Stochastic Event-Based Control and Estimation

Henningsson, Toivo

2012

Document Version:
Publisher's PDF, also known as Version of record

Link to publication

Citation for published version (APA):
Henningsson, T. (2012). *Stochastic Event-Based Control and Estimation*. [Doctoral Thesis (compilation), Department of Automatic Control]. Department of Automatic Control, Lund Institute of Technology, Lund University.

Total number of authors:
1

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Stochastic Event-Based Control and Estimation

Stochastic Event-Based Control and Estimation

Toivo Henningsson

Department of Automatic Control
Lund University
Lund, December 2012

Till Maria, Arthur och Alice

Department of Automatic Control
Lund University
Box 118
SE-221 00 LUND
Sweden

ISSN 0280–5316
ISRN LUTFD2/TFRT--1095--SE

Printed in Sweden by Media-Tryck.
Lund 2012

Abstract

Digital controllers are traditionally implemented using periodic sampling, computation, and actuation events. As more control systems are implemented to share limited network and CPU bandwidth with other tasks, it is becoming increasingly attractive to use some form of *event-based control* instead, where precious events are used only when needed.

Forms of event-based control have been used in practice for a very long time, but mostly in an ad-hoc way. Though optimal solutions to most event-based control problems are unknown, it should still be viable to compare performance between suggested approaches in a reasonable manner.

This thesis investigates an event-based variation on the stochastic linear-quadratic (LQ) control problem, with a fixed cost per control event. The *sporadic constraint* of an enforced minimum inter-event time is introduced, yielding a mixed continuous-/discrete-time formulation. The quantitative trade-off between event rate and control performance is compared between periodic and sporadic control. Example problems for first-order plants are investigated, for a single control loop and for multiple loops closed over a shared medium.

Path constraints are introduced to model and analyze higher-order event-based control systems. This component-based approach to stochastic hybrid systems allows to express continuous- and discrete-time dynamics, state and switching constraints, control laws, and stochastic disturbances in the same model. Sum-of-squares techniques are then used to find bounds on control objectives using convex semidefinite programming.

The thesis also considers state estimation for discrete-time linear stochastic systems from measurements with convex set uncertainty. The Bayesian observer is considered given *log-concave* process disturbances and measurement likelihoods. Strong log-concavity is introduced, and it is shown that the observer preserves log-concavity and propagates strong log-concavity like inverse covariance in a Kalman filter. A recursive state estimator is developed for systems with both stochastic and set-bounded process and measurement noise terms. A time-varying linear filter gain is optimized using convex semidefinite programming and ellipsoidal over-approximation, given a relative weight on the two kinds of error.

Acknowledgments

First, I would like to thank my supervisor Anton Cervin for many rewarding discussions about the proper way to do event-based control. Anton has also been a great support, available to give good advice and pitching in when it was most needed (such as close to deadlines). I would also like to thank Karl Johan Åström for introducing me to the field of event-based control, and Bo Bernhardsson for feedback, collaboration, and many insights into hairy mathematical issues.

I have enjoyed many interesting discussions with my coworkers about both event-based control and other things; special thanks go to Peter Alriksson, Mikael Lindberg, and Kristian Soltesz who got to listen to my ideas and tried to understand what I was doing. Thanks also to Erik Johannesson for the work we did together on Paper III. Proofreading and feedback on the material in this thesis was kindly given by Karl Johan Åström, Bo Bernhardsson, Anders Rantzer, Anton Cervin, Peter Alriksson, Ather Gattami, and Maria Henningsson.

Leif Andersson deserves credit for computer support and help with the typesetting of this thesis. Thanks go also to Eva Westin and Lizette Borgeram for taking care of so many of the practicalities involved with the dissertation. Many thanks go to Anders Blomdell and Rolf Braun, for patient help with many experimental and practical matters.

This work has been supported by the Swedish Research Council, the LCCC Linnaeus Center, and the ELLIIT Excellence Center.

I am grateful to Lasse, Kerstin, Rasmus, Harald, and Maja for all their support and for valuable breaks from the world of research. Last, and most importantly, I want to thank Maria for all her love and support and Arthur and Alice for being great.

Contents

Preface

The field of event-based control is a rich subject area that spans many diverse and challenging control and estimation problems. Classical control rests on the assumption that the controller is connected to sensors and actuators through perfect, dedicated communication channels. This assumption breaks down when one or more of the channels are constrained, so that transmission events are costly or limited in number or timing. A corresponding form of event-based control can then provide sizable performance benefits compared to periodic control. These kinds of control problems are becoming of increasing interest as more control systems are implemented to share limited communication and computation resources with other tasks.

While there has been a recent surge of interest in the field, the available theory is still limited, mainly because of the mathematical difficulties involved. Event-based constraints can be seen as a form of nonlinearity or hybrid element in the control loop, and they invariably seem to cause classical closed-form solutions such as employed in linear-quadratic Gaussian (LQG) control to break down. It is not surprising that these problems have been approached from a number of different perspectives, including, e.g., nonlinear, hybrid, and optimal control.

Since any definitive answer still seems to be well out of reach, there remains a need to experiment with different approaches, and to develop tools to compare the various approaches suggested in a quantitative manner. The study of prototype problems is also important to gain insight into the problem structure and the characteristics of event-based control.

The focus of this thesis is on event-based control in the face of stochastic disturbances. This setting is particularly natural with communication and sensing constraints, where it is the disturbances that create the need for communication and sensing, respectively. Two kinds of problems are studied: The first is an LQ-like control problem with added cost per communication or actuation event, and the *sporadic constraint* of a minimum inter-event time. The second kind of problem is state estimation from measurements with set-bounded uncertainty, such as with coarsely quantized measurements. This case also occurs with intelligent sensor nodes or with state estimators for event-based control, which only transmit a state estimate when it deviates by some tolerance from what the receiving node can predict.

Sporadic Control

The event in event-based control can stand for many things, among them the transmission of a message across a data network and the execution of a controller update. In reality, two such events can almost never occur arbitrarily close in time. A major objective of this thesis is to study event-based control subject to this almost universal limitation, expressed by the sporadic constraint of an enforced minimum inter-event time.

Most sporadic control problems in the thesis are posed in a continuous-time setting, such that an event may occur at any time after the minimum inter-event time has elapsed. To be able to guarantee a minimum inter-event time, other authors have considered event-based control in a discrete-time setting, or in continuous time where they derive a posteriori bounds on the minimum inter-event time given some triggering law.

Motivations to study continuous-time sporadic control include a potential for better performance than a discrete-time formulation, and the possibility to make good use of the available resources by enforcing the inter-event time explicitly so that the controller can adapt to it. A continuous-time formulation also precludes the need for synchronized clocks (for, e.g., control over networks), and though it brings some technical complications, it actually simplifies the control problems in some respects. To the best of the author's knowledge, this thesis contains the first investigation of stochastic event-based control in a continuous-time sporadic setting.

Path Constraints

Though a good number of event-based controllers are described in the literature, it is often difficult to evaluate and compare them. Many papers focus on some particular mixture of architectural assumptions, analyzed with some particular method. Few papers compare performance to other methods; even a fair comparison to periodic control is often missing.

While it is probably too hard to find optimal event-based controllers in most cases, quantitative evaluation of controllers for a given process need not be. To this end, development of analysis methods on one hand, and investigation of problem formulations and architectures for event-based control on the other, must be separated.

A common framework that captures many event-based control problems is given by hybrid systems. Though extremely expressive, hybrid modeling is in many frameworks not unlike filling out a form, and often quite cumbersome. The framework of *path constraints* is an attempt at a minimal component-oriented framework for stochastic hybrid systems, where the user needs only bring in exactly those kinds of components that are needed in his or her model. Any such model is amenable to analysis of upper and lower bounds on control objectives, a procedure which can (and partly has) been automated (though numerical issues do create a need for human intervention).

Outline and Contributions of the Thesis

The thesis is composed of six introductory chapters and seven papers. Papers I–V are on the topic of event-based control, while Papers VI and VII consider event-based state estimation. This section outlines the contents of the introductory chapters and the contributions of each paper.

Chapter 1 – Introduction

Event-based control can mean many things. An attempt is made to describe what it is, or can be, by way of example.

Chapter 2 – Related Work

This chapter gives a general overview of the literature on event-based control, with a focus on topics more closely related to the thesis.

