TAMPERE UNIVERSITY OF TECHNOLOGY

OLLI SUOMINEN

OPTIMAL MEASUREMENT SCHEDULING IN LQG SYSTEMS

Master of Science Thesis

Examiners: Professor Risto Ritala
Professor Robert Piché
Examiners and topic approved in Faculty of Automation, Mechanical and Materials Engineering Council meeting 8.6.2011

ABSTRACT

TAMPERE UNIVERSITY OF TECHNOLOGY
Master's Degree Programme in Automation Technology
SUOMINEN OLLI: Optimal Measurement Scheduling in LQG Systems
Master of Science Thesis, 36 pages
- 2011
Major: Measurement and Information Technology
Examiners: Professor Risto Ritala, Professor Robert Piché
Keywords: Optimal control, measurement, linear quadratic Gaussian, Kalman filter, dynamic programming, value iteration

Optimal control mathematically defines a policy so that the state of a system fulfills a certain performance criterion. Measurements provide an often incomplete method of observing the state of a stochastic process. When presented with a set of measurement alternatives the measurements must be chosen to enable optimal control of the process. In this thesis we seek to find a way of choosing optimal measurements from a set of discrete measurement alternatives in discrete time linear quadratic Gaussian systems (LQG).

The first part of the thesis presents the Kalman filter used to update information on the state of the system recursively. The Kalman filter relies on conditional probabilities and Bayes' rule. Information on the state is represented with the Gaussian distribution. The second part of the thesis derives the optimal control policies by first applying dynamic programming to a system where the state is known without uncertainty. This is expanded on by considering a system where measurements are used to provide information on the state. The derivation of the optimal control policy provides the objective function for the measurement problem, which may be analyzed separately from the control problem.

Two procedures for choosing the optimal measurement are presented. The first chooses a short prediction time for which an optimal measurement sequence is found by evaluating all possible sequences. The first measurement is performed and this procedure is repeated at each time step. This provides a local on-line solution to the measurement problem. The second procedure uses value iteration to optimize the measurement decisions. The space of the measurement problem is approximated with a finite grid. In this case the value function converges to the infinite horizon optimal value function. This method provides a global measurement policy off-line. Numerical studies are performed on three different systems to show the applicability of the optimization procedures.

TIIVISTELMÄ

TAMPEREEN TEKNILLINEN YLIOPISTO
Automaatiotekniikan koulutusohjelma
SUOMINEN, OLLI: Mittausten optimaalinen ajoittaminen LQG systeemeissä
Diplomityö, 36 sivua
- 2011
Pääaine: Mittausinformaatiotekniikka
Tarkastaja: professori Risto Ritala, professori Robert Piché
Avainsanat: Optimisäätö, mittaus, lineaari-kvadraatti-gaussinen, Kalman suodatus, dynaaminen ohjelmointi

Optimisäädössä määritetään säätöpolitiikka matemaattisesti täyttämään tietty tavoitekriteeri. Stokastisen prosessin tilaa mitataan säätöä varten. Jos mittausvaihtoehtoja on useampi on optimaalinen mittaus määritettävä säädön parhaan suorituskyvyn takaamiseksi. Tässä työssä pyritään valitsemaan optimaalinen mittaus joka hetkellä diskreetistä joukosta mittausvaihtoehtoja diskreettiaikaisessa lineaari-kvadraatti-gaussisessa järjestelmässä.

Työn ensimmäisessä osassa esitetään Kalman suodatus, jonka avulla tietoa prosessin tilasta voidaan rekursiivisesti päivittää. Kalman suodatus johdetaan esittämällä sen perustana toimivat ehdollinen todennäköisyys ja Bayesin sääntö. Tietoa prosessin tilasta kuvataan Gaussin jakauman avulla. Työn toisessa osiossa johdetaan säätöpolitiikka prosessin tilalle ensin tilanteessa, jossa prosessin tila tiedetään ilman epävarmuutta. Samaa ongelmaa tarkastellaan tilanteessa, jossa prosessin tilasta saadaan tietoa epävarmojen mittausten avulla. Tämän johdon avulla saadaan määritettyä tavoitefunktio mittauksen valinnalle. Tilan ja mittauksen säätöä voidaan tässä tapauksessa tarkastella erikseen.

Työssä esitetään kaksi optimointimenetelmää mittauksen valinnalle. Ensimmäisessä valitaan lyhyt ennustusaika, jonka kaikkia mahdollisia mittaussekvenssejä tarkastelemalla valitaan optimaalinen. Tästä sekvenssistä ensimmäinen mittaus suoritetaan ja seuraavalla ajanhetkellä optimointi toistetaan. Tämän menetelmän avulla saadaan lokaalisti optimaalinen mittaus valittua prosessin ajon aikana. Vaihtoehtoisesti mittausongelman optimaalinen arvofunktio määritetään koko mittausongelman avaruudessa. Arvofunktio konvergoituu tässä tapauksessa äärettömän horisontin tapaukseen. Tämän menetelmän avulla saadaan globaali mittauspolitiikka määritettyä ennen ajoa. Menetelmien toimivuus varmistetaan tarkastelemalla numeerisesti kolmea erilaista järjestelmää.

PREFACE

This thesis was made at Tampere University of Technology for the Department of Automation Science and Engineering during 2010-2011 with contributions from the Tampere University of Technology foundation. I would like to thank Professor Risto Ritala for providing the basis for the ideas and information in this work. His instructions and teaching has provided invaluable help during the process. I would also like to thank Professor Robert Piché for acting as an examiner.

The Department of Automation Science and Engineering has provided an interesting and encouraging workplace.

I appreciate all the support provided by my family and friends.

Tampere 14.6.2011
Olli Suominen

CONTENTS

1 INTRODUCTION

1.1 Background of the Thesis

The use of automation systems is prevalent throughout all facets of modern life. A central question in systems research has always been how to control a system. Ever since the 1950s optimal control theory has been an important part of control systems theory. Optimal control constitutes a way of deriving mathematically the optimal control policies when considering a certain performance criterion.

Consider the task of navigating a ship from one port to the next. The ship's starting location may be considered its initial state in the system and reaching the destination port the goal. The controller is able to determine the ship's motion by manipulating its rudder and engine output. There exist an infinite number of routes which realize the end goal. Some are, however, not realistic due to technical and physical limitations. The available routes are restricted to areas with water deep enough for the ship. The ship's rudder has a maximum angle and the engine's power is limited which present limitations on the control resources available to the controller. The goal for the controller is to determine the best plan of action whose optimality is defined by for example, the shortest travel time, least fuel expenditure or a combination of these.

The dynamic properties of the system determine how control influences the evolution of the system state. A chemical process is influenced by concentrations of the ingredients, temperature and catalysts among other variables. The controller uses these properties with a physical model to predict the state of the system under each plan of action. The time horizon of the prediction affects the optimization and depends on the problem. A chemical process may reach a terminal state at some time horizon or one may want to continue the process indefinitely.

Measurements provide a way for the controller to determine the state of a system and to compensate for unpredictable circumstances. The cruise control of a car knows the appropriate throttle level which sustains the given velocity on level ground. An uphill or downhill incline will however affect the speed. By measuring the speed the controller is able to modify the throttle level to achieve the desired result. Measurements often provide only incomplete information about the true state of the system and the controller must form an estimate of the system state.

Making a measurement may have a cost, and the control must decide if the measurement is worth performing. Measurement cost may arise if personnel are required or the measurement method consumes some chemical or other resources. For example, X-ray diffraction is a common laboratory measurement in metalworking

which determines stress levels in processed metals, unobservable otherwise. Most x-ray devices require a technician to prepare the sample and to manipulate the device. Etching is required if information on the internal stress levels is to be obtained thus destroying the sample.

Many systems contain multiple mutually exclusive measurement options amongst which one must be chosen to support control actions. For example, a driver chooses which direction to look when driving a car. Additionally the driver may wish to look at the speedometer or fuel gauge. An additional example can be found in the process industries where measurement systems exist that analyze multiple sampling lines and only one analysis may be made at a time.

One of the fundamental optimal control problems is the problem of a linear system perturbed by additive Gaussian noise, in which information on the true state of the system is incomplete, and in which the objective is to minimize a quadratic function of the system state. Control of these linear-quadratic Gaussian (LQG) systems under rather general assumptions separates into an estimation and regulation problem. This property of LQG systems is referred to as the separation principle.

The optimal control in LQG systems is derived with dynamic programming. Dynamic programming was first formulated by Richard E. Bellman [1]. Dynamic programming solves a complex multistage control problem by analyzing the problem one stage at a time. The optimal solutions are combined to form the complete solution.

The optimal control for time-invariant systems where no optimization horizon is determined may be solved with dynamic programming. The value function determining the value of the optimal objective function over the whole state space is iterated until it converges. This method is referred to as value iteration.

1.2 Objectives of the Thesis

The problem of a discrete time LQG system which has a discrete set of measurement choices to choose between is the core problem considered in this work. The optimal control problem pertaining to the regulation of the system has been studied comprehensively [2]. We show the derivation of the optimal control policies with regard to the control and measurement choice.

The optimal measurement policy is fundamentally different from the regulation problem as it is not analytically solvable. We show that the objective function relating to the measurement problem seeks to minimize a combination of the estimation error and the incurred costs.

Two methodologies to solve the measurement problem are presented. Because of its structure both methods are approximate. The first method works by considering only a limited time horizon. It works by predicting the evolution of the system at each time step to find the optimal measurement sequence. The first control is used from the sequence and the optimization is repeated at the next time step. Thus this method

provides a control valid only in the present system state and requires the control to be calculated on-line. This is the solution procedure as presented by Meier et al. [3,4].

In the second method we apply value iteration to the measurement problem. This allows us to calculate the policy before any application, i.e. off-line. The value iteration method seeks to provide a measurement policy for all possible cases. The space of the measurement problem is approximated with a finite grid.

1.3 Contents of the Thesis

The basic theory pertaining to the estimation of a system state through use of measurements needs to be shown before deriving the optimization procedures. Chapter 2 reviews the estimation procedure applied in linear Gaussian systems. Bayes' theorem and conditional probabilities are used to update information on the state of a system. The estimation uses the Gaussian probability distribution to represent the information on the state. These lead to the recursive estimation procedure, the Kalman filter.

Chapter 3 derives the optimization procedure by applying dynamic programming. We start by applying dynamic programming to a system where the state is known without uncertainty. This is expanded to the case where measurements are used to observe the state. Finally we present how value iteration may be applied to the measurement problem.

The optimization is applied to three different process control problems in Chapter 4 with numerical simulation in Matlab [5]. We start with a simple problem with a one-dimensional state space. The structure of the optimal measurement policy is then shown in a two-dimensional problem. Finally the on-line solution is applied to a four-dimensional control problem.

Chapter 5 summarizes and concludes the thesis with a final comparison of the optimization procedures. Possible future work and improvements are also discussed.

2 BACKGROUND: ESTIMATION

This chapter derives the recursive estimation procedure used to update the information on the process state as new measurements become available. The process is defined in discrete time with the linear relation:

$$x_{k+1} = A_k x_k + B_k u_k + w_k$$

where x_k denotes the state at the kth time step, u_k the control and w_k additive disturbances, distributed identically and independently. The state is not necessarily directly observable. The measurements y_k are defined with the linear relation:

$$y_k = C_k x_k + v_k$$

where v_k are additive disturbances, distributed identically, mutual independently and independently of w_k. Here the disturbances are modeled with Gaussian distributions. This allows the information on the state to be presented with the Gaussian distribution. The information prior to a measurement and the information provided by a measurement are combined with Bayes' theorem.

