for Optimal and Baseline Sampling Strategies - Unstable System and State output matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$

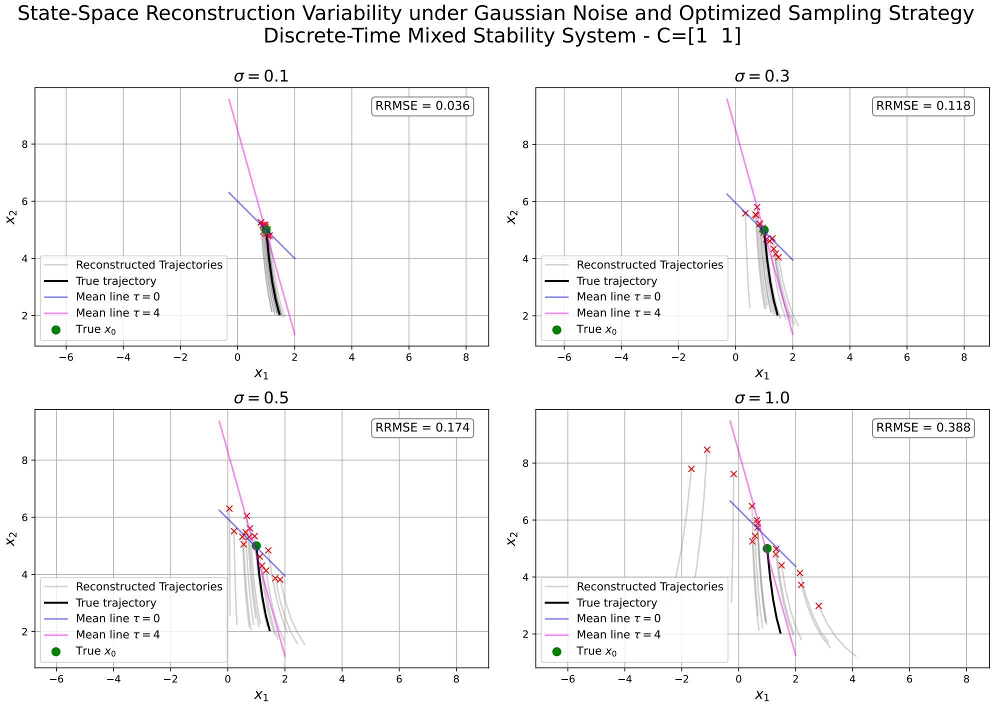

(a) Optimal Sampling Strategy

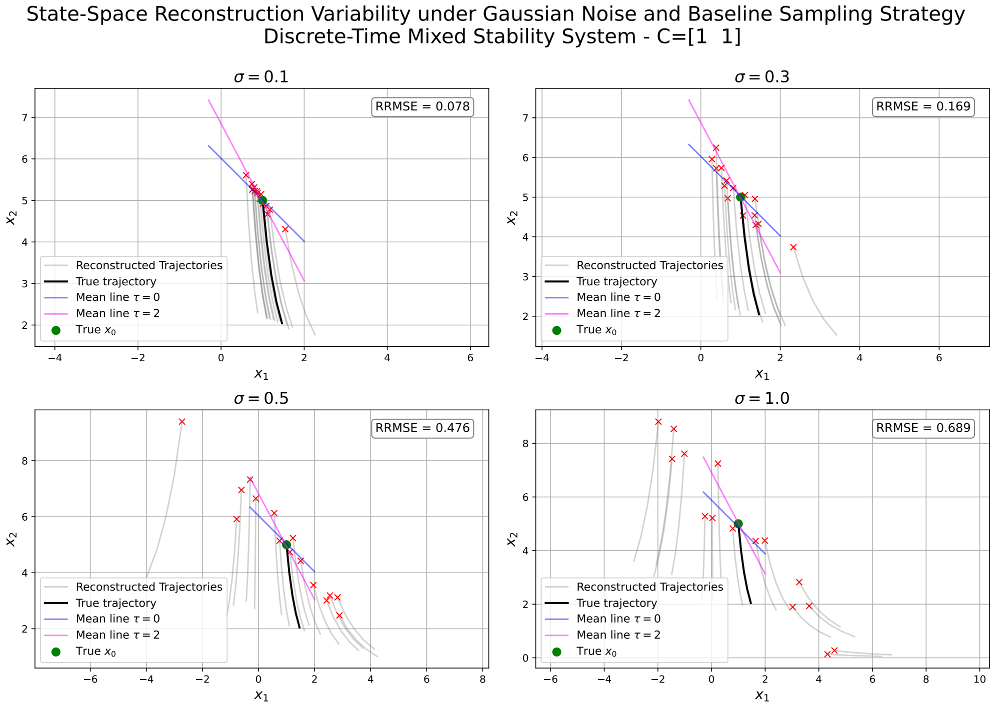

(b) Baseline Sampling Strategy

Figure 4.6: State Space Reconstruction Comparison for Optimal and Baseline Sampling Strategies - Mixed System and State output matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$

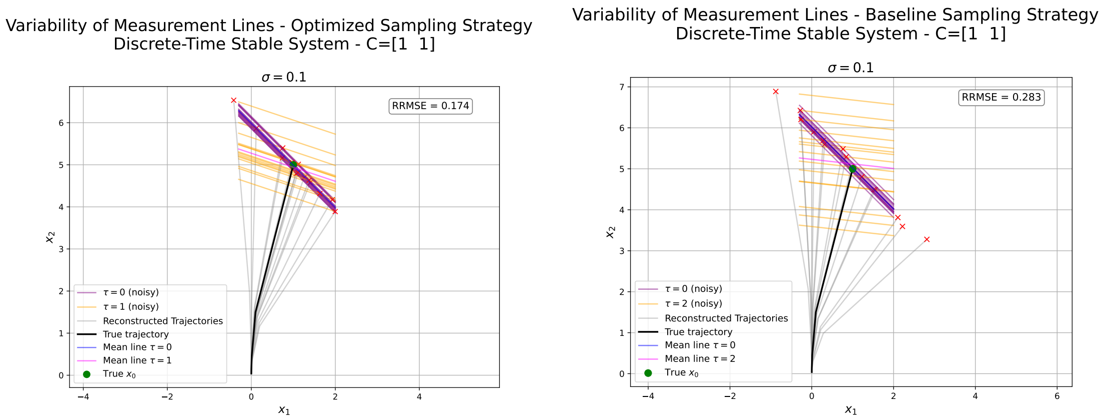

(a) Optimal Sampling Strategy

(b) Baseline Sampling Strategy

Figure 4.7: Measurement line variance for Optimal and Baseline Sampling Strategies - Stable System and State output matrix $C = \begin{pmatrix} 1 & 1 \end{pmatrix}$

Complex Eigenvalue Case - Periodic system (Figure 4.8): Finally, we consider a system with purely imaginary eigenvalues, where the state matrix A is

$$A = \begin{pmatrix} \cos(\frac{\pi}{20}) & -\sin(\frac{\pi}{20}) \\ \sin(\frac{\pi}{20}) & \cos(\frac{\pi}{20}) \end{pmatrix}, \tag{4.4}$$

this corresponds to a discrete-time rotation of the state vector by $\frac{\pi}{20}$ rad. at each time step, resulting in a circular trajectory in the state space.

Although this system has complex eigenvalues and a non-diagonal state matrix, it presents the same properties of solutions as in the previous cases. The optimization offers sampling times that induce highly orthogonal measurement lines. Interestingly, the sampling times are separated by an angle of $\frac{\pi}{2}$ and due to the periodic nature of the dynamics, these optimal sampling schedules also appear periodically as the time horizon increases. This case validates the idea that the optimization framework adapts to the system dynamics, even when the matrix is not diagonal, by selecting sampling times that maximize observability and minimize covariance of estimate.