Chapter 3 – Event-Based LQ Problems

An event-based linear-quadratic (LQ) control problem is formulated, which roughly encompasses the problems studied in Papers I–V. The problem is discussed, as well as how the papers relate to this formulation.

Chapter 4 – State-Space Methods for Stochastic Optimal Control

The primary tools employed in Papers I–V are various forms of state-space methods. This chapter presents a brief overview of related methods and concepts such as dynamic programming, stationary state distributions and value functions, and outlines how the papers make use of different parts of this toolbox.

Chapter 5 – Event-Based Estimation under Sensing Constraints

Papers VI and VII treat event-based estimation problems that arise with sensing constraints that cause set membership uncertainty at the estimator (in combination with stochastic uncertainty). A prototypical event-based estimation problem is formulated and discussed, and the approaches of the papers are outlined.

Chapter 6 – Outlook

This chapter outlines some interesting directions for future work.

Paper I

Henningsson, T. (2011): "Sporadic event-based control using path constraints and moments." In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*. Orlando, Florida, USA.

This paper treats an event-based variation on a general LQ control problem, in mixed continuous and discrete time with a fixed cost and minimum waiting time after each actuation event. *Path constraints* are introduced as a means to model this kind of stochastic hybrid systems in a component-oriented manner. Constraints are provided to model continuous-time and discrete-time dynamics, random and controlled mode switching, and stochastic disturbances. Constraints on trajectories' moments are extracted from the path constraints and used to bound control objectives over an infinite horizon, with the aid of sum-of-squares techniques and convex semidefinite programming. Bounds that can be derived include lower bounds on achievable cost for any controller, and upper bounds on cost with a given controller. The bounds become monotonically tighter as the order of trajectory moments is increased.

Paper II

Bernhardsson, B. and T. Henningsson (2012): "A Riccati-like equation for finding optimal elliptical triggering rules." Manuscript in preparation.

This paper builds on the methods developed in Paper I to derive explicit solutions to the same control problem when the process is an n-dimensional integrator, with arbitrary positive definite process noise covariance and state penalty matrices. It turns out that in this case there exists an analytic expression for the value function, whose parameters can be determined from a novel Riccati-like equation. The value function is used in two cases to find the optimal threshold, which is ellipsoidal, and to express the trade-off between state variance and event rate.

The Riccati-like equation for the optimal value function was suggested by B. Bernhardsson and analyzed by both authors. T. Henningsson developed the value-function-based proof machinery. B. Bernhardsson investigated the optimal state distributions and the scalar integrator plant with minimum inter-event time constraint. Both authors worked on the general formulation in $\mathbb{R}^n$.

Paper III

Henningsson, T., E. Johannesson, and A. Cervin (2008): “Sporadic event-based control of first-order linear stochastic systems.” *Automatica*, **44:11**, pp. 2890–2895.

This paper introduces the concept of sporadic control to model the practical constraint of a minimum time between any two control events, in contrast to aperiodic control where events may occur arbitrarily close. Two variants of sporadic controllers are proposed, using continuous-time and discrete-time measurements. It is argued for first-order systems with impulse control and white noise disturbances that the optimal sporadic controller will use a threshold strategy, generating a control event whenever it is allowed and the state becomes big enough. Ways of computing this optimal threshold and the associated performance in terms of regulation and event rate are described. The best achievable trade-off between regulation error and event rate is characterized for the sporadic, periodic and aperiodic controllers, showing that many of the benefits of aperiodic control are retained even with the practical constraint of a minimum inter-event time. It is also shown that sporadic control can not only greatly reduce the required communication rate, but also reduce the regulation error somewhat, compared to periodic control. Different possibilities for the generalization of the control problem at hand to higher dimensional systems are discussed.

T. Henningsson wrote about and did the simulations for the continuous-time case, and rewrote the paper for *Automatica*. E. Johannesson wrote about and did the simulations for the discrete-time case. A. Cervin wrote the introduction and assisted in the structuring and editing of the manuscript.

This paper is an extension of

Johannesson, E., T. Henningsson, and A. Cervin (2007): “Sporadic control of first-order linear stochastic systems.” In *Proceedings of the 10th International Conference on Hybrid Systems: Computation and Control*, Pisa, Italy.

Paper IV

Cervin, A. and T. Henningsson (2008): "Scheduling of event-triggered controllers on a shared network." In *Proceedings of the 47th IEEE Conference on Decision and Control*. Cancún, Mexico.

This paper treats the control problem when several loops (for integrator plants) of the same type as in Paper III are closed over a shared network. Models for the medium access protocols TDMA, FDMA, and CSMA, the latter with three different prioritization mechanisms, are stated. A suitable control setup is described for each case, and ways to evaluate the expected regulation error and to choose optimal parameters for the controllers are described, some based on Monte Carlo simulations. The performance when controlling N integrator plants is compared for the protocols, showing that CSMA quickly yields superior performance, asymptotically requiring only a third of the bandwidth or giving a third of the error compared to any of the other protocols. A case study for control of one stable, one integrator, and one unstable plant is also presented, yielding similar conclusions.

A. Cervin wrote most of the paper and performed most of the simulations. T. Henningsson derived the optimal control policy for the case of two integrators and shared information between the controllers, using dynamic programming for controlled Markov processes.

Paper V

Henningsson, T. and A. Cervin (2010): "A simple model for the interference between event-based control loops using a shared medium." In *Proceedings of the 49th IEEE Conference on Decision and Control*. Atlanta, GA.

This paper contains an approximate analytical analysis of the sporadic control problem treated in Paper IV, with identical control loops for integrator plants closed over a shared medium. Each transmission takes a fixed amount of time, during which no other loop may transmit. A simple approximate system model is formulated by partitioning the system into one foreground loop and $N-1$ background loops, and replacing the fixed minimum inter-event time with an exponentially distributed one. A Markov process model is formed with a discrete channel busy state, and the continuous state of the foreground plant. An analytic expression for the model's stationary state distribution is derived from a boundary value problem formulation, and used to investigate the approximate aggregate behavior of the control loops. The results are compared to the Monte Carlo simulations from Paper IV.

T. Henningsson wrote most of the paper. A. Cervin assisted with reviewing and performed the Monte Carlo simulations.

Paper VI

Henningsson, T. and K. J. Åström (2006): "Log-concave observers." In *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems*. Kyoto, Japan.

This paper investigates the problem of Bayesian state estimation with log-concave measurement likelihoods and process noise — a considerable generalization of the Kalman filter setting — and the properties of log-concave functions that can be used to say something about the state estimation problem. The optimal Bayesian state estimator for discrete-time linear systems is described, and split up into dynamics, process noise, and measurement updates, each acting on the probability distribution of the state conditioned on the measurements. It is shown that the conditional state density will be log-concave. Using the concept of strongly log-concave functions, theorems are derived that allow to upper bound the estimation error covariance of the Bayesian estimator. The upper bound is compared in simulations to a grid-based approximation of the Bayesian estimator and a Kalman filter designed using insight gained on the estimation problem, for a double integrator with quantized measurements.

T. Henningsson wrote most of the paper. K. J. Åström did extensive reviewing.

Paper VII

Henningsson, T. (2008): "Recursive state estimation for linear systems with mixed stochastic and set-bounded disturbances." In *Proceedings of the 47th IEEE Conference on Decision and Control*. Cancún, Mexico.

This paper presents a recursive state estimator for linear systems that are subject to both stochastic and uncertain (set-bounded) process and measurement disturbances. The structure for a recursive state estimator is proposed, allowing to model general state estimation problems with combined stochastic and set-bounded measurement and process disturbances. An optimization procedure for selecting the filter gain considering both the resulting stochastic and set-bounded error is described. For the case of ellipsoidal uncertainty sets, an LMI is formulated to optimize the feedback gain and fit the best one-step optimal ellipsoidal over-bound available using the S-procedure onto the set-bounded uncertainty. The estimator is compared in simulations to a grid filter and a time-varying Kalman Filter for the case of a double integrator with quantized measurements, where it comes quite close to the almost optimal grid filter.

Other Publications

Henningsson, T. (2008): "Event-based control and estimation with stochastic disturbances." Licentiate Thesis ISRN LUTFD2/TFRT--3244--SE. Department of Automatic Control, Lund University, Sweden.