2.1 Conditional Probability and Bayes' Theorem

The conditional probability $p(A|B)$ is the probability of event A occurring given the occurrence of some other event B. Conditional probability can be defined through the joint probability $p(A, B)$ of A and B occurring as:

$$p(A|B) = \frac{p(A, B)}{p(B)} \quad \text{where } p(B) > 0 \tag{1}$$

Consider now that we receive a measurement y and wish to infer information on the state x. Using conditional probabilities this may be done with Bayes' theorem:

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} \tag{2}$$

The prior distribution $p(x)$ represents the information on the state before receiving the observations. The conditional distribution $p(y|x)$ is referred to as the measurement model. It models the causal but noisy or inaccurate relationship between the state and measurements. The posterior distribution $p(x|y)$ represents the information when all the data in the observations and information prior to the observations is used.

The divisor $p(y)$ is a normalization factor which ensures that the total probability is 1. It may be written as:

$$p(y) = \int p(y|x)\, p(x) dx \tag{3}$$

As an example consider the problem of estimating today's temperature. From the previous years' temperature data a prior distribution $p(x)$ for temperature may be formed. The process of taking a temperature reading can be modeled with a measurement model $p(y|x)$. The knowledge of today's temperature $p(x|y)$ is updated with Bayes' theorem when a temperature reading is available. Furthermore, the new posterior distribution can be used to form the prior for the problem of estimating tomorrow's temperature.

2.2 The Gaussian Distribution

Bayesian inference provides the method for computing the posterior distribution for any model. However, often the calculations are very demanding and unsolvable analytically. The relevant distributions may be approximated with some simpler and more easily handled distributions. A typical choice is the Gaussian distribution.

A justification for the choice of the Gaussian distribution is that it maximizes entropy thus minimizing the prior information built into the choice of distribution, if only mean and covariance of the random variable are known. Choosing the Gaussian distribution is also convenient as it allows for comparatively easy calculations. Furthermore the Central Limit Theorem states that under very general conditions given n independent random variables the distribution of their sum approaches a Gaussian distribution as n grows without limit [6, pp.214-219]. Thus the Gaussian distribution is a natural choice for modeling many different phenomena.

A k-dimensional vector random variable X is multivariate Gaussian distributed, denoted $X \sim N(\mathrm{x}; \mu, \Sigma)$, if its probability density function is given as

$$f_X(x) = \frac{1}{(2\pi)^{k/2}|\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)\right) \tag{4}$$

A multivariate Gaussian distribution is fully parameterized by its mean vector $\mu \in \mathbb{R}^k$ and covariance matrix $\Sigma \in \mathbb{R}^{kxk}$. The covariance matrix must be positive-semidefinite. $|\Sigma|$ denotes the determinant of the covariance matrix.

The multivariate Gaussian distribution is a continuous probability distribution and identifies the probability of the value falling in an infinitesimal interval $(x, x + dx)$. In the same way as any probability distribution the sum of all possible states, here the integral of the multivariate Gaussian distribution over the whole sample space, equals 1. This is ensured by the normalization factor, the prefactor to the exponential expression in Equation (4).

By denoting the ith element of X as X_i, the i,jth element of the covariance matrix may be expressed as:

$$\Sigma_{i,j} = E\{(X_i - E\{X_i\})(X_j - E\{X_j\})\} \tag{5}$$

Equation (5) with $i = j$ is referred to as the variance σ_i^2 of the ith element and describes the uncertainty of the element. The greater the variance is, the wider the probability distribution. Figure 1 illustrates how a change in the parameters influences the form of the probability density function in a univariate case.

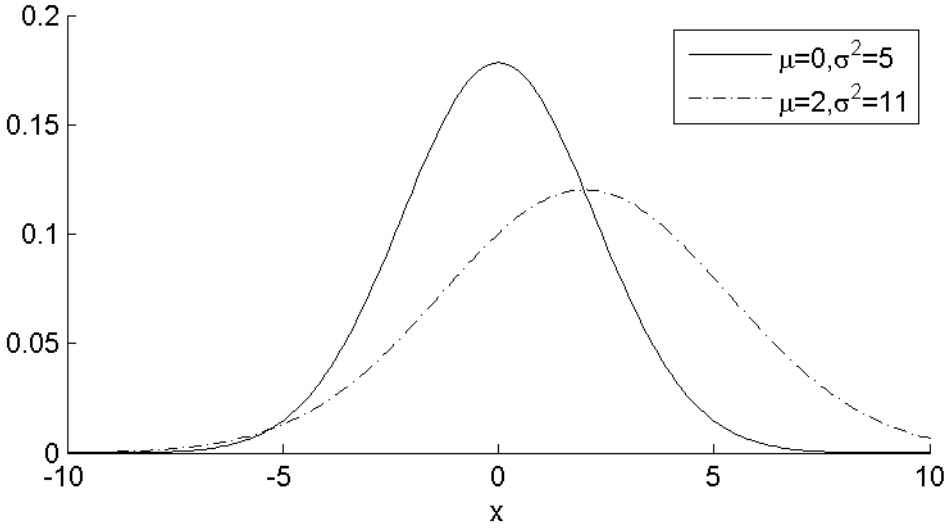

Figure 1: Two univariate Gaussian probability density functions with different parameters.

Joint probabilities are useful when combining information. Some of their properties are expressed as in the following two results. Let the joint probability of the random variables X and Y be:

$$\begin{bmatrix} X \\ Y \end{bmatrix} \sim N\left(\begin{bmatrix} x \\ y \end{bmatrix}; \begin{bmatrix} a \\ b \end{bmatrix}, \begin{bmatrix} A & C \\ C^T & B \end{bmatrix}\right)$$

Assuming B is nonsingular, the marginal and conditional distributions of X are:

$$\begin{gathered} X \sim N(x; a, A) \\ X|Y \sim N(x; a + CB^{-1}(y - b), A - CB^{-1}C^T) \end{gathered} \tag{6}$$

Additionally let $X \sim N(x; a, A)$ and the conditional distribution $Y|X \sim N(y; Qx + b, B)$ then the joint probability of X and Y is:

$$\begin{bmatrix} X \\ Y \end{bmatrix} \sim N\left(\begin{bmatrix} x \\ y \end{bmatrix}; \begin{bmatrix} a \\ Qa + b \end{bmatrix}, \begin{bmatrix} A & AQ^T \\ QA & QAQ^T + B \end{bmatrix}\right) \tag{7}$$

2.3 Linear Gaussian Processes

To carry out the intended analysis and optimization of the system we must present the definitions and assumptions related to the behavior of the system. This subsection presents the assumed transition, control and observation models.

Any analysis of a system relies on a definition of the process state. The state may be a number, collection of numbers or a function that describe the behavior of the system. The state space consists of all values of the system state. For example the position of an object may be described by its three-dimensional coordinates. In some cases this may not be sufficient and can be expanded, for example, by including the orientation angles of the object to the state space. To describe the motion of the object, the state space would also include the linear and angular velocities of the object. Such a definition would constitute a continuous state space as the state variables are real-valued.

A discrete-time dynamic system can be defined in terms of states $x_k \in \mathbb{R}^n$, controls $u_k \in \mathbb{R}^p$ and random disturbances $w_k \in \mathbb{R}^n$ referred to as the process noise. The subscript denotes the time instant and runs from 0 to a terminal time horizon N. The special case of a linear system is considered:

$$x_{k+1} = A_k x_k + B_k u_k + w_k \tag{8}$$

The matrices A_k of size nxn and B_k of size nxp are assumed given. The matrices will be assumed constant if the time index is omitted. The disturbances w_k are assumed to be independent random vectors with a known probability distribution. Here we assume a Gaussian n-dimensional distribution $N(w_k; 0, \Sigma_W)$. The process noise is assumed to be independent of the state x_k and control u_k.

We consider the case in which the state is not directly observable. The measurement model of Equation (9) specifies the observation's relation to the state and the uncertainty in the observation. Any observation on the state x_k provides data y_k which is on average linearly related to the state by the matrix $C_k^{(m_k)}$ and perturbed by an additive Gaussian white noise $v_k^{(m_k)}$. The superscript (m_k) is used to differentiate between different measurements. Thus $m_k = 1,2,\ldots$ depending on how many different measurements there are to choose from.

$$y_k = C_k^{(m_k)} x_k + v_k^{(m_k)} \tag{9}$$

The measurement disturbances $v_k^{(m_k)}$ are assumed to be independent random vectors with the distribution $N\left(v_k^{(m_k)}; 0, \Sigma_M^{(m_k)}\right)$. Thus the distribution of the measurement noise depends on which of the measurements is chosen.

With these assumptions the transition and measurement models can be described by the distributions:

$$p_S(x_{k+1}|x_k, u_k) = N(x_{k+1};\ A_k x_k + B_k u_k, \Sigma_W) \tag{10}$$

$$p_M^{(m_k)}(y_k|x_k) = N(y_k;\ C_k^{(m_k)} x_k, \Sigma_M^{(m_k)}) \tag{11}$$

2.4 The Kalman Filter

The Kalman filter is an efficient, optimal and recursive state estimation method for systems described by Equations (10) and (11) [7]. The Kalman filter works by predicting a value, estimating the uncertainty of the predicted value, and combining the prediction with the measured value. In essence it provides a weighted average where the weights are derived from the uncertainty. Though not originally defined with the Bayesian approach, the derivation of the filtering equations may be performed with application of Bayes' theorem [8].

First we note that the linear Gaussian system defines a Markov process. A Markov process is a sequence of random variables $\{X_1, X_2, X_3, \dots\}$ in which the probability of a transition to state $X_{n+1} = x_{n+1}$ depends only on the state $X_n = x_n$, not on earlier states. The fact that the transition does not depend on any history of the system before its arrival in x_n is referred to as the Markov property. The Markov property is defined as:

$$p(x_{n+1}|x_1, x_2, \dots, x_n) = p(x_{n+1}|x_n) \tag{12}$$

Further the measurements y_k are assumed conditionally independent given x_k. Thus the measurements are independent of past measurements $\{y_0, y_1, \dots, y_{k-1} = y_{1:k-1}\}$ or past values of the state $\{x_0, x_1, \dots, x_{k-1} = x_{0:k-1}\}$:

$$p(y_k|x_{0:k}, y_{0:k-1}) = p(y_k|x_k) \tag{13}$$

To derive the Kalman filter formulation we start by assuming that at time k the available control and measurement history results in an estimate for the state of the form:

$$p(x_k|y_{0:k}, u_{1:k-1}) = N\left(x_k; \mu_{k|k}, \Sigma_X^{(k|k)}\right) \tag{14}$$

The control u_k is applied but before the measurement y_{k+1} is taken we wish to form information on the state x_{k+1}. Due to the Markov property of the system we may write the joint probability of x_k and x_{k+1} as:

$$p(x_{k+1}, x_k|y_{0:k}, u_{1:k}) = p(x_{k+1}|x_k, u_k)p(x_k|y_{0:k}, u_{1:k-1}) \tag{15}$$

The first probability distribution $p(x_{k+1}|x_k, u_k)$ is provided as Equation (10) and the second $p(x_k|y_{1:k}, u_{1:k-1})$ is the assumed information. Using Equation (7) the joint probability is:

$$p(x_k, x_{k+1}|y_{0:k}, u_{1:k}) = N\left(\begin{bmatrix} x_k \\ x_{k+1} \end{bmatrix}; \begin{bmatrix} \mu_{k|k} \\ A\mu_{k|k} + Bu_k \end{bmatrix}, \begin{bmatrix} \Sigma_X^{(k|k)} & \Sigma_X^{(k|k)} A_k^T \\ A_k \Sigma_X^{(k|k)} & A_k \Sigma_X^{(k|k)} A_k^T + \Sigma_W \end{bmatrix}\right) \tag{16}$$

The distribution for x_{k+1} is the marginal distribution given by integrating over x_k. By Equation (6):

$$\begin{aligned} p(x_{k+1}|y_{0:k}, u_{1:k}) &= \int p(x_{k+1}, x_k|y_{0:k}, u_{1:k})\, dx_k \\ &= N\left(x_{k+1}; A_k\mu_{k|k} + B_k u_k, A_k\Sigma_X^{(k|k)}A_k^T + \Sigma_W\right) \\ &= N\left(x_{k+1}; \mu_{k+1|k}, \Sigma_X^{(k+1|k)}\right) \end{aligned} \tag{17}$$

The last row defines the notation for the predicted distribution parameters. The above is referred to as the prediction step as the state is predicted with information prior to time $k+1$.