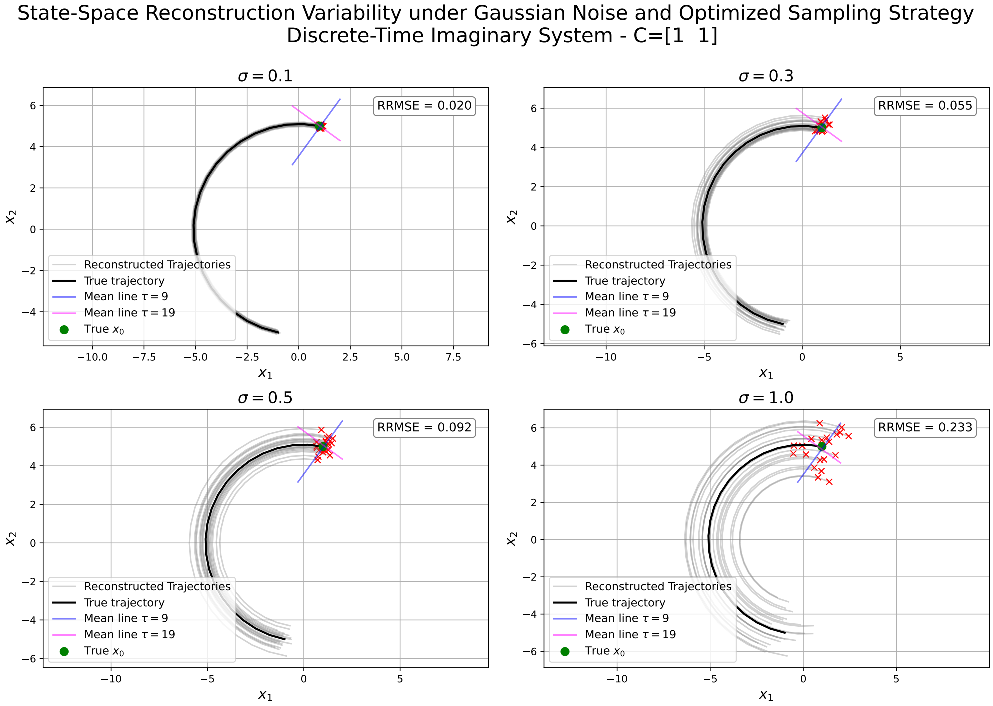

(a) Optimal Sampling Strategy

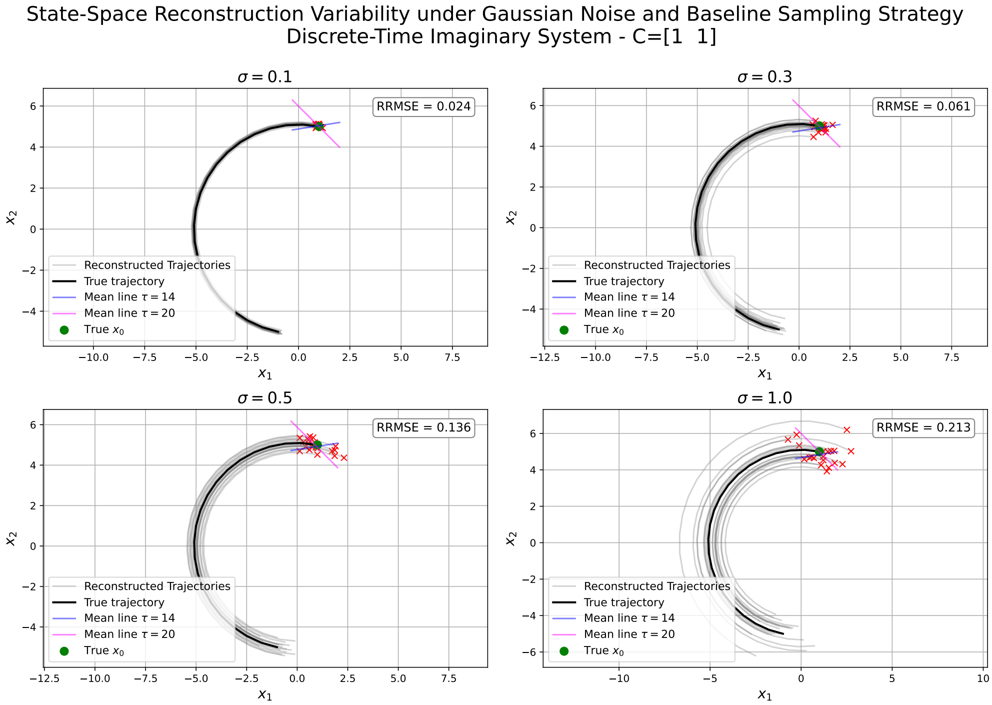

(b) Baseline Sampling Strategy

Figure 4.8: State Space Reconstruction Comparison for Optimal and Baseline Sampling Strategies - Periodic System and State output matrix $C = I_2$

Chapter 5

Conclusion

The purpose of this master thesis was to investigate the problem of reconstructing the initial state of a linear dynamical system from a finite set of measurements. In particular, we have formulated and studied the problem of optimally selecting a finite number of sampling times to improve the quality of the initial-state estimate, within the framework of LTI systems. Our work was structured in three steps: first, we laid out the theoretical foundations of the problem and formulated the optimization problem; second, we derived and analyzed the objective function in closed form for one- and two-dimensional specific systems; and third, we validated our approach through numerical simulations that explored the geometric and statistical implications of the solutions in the state space.

Summary

We began by laying out the theoretical background required to formulate the optimization problem. We discussed the concepts of observability, least-squares estimation, and the observability Gramian. We then derived an objective function that quantifies the quality of the reconstruction: the trace of the inverse of the observability Gramian. This quantity is directly proportional to the trace of the covariance matrix of the least-squares estimator, which is also the minimum-variance unbiased estimator in the Gaussian noise setting. We concluded that minimizing this objective corresponds to minimizing the average variance along the principal axes of the uncertainty ellipse.

We then analyzed the closed form of the objective function in low-dimensional systems to gain insight into the optimization landscape and optimal solutions. In the one-dimensional case, we derived explicit analytical solutions and showed that the choice of optimal sampling times is entirely governed by the stability of the system: for stable systems early sampling is optimal as information about the initial state decays over time and for unstable systems late sampling is optimal, since the state grows and becomes more observable. In the two-dimensional case, we considered diagonal systems with both identity and unit vector output matrices. We derived