The licentiate thesis includes Papers III, IV, VI, VII, as well as the paper [Henningsson and Cervin, 2009] below.

Henningsson, T. and A. Cervin (2009): "Comparison of LTI and event-based control for a moving cart with quantized position measurements." In *Proceedings of the European Control Conference*. Budapest, Hungary.

Henningsson, T. and A. Rantzer (2007): "Scalable distributed Kalman filtering for mass-spring systems." In *Proceedings of the 46th IEEE Conference on Decision and Control*. New Orleans, LA.

1

Introduction

This chapter attempts to describe what event-based control is, why we might want to use it, and some of the general challenges involved. We will not try to give a rigorous definition of the term, since it applies to many problems that are very different to each other. Instead, we will try to define it by example (and with periodic control as the counterexample).

1.1 What is Event-Based Control?

Control loops are usually implemented using periodic sensing, computation and actuation. When some of these *events* come with a cost or limitation, however, it makes sense to use a control strategy that takes this into account, e.g. by triggering events only when necessary. In *event-based* control problems, such event limitations are incorporated in the formulation.

This broad class of event-based control problems includes problems with very different kinds of limitations, that arise in different scenarios, and will typically need correspondingly different kinds of solutions (for an overview of possible scenarios, see also [Åström, 2007]). Depending on the limitation, it might e.g. be addressed with

- sensors and controllers connected over a data network, that only transmit when the message differs from what the recipient expects.
- a control law that is only updated when measurements have changed by a certain amount, to save processing power.
- a state estimator with updates triggered by changes in a quantized measurement signal.
- a controller that plans the trajectory of an on/off actuator to optimize the process response using few switches.

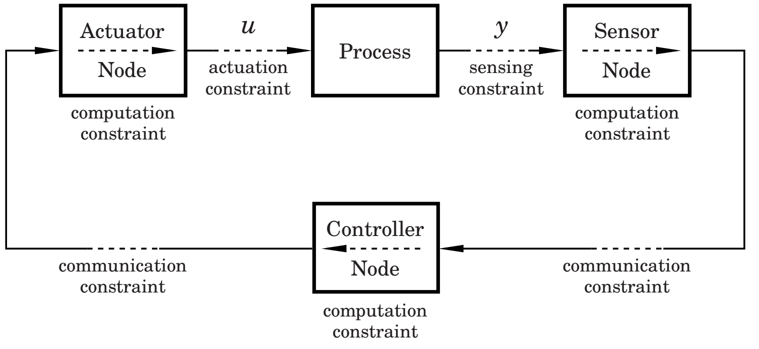

Figure 1.1 Control loop with various constraints motivating event-based control.

Possible advantages with event-based control include

- better use of limited or costly resources for communication/actuation/computation/measurement, etc.
- reduced need for centralized coordination and synchronized clocks.
- simplified control laws in some cases, due to the fact that they only need to be evaluated under event trigger conditions.
- more natural treatment of components with discrete characteristics, such as on/off actuators, sensors with quantized measurements, etc.

On the other hand, event-based control often implies

- harder analysis and design compared to periodic control, since the sampling operator becomes nonlinear.
- less predictable event patterns, which may make it harder to schedule the use of shared resources.

1.2 Scenarios for Event-Based Control

We will now describe some possible event-based limitations, and how they can enter a control loop such as in Figure 1.1. This will give an important way to classify event-based control problems. Practical problems will often have multiple kinds of limitations.

Communication constraints. The smallest unit of communication in a modern data network is a *data packet*. Transmission of a packet takes some time, and often requires exclusive transmit access to the link. In wireless networks it is harder to avoid packet losses, and battery-powered nodes must make do with limited power to spend on sending and listening to the channel.

The overhead associated with each packet and the relatively low bit count needed in most control problems means that it is usually the timing of transmission events that is restrictive, rather than the bit rate per se. While there is a trend of increasing bit rate in data networks, this seems to be achieved at the expense of longer packets and lower packet rate.

Actuation constraints. Some actuators can not respond to a continuously varying control signal. Examples include on/off actuators such as relays and some satellite thrusters [Dodds, 1981], and medical injections that naturally take the form of impulses. Wear and tear can sometimes be reduced by limiting the number of changes in control signals. Actuators driven by a piecewise constant control signal (zero-order hold) or generalized hold (see e.g. [Åström, 2007]) also give rise to actuation constraints.

Sensing constraints. Many common kinds of sensors produce measurements at ahead of time unknown events. Examples include rotary motion encoders that give pulses at fixed angular increments and A/D-converters with coarse resolution, where each change in the measured value can be seen as an event. Another example is the *send on delta* protocol [Neugebauer and Kabitzsch, 2004], where a sensor will only transmit measurements that differ by more than some given tolerance from the last one transmitted. These scenarios all have in common that both the occurrence of an event and its *absence* carries some information.

Computation constraints. Almost all controllers deployed today are implemented using computers. In cost-sensitive applications, it is often desirable to keep down the amount of processing power available to the control task, and to make it share these resources with other tasks. Energy and hardware costs can be saved if the controller can do fewer computations, only do them when necessary, and preferably at any time that is opportune for the scheduler. Computation constraints can be present to a greater or lesser degree in each of the sensor, controller, and actuator nodes to be designed in Figure 1.1.

There are some interrelations between the different kinds of constraints. Examples:

- If the actuator node's strategy in Figure 1.1 is known or fixed, a communication constraint between controller and actuator will be a special case of actuation constraints, since the actuator node will act like a generalized hold.

- If, on the other hand, the actuator node is to optimize its strategy freely based on the information that it has, and the controller uses a threshold to decide when to transmit, the actuator node will experience this threshold rule as a sensing constraint.

- Without computation constraints, it can be desirable for the sensor node in Figure 1.1 to implement an observer and transmit filtered state estimates, and for the actuator node to keep a running open-loop prediction of the state between control updates to recalculate its control signal. A fixed strategy in one of these nodes can be seen as an extreme computation constraint, but also acts like a sensing or actuation constraint, respectively.

As is often the case in optimization, adding constraints might make a problem either easier or harder to solve. An example of the latter is that the separation principle in [Molin and Hirche, 2009] ceases to hold when the actuator node is forced to use a zero order hold to generate its output.

1.3 Event-Based Control over Data Networks

One important motivation for event-based control in general, and for many of the papers in this thesis in particular, is control over data networks. While the same network should probably be shared by many nodes in order to gain practical benefits, it may be easier to first investigate a single control loop in isolation to build up an understanding of the phenomena involved.

Figure 1.2 illustrates one possible architecture for event-based control over networks, seen from the perspective of a single control loop. An event generator at the process receives the raw sensor signal, filters it, and decides when to try inform the controller of what it knows. The message is sent to the event-based observer at the controller, which updates its state estimate, and then predicts the process state in open loop until it receives new information. The control event generator uses the state estimate to decide whether to send a new command to the control signal generator at the process input, which acts as a *generalized hold* (see e.g. [Åström, 2007]) that generates a control signal waveform based on its the most recently received message.

The setup can be varied in many ways. The network could be shared with other control loops and with other, non-control related tasks, usually with longer deadlines. There could be many sensor and actuator nodes

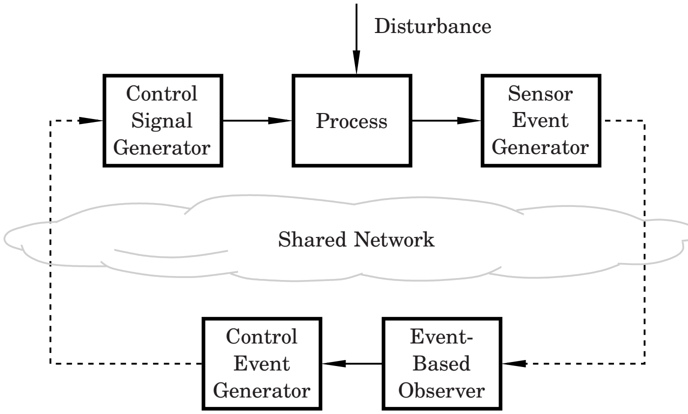

Figure 1.2 Event-based control over a shared network. Solid lines represent continuous signal transmission, while dashed lines represent event-based transmission.

attached to the process in different places, orchestrated by the controller. Some of the signals transmitted across the network in Figure 1.2 may instead use dedicated paths.