At time $k+1$ a measurement y_{k+1} is recieved and we wish to update the estimate of the state given all the information now available. We may apply Bayes' theorem and recalling that the measurements were assumed conditionally independent write:

$$\begin{aligned} p(x_{k+1}|y_{k+1}, y_{0:k}, u_{1:k}) &= Dp(y_{k+1}|x_{k+1})p(x_{k+1}|y_{0:k}, u_{1:k}) \\ &= Dp(x_{k+1}, y_{k+1}|y_{0:k}, u_{1:k}) \end{aligned} \tag{18}$$

where D is the appropriate normalization factor. The distributions are the measurement model, Equation (11), and the predictive estimate for the state specified by Equation (17). We may again form the joint probability of y_{k+1} and x_{k+1} which by Equation (7) is:

$$\begin{aligned} p(x_{k+1}, y_{k+1}|y_{0:k}, u_{1:k}) = N\Bigg(&\begin{bmatrix} x_{k+1} \\ y_{k+1} \end{bmatrix}; \begin{bmatrix} \mu_{k+1|k} \\ C_{k+1}^{(m_{k+1})}\mu_{k+1|k} \end{bmatrix}, \\ &\begin{bmatrix} \Sigma_X^{(k+1|k)} & \Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T \\ C_{k+1}^{(m_{k+1})}\Sigma_X^{(k+1|k)} & C_{k+1}^{(m_{k+1})}\Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T + \Sigma_M^{(m_{k+1})} \end{bmatrix}\Bigg) \end{aligned} \tag{19}$$

Using the notation:

$$S_{k+1}^{(m_{k+1})} = C_{k+1}^{(m_{k+1})}\Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T + \Sigma_M^{(m_{k+1})}$$

$$K_{k+1}^{(m_{k+1})} = \Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T {S_{k+1}^{(m_{k+1})}}^{-1}$$

By Equation (6) the conditional distribution is given as:

$$\begin{aligned} &p(x_{k+1}|y_{k+1}, y_{0:k}, u_{1:k}) \\ &= N\Big(x_{k+1}; \mu_{k+1|k} + K_{k+1}^{(m_{k+1})}\left(y_{k+1} - C_{k+1}^{(m_{k+1})}\mu_{k+1|k}\right), \Sigma_X^{(k+1|k)} \\ &\quad - K_{k+1}^{(m_{k+1})}S_{k+1}^{(m_{k+1})}{K_{k+1}^{(m_{k+1})}}^T\Big) \end{aligned} \tag{20}$$

This provides the updated information about the state as D is independent of x_{k+1}. Thus we may write the prediction step as:

$$\begin{aligned} \mu_{k+1|k} &= A_k\mu_{k|k} + B_k u_k \\ \Sigma_X^{(k+1|k)} &= A_k\Sigma_X^{(k|k)}A_k^T + \Sigma_W \end{aligned} \tag{21}$$

And the update step:

$$\begin{aligned} \mu_{k+1|k+1} &= \mu_{k+1|k} + K_{k+1}^{(m_{k+1})}\left(y_{k+1} - C_{k+1}^{(m_{k+1})}\mu_{k+1|k}\right) \\ \Sigma_X^{(k+1|k+1)} &= \Sigma_X^{(k+1|k)} - K_{k+1}^{(m_{k+1})} S_{k+1}^{(m_{k+1})} K_{k+1}^{(m_{k+1})^T} \end{aligned} \tag{22}$$

where $\mu_{k+1|k+1}$ is the resulting state estimate and $\Sigma_X^{(k+1|k+1)}$ its uncertainty. If a measurement is not available, the update step may be left out. $K_{k+1}^{(m_{k+1})}$ is referred to as the Kalman gain as it governs how the difference between actual measurement data and the expected measurement prior to the actual measurement is incorporated into the estimate.

These two steps provide the optimal filtering equations of the Kalman filter provided that initial information on the state is available.

It is noteworthy that the ensuing covariance $\Sigma_X^{(k+1|k+1)}$ of the estimate is independent of the control u_k. This may be shown by examining the prediction step. Let x_{k+1} and $\tilde{x}_{k+1}$ have the following linear relation to x_k:

$$\begin{aligned} x_{k+1} &= A_k x_k + B_k u_k + w_k \\ \tilde{x}_{k+1} &= A_k x_k + w_k \end{aligned}$$

By the joint distribution and marginal distribution the predictive covariance $\Sigma_X^{(k+1|k)}$ of x_{k+1} and $\Sigma_{\tilde{X}}^{(k+1|k)}$ of $\tilde{x}_{k+1}$ are the same:

$$\Sigma_X^{(k+1|k)} = \Sigma_{\tilde{X}}^{(k+1|k)} = A\Sigma_X^{(k|k)}A^T + \Sigma_W^{(k)}$$

In the update step only the measurement covariance influences $\Sigma_X^{(k+1|k+1)}$ thus the value of u_k does not influence the uncertainty of the estimate. The covariance is also independent of the realized value of y_{k+1}.

3 BACKGROUND: OPTIMAL CONTROL

The performance of a process is evaluated and optimized with regard to an objective function. The objective may be, for example, to maximize utility, minimize costs or keep the process state close to a given setpoint. In optimal control the goal is to derive an optimal policy which maps any value of the state to the control which best fulfills the given objective.

The linear Gaussian system presented above is considered with a quadratic cost function. The optimal control law may be solved by applying dynamic programming. The control is solved analytically. The measurement selection and control problems separate and may be analyzed independently.

3.1 Dynamic Programming

Dynamic programming is used to solve multistage planning problems. The dynamic programming approach breaks the multistage problem into a recursive algorithm using functional equations.

The goal in dynamic programming is to minimize the accumulating cost defined by

$$V(x_0) = \min_{\{u_k\}_{k=0}^{T}} E\left\{\sum_{k=0}^{T} \alpha^k F_k(x_k, u_k, w_k)\right\} \quad \text{subject to } x_{k+1} = T_k(x_k, u_k, w_k) \tag{23}$$

$F_k(x_k, u_k, w_k)$ denotes the instantaneous cost functions. The complete sequence of controls $\{u_k\}_{k=0}^{T}$ is to be minimized. The transition from one state to the next is restricted by the transition function $x_{k+1} = T_k(x_k, u_k, w_k)$ defined by the dynamics of the process. Here the cost is considered up to a terminal horizon T. Because of uncertainty in the future state of the system or as a modeling of impatience, immediate costs may be valued more than future costs. This devaluation of future costs is modeled by the parameter α referred to as the discount factor, with $0 < \alpha \leq 1$. The minimization task is expressed by the following sum where in the case of a stochastic process the expectation must be considered.

The founding block for the optimization procedure is provided by Bellman's Principle of Optimality [9,p.15]. It states: "Whatever the initial state and decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the initial decision."

The dynamic programming approach works by defining a sequence of functions $\{V_k(x_k)\}_{k=0}^{T}$. First $V_T(x_T)$ is defined as the value function for all states x_T which minimizes the expected immediate cost at the last time instant:

$$V_T(x_T) = \min_{u_T} E\{F_T(x_T, u_T, w_T)\}$$

Next the problem is considered at time $T-1$:

$$V_{T-1}(x_{T-1}) = \min_{u_T, u_{T-1}} E\{F_{T-1}(x_{T-1}, u_{T-1}, w_{T-1}) + \alpha E\{F_T(x_T, u_T, w_T)\}\}$$

$$= \min_{u_{T-1}} E\left\{F_{T-1}(x_{T-1}, u_{T-1}, w_{T-1}) + \alpha \min_{u_T} E\{F_T(x_T, u_T, w_T)\}\right\}$$

Written in this form we may note that the latter minimization has already been solved and write:

$$V_{T-1}(x_{T-1}) = \min_{u_{T-1}} E\{F_{T-1}(x_{T-1}, u_{T-1}, w_{T-1}) + \alpha E\{V_T(x_T)\}\}$$

$V_T(x_T)$ provides the optimal cost-to-go. The general iteration for the functions $V_k(x_k)$ is:

$$V_{k-1}(x_{k-1}) = \min_{u_{k-1}} E\{F_{k-1}(x_{k-1}, u_{k-1}, w_{k-1}) + \alpha E\{V_k(x_k)\}\} \quad (24)$$

Finally the minimum value of $V(x_0)$ is given by $V_0(x_0)$.

Equation (24) is the basic functional equation used to solve dynamic programming problems. By considering the multistage problem one stage at a time the optimization is clearly less complex. Alternatively the functional equation can be written with forward induction where rather than the optimal cost-to-go the optimal prior cost is considered. A more thorough derivation and discussion of the previous equation may be found, for example, in [9,pp.1-15]. The solution provided by the dynamic programming approach is not guaranteed to be unique but its optimality is ensured.

The major problem now is how to exactly compute the iterative step between the functions $V_k(x_k)$. It is clearly impossible or very demanding to tabulate values for every x_k in each step even for a finite set of values. The simplest solution is to define a finite grid from which values are interpolated in some manner. However, some problems exhibit properties that make an analytical solution possible.

The main source of complexity is in the dimension of the state variable where adding a state variable makes the problem exponentially more complex. This is referred to as the curse of dimensionality.

The complexity of the problem also increases with the length of the time horizon. When only limited resources are available in optimizing a system to some time horizon T_1, it is possible to form an approximate procedure which at some time d performs the optimization up to time horizon $T_2 < T_1$. At time $d+1$ the same optimization is performed up to $T_2 + 1$. At each time instant only the first control is used. In process control this idea forms the basis for Model Predictive Control [10,11].

3.2 Linear Quadratic Gaussian Systems and Dynamic Programming

In LQG systems the cost function is defined as a quadratic function. The goal is to minimize the sum:

$$E\left\{\sum_{k=0}^{T}\alpha^k\left((x_k-\bar{x}_k)^T Q_k(x_k-\bar{x}_k)+u_k^T R_k u_k\right)\right\} \tag{25}$$

The matrices Q_k and R_k are weighting matrices of appropriate size which are assumed positive semidefinite symmetric and positive definite symmetric respectively. The weighting matrices are specified when designing the system and define which system state dimensions are seen as important in relation to each other. The quadratic function is often used in the formulation of regulator problems [12,p.130]. For example a regulator may be used to control the voltage to govern engine speed. It expresses the desire to keep the state close to a given trajectory $\{\bar{x}_0, \bar{x}_1, \dots\}$. The quadratic cost function penalizes large deviations from the set point more than smaller deviations.

We begin by deriving the solution in terms of a system where the current state is known exactly. Here we assume that the goal is to keep the state close to the origin.

$$\begin{gathered}\min_{\{u_k\}_{k=0}^{T}} E\left\{\sum_{k=0}^{T}\alpha^k\left(x_k^T Q_k x_k + u_k^T R_k u_k\right)\right\} \\ \text{subject to } x_{k+1} = A_k x_k + B_k u_k + w_k \\ \text{where } N\left(w_k; 0, \Sigma_W^{(k)}\right),\ E\{w_i w_j^T\} = 0\ \forall\ i \neq j\end{gathered} \tag{26}$$

Equation (26) reiterates the complete optimization problem where all the matrices defining the system dynamics $\{A_i, B_i\}_{i=1}^T$ and cost function $\{Q_i, R_i\}_{i=1}^T$ are known, in addition to the covariance $\Sigma_W^{(k)}$ of the process noise. The process noise is assumed to be independent of the state and control.