the closed form of the objective functions and their gradients and explored the optimization landscape. While the intuition of the one-dimensional case remained, the solutions presented were more complicated to interpret: the distance between the chosen sampling times is not easily explained, and we explored the mixed stability case where the optimizer selected sampling times that seemed to reflect a compromise between the two modes of the system.

To investigate these cases, we performed simulations in the two-dimensional state space for different sampling strategies and their performance. We compared the initial-state reconstruction using optimal sampling strategies and baseline sampling strategies. These experiments showed a tradeoff role played by the objective function: measurement independence; by maximizing independence of the measurement direction, consequently improving the conditioning of the observability Gramian, and measurement robustness; by minimizing the variance across trials and making estimation more stable. Finally, we extended the analysis to a system with imaginary eigenvalues. Despite its non-diagonal structure, the optimization framework naturally selected sampling times spaced $\frac{\pi}{2}$ apart, corresponding to orthogonal directions in the circular dynamics, and showing once again that the method adapts to the system's structure and exploits geometric features to enhance observability.

Research perspectives

While this master's thesis focused on one- and two-dimensional linear systems and provided a detailed study of the sampling optimization problem in specific settings, several points raise future exploration directions.

First, our approach relied on grid search to solve the optimization problem, which is not computationally scalable. Exploring more advanced optimization techniques such as heuristics or evolutionary algorithms [16] could make the solutions applicable to larger systems. This naturally leads us to the next direction: extending the analysis to higher-dimensional systems, which better reflect real-world scenarios. It would be interesting to compare optimal sampling schedules and baseline ones to evaluate the trade-off between performance and computational cost.

Another extension is to consider stochastic systems and systems with nonlinear dynamics. Such systems allow to model realistic cases, they offer a lot of complexity to the problem but also broaden the applicability.