The kind of medium access scheme used for a network interacts strongly with the choice of control strategy. In a *time-division multiple access* (TDMA) scheme, medium access for different nodes is scheduled in advance, which is suitable for periodic control. For event-based control, it is more natural to use a *random access* scheme instead, where multiple nodes may attempt to transmit packets at the same time. When they do, the packets are said to *collide*.

While TDMA can be virtually collision free, random access schemes are more flexible when it comes to redistributing communication resources, due to added/removed nodes or changing needs. Since there will usually be collisions, event-based control over networks is more dependent on network characteristics than periodic control.

Modeling of Data Networks for Event-Based Control

To formulate a control problem with a data network in the feedback loop, a model of the corresponding communication constraints is needed. While accurate modeling of data networks is an intricate task, simplified models are usually sought for control design.

Various characteristics of a real network may be included in the model. *Slotted* networks divide time into transmission slots, whereas packet initiation may take place at any time in an *unslotted* network. Networked communications often suffer from *packet dropouts* and *delay*. In many network protocols, a node that receives a message will transmit an *acknowledgment* message back to the sender, which will also be subject to loss and delay; the sender will often attempt to *retransmit* if the acknowledgment is not received.

Simple network constraints that have been considered in the literature include

- to use a *limited number of packets* over a finite horizon,
- to use a *limited average packet rate* over an infinite horizon,
- to include a *fixed cost per packet* as a term in the control objective,
- to enforce a *minimum inter-event time* (sporadic control).

Looking closer, a major distinction in network characteristics is between wired and wireless networks.

Wired networks. In wired networks, it is usually possible to arbitrate access to the medium upon collisions so that all nodes except one will back off, without disturbing the transmission. Simple models of arbitration include

Static priority with a fixed ranking between nodes.

Dynamic priority, where each node ranks the importance of its packets.

Random priority, where any node has equal probability to gain access.

The *Controller Area Network* (CAN) standard uses a shared bus in a wired AND configuration and a priority field in the packet header to give access to the highest priority packet; the priority may be chosen dynamically by the sending node from a limited set of priorities.

Wired networks may consist of one shared link or many point-to-point links connected in a network graph; collisions can be avoided in the latter case, but packets may still have to compete for access to the same link.

Wireless networks. In wireless networks, the topology is dictated by the radio fading conditions between different nodes. Over modest distances and without obstructions it seems reasonable to assume that all nodes will interfere with each other. Unlike wired networks, it is usually not possible to guarantee mutual exclusive access to medium, and packets that overlap in time will typically be lost. In the simple ALOHA protocol, senders make no attempt to avoid collisions. In *carrier-sense multiple access* (CSMA) protocols, transmitters attempt to avoid collisions by listening to the channel to see if it is free before sending.

1.4 General Challenges

One fundamental difficulty with event-based control problems is that the sampling operator becomes nonlinear. This typically means that the kind of closed form solutions used in e.g. linear-quadratic (LQ) control break down; most event-based control problems do not have a known closed-form solution. Whereas e.g. the optimal LQ control law is independent of the presence and intensity of stochastic disturbances, optimal event-based controllers will typically not be.

Event-based control problems also often include elements of other problems that are known to be hard. Examples are:

Dual control. In dual control (see [Feldbaum, 1960 1961]), the controller might use its control signal to influence the information gained from measurements. Certainty equivalence controllers will generally not be optimal since they lack any probing effect, and tractable optimal design procedures are generally not known.

The most obvious case where dual control shows up is with control under sensing constraints, but these kind of issues also have to be considered in cases with communication constraints and several agents.

Distributed control with non-global knowledge. With control under communication constraints, control signals will often be issued by multiple agents that have different information. Such problems are known to be hard even in an otherwise traditional setting with linear dynamics, quadratic costs, and Gaussian disturbances (see [Witsenhausen, 1968]), and the optimal controller will typically not be linear.

The information pattern is *partially nested* if, for any decision that affects another decision, the second decision maker knows all that the first one did. In this case there is no incentive to use the control signal for signaling (information transmission).

Optimization under complexity constraints. Optimization problems where the complexity of program code is part of the objective are in general extremely hard to solve exactly. This makes realistic computation constraints hard to quantify.

Approach to the Challenges

The papers in this thesis take a rather practical approach to circumvent the challenges outlined above. The main focus lies on how to deal with the nonlinear sampling operator. For the other challenges;

- Dual control is avoided by either having a single agent with complete state information (control case) or not closing the loop (estimation case).
- Distributed control problems with multiple loops closed over a shared network are handled by restricting attention to a single-parameter family of control laws, which are based on local information.
- Complexity constraints are avoided in the control case by not posing any computation constraints a priori; the resulting control laws turn out to have modest computation requirements anyway.
- The only direct measure to reduce online computation complexity is the ellipsoidal over-approximation of uncertainty sets used in Paper VII.

2

Related Work

This chapter gives an overview of the literature on event-based control, with a focus on topics more closely related to the contributions of the thesis. For an overview of early work on event-based control, see [Åström, 2007] and [Heemels *et al.*, 2012], and references therein.

2.1 Comparison between Event-Based and Periodic Control

Stochastic Case

The trade-off between event rate and control performance for event-based control as compared to periodic is investigated in [Åström and Bernhardsson, 2002]. A continuous-time first-order plant with white noise input disturbance is considered, with impulse control at events. It is shown that event triggering using a simple threshold rule outperforms periodic control with a factor of three regarding either stationary state variance or average event rate, when the other is fixed. Favorable trade-off between state variance and number of events compared to periodic control is also demonstrated in various settings for integrator plants in [Imer and Basar, 2006a], [Molin and Hirche, 2009], [Molin and Hirche, 2010c], [Molin and Hirche, 2010a], for first-order plants in [Xu and Hespanha, 2004b], [Xu and Hespanha, 2006], [Rabi *et al.*, 2006], [Rabi and Baras, 2007], [Rabi *et al.*, 2012], [Cervin and Johannesson, 2008], and for a second-order plant in [Li *et al.*, 2010].

Performance compared to periodic control has also been investigated for multiple control loops closed over a shared medium. [Molin and Hirche, 2011a] demonstrates favorable state variance compared to periodic control for a number of integrator plants, closed over a shared medium with random arbitration. [Blind and Allgöwer, 2011a], [Blind and Allgöwer, 2011b] consider instead a wireless network using the simple ALOHA protocol, and compare state variance between periodic control and a heuristic event-based strategy proposed in [Rabi and Johansson, 2009b]. They find that periodic control performs better in this case. [Rabi and Johansson, 2009b] come to the opposite conclusion; partly due to different assumptions and partly because they care about reducing network traffic as well as state variance.

Lower state variance with event-based control compared to state independent scheduling is also demonstrated in [Rabi and Johansson, 2009a], [Rabi *et al.*, 2006], [Imer and Basar, 2006b], [Ramesh *et al.*, 2009], and for transmission of independent identically distributed (IID) measurements in [Imer and Basar, 2005] and [Imer and Basar, 2010].

Non-Stochastic Case

[Wan and Lemmon, 2009] demonstrates favorable convergence for distributed optimization as a function of the number of messages transmitted for an event-triggered strategy compared to periodic transmissions. [Wang and Lemmon, 2008] demonstrates favorable average period to stabilize an inverted pendulum for event-based control compared to self-triggered control and periodic control with a guaranteed stabilizing period. The case study [Sandee *et al.*, 2007a] investigates event-based control applied to velocity control of the paper feed in a printer, and demonstrates that it allows to use a cheaper encoder with coarser quantization, reduced processor load, and fewer controller tuning parameters.

2.2 Stochastic Event-Based Control

[Xu and Hespanha, 2004a] introduces an event-based *remote estimation problem*, with an intelligent sensor that decides when to transmit its information to a remote observer, to achieve a good remote state estimate using few transmission events. A continuous-time setting with a linear plant and Gaussian noise is considered, with an open-loop predictor at the receiver. Stochastic triggering rules are proposed with an estimation error dependent Poisson intensity of the form $\lambda = (e^T Pe)^k$, and are shown to give moment bounded estimation error with for any positive definite P. Conditions for moment stability with constant Poisson triggering intensity $\lambda_{\max}$ are also given; [Xu and Hespanha, 2005] shows that these kind of conditions hold equally well with a saturated error-dependent triggering intensity $\lambda = \min(\lambda_{\max}, (e^T Pe)^k)$.