The goal is to find an optimal control policy $\pi_k(x_k)$ which maps values of the state x_k to values of u_k:

$$u_k = \pi_k(x_k)$$

We will prove inductively that any $V_k(x_k)$ function of the dynamic programming approach is a quadratic function of the form:

$$V_k(x_k) = x_k^T S_k x_k + s_k \tag{27}$$

where S_k is symmetric positive semidefinite and the scalar s_k independent of the state and control.

For $k = T$, $V_T(x_T)$ is:

$$V_T(x_T) = \min_{u_T} E\{x_T^T Q_T x_T + u_T^T R_T u_T\} \tag{28}$$

The quadratic form $u_T^T R_T u_T$ is strictly positive for all $u_T \neq 0$ as R_T was assumed positive definite symmetric. Thus the optimal control u_T^* is simply $u_T^* = 0$.

This means that the function $V_T(x_T)$ is quadratic:

$$V_T(x_T) = x_T^T Q_T x_T$$

For the inductive step we assume $V_{k+1}(x_{k+1})$ is of the quadratic form of Equation (27) such that for $V_k(x_k)$ we have the form:

$$V_k(x_k) = \min_{u_k} E\left\{x_k^T Q_k x_k + u_k^T R_k u_k + \alpha E\{V_{k+1}(x_{k+1})\}\right\}$$
$$= \min_{u_k} E\{x_k^T Q_k x_k + u_k^T R_k u_k + \alpha E\{x_{k+1}^T S_{k+1} x_{k+1}\} + \alpha s_{k+1}\}$$

For the inner term:

$$E\{x_{k+1}^T S_{k+1} x_{k+1}\}$$
$$= E\{x_{k+1}\}^T S_{k+1} E\{x_{k+1}\}$$
$$+ E\{(x_{k+1} - E\{x_{k+1}\})^T S_{k+1}(x_{k+1} - E\{x_{k+1}\})\}$$

As $E\{w_k\} = 0$ and given the value of x_k:

$$E\{x_{k+1}\} = A_k x_k + B_k u_k$$

Thus:

$$E\{(x_{k+1} - E\{x_{k+1}\})^T S_{k+1}(x_{k+1} - E\{x_{k+1}\})\} = E\{w_k^T S_{k+1} w_k\}$$

By differentiating with respect to u_k and setting the derivative to zero, gives:

$$(R_k + \alpha B_k^T S_{k+1} B_k) u_k = -\alpha B_k^T S_{k+1} A_k x_k$$

Thus the optimal control u_k^* is:

$$u_k^* = -\alpha(R_k + \alpha B_k^T S_{k+1} B_k)^{-1} B_k^T S_{k+1} A_k x_k$$

Substituting the expression of u_k^* for u_k gives the required quadratic form:

$$V_k(x_k) = x_k^T S_k x_k + s_k$$

$$\text{where } S_k = Q_k + \alpha A_k^T S_{k+1} A_k - \alpha^2 A_k^T S_{k+1} B_k \left(R_k + \alpha B_k^T S_{k+1} B_k\right)^{-1} B_k^T S_{k+1} A_k \tag{29}$$

$$s_k = \alpha\left(E\left\{w_k^T S_{k+1} w_k\right\} + s_{k+1}\right)$$

Thus the assumed quadratic form holds for any k.

This provides the optimal control law $\pi_k^*(x_k)$ as the linear relation:

$$\pi_k^*(x_k) = L_k x_k$$

$$\text{where } L_k = -\alpha\left(R_k + \alpha B_k^T S_{k+1} B_k\right)^{-1} B_k^T S_{k+1} A_k \tag{30}$$

The recursion for matrices S_k:

$$S_T = Q_T$$

$$S_k = Q_k + \alpha A_k^T S_{k+1} A_k - \alpha^2 A_k^T S_{k+1} B_k \left(R_k + \alpha B_k^T S_{k+1} B_k\right)^{-1} B_k^T S_{k+1} A_k \tag{31}$$

We shall use the trace function denoted $\mathrm{tr}(\cdot)$ to further simplify the expression. The trace function has the property: $\mathrm{tr}(UVX) = \mathrm{tr}(XUV) = \mathrm{tr}(VXU)$.

As S_k is clearly a symmetric square matrix, we may write:

$$E\{w_k^T S_{k+1} w_k\} = E\{\mathrm{tr}(w_k^T S_{k+1} w_k)\} = E\{\mathrm{tr}(S_{k+1} w_k w_k^T)\} = \mathrm{tr}\left(S_{k+1} \Sigma_W^{(k)}\right)$$

The optimal cost is provided by the last step, if the initial state and its covariance are distributed as $N\left(x_0; \mu_0, \Sigma_X^{(0)}\right)$, and is:

$$V_0(x_0) = \mu_0^T S_0 \mu_0 + \mathrm{tr}\left(S_0 \Sigma_X^{(0)}\right) + \sum_{k=0}^{T-1} \alpha^{k+1} \mathrm{tr}\left(S_{k+1} \Sigma_W^{(k)}\right)$$

In this case the final cost function in itself is not required as the calculation of the matrices S_k provides all the information needed to determine the control. The optimal control law may be seen as a feedback controller. This is illustrated in Figure 2. The current state determines the control and thus provides the next state.

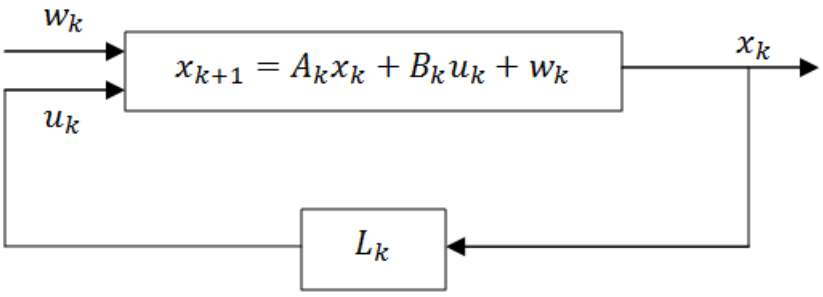

Figure 2: Linear feedback structure of the optimal control

The following generalizes this approach to deal with the case where the current state is not exactly known.

3.3 Imperfect State Information

In the previous section it was assumed that the controller has direct access to the state. A common problem is to find a similar derivation for the problem where state information is not perfect. As previously mentioned in the linear Gaussian model these observations are modeled with a linear measurement model. The observations y_k are perturbed by the zero mean Gaussian noise $v_k^{(m_k)}$ whose covariance matrix $\Sigma_M^{(m_k)}$ is known. This noise is also assumed independent of the state, control and process noise. Each measurement choice is associated with a positive scalar cost given by the function $h(m_k)$. The linear relation is:

$$y_k = C_k^{(m_k)} x_k + v_k^{(m_k)}, \quad k = 0, \ldots, T-1$$

$$\text{where } v_k^{(m_k)} \sim N\left(0, \Sigma_M^{(m_k)}\right), \qquad E\left\{v_i^{(m_i)} v_j^{(m_j)^T}\right\} = 0 \;\forall\; i \neq j \tag{32}$$

$$h(m_k) > 0$$

We seek to minimize the sum:

$$\min_{\{u_k\}_{k=0}^{T},\{m_k\}_{k=1}^{T}} E\left\{\sum_{k=0}^{T}\alpha^k\left(x_k^T Q_k x_k + u_k^T R_k u_k + h(m_{k+1})\right)\right\} \quad (33)$$

$$\text{subject to } x_{k+1} = A_k x_k + B_k u_k + w_k, \quad N\left(w_k; 0, \Sigma_W^{(k)}\right)$$

With these assumptions the state information is obtained recursively with the Kalman filter. We assume that at each time k the measurement data $\{y_k\}_{i=0}^{k}$ is available. At each time instant k the measurement to be performed at time $k+1$ is chosen. The estimate for the state at time k $E\{x_k\}$ is denoted $\mu_{k|k}$ and its covariance $\Sigma_X^{(k|k)}$. We assume the initial state information is $N\left(x_0; \mu_{0|0}, \Sigma_X^{(0|0)}\right)$. Thus the value function, in general depending on current state information is given as a function of the mean and covariance, $V_k\left(\mu_{k|k}, \Sigma_X^{(k|k)}\right)$.

We use the following result where we assume the measurement data at time T has been obtained.

$$\begin{aligned} E\{x_T^T Q_T x_T\} &= E\{x_T\}^T Q_T E\{x_T\} + E\{(x_T - E\{x_T\})^T Q_T (x_T - E\{x_T\})\} \\ &= E\{x_T\}^T Q_T E\{x_T\} + \text{tr}(Q_T E\{(x_T - E\{x_T\})(x_T - E\{x_T\})^T\}) \\ &= \mu_{T|T} Q_T \mu_{T|T} + \text{tr}\left(Q_T \Sigma_X^{(T|T)}\right) \end{aligned}$$

The first function for the dynamic programming approach is given as follows if we assume the measurement at time T is taken and no more measurements are to be performed:

$$\begin{aligned} V_T\left(\mu_{T|T}, \Sigma_X^{(T|T)}\right) &= \min_{u_T}(E\{x_T^T Q_T x_T + u_T^T R_T u_T\}) \\ &= \min_{u_T}\left(\mu_{T|T}^T Q_T \mu_{T|T} + u_T^T R_T u_T + \text{tr}\left(Q_T \Sigma_X^{(T|T)}\right)\right) \end{aligned} \quad (34)$$

The optimal control is clearly $u_T^* = 0$. The value is quadratic in terms of the estimate and has an additional scalar term that results from the estimation error.

We seek to prove that the functions $V_k\left(\mu_{k|k}, \Sigma_X^{(k|k)}\right)$ remain quadratic for all k, as in the case of perfect state information. Thus we assume:

$$V_{k+1}\left(\mu_{k+1|k+1}, \Sigma_X^{(k+1|k+1)}\right) = \mu_{k+1|k+1}^T S_{k+1} \mu_{k+1|k+1} + s_{k+1}\left(\Sigma_X^{(k+1|k+1)}\right) \quad (35)$$

Where S_{k+1} is symmetric semidefinite and the function $s_{k+1}\left(\Sigma_X^{(k+1|k+1)}\right)$ scalar and independent of the control.

The iteration is:

$$\begin{aligned} V_k\left(\mu_{k|k}, \Sigma_X^{(k|k)}\right) & \\ &= \min_{u_k, m_{k+1}} E\left\{x_k^T Q_k x_k + u_k^T R_k u_k + h(m_{k+1})\right. \\ &\quad \left. + \alpha E\left\{V_{k+1}\left(\mu_{k+1|k+1}, \Sigma_X^{(k+1|k+1)}\right)\right\}\right\} \end{aligned} \quad (36)$$

where $E\left\{V_{k+1}\left(\mu_{k+1|k+1}, \Sigma_X^{(k+1|k+1)}\right)\right\}$ needs to be evaluated.