Bibliography

[1] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

[2] P.C. Müller and H.I. Weber. Analysis and optimization of certain qualities of controllability and observability for linear dynamical systems. *Automatica*, 8(3):237–246, 1972.

[3] Mohammed Dahleh, Munther A. Dahleh, and George Verghese. Dynamic systems and control. `https://ocw.mit.edu/courses/6-302-feedback-systems-spring-2007/`, 2007. MIT OpenCourseWare, Massachusetts Institute of Technology.

[4] Young-Gu Lee and Hyochoong Bang. Relative state estimation of satellite formation flying using kalman filter. *IFAC Proceedings Volumes*, 41(2):2111–2116, 2008. 17th IFAC World Congress.

[5] Wei Zhang, Michael P. Vitus, Jianghai Hu, Alessandro Abate, and Claire J. Tomlin. On the optimal solutions of the infinite-horizon linear sensor scheduling problem. In *49th IEEE Conference on Decision and Control (CDC)*, pages 396–401, 2010.

[6] Michael P. Vitus, Wei Zhang, Alessandro Abate, Jianghai Hu, and Claire J. Tomlin. On efficient sensor scheduling for linear dynamical systems. *Automatica*, 48(10):2482–2493, 2012.

[7] Le Yi Wang, Chanying Li, G. George Yin, Lei Guo, and Cheng-Zhong Xu. State observability and observers of linear-time-invariant systems under irregular sampling and sensor limitations. *IEEE Transactions on Automatic Control*, 57(12):2996–3010, 2012.

[8] Haotian Zhang, Raid Ayoub, and Shreyas Sundaram. Sensor selection for optimal filtering of linear dynamical systems: Complexity and approximation. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 5002–5007, 2015.

[9] Morgan Henrie, Philip Carpenter, and R. Edward Nicholas. Chapter 4 - real-time transient model–based leak detection. In Morgan Henrie, Philip Carpenter, and R. Edward Nicholas, editors, *Pipeline Leak Detection Handbook*, pages 57–89. Gulf Professional Publishing, Boston, 2016.

[10] Abderrahman Iggidr and Max O. Souza. State estimators for some epidemiological systems. *Journal of Mathematical Biology*, 78(1):225–256, 2019.

[11] Giacomo Baggio, Fabio Pasqualetti, and Sandro Zampieri. Energy-aware controllability of complex networks. *IEEE Transactions on Control of Network Systems*, 7(2):1104–1114, 2020.

[12] Arman Rajaei, Mahsa Raeiszadeh, Vahid Azimi, and Mojtaba Sharifi. State estimation-based control of covid-19 epidemic before and after vaccine development. *Journal of Process Control*, 102:1–14, 2021.

[13] Julian Staiger, Lorenzo Mazzanti, and Frank Naets. State-oriented evaluation of observability and sensor placement for mechanical estimation applications. In *2023 IEEE International Conference on Mechatronics (ICM)*, pages 1–6, 2023.

[14] Julian Staiger, Lorenzo Mazzanti, and Frank Naets. State-oriented evaluation of observability and sensor placement for mechanical estimation applications. In *2023 IEEE International Conference on Mechatronics (ICM)*, pages 1–6, 2023.

[15] Lei Wang, Romeo Ortega, and Alexey Bobtsov. Observability is sufficient for the design of globally exponentially stable state observers for state-affine nonlinear systems. *Automatica*, 149:110838, 2023.

[16] Antoine Aspeel. Optimal sampling for state estimation of stochastic dynamical systems. Master's thesis, Université catholique de Louvain, 2024. Promoters: Benoît Macq and Raphaël Jungers.

[17] Gianluca Bianchin. Stochastic processes: Estimation and prediction. `https://uclouvain.be/en-cours-2025-linma1731`, 2025. Course at UCLouvain, LINMA1731.

[18] Chi-Tsong Chen. *Linear System Theory and Design.* Oxford University Press, New York, 3rd edition, 2004.

[19] Gilbert Strang. *Linear Algebra and Its Applications.* Brooks/Cole, 4 edition, 2006.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain
Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | **www.uclouvain.be/epl**