[Ramesh *et al.*, 2009] compares a number of medium access schemes for a single control loop closed over a wireless network. Control of multiple loops for integrator plants over a wireless network is considered in

[Rabi and Johansson, 2009b], [Blind and Allgöwer, 2011a], and [Blind and Allgöwer, 2011b]. A Markov chain model for the medium access protocol in wireless networks is proposed in [Ramesh *et al.*, 2011b] and [Ramesh *et al.*, 2012], and used to analyze control of multiple loops over a wireless network. [Weimer *et al.*, 2012] considers a centralized state estimation problem with multiple sensors that communicate over a wireless network, with energy cost for both sending and listening.

Stochastic Optimal Control

In optimal control, assumptions are essential. A control policy that is optimal under one set of assumptions may be suboptimal under slightly different ones.

Most results on event-based stochastic optimal control in the literature seem to be derived for single loop control problems. Most are also derived under communication constraints, where each event is a transmission of the state (estimate) from sensor to controller, and under a *nested information pattern*, where the sensor knows everything that the controller does. A few optimality results are derived under actuation constraints, where each event effects a change in a piecewise constant control signal.

Analytic solutions. Analytic solutions to event-based optimal control problems in the literature seem to be limited to first-order plants with white noise input. In [Rabi, 2006] and [Rabi *et al.*, 2012], continuous-time optimal remote estimation problems are solved on a finite horizon given a limited number of samples. [Rabi, 2006] also solves a corresponding infinite horizon problem given a limit on the average sample rate. [Rabi *et al.*, 2008] solves a single sample optimal control problem with a zero-order hold actuator.

Structural results. [Xu and Hespanha, 2004a] considers a single-loop discrete-time remote state estimation problem. It is shown how to formulate the problem using dynamic programming, and that value iteration converges to the optimal solution under mild technical assumptions. The problem is posed for a general linear plant with Gaussian process noise, with objective to minimize the infinite horizon mean square estimation error plus a linear penalty on event rate.

[Molin and Hirche, 2009] derives a separation principle for single-loop event-based control with communication constraints and a nested information pattern. (For a more accessible account, see [Molin and Hirche, 2012b].) The problem is a finite horizon discrete-time optimal control problem with general linear plant and Gaussian process noise. There are two players: the controller, which decides the control signal at each time step, and the sensor, which observes the state of the process and decides when to transmit a message to the controller. The objective is to minimize the sum of an LQ cost and a fixed cost per event.

Under these assumptions, it is shown that the optimal controller uses standard LQ feedback from its mean square plant state estimate, which it under symmetry assumptions can be form using a linear state predictor that runs in open loop between events. At events, it is sufficient for the sensor to transmit the current state. The optimal trigger problem at the sensor is reduced to a remote estimation problem, with quadratic weighting on the estimation error derived from the LQ control solution.

[Molin and Hirche, 2010a] derives corresponding results in continuous time, while [Molin and Hirche, 2010b] extends the results to noisy measurements at the sensor, which then needs to run a stationary Kalman filter. [Lipsa and Martins, 2009] and [Lipsa and Martins, 2011] prove similar separation results for the special case of a first-order plant. Similar problems are also considered for first-order systems in the context of estimation in [Imer and Basar, 2005] and [Imer and Basar, 2010], control in [Imer and Basar, 2006a] and [Bommannavar and Basar, 2008], and a case where the controller is to decide at each sample whether to measure or control in [Imer and Basar, 2006b]. For an argument that this is a hard problem, see [Ramesh *et al.*, 2011a].

Suboptimal Control

[Cogill *et al.*, 2007] proposes an easily computable ellipsoidal threshold for a remote estimation problem with stochastic disturbances in discrete time over an infinite horizon, and shows that it realizes a cost within a factor six of the optimal. [Li and Lemmon, 2011] proposes a similar threshold policy that works for unstable systems and analog transmission noise, but appears to give slightly worse performance. [Li *et al.*, 2010] considers a similar remote estimation problem, but with limited event budget over a finite horizon. The value function is overestimated as a minimum of quadratic functions given the number of samples and transmissions left, and used to approximate the optimal triggering rule.

[Cogill, 2009] applies similar techniques as in [Cogill *et al.*, 2007] to an actuation constrained single-loop discrete-time impulse control problem, to find jointly an ellipsoidal triggering rule and a linear control law that minimizes an upper bound on the cost. The procedure is based on optimization over quadratic approximate value functions using modest size linear matrix inequalities (LMI:s).

[Molin and Hirche, 2011a] considers multiple loops closed over a shared medium with random arbitration. A two level design procedure is proposed, with a convex packet rate allocation problem based on the controllers designed in isolation. The design is shown to be asymptotically optimal as the number of control loops grows, under assumptions of uncorrelated disturbances between the loops. Sufficient conditions are given for stability with a finite number of loops.

The separation principle in [Molin and Hirche, 2009] relies on a nested information pattern; in particular that the sensor knows immediately whether a packet was delivered successfully to the controller. [Molin and Hirche, 2010c] considers the case when nestedness breaks down due to a fixed acknowledgment delay from controller to sensor. A heuristic strategy is proposed where the sensor estimates delivery success from measurements, while waiting for acknowledgments.

[Molin and Hirche, 2011b] exploits the same separation results to propose an order reduction heuristic for event triggering rule design, based on projection of the process dynamics onto the subspace that determines the current LQ control signal, and shows that it preserves stability. [Molin and Hirche, 2012a] builds on the same results to illustrate that a non-symmetric state predictor at the controller can beat a symmetric one, for a symmetric problem with multimodal process noise.

[Ramesh *et al.*, 2011a] restricts admissible control laws to a subset, in order to derive a separation principle similar to [Molin and Hirche, 2009]. [Antunes *et al.*, 2012b] considers a plant with a single controller that schedules access to a shared medium for a number of sensor and actuator nodes, in discrete time with linear dynamics and Gaussian noise. The controller tries to optimize the schedule online to minimize a quadratic cost given a prescribed control law, and is able to trigger actuation events in an event-based fashion.

2.3 Event-Based Control under Bounded Disturbances

[Heemels and Sandee, 2006] and [Heemels *et al.*, 2008] consider event-based control of a linear system with bounded disturbance, where the control signal is held constant between events. Two sampling mechanisms are considered: *uniform* (later called *periodic event-triggered control* (PETC)), where the controller decides at periodic instants whether to trigger an event, and *non-uniform* (or sporadic) with a minimum inter-event time constraint. It is shown how properties of ultimate boundedness for the event-triggered system can be shown from ultimate boundedness of discrete-time piecewise linear (uniform case) or linear system (non-uniform case). [Heemels *et al.*, 2011] considers analysis methods for PETC based on LMI:s, and shows how to bound exponential convergence rate, disturbance gain, and minimum inter-event time for a given linear system, control law, and quadratic trigger threshold.

[Lunze and Lehmann, 2010] considers continuous-time event-based control with a network link between the sensor, which has full state information, and the controller, with a state predictor. Bounds are shown on disturbance gain and minimum inter-event time. The analysis is extended to output feedback with an observer at the sensor in [Lehmann and Lunze, 2011a], to short communication delays and packet losses in [Lehmann and Lunze, 2012], with an integral part in [Lehmann and Lunze, 2011b], and adapted to input-output linearizable nonlinear plants in [Stoecker and Lunze, 2011]. The setup is compared to to periodic control, and event-based control with a zero-order hold at the actuator in [Lehmann *et al.*, 2012], which finds better bounds on event rate given the same performance with a predictor at the actuator as opposed to zero-order hold. A discrete-time formulation is given in [Grüne *et al.*, 2010], together with a global control approach for nonlinear systems based on coarsely quantized state information and dynamic programming. A continuous-time version with zero-order hold at the actuator is analyzed in [Wang and Lemmon, 2011a].