From the Kalman filter equations it is known that:

$$\mu_{k+1|k+1} = \mu_{k+1|k} + K_{k+1}^{(m_{k+1})}\left(y_{k+1} - C_{k+1}^{(m_{k+1})}\mu_{k+1|k}\right)$$

The distribution for y_{k+1} is known by Equation (19) to be:

$$N\left(y_{k+1}; C_{k+1}^{(m_{k+1})}\mu_{k+1|k}, C_{k+1}^{(m_{k+1})}\Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T + \Sigma_M^{(m_{k+1})}\right)$$

$$\text{where } \Sigma_X^{(k+1|k)} = A_k \Sigma_X^{(k|k)} A_k^T + \Sigma_W^{(k)}$$

This means that given a fixed $\mu_{k|k}$ and taking y_{k+1} as a random variable:

$$E\{\mu_{k+1|k+1}\} = E\{\mu_{k+1|k}\} = A_k\mu_{k|k} + B_k u_k$$

$$\begin{aligned}
&E\left\{S_{k+1}\left(\mu_{k+1|k+1} - E\{\mu_{k+1|k+1}\}\right)\left(\mu_{k+1|k+1} - E\{\mu_{k+1|k+1}\}\right)^T\right\} \\
&\quad = E\left\{S_{k+1}\left(K_{k+1}^{(m_{k+1})}\left(y_{k+1}\right.\right.\right. \\
&\quad \left.\left.\left. - C_{k+1}^{(m_{k+1})}\mu_{k+1|k}\right)\right)\left(K_{k+1}^{(m_{k+1})}\left(y_{k+1} - C_{k+1}^{(m_{k+1})}\mu_{k+1|k}\right)\right)^T\right\} \\
&\quad = \text{tr}\left(S_{k+1}K_{k+1}^{(m_{k+1})}\left(C_{k+1}^{(m_{k+1})}\Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T + \Sigma_M^{(m_{k+1})}\right)K_{k+1}^{(m_{k+1})^T}\right)
\end{aligned}$$

Thus:

$$\begin{aligned}
&V_k\left(\mu_{k|k}, \Sigma_X^{(k|k)}\right) \\
&\quad = \min_{u_k, m_{k+1}} E\left\{x_k^T Q_k x_k + u_k^T R_k u_k + h(m_{k+1})\right. \\
&\quad + \alpha E\{\mu_{k+1|k+1}^T S_{k+1}\mu_{k+1|k+1}\} + \alpha s_{k+1}\left(\Sigma_X^{(k+1|k+1)}\right)\Big\} \\
&\quad = \min_{u_k}\left(\mu_{k|k}^T Q_k \mu_{k|k} + \alpha\left(A_k\mu_{k|k} + B_k u_k\right)^T S_{k+1}\left(A_k\mu_{k|k} + B_k u_k\right)\right. \\
&\quad \left. + u_k^T R_k u_k\right) + \text{tr}\left(Q_k\Sigma_X^{(k|k)}\right) \qquad (37) \\
&\quad + \min_{m_{k+1}}\left(h(m_{k+1})\right. \\
&\quad + \alpha\text{tr}\left(S_{k+1}K_{k+1}^{(m_{k+1})}\left(C_{k+1}^{(m_{k+1})}\Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T\right.\right. \\
&\quad \left.\left.\left. + \Sigma_M^{(m_{k+1})}\right)K_{k+1}^{(m_{k+1})^T}\right) + \alpha s_{k+1}\left(\Sigma_X^{(k+1|k+1)}\right)\Big)
\end{aligned}$$

By differentiating with respect to u_k and setting the derivative to zero:

$$(R_k + \alpha B_k^T S_{k+1} B_k)u_k = -\alpha B_k^T S_{k+1} A_k \mu_{k|k}$$

Thus the optimal control u_k^* is:

$$u_k^* = -\alpha(R_k + \alpha B_k^T S_{k+1} B_k)^{-1} B_k^T S_{k+1} A_k \mu_{k|k}$$

The resulting uncertainty $\Sigma_X^{(k+1|k+1)}$ depends on the measurement to be chosen:

$$\begin{aligned}
\Sigma_X^{(k+1|k+1)} &= \Sigma_X^{(k+1|k)} \\
&\quad - \Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T\left(C_k^{(m_{k+1})}\Sigma_X^{(k+1|k)}{C_{k+1}^{(m_{k+1})}}^T\right. \\
&\quad \left. + \Sigma_M^{(m_{k+1})}\right)^{-1} C_{k+1}^{(m_{k+1})}\Sigma_X^{(k+1|k)} \\
&\text{where } \Sigma_X^{(k+1|k)} = A_k \Sigma_X^{(k|k)} A_k^T + \Sigma_W^{(k)}
\end{aligned}$$

Substituting the optimal control into Equation (37) and rearranging the terms gives:

$$\begin{aligned} V_k\left(\mu_{k|k}, \Sigma_X^{(k|k)}\right) &= \mu_{k|k}^T\left(Q_k + \alpha A_k^T S_{k+1} A_k\right)\mu_{k|k} \\ &- \alpha^2 \mu_{k|k}^T A_k^T S_{k+1} B_k \left(R_k + \alpha B_k^T S_{k+1} B_k\right)^{-1} B_k^T S_{k+1} A_k \mu_{k|k} \\ &+ \mathrm{tr}\left(Q_k \Sigma_X^{(k|k)}\right) \\ &+ \min_{m_{k+1}}\left(\alpha \mathrm{tr}\left(S_{k+1} K_{k+1}^{(m_{k+1})}\left(C_{k+1}^{(m_{k+1})} \Sigma_X^{(k+1|k)} C_{k+1}^{(m_{k+1})^T}\right.\right.\right. \\ &\left.\left.\left. + \Sigma_M^{(m_{k+1})}\right) K_{k+1}^{(m_{k+1})^T}\right) + h(m_{k+1}) + \alpha s_{k+1}\left(\Sigma_X^{(k+1|k+1)}\right)\right) \end{aligned} \tag{38}$$

Finally the recursion is:

$$\begin{aligned} V_k\left(\mu_{k|k}, \Sigma_X^{(k|k)}\right) &= \mu_{k|k}^T S_k \mu_{k|k} + s_k\left(\Sigma_X^{(k|k)}\right) \\ \text{where } S_k &= Q_k + \alpha A_k^T S_{k+1} A_k \\ &- \alpha^2 A_k^T S_{k+1} B_k (R_k + \alpha B_k^T S_{k+1} B_k)^{-1} B_k^T S_{k+1} A_k \\ S_T &= Q_T \\ \text{and } s_k &= \mathrm{tr}\left(Q_k \Sigma_X^{(k|k)}\right) \\ &+ \min_{m_{k+1}}\left(h(m_{k+1})\right. \\ &+ \alpha \mathrm{tr}\left(S_{k+1} K_{k+1}^{(m_{k+1})}\left(C_{k+1}^{(m_{k+1})} \Sigma_X^{(k+1|k)} C_{k+1}^{(m_{k+1})^T}\right.\right. \\ &\left.\left.\left. + \Sigma_M^{(m_{k+1})}\right) K_{k+1}^{(m_{k+1})^T}\right) + \alpha s_{k+1}\left(\Sigma_X^{(k+1|k+1)}\right)\right) \\ s_T &= \mathrm{tr}\left(Q_T \Sigma_X^{(T|T)}\right) \end{aligned} \tag{39}$$

We have shown that the evolution of the estimation covariance $\Sigma_X^{(k|k)}$ is independent of the control thus allowing us to conclude the inductive step.

The optimal control law $\pi_k^*(x_k)$ is the linear relation:

$$\begin{aligned} \pi_k^*(x_k) &= L_k \mu_{k|k} \\ \text{where } L_k &= -\alpha\left(R_k + \alpha B_k^T S_{k+1} B_k\right)^{-1} B_k^T S_{k+1} A_k \end{aligned} \tag{40}$$

The relation to the optimal control law of the previous chapter, Equation (30), is apparent. The difference is in that instead of the exact value of the state variable, its estimate is used.

As previously if we assume the initial state estimate distributed as $N\left(x_0; \mu_{0|0}, \Sigma_X^{(0|0)}\right)$ the optimal cost is:

$$\begin{aligned} V_0\left(\mu_{0|0}, \Sigma_X^{(0|0)}\right) = {} & \mu_{0|0} S_0 \mu_{0|0} + \mathrm{tr}\left(Q_0 \Sigma_X^{(0|0)}\right) + \alpha^T \mathrm{tr}\left(Q_T \Sigma_X^{(T|T)}\right) \\ & + \sum_{k=0}^{T-1} \left(\alpha^{k-1} h(m_{k+1}) + \alpha^{k-1} \mathrm{tr}\left(Q_k \Sigma_X^{(k|k)}\right) \right. \\ & + \alpha^k \mathrm{tr}\left(S_{k+1} K_{k+1}^{(m_{k+1})} \left(C_{k+1}^{(m_{k+1})} \Sigma_X^{(k+1|k)} {C_{k+1}^{(m_{k+1})}}^T \right. \right. \\ & \left. \left. + \Sigma_M^{(m_{k+1})} \right) {K_{k+1}^{(m_{k+1})}}^T \right) \Bigg) \end{aligned} \tag{41}$$

The structure of the optimal controller is summarized in Figure 3. This is similar to Figure 2 except it contains an additional layer in the measurement and estimation. The estimator uses the previous control to perform the prediction and the received measurement to update the estimate.

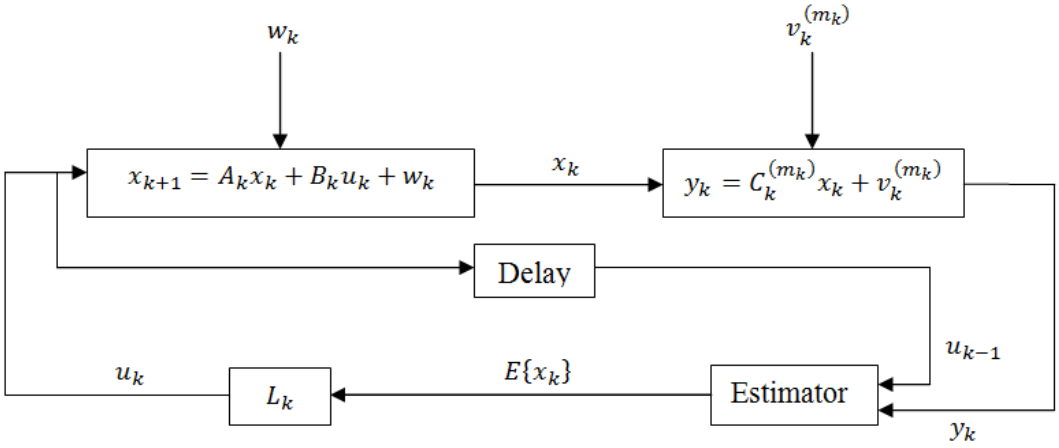

Figure 3: Structure of the optimal controller when an estimator is used to determine the state through measurements

That the choice of measurement does not affect the optimal control law and that the estimation error is independent of the control is an important property of LQG systems.

3.4 Separation Principle

The fact that the estimation and control problems may be viewed independent of each other in the LQG system is referred to as the separation principle. The first proof was given in [13], soon after the formulation of the Kalman filter.

This result implies that the optimal structure of the controller may be seen as a combination of the optimal filter and a deterministic controller. The two parts can be designed independently. The optimal estimator then feeds the estimate of the state to the optimal controller which may be designed as in the perfect information case. Thus the estimator forecasts the evolution of the state and the controller uses the estimate without taking into account any risks associated with it.

The separation of the control and estimation problems also leads to the result that the measurement optimization i.e. minimization of the estimation error may be performed offline. This is apparent if one examines the evolution of the estimation covariance which is a deterministic problem, i.e. independent on the measurement data, with regard to the choice of measurement.

3.5 Value Iteration

The problem structure where the number of decision stages is infinite or at least a very large number for which an infinite-horizon approximation is appropriate can be analyzed with value iteration. This means that the optimization is to be performed into the distant future where the exact decision horizon may not be known. Contrary to the previous system models, here we assume that the model, cost per stage and all random disturbance statistics are time-invariant. This does not mean that the process itself is a stationary stochastic process.

The minimization task is of the following form:

$$V(x_0) = \min_{\{u_k\}_{k=0}^{\infty}} E\left\{\sum_{k=0}^{\infty} \alpha^k F(x_k, u_k, w_k)\right\} \quad \text{subject to } x_{k+1} = T(x_k, u_k, w_k) \tag{42}$$

The value function iteration $\{V_k(\cdot)\}_{k=0}^{\infty}$, which describe the expected optimal value to be gained from optimal actions in the future, converges as $k \to \infty$. However, it is clear that the analysis cannot be started as before from the last time instant. The assumption that the system equations are time-invariant means that first only the immediate cost is considered. Then considering a value function one step into the future means that the previous value function describes the value from which the optimum for the future rewards can be calculated. The iteration is:

$$V_0(x_0) = \min_{u_0} E\{F(x_0, u_0, w_0)\}$$

$$V_1(x_0) = \min_{u_0} E\{F(x_0, u_0, w_0) + \alpha E\{V_0(x_1)\}\}$$

...

$$V_{k+1}(x_0) = \min_{u_0} E\{F(x_0, u_0, w_0) + \alpha E\{V_k(x_1)\}\}$$

The value function considered at each iteration takes into consideration a time horizon one step longer than the previous one. This is continued until a sufficient approximation of the infinite horizon case is obtained.

Often the value function is formed as a tabulation of the optimal values for the state space. A one step look-a-head approach may be used if the control is not directly saved when iterating the value functions. When a state is reached at a time instant the control is found by solving the value function iteration once with the optimal cost-to-go solved earlier.

The LQG system where the system dynamics $\{A, B\}$, the weighting matrices $\{Q, R\}$ and the distribution parameters of the process covariance $\{\Sigma_W\}$ are fixed for all k may be solved using value iteration. This allows the measurement problem to be solved a priori.

The iteration may be directly taken from the previous solution of the dynamic programming approach to the imperfect state knowledge problem. Here we assume no measurement is to be performed in the future when no time steps are remaining. For the first step:

$$V_0\left(\mu_{0|0},\, \Sigma_X^{(0|0)}\right) = \mu_{0|0} Q_1^{(V)} \mu_{0|0} + \mathrm{tr}\left(Q\, \Sigma_X^{(0|0)}\right), \qquad \text{where } u_0^* = 0 \text{ and } Q_1^{(V)} = Q$$

A set of functions $\left\{G_k\left(\Sigma_X^{(0|0)}\right)\right\}_{k=0}^{\infty}$ is defined that provides the value related to the measurement optimization. For the first iteration $G_0\left(\Sigma_X^{(0|0)}\right)$ is:

$$G_0\left(\Sigma_X^{(0|0)}\right) = \mathrm{tr}\left(Q\Sigma_X^{(0|0)}\right)$$

Then any value function is given as:

$$V_k\left(\mu_{0|0}, \Sigma_X^{(0|0)}\right) = \mu_{0|0} Q_{k+1}^{(V)} \mu_{0|0} + G_k\left(\Sigma_X^{(0|0)}\right)$$

$$where\ Q_{k+1}^{(V)} = Q + \alpha A^T Q_k^{(V)} A - \alpha^2 A^T Q_k^{(V)} B\left(R + \alpha B^T Q_k^{(V)} B\right)^{-1} B^T Q_k^{(V)} A \tag{43}$$

$$u_0^* = -\alpha\left(R + \alpha B^T Q_k^{(V)} B\right)^{-1} B^T Q_k^{(V)} A \mu_{0|0} = -L_k \mu_{0|0}$$

$Q_{k+1}^{(V)}$ will with a few assumptions approach a stationary value as $k \to \infty$ [12, pp.202-203]. $Q_k^{(V)}$ may then be substituted with the stationary value $Q^{(V)}$. The algebraic Riccati equation may be used to solve $Q^{(V)}$ from the equality:

$$Q^{(V)} = Q + \alpha A^T Q^{(V)} A - \alpha^2 A^T Q^{(V)} B\left(R + \alpha B^T Q^{(V)} B\right)^{-1} B^T Q^{(V)} A$$

The measurement problem is iterated as:

$$\begin{aligned} G_{k+1}\left(\Sigma_X^{(0|0)}\right) = \mathrm{tr}\left(Q\Sigma_X^{(0|0)}\right) \\ + \min_{\mathrm{m}}\Bigg(h(m) \\ + \alpha \mathrm{tr}\Big(Q_k^{(V)} \Sigma_X^{(1|0)} C^{(m)^T}\Big(C^{(m)} \Sigma_X^{(1|0)} C^{(m)^T} \\ + \Sigma_{\mathrm{M}}^{(\mathrm{m})}\Big)^{-1} C^{(m)} \Sigma_X^{(1|0)}\Big) + \alpha G_k\left(\Sigma_X^{(1|1)}\right)\Bigg) \end{aligned} \tag{44}$$

$$\text{where } \Sigma_X^{(1|0)} = A\, \Sigma_X^{(0|0)} A^T + \Sigma_W$$

$$\Sigma_X^{(1|1)} = \Sigma_X^{(1|0)} - \Sigma_X^{(1|0)} C^{(m)^T}\left(C^{(m)} \Sigma_X^{(1|0)} C^{(m)^T} + \Sigma_{\mathrm{M}}^{(\mathrm{m})}\right)^{-1} C^{(m)} \Sigma_X^{(1|0)}$$

Thus the measurement problem presents a functional equation where an analytical solution is not possible in this form. In this work we approximate the functions $G_k(\cdot)$ by a grid over the space of the covariance. Each element of the state space has a variance and a correlation with all other elements. An n-dimensional state space will have $\frac{n(n+1)}{2}$

variances and correlations that need to be discretized. Any value that falls between points may be approximated to the nearest value or interpolated.

Value iteration is often used in conjunction with partially observable Markov decision processes [14,15]. There the state space is discrete which allows for the state information, also called the state belief, to be presented by a piecewise linear function. In continuous state spaces the problem is substantially more difficult.

4 RESULTS

This chapter first describes some of the issues in the implementation of the above optimization methods. These are then applied to the optimization of measurement sequences in a simple one-dimensional system. We expand on this by optimizing the measurement sequence of a more complex two-dimensional system. Finally the on-line solution method is applied to a system with a four-dimensional state space and two measurement alternatives.

4.1 On-line Solution

The on-line solution to the control and measurement problem works by calculating the value of all measurement sequences from any time k to $k + T$ where the estimate and uncertainty are known at time $k - 1$. The control is simply found by iterating the matrices S_k and using Equation (40). The measurement problem, however, is solved by calculating all values of the measurement sequence up to the specified time horizon. From the solution the first measurement is performed and at the next time $k + 1$ the same optimization is performed up to $k + 1 + T$. This is the solution structure presented by Meier et al [3,4].

This solution is computationally hard as there exist H^T different possibilities for the measurement sequence where H denotes the amount of different measurement choices. However, in calculating the value of a new measurement sequence we may use previous results to somewhat reduce the quantity of computations to be performed. Only the last measurement choice which leads to a new measurement sequence is changed. Denoting the time instance of the changed measurement choice as k_1, the covariance at $k + k_1 - 1$ remains unchanged and only the new covariances up to $k + T$ need to be calculated. Thus we apply $\sum_{i=1}^{T} H^i$ updates to the covariances to find all values of the sequences in the specified time horizon. This is in comparison to TH^T updates to the covariance if all sequences were to be cycled through from beginning to end.

The length of the prediction horizon is restricted by the computation time available between measurement instances. Thus this is an approximate solution which however approaches the infinite horizon solution when the discount factor is suitably small. Additionally, the measurement problem leads to a repeating sequence of optimal measurement choices as the updating of the covariance is in fact a deterministic problem.

4.2 Off-line Solution

The measurement problem is alternatively solved with the described value iteration method. This method is also approximate as the resulting value function is nonlinear and the space defined by the covariance must be approximated by a finite grid. In comparison to the method above the approximation is not in the horizon length but in the discretization of the covariance. For example, in a 2-dimensional state space the covariance associated with state information may be written as

$$\Sigma_X^{(k|k)} = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}.$$

σ_1^2 denotes the variance of the first state variable and ρ the correlation between the two variables with $\rho \in [-1,1]$. The maximum and minimum variances to be taken into account depend on the system specifications. These are given by finding values of the state information covariance $\Sigma_X^{(k|k)}$ as $k \to \infty$. These values are solved by simple iteration or using the Riccati equation where only one of the measurement choices is used. Values must be extrapolated if the system is defined such that the uncertainty does not converge.

The choice of how grid points are spaced is affected by calculation and memory restrictions. Let D denote the amount of grid points in the space of each variable of an n-dimensional state space. Altogether there are $D^{n(n+1)/2}$ points at which the minimization task must be solved. At any iteration step the previous grid has to be saved as well as the new resulting grid. These extreme memory requirements are the main hindrance in the implementation of the value iteration method.

The main benefit in comparison to the on-line solution is that the value function may be calculated a priori. There is no need to do costly optimizations between measurements as the algorithm provides a look-up-table of optimal measurement decisions.

4.3 One-dimensional State Space

The two solution procedures are used in the control of the following one-dimensional system. The system dynamics and costs are:

$$A = 0.9, B = 0.6, \Sigma_W = 0.1, Q = 1, R = 1, \alpha = 0.9$$

The initial state, estimate and uncertainty are:

$$x_0 = x_{0|0} = 10, \Sigma_X^{(0|0)} = 0.01$$

The choice of measurement is made between the following two alternatives:

$$\Sigma_M^{(1)} = 0.5, C^{(1)} = 1, h(1) = 0.1$$

$$\Sigma_M^{(2)} = 0.03, C^{(2)} = 1, h(2) = 0.19$$

The cheaper less exact measurement would be expected to be used when the state information uncertainty is low and the costlier measurement when information must be

improved more. The on-line method is used with a prediction horizon of 10. In this case the two methods lead to the very similar results as the discount factor is quite low.

The minimum and maximum values for the discretization of the estimation uncertainty are given by solving the following for Σ_X:

$$\Sigma_X^+ = A\Sigma_X A^T + \Sigma_W$$

$$\Sigma_X = \Sigma_X^+ - \Sigma_X^+ C^{(m)^T}\left(C^{(m)}\Sigma_X^+ C^{(m)^T} + \Sigma_{\mathrm{M}}^{(\mathrm{m})}\right)^{-1} C^{(m)}\Sigma_X^+$$

The minimum is given by repeated use of measurement 2: $\Sigma_X^{(min)} = 0.0240$. The maximum is given by repeated use of measurement 1: $\Sigma_X^{(max)} = 0.1557$. This interval is discretized into 1000 equally distanced points. The value function for the measurement problem is initialized with:

$$G_0(\Sigma_X) = \mathrm{tr}(Q\Sigma_X)$$

The value function converges in under 200 iterations. We use the sum of the differences between $G_{k+1}(\Sigma_X)$ and $G_k(\Sigma_X)$ to gauge the convergence. The value of the resulting $G_k\left(\Sigma_X^{(1|1,m)}\right)$ is interpolated linearly at each grid point when calculating $G_{k+1}(\Sigma_X)$. The decision boundary above which measurement 2 is used is $\Sigma_X = 0.1490$.

Figure 4 shows the result of a simulation of the system over 50 time steps. Measurement 1 is used until the estimate uncertainty exceeds the decision boundary after which measurement 2 is used once. The dotted line in the second graph of Figure 4 shows the boundary above which measurement 2 is used.