[Donkers and Heemels, 2012] considers a similar case, where both control and sensor signals are transmitted over a collision free network, and held constant between their respective transmission events. Asymptotic stability and disturbance gain are analyzed using LMI:s. [Donkers and Heemels, 2013] derives similar results (but in discrete time) with predictors for control and sensor signals between events, and demonstrates reduced communication frequency with maintained performance compared to zero order hold in examples. [Wang and Hovakimyan, 2010] uses $\mathcal{L}_1$-adaptive control to bound the difference between actual and ideal state trajectories in a continuous-time setting where nodes sample their own state continuously, and use events to broadcast it.

2.4 Event-Based Control for Deterministic Systems

Single Control Loop

Control of nonlinear continuous-time systems with zero-order hold is considered in an input-to-state stability (ISS) framework in [Tabuada, 2007b], [Wang and Lemmon, 2008], and [Wang and Lemmon, 2011c], where the latter two prescribe an exponentially decaying reference value that the Lyapunov function must stay below. [Durand *et al.*, 2011] investigates a similar problem for linear systems. [Yu and Antsaklis, 2011] presents a simple event-triggered static output feedback strategy for passive nonlinear systems. Patterns of sampling interval sequences for control of continuous-time linear systems using zero-order hold are investigated in [Velasco *et al.*, 2009], where it is shown that the sequences can become chaotic. Event-based control is applied to consensus-like problems in e.g.

[Speranzon *et al.*, 2007], [Seyboth *et al.*, 2011].

A related approach to deterministic event-based control is *self-triggered control* (see e.g. [Lemmon *et al.*, 2007], [Anta and Tabuada, 2010]), where the controller calculates both a control signal and a suitable deadline for its next activation; essentially adapting the sampling period depending on the process state. While self-triggered control is often based on a pre-designed continuous-time control law, [Donkers *et al.*, 2011] formulates an online strategy to jointly optimize control signal and deadline using linear programming.

Control over Networks

Control of a number of systems closed over a shared network is considered in [Wang and Lemmon, 2010], [Wang and Lemmon, 2011b], [Mazo and Tabuada, 2011], and [Guinaldo *et al.*, 2012], with continuous-time linear dynamics and zero-order hold. [Antunes *et al.*, 2012a] considers a plant with a single controller that schedules access to a shared medium for sensor and actuators, in discrete time with linear dynamics. The controller tries to optimize the schedule online to minimize a quadratic cost given a prescribed control law. [Wan and Lemmon, 2009] considers distributed optimization with event-based updates.

2.5 Estimation and Control under Sensing Constraints

[Heemels *et al.*, 1999] and [Sandee *et al.*, 2007a] consider velocity control problems for rotating servos with coarsely quantized angle measurements. The control problems are parametrized in distance instead of time, yielding a discrete-time nonlinear plant, which is controlled using gain scheduling. [Marchand, 2008] proposes a static output feedback law based on quantized state measurements to stabilize a chain of integrators with input saturation. [Rabi and Baras, 2007] investigates feedback from quantized measurements for a first-order system with white noise disturbance. [Sijs and Lazar, 2009] considers a state estimator for linear systems with Gaussian noise, where measurements are triggered when the state reaches the boundary of a space-time tube. The likelihood of the state given that no event has been triggered is applied periodically to update the estimate, using a sum-of-Gaussians approximation of the indicator function. [Sijs *et al.*, 2010] uses the estimator in a model predictive control (MPC) formulation.

3

Event-Based LQ Problems

Linear-quadratic control is a central part of classic linear control theory, and a popular starting point to formulate event-based control problems. Most of the literature on stochastic optimal and suboptimal event-based control cited in Section 2.2 is based on LQ formulations, in either continuous or discrete time.

In contrast, the focus of the current chapter and of this thesis is on mixed time formulations that arise from continuous-time *sporadic control*, where the system evolves in continuous time waiting for the next event, and takes a discrete-time step over the minimum inter-event time after each event. This chapter formulates an event-based LQ control problem that roughly encompasses the problems studied in Papers I–V, and describes how it relates to these problems.

The treatment is restricted to stationary control problems, where the optimal control law is time independent, and in particular to average cost objectives. This setup is commonly used in practice, and is in some senses less complex than time-varying formulations, although the average cost formulation does necessitate some extra technicalities.

3.1 Linear-Quadratic Control

The classic LQ control problem considers a linear plant and quadratic cost on the plant state x and control signal u. A stationary LQ control problem to search for a state feedback control law f that minimizes the average cost J can be formulated in continuous time as

$$\begin{aligned}
\min_f J &= \limsup_{T\to\infty} \mathrm{E}\,\frac{1}{T}\int_0^T \left(x^T Q_x x + u^T Q_u u\right) dt \\
\text{s.t. } dx &= Axdt + Budt + dw, \quad \mathrm{E}(dwdw^T) = Rdt, \\
u &= f(x)
\end{aligned} \tag{3.1}$$

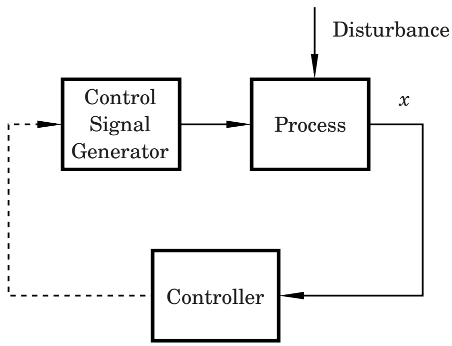

Figure 3.1 Event-based state feedback.

where E denotes expectation, x, u, and the disturbance w are functions of time, w is a Wiener process with incremental covariance $R \succeq 0$, $Q \succeq 0$ specifies the cost function, A, B are model matrices, and $\succeq 0$ means positive semidefinite. A similar LQ problem can be formulated in discrete time as

$$\begin{aligned} \min_f J &= \limsup_{N\to\infty} \mathrm{E}\,\frac{1}{N}\sum_{k=0}^{N} x_k^T Q_x x_k + u_k^T Q_u u_k \\ \text{s.t. } x_{k+1} &= \Phi x_k + \Gamma u_k + w_k, \quad \mathrm{E}(w_k w_k^T) = P, \\ u_k &= f(x_k) \end{aligned} \tag{3.2}$$

where x, u, and w are now functions of the discrete time k, w_k are independent zero mean Gaussian random variables, and Φ and Γ are discrete-time model matrices.

The solutions to both (3.1) and (3.2) are well known, see e.g. [Åström, 2006]. The optimal objective value J^* is finite under suitable assumptions, such as stabilizability. A quadratic value function can then be found for the system with optimal controller, by solving a continuous or discrete-time algebraic Riccati equation. The optimal value function in turn gives an optimal linear state feedback controller f that realizes the minimum $J = J^*$.

3.2 A Sporadic Linear-Quadratic Control Problem

We will now modify the LQ control problem to a form that might be suitable for the control loop in Figure 3.1, where

- there is a network link between the controller and control signal generator, with minimum inter-event time $\Delta T > 0$.

- the control signal generator is a fixed linear dynamic system, perhaps a zero-order hold, or a state predictor with an LQ feedback law at the output.
- each event results in a Dirac impulse at the input of the control signal generator.
- the controller causally decides the value and timing of events, given continuous measurements of the state x.

Since the formulation is in continuous time, we will take problem (3.1) as our starting point.

The plant dynamics

$$dx = Axdt + Budt + dw, \quad \mathrm{E}(dwdw^T) = Rdt,$$

take the same form as in (3.1), but now represent the *interconnection of the control signal generator and process* in Figure 3.1. The control signal

$$u = \sum_k u_k \delta(t - t_k), \qquad t_{k+1} - t_k \geq \Delta T, \quad \forall k, \tag{3.3}$$

represents the stream of events from controller to control signal generator in Figure 3.1, where $\{t_k\}_k$ are the *event times* and we have added the *sporadic constraint* of minimum inter-event time $\Delta T > 0$. In most realistic examples, the Dirac impulses will act directly only on the states of the control signal generator, which will in turn affect the process through the control signal.

We seek an admissible control law, i.e. one that causally decides

- the event times t_k, and
- the control signals u_k at events.

To capture both classic control performance and communication cost, we consider two objectives: the quadratic cost J_x and the event rate f_u,

$$\begin{aligned} J_x &= \limsup_{T\to\infty} \mathrm{E}\, \frac{1}{T} \left(\int_0^T x^T Q_x x\, dt \; + \sum_{\{k; t_k \leq T\}} u_k^T Q_u u_k \right), \\ f_u &= \limsup_{T\to\infty} \mathrm{E}\, \frac{1}{T} \sum_{\{k; t_k \leq T\}} 1. \end{aligned} \tag{3.4}$$

Note that the state cost $x^T Q_x x$ is integrated in continuous time as in (3.1), while the control signal cost $u_k^T Q_u u_k$ is summed over discrete events, as in (3.1). Since the control signal generator in Figure 3.1 will be part of the plant model in this formulation, a cost on the control signal that enters the process can be modeled as a term in Q_x.

The sporadic LQ control problem can now be formulated: find all admissible Pareto optimal state feedback controllers that realize a minimal pair (f_u, J_x). Another formulation is to minimize $J = J_x + \rho f_u$ for some event cost ρ.

3.3 Control Problems in Papers I–V

The sporadic constraint in (3.3) is accommodated in the papers by sampling the system for a time step ΔT, which is applied after each event. The more realistic case that the control impulse arrives at the plant with a delay equal to the minimum inter-event time can be incorporated into the discrete-time step without changing the model order.

Paper I considers a single control loop that takes a discrete-time jump akin to the discrete-time dynamics in (3.2) at each event, and evolves in continuous time in between, with arbitrary model matrices A, B, R, $\Phi, \Gamma, P, Q_{\text{flow}}, Q_{\text{jump}}$. Upper and lower bounds on the cost $J = J_x + \rho f_u$ are derived. This formulation is a generalization of the sporadic LQ problem outlined above.

Paper II derives explicit solutions to the event-based LQ problem when $A \in \mathbb{R}^n = 0$, in the two cases when either the minimum inter-event time $\Delta T = 0$, or when the state dimension $n = 1$, respectively.

Paper III investigates the trade-off curve (f_u, J_x) for a single control loop for a first-order plant (scalar x). Beside the continuous-time formulation, sporadic control is also considered in discrete time, where the controller samples the state with period ΔT and decides at each sample whether to trigger an event.

Papers IV and V treat the case with a number of identical integrator plants closed over a shared medium; the former using Monte Carlo simulations and the latter using analytical solutions for an approximate model. Paper IV also investigates a case with one stable, one marginally stable, and one unstable first-order plant closed over a shared medium.

4

State-Space Methods for Stochastic Optimal Control

Linear-quadratic (LQ) control problems, as e.g. described in the previous chapter, are usually solved using state-space methods. The most common approach is based on value functions, though a dual formulation using state covariance is also possible, as in e.g. [Rantzer, 2006].

Papers I–V also predominantly make use of state-space methods in various forms to treat event-based LQ problems. This chapter gives a brief and informal overview of related methods and concepts used, such as dynamic programming, stationary state distributions, and value functions, and outlines how the papers employ different parts of this toolbox.

Notation and Scope

We will limit the discussion in this chapter to (controlled) Markov processes with a finite state space $\mathcal{X}$, composed of N states. Let the probability distribution over the state at some given time be described by $p \in \mathbb{R}^{\mathcal{X}}$, such that $\mathbf{1}^T p = 1, p \geq 0$, where $\mathbf{1}$ is a vector of all ones, and $\geq$ is taken elementwise. We will take $V(x)$ to denote the element of the vector V at index x, to emphasize the similarity to the case when V is a function. Subscripts will be reserved to distinguish between e.g. vectors V_k.

The finite state formulation is directly applicable to control problems with a discretized state space. The same ideas as presented in this chapter can also be applied with an uncountable state space $\mathcal{X}$ under various technical assumptions, as is done in several of the papers. The reader may imagine the material in this chapter with probability vectors replaced by measures, cost vectors by functions, matrices by linear operators, etc.

4.1 Modeling

We first consider an uncontrolled Markov process; this includes the case of a controlled Markov process with a fixed (Markovian) controller.

The dynamics of the probability vector p can be modeled in either continuous or discrete time as

$$\dot{p} = \mathcal{A}p + \kappa, \tag{4.1}$$
$$p^+ = \mathcal{H}p + \Delta T\kappa, \tag{4.2}$$

respectively, where p^+ is the probability vector p at the next time step for the discrete-time dynamics (4.2), with step length ΔT. The source density per unit time κ can be used to model initiation and termination of trajectories. The dynamics matrices $\mathcal{A}, \mathcal{H} \in \mathbb{R}^{X\times X}$ must be such that

$$\begin{aligned} \mathbf{1}^T\mathcal{A} &= 0, & p(x) = 0, p \geq 0 &\implies (\mathcal{A}p)(x) \geq 0, \\ \mathbf{1}^T\mathcal{H} &= \mathbf{1}^T, & p \geq 0 &\implies \mathcal{H}p \geq 0, \end{aligned} \tag{4.3}$$

where $p(x)$ denotes the probability to be in state x.

Any control objective J that can be expressed as an expectation over the state x will be linear in p, e.g.

$$J = \mathrm{E}(c(x)) = c^T p, \tag{4.4}$$

for some suitable state dependent cost $c \in \mathbb{R}^X$. We will mostly be concerned with the case when J is an average cost over an infinite time horizon. This is achieved when (4.4) is evaluated with p as a stationary probability.

4.2 Stationary Distributions

Stationary distributions are stationary solutions to the dynamics (4.1) or (4.2), and so must satisfy

$$\mathcal{A}p + \kappa = 0, \tag{4.5}$$
$$(\mathcal{H} - I)p + \Delta T\kappa = 0, \tag{4.6}$$

respectively, with $p \geq 0, \mathbf{1}^T p = 1$ and $\mathbf{1}^T\kappa = 0$.

The similarity between the stationarity conditions (4.5) and (4.6) can be exploited to translate a model between continuous and discrete time, as long as we are interested only in the stationary distribution. Using these conditions, we see that the dynamics (4.1) and (4.2) will have the same stationary distributions as

$$p^+ = (I + \Delta T \mathcal{A})p + \Delta T \kappa, \tag{4.7}$$

$$\dot{p} = \frac{1}{\Delta T}(\mathcal{H} - I)p + \kappa, \tag{4.8}$$

respectively. With (4.7), we must make sure to pick ΔT small enough that $I + \Delta T \mathcal{A}$ is nonnegative, which is not always possible when the state space $\mathcal{X}$ is infinite.

4.3 Dynamic Programming

Dynamic programming as such is not used much in the thesis, but it is instructive to see how it relates to the techniques that are used. For ease of exposition, we consider dynamic programming in discrete time.

Consider a discrete-time controlled Markov process. Instead of a single set of dynamics and cost $(\mathcal{H}, c)$, we have a set indexed by the control variable a belonging to some finite action set $\mathcal{U}$, such that $(\mathcal{H}, c) = (\mathcal{H}_a, c_a)$. With a state dependent control $a = u(x), u \in \mathcal{U}^{\mathcal{X}}$, we can define $\mathcal{H}_u$ and c_u according to

$$\mathcal{H}_u p = \sum_{x \in \mathcal{X}} \mathcal{H}_{u(x)} p(x),$$

$$c_u^T p = \sum_{x \in \mathcal{X}} c_{u(x)}(x) p(x).$$

Consider the optimal expected cost

$$V_n^T p_n = \min_U \sum_{k=n}^{n_f} c_{u_k}^T p_k$$
$$\text{s.t. } p_{k+1} = \mathcal{H}_{u_k} p_k$$

over the time sequence of control laws $U = \{u_k\}_n^{n_f}$. We see that

$$V_n^T p_n = \min_u c_u^T p_n + V_{n+1}^T p_{n+1} = \min_u c_u^T p_n + V_{n+1}^T \mathcal{H}_u p_n,$$

so that V_n satisfies the recurrence

$$V_n(x) = \min_u c_u(x) + (\mathcal{H}_u^T V_{n+1})(x), \tag{4.9}$$

which is Bellman's equation for a discrete-time Markov process. If we can keep track of the value function V_n and carry out the minimization

(4.9), we can solve finite horizon optimal control state feedback problems for our Markov process. Under ergodicity assumptions, the iteration (4.9) converges except for the average value of V. Such a fixed point will satisfy

$$V(x) + J = \min_u c_u(x) + \left(\mathcal{H}_u^T V\right)(x). \tag{4.10}$$

4.4 Cost Bounding Using Value Functions

Dynamic programming allows to compute optimal costs and controllers for state feedback, but it quickly becomes intractable with increasing problem complexity. One difficulty with dynamic programming is that it is hard to approximate the minimization and the value function, and still have some kind of convergence to a reasonable solution. We will now consider related methods where approximations enter more naturally.

Consider the discrete-time dynamics (4.2). We want to evaluate and minimize the cost (4.4) in stationarity (4.6). Using the equality (4.6), we see that for any $V \in \mathbb{R}^{\mathcal{X}}$, the cost J satisfies

$$\begin{aligned} J = c^T p &= c^T p + V^T \Big((\mathcal{H} - I)p + \Delta T \kappa\Big) \\ &= (c + (\mathcal{H} - I)^T V)^T p + \Delta T V^T \kappa. \end{aligned}$$

Here, V can be interpreted as a value function.

Exploiting the fact that $p \geq 0, \mathbf{1}^T p = 1$, we see that

$$\min(c + (\mathcal{H} - I)^T V) \leq J - \Delta T V^T \kappa \leq \max(c + (\mathcal{H} - I)^T V). \tag{4.11}$$

The minimum and maximum above will coincide if we find a $V = V^*$ such that $c + (\mathcal{H} - I)^T V$ is constant across the states, which corresponds to a stationary solution (4.10), and will give an exact value for J:

$$c + (\mathcal{H} - I)^T V = \hat{J}\mathbf{1} \qquad \Longrightarrow \qquad J = \hat{J} + \Delta T V^T \kappa. \tag{4.12}$$

The usefulness of the current formulation is that we can get bounds on the objective J even without such an optimal value function V^*, but which will generally be tighter the closer we come to V^*. We also note that it is a convex problem to find the tightest upper or lower bound on J over a convex set of value functions $\mathcal{V}$, e.g.

$$\begin{aligned} &\max_{V \in \mathcal{V}} \underline{J} \\ &\text{s.t. } c + (\mathcal{H} - I)^T V \geq (\underline{J} - \Delta T V^T \kappa)\mathbf{1}. \end{aligned}$$

On the other hand, it is a non-convex problem to jointly optimize the control law and upper bound (not shown here). Dynamic programming has an advantage in this respect, when it can be applied.

4.5 State-Space Methods in Papers I–V

Paper I uses a continuous-time operator formulation akin to (4.1), where the source term κ represents switching between different modes. Finite duration jumps are incorporated using the reformulation (4.8). Control signals $u \in \mathbb{R}^n$ are represented as additional states with unspecified dynamics. To evaluate a given controller, a control law is added as an algebraic constraint $u = Kx$. Inequalities akin to (4.11) are used to find upper and lower bounds on the control objective. Search for the tightest bounds is a convex problem in the value function V. A sum-of-squares restriction is used to parametrize V as a polynomial of specified degree, and yields a finite dimensional convex semidefinite program to optimize the bounds.

Paper II uses the exact value function condition (4.12) to to solve for the value function V when the dynamics matrix $A \in \mathbb{R}^n = 0$, and proceeds to characterize the full solution in two cases. It also uses a formulation based on the stationarity condition (4.5) to derive the optimal stationary state density in some cases.

Paper III uses a recurrent state formulation where the discrete-time evolution after an event serves as boundary conditions for the entry into and exit from continuous-time evolution; effectively coupling them by making the κ variables in (4.1) and (4.2) sum up to zero. Analytical expressions for the stationary distribution (by (4.5)) and value function (by (4.12)) are derived in the form of double integrals, and it is shown how control cost and event rate can be derived from either. For the discrete-time case, the stationary distribution is computed by gridding the state space and using the discrete-time dynamics (4.2) as a fixed point iteration.

Paper IV relies mainly on Monte Carlo simulation to evaluate event-based control. There is, however, one example that is solved using dynamic programming as in (4.9), after gridding the state space.

Paper V formulates an approximate system model as a Markov process with one discrete available/busy state for the channel, and one continuous state for one of the plants. In this particular case, the continuous-time stationarity condition (4.5) reduces to a linear space invariant ODE boundary value problem, from which an analytical expression for the stationary state distribution is derived.

5

Event-Based Estimation under Sensing Constraints

This chapter discusses event-based state estimation constrained by sensors that only deliver measurements at certain events, such as with coarse quantization. Optimal solutions to such problems usually require much heavier online computations than in the case of communication constrains discussed previously, which might be explained by the need to keep track of e.g. the whole conditional state distribution, instead of just evaluating the control strategy for a given state. This chapter is based on Chapter 3 in [Henningsson, 2008a].

The practical difficulty to obtain a good state estimate under sensing constraints varies greatly with the regularity of events. If long event-free periods are never expected to occur, a time-varying Kalman filter that uses only the information at events will be able to extract most of the state information from the measurements. The main challenge lies in how to exploit the information contained in the *absence* of events, when this is needed.

Optimal control problems with sensing constraints in the loop generally contain an element of *dual control* [Feldbaum, 1960 1961], where the controller may choose to excite the process to extract more information. Since optimal dual control is known to be very hard, we will restrict our discussion in this chapter to the state estimation problem, as depicted in Figure 5.1.

5.1 Problem Formulation

Let us consider a process with linear dynamics and a white noise disturbance,

$$dx = Axdt + dw, \tag{5.1}$$

where x is the state, w is a Wiener process with incremental variance $\mathrm{E}(dwdw^T) = Rdt$, and A is the dynamics matrix. (We have left out any exogenous inputs for ease of exposition.) Measurements provide set membership information

$$x(t) \in X(t), \qquad \forall t. \tag{5.2}$$

Given the dynamics (5.1) and the stream of measurements (5.2), we want to create an online estimator for the state x.

Events enter the problem if we take the measurement $X(t)$ to be a piecewise continuous function of time (in e.g. the Hausdorff metric), with jumps at *events* $t = t_k$. Assuming that the state $x(t)$ itself is continuous, we can take X at an event t_k as

$$X(t_k) = X^-(t_k) \cap X^+(t_k),$$

where X^- and X^+ are left and right limits, respectively. Given that (5.2) holds at all times, the intersection $X(t_k)$ must be nonempty.

We are interested in estimation problems where the measurements $x(t) \in X(t)$ contain significantly more information at events t_k than otherwise. For instance, with quantized measurements or send-on-delta transmission, $X(t)$ will at most times describe an interval on some state variable, but at events, the value will be given exactly. We note that all X are convex in these examples.

It is not always realistic to monitor the measurement $X(t)$ in continuous time. Another option is to sample $X(t)$ periodically, and use an assumption of small state change between two samples when $X(t)$ changes abruptly.

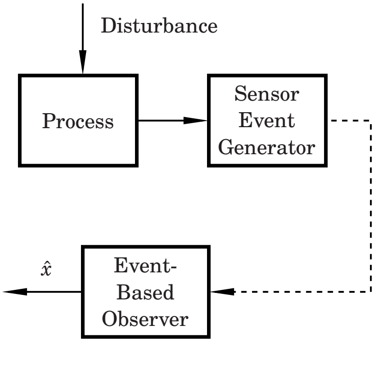

Figure 5.1 State estimation with event-based measurements.

5.2 Optimal Estimation

The optimal solution to the state estimation problem is given by the Bayesian state estimator, described for the case of quantized measurements in [Curry, 1970], and in Paper VI. The state of the estimator is the conditional probability density $f_{X|Y}$ of the plant state x conditioned on the available measurements y, and is generally infinite-dimensional. The estimator's updates are described by partial differential equations in continuous time, or convolutions, affine transformations and pointwise multiplications in discrete time (see Paper VI).

The Kalman filter [Kalman, 1960; Kalman and Bucy, 1961] can be derived as a special case of the Bayesian estimator, where the conditional state density $f_{X|Y}$ is known to be Gaussian, and can thus be summarized by its mean and covariance. Very few other state estimation problems seem to admit such a finite state representation, making an exact implementation untractable.

5.3 Approximate Estimation

A number of different approaches are possible in order to find an approximate solution to sensing constrained state estimation problems, including:

Linear time-invariant estimation. It is often possible to express or relax the measurement condition (5.2) as $||Cx - y|| \leq a$, for some C, a, and y. If a relaxation can be found using constant C and a, then x can be estimated from y using linear time-invariant (LTI) filtering, by trying to keep down the gain from measurement disturbance $\tilde{y