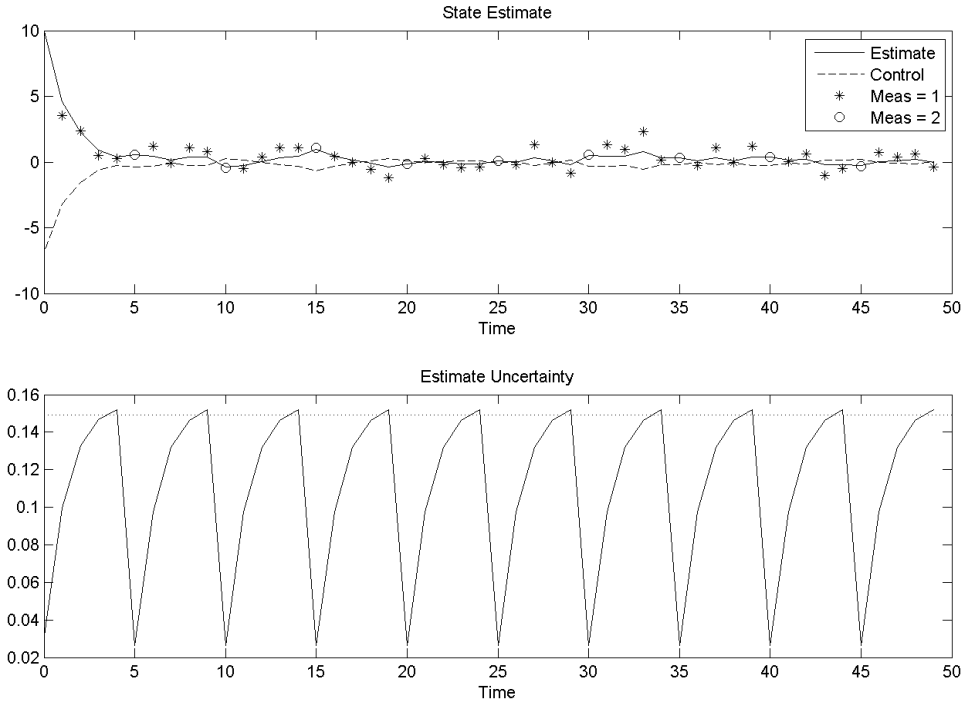

Figure 4: Simulation result. The first figure shows the state estimate and received measurement values. The second shows the estimation uncertainty. The dotted line shows the decision boundary.

The repeating structure of the optimal measurement sequence is apparent from the above figure. This structure could be exploited in cases where computational resources are very limited in implementation. However, during application of the optimization procedures the value iteration method already provides a computationally less expensive method. The following example provides an example of the off-line method applied to a more complex system.

4.4 Two-dimensional State Space

The second example system where measurement optimization is applied has two state variables. The system dynamics are:

$$A = \begin{bmatrix} 0.7 & -0.01 \\ 0 & 0.4 \end{bmatrix}, \qquad B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \qquad \Sigma_W = \begin{bmatrix} 0.2025 & 0 \\ 0 & 0.15 \end{bmatrix}$$

The weighting matrices:

$$Q = \begin{bmatrix} 1 & 0 \\ 0 & 5 \end{bmatrix}, \qquad R = \begin{bmatrix} 20 & 0 \\ 0 & 12 \end{bmatrix}$$

The initial state and estimate:

$$x_0 = x_{0|0} = \begin{bmatrix} 10 \\ 15 \end{bmatrix}, \qquad \Sigma_{\mathrm{X}}^{(0|0)} = \begin{bmatrix} 0.1 & 0 \\ 0 & 0.1 \end{bmatrix}$$

The measurement choice is made between three alternatives. The first two measurements are used to measure only one variable at a time whereas the third is able to measure both at once. In this system the correlation between the variables remains small and the first two measurements provide very little information on the variable that they are not directly related to. The three measurement alternatives are defined as follows:

$$\Sigma_M^{(1)} = 0.03, \qquad C^{(1)} = [1 \quad 0], \qquad h(1) = 0.005$$

$$\Sigma_M^{(2)} = 0.055, \qquad C^{(2)} = [0 \quad 1], \qquad h(2) = 0.028$$

$$\Sigma_M^{(3)} = \begin{bmatrix} 0.04 & 0.02 \\ 0.02 & 0.04 \end{bmatrix}, \qquad C^{(3)} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \qquad h(3) = 0.055$$

Figure 5 depicts the resulting measurement policy where the correlation coefficient is 0. This policy has been calculated with a discount factor of $\alpha = 0.9$ and is the result after 25 iterations. The figure depicts a slice of the three-dimensional grid where each variable in the uncertainty space was discretized into 101 equidistant points. Value function interpolation was performed with linear interpolation. The variances have been calculated from values of 0.01 to 1. The third measurement is used when information is needed on both variables. The third measurement choice is not present if the variables correlate because then a measurement on one variable gives information on the other.

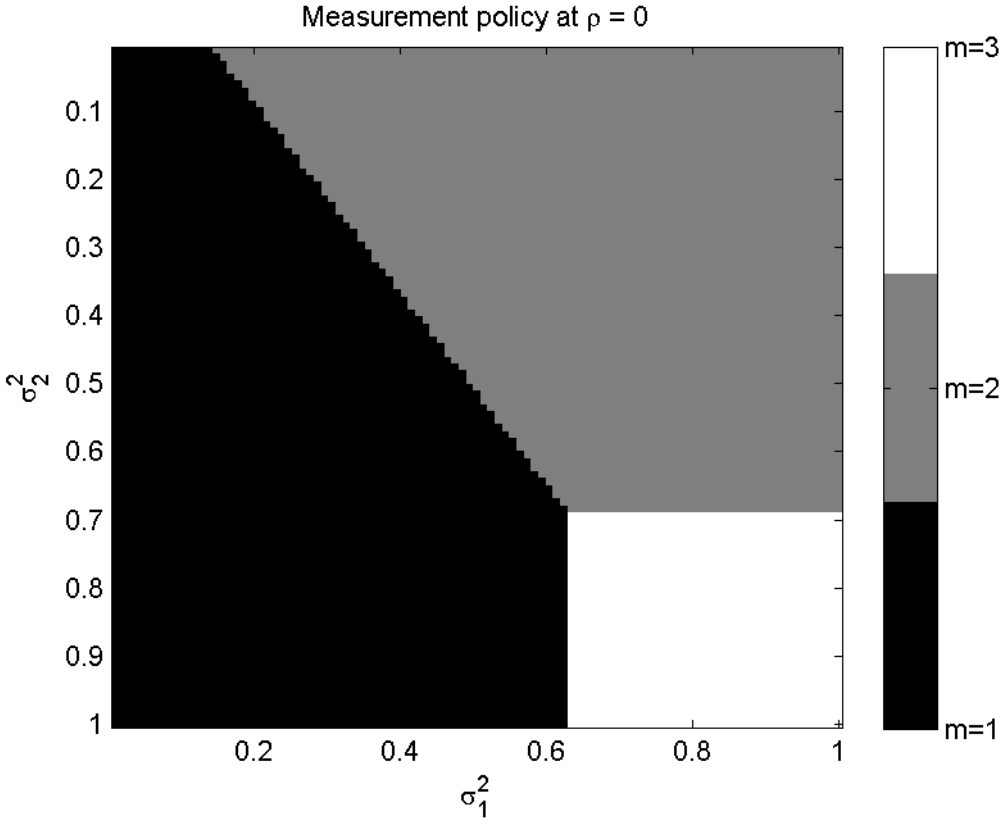

Figure 5: Optimal measurement policy as calculated after 25 iterations

The optimal measurement sequence in fact only makes use of the first two measurement alternatives. The second alternative is used at every third time step. The resulting measurement sequence is shown in Figure 6.

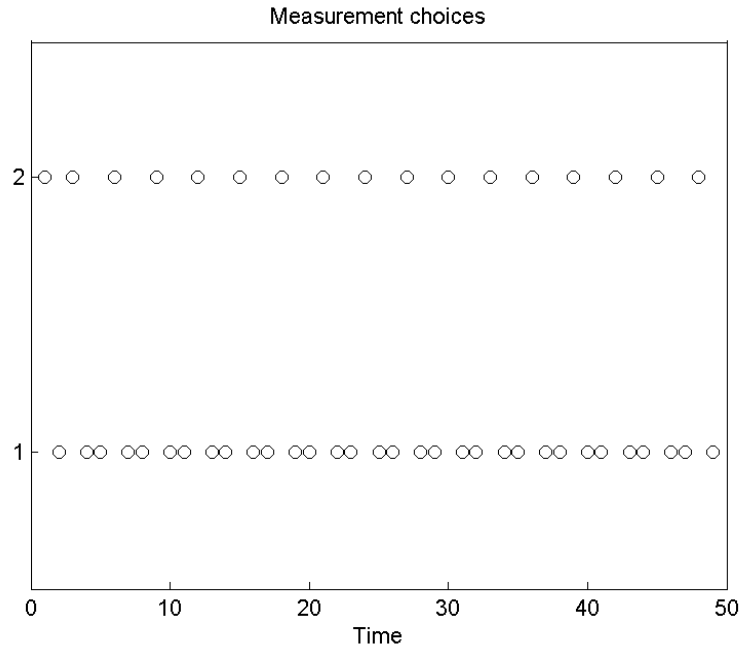

Figure 6: Measurement choices

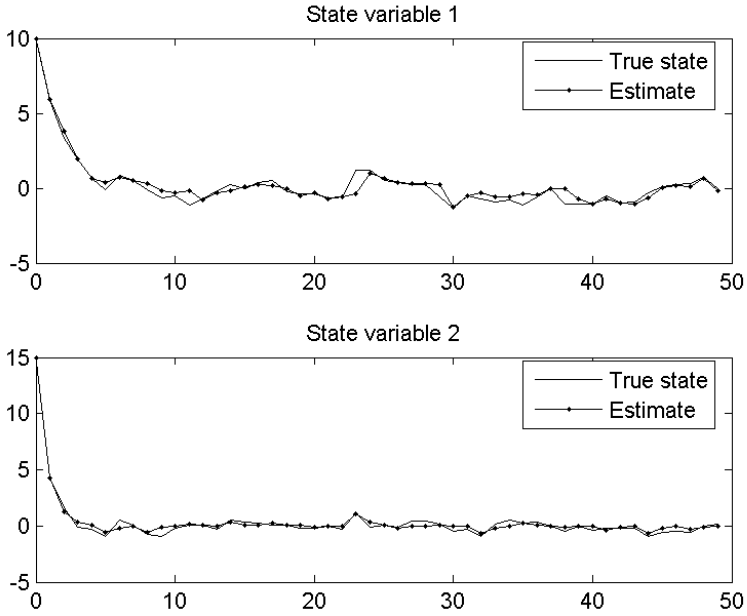

Figure 7: Simulation result. True state and estimate over 50 time steps

Figure 7 shows the simulated system over 50 time steps. The state reaches the goal quickly. However, the weighting matrices lead to the control of the second variable being quicker as any deviation in the second variable leads to a higher penalty.

The same results are reached with the on-line solution. Precalculation of the measurement policy provides information on the exact structure of the measurement problem in many cases. This could be used to aid in the tuning of control systems. The main hindrance is presented by the time needed to perform the calculations.

4.5 Four-dimensional State Space

The off-line solution becomes impossible to calculate with a gridded solution in higher dimensions. However, the on-line solution performs well if the amount of measurements is kept low. Let us define a discrete time four-dimensional process where the first two state variables describe the position $(x_{1,k}, x_{2,k})$ of a moving agent in a two-dimensional coordinate system at time k. The remaining two variables define the velocities $(\dot{x}_{1,k}, \dot{x}_{2,k})$ in the direction of the two axes such that the state vector is: $x_k = [x_{1,k}\ x_{2,k}\ \dot{x}_{1,k}\ \dot{x}_{2,k}]^T$.

The dynamic model is defined as:

$$x_{k+1} = Ax_k + Bu_k + w_k$$

$$\text{where } A = \begin{bmatrix} 1 & 0 & 0.5 & 0 \\ 0 & 1 & 0 & 0.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix},$$

$$w_k \sim N\left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0.0008 & 0 & 0.0025 & 0 \\ 0 & 0.0021 & 0 & 0.0063 \\ 0.0025 & 0 & 0.0100 & 0 \\ 0 & 0.0063 & 0 & 0.0250 \end{bmatrix} \right)$$

The control affects the velocity in the directions of each axis in the two-dimensional space. The system is to choose between two measurement choices. The first measurement measures the position of the agent in the first coordinate axis and the second measurement in the second coordinate axis.

$$\Sigma_M^{(1)} = 0.5, \qquad C^{(1)} = [1\ 0\ 0\ 0], \qquad h(1) = 0.1$$

$$\Sigma_M^{(2)} = 0.5, \qquad C^{(2)} = [0\ 1\ 0\ 0], \qquad h(2) = 5$$

The weights for the cost function are defined as:

$$Q = \begin{bmatrix} 0.5 & 0 & 0 & 0 \\ 0 & 0.5 & 0 & 0 \\ 0 & 0 & 5 & 0 \\ 0 & 0 & 0 & 5 \end{bmatrix}, \qquad R = \begin{bmatrix} 20 & 0 \\ 0 & 20 \end{bmatrix}$$

The measurement problem is solved with a prediction horizon of 10 and the discount factor is $\alpha = 0.9$. The initial state and estimate is defined as follows with the goal being to reach and to stay at the origin.

$$x_0 = x_{0|0} = \begin{bmatrix} 50 \\ 50 \\ 75 \\ -30 \end{bmatrix}, \qquad \Sigma_X^{(0|0)} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Figure 8 shows the first two components of the state vector after a simulation of 75 time steps. The estimate is able to follow the true trajectory quite well. The control works well as there is no overshoot after the origin is reached.

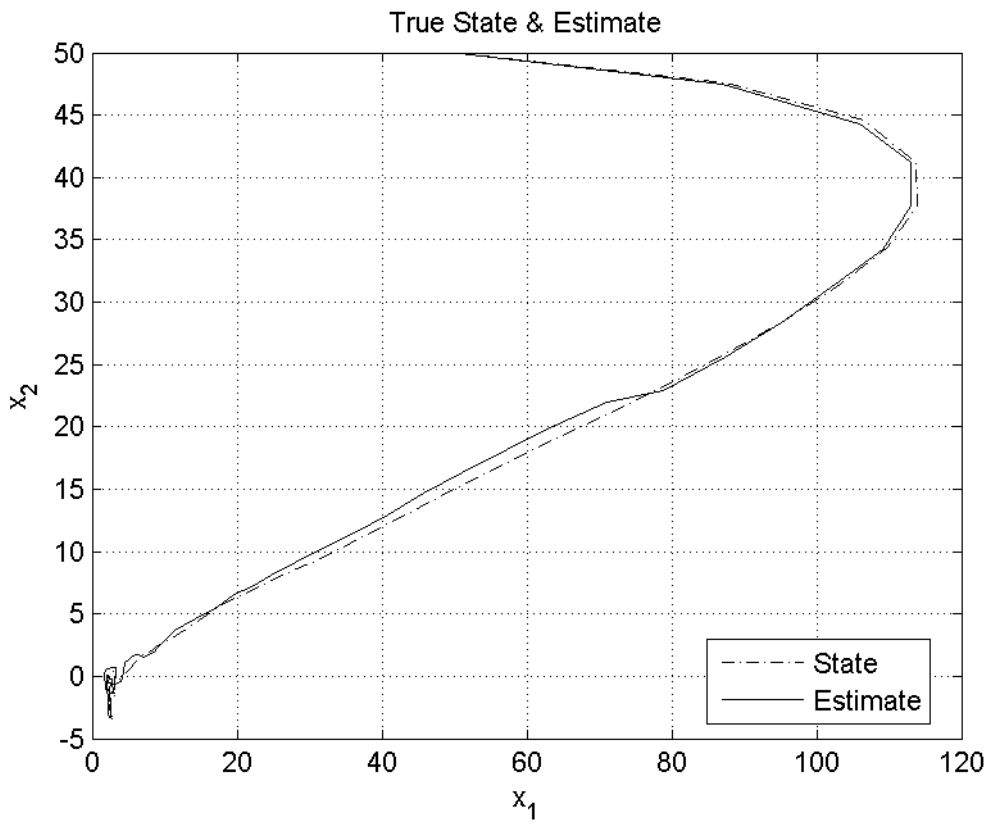

Figure 8: Simulation result. The first two components of the true state and state estimate after a simulation of 75 time steps.

The measurement sequence reaches a stationary policy after a short time. The measurement choices are shown in Figure 9. After an initial settling period, the second measurement is used only every sixth time step.

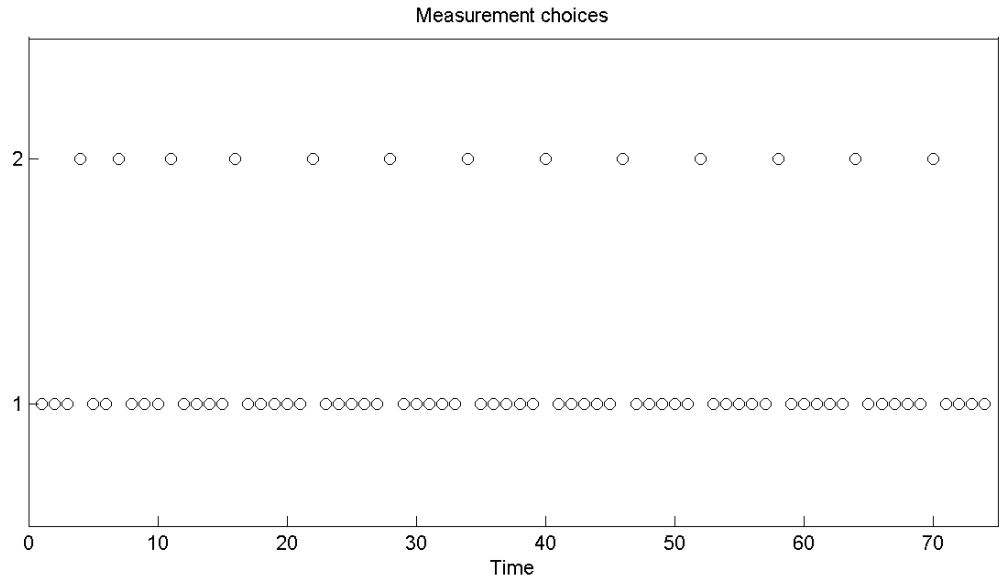

Figure 9: Measurements as chosen by the on-line solution

Though the on-line solution does not provide any insight into the structure of the measurement problem it provides a potential practical implementation. The on-line solution proves practical in systems with a large state space.

5 CONCLUSION AND FUTURE WORK

5.1 Conclusion

This thesis analyzed the problem of how to find an optimal measurement sequence in a discrete-time linear quadratic Gaussian system when presented with a discrete set of measurement alternatives. LQG-systems are popular in process control applications and often have many alternative measurements available. The optimization procedures presented provide a basis for optimal control of the measurement process in conjunction with traditional optimal control.

The problem was presented by first deriving the Kalman filter which is used to recursively update the information on the state of the system. The Gaussian distribution is used to represent the information on the system state. By repeated use of Bayes' theorem and conditional probabilities new information is combined with any old information.

The quadratic objective function was presented in the linear Gaussian case. Dynamic programming shows that this leads to a control policy as a linear function of the state estimate. The derivation further provides the objective for the measurement decision. The separation principle which provides for the separate examination of the control and measurement problems was explained. The measurement problem was solved first by considering all possible resulting measurement sequences and finding the optimal sequence. This approach approximates the optimal sequence in considering only a short prediction horizon. Alternatively the measurement problem was considered with value iteration which uses a discrete grid to represent the relevant uncertainty space. The first solution needs to be used on-line whereas the value iteration method provides a policy off-line.

The two different approaches in finding an optimal measurement sequence were applied to three different process control problems. The first problem was a process with a one-dimensional state and two measurement alternatives. With the studied parameters the two approaches lead to identical measurement sequences.

The advantages of the off-line value iteration method were illustrated with an example having a two-dimensional state space. The value iteration provides insight into the structure of the measurement policy. This can be seen as a potential tool in designing control and measurement systems. Most importantly the policy may be calculated in advance. Thus the computational load during application is very low.

The problem with implementation of the value iteration method is that its complexity grows quickly with the dimensionality of the state space. This was evident in the last example in which the on-line solution was used in the control of a system with a four-dimensional state space. This example showed the versatility of the on-line solution. The computational complexity of the on-line solution does not grow as quickly with the state space dimensionality as that of value iteration method. The solution is found for a short prediction horizon from which only the first decision is implemented. This is repeated at each time step. Often the optimal measurement sequence repeats itself after an initial settling period and this can be used to decrease the computational load of the method during application.

The procedures presented provide a theoretical basis for measurement optimization in LQG-systems.

5.2 Future Work

The assumption of the system being LQG in this thesis is rather restrictive and more complex system structures need to be considered in future. These include systems in which constraints are set for the control or state space. As a result the separation principle will no longer be valid so that the measurement and control problems need to be considered jointly. Nonlinear systems and/or non-quadratic objective functions present potential research areas for measurement optimization.

The measurement optimization procedure is a subproblem within measurement system design. This pertains not to the problem of finding the best measurement sequence but to the problem of defining the measurements to best measure the system state.

Efficient methods for solving optimal value functions can also be found in literature about partially observable Markov decision processes. These methods deal with highly complex structures very close to the considered problem. However, when applying them the need for discretization may be a limiting factor.

REFERENCES

[1] Bellman, R.E., 1957. Dynamic Programming. 4th reprint (1965). Princeton, New Jersey, Princeton University Press. 342 p.

[2] IEEE Transactions on Automatic Control, 1971. Special Issue on Linear-Quadratic Gaussian Problem, Vol. AC-16.

[3] Meier, L., III, Peschon, J. & Dressler, R.M. Optimal Control of Measurement Subsystem, IEEE Transactions on Automatic Control, vol. AC-12(1967)5, pp. 528-536.

[4] Meier, L., III, Larson, R.E. & Tether, A.J. Dynamic Programming for Stochastic Control of Discrete Systems, IEEE Transactions on Automatic Control, vol. AC-16(1971)6, pp. 767-775.

[5] MATLAB – The Language of Technical Computing [WWW]. [Referenced 3.12.2010]. Available: http://www.mathworks.com/products/matlab/

[6] Papoulis, A., 1965. 3rd international edition (1991). Probability, Random Variables and Stochastic Processes. New York, USA, McGraw Hill. 666 p.

[7] Kalman, R.E. A new approach to linear filtering and prediction problems, ASME Journal of Basic Engineering, series D Vol. 82(1960), pp. 34-45.

[8] Ho, Y.C. & Lee, R.C.K. 1964. A Bayesian Approach to Problems in Stochastic Estimation and Control. IEEE Transactions on Automatic Control, vol. 9(1964)4. pp. 333-339.

[9] Bellman, R.E. & Dreyfus, S.E. 1962. Applied Dynamic Programming, 3rd reprint (1966). Princeton, New Jersey, Princeton University Press. 363 p.

[10] Richalet, J., Rault, A., Testud, J.L. & Papon, J. Model Predictive Control: Applications to Industrial Processes. Automatica, vol. 14(1978)5. pp. 413-428.

[11] Camacho, E.F. & Bordons, C., 1999. Model Predictive Control. In: Grimble, M.J. & Johnson, M.A., editors. Advanced Textbooks in Control and Signal Processing. London, UK, Springer. 280 p.

[12] Bertsekas, D 1995. Dynamic Programming and Optimal Control. Belmont (Mass.), Athena Scientific. 387 p.

[13] Joseph, P.D., & Tou, J.T.. On linear control theory, Transactions of the AIEE, pt. II 80(1961), pp. 193-196.

[14] Smallwood, R.D. & Sondik, E.J., The Optimal Control of Partially Observable Markov Processes Over a Finite Horizon. Operations Research, vol. 21(1973)5, pp. 1071-1088

[15] Kaelbling, L.P., Littman, M.L. & Cassandra, A.R. Planning and Acting in Partially Observable Stochastic Domains, Artificial Intelligence, vol. 101(1998)1-2, pp. 99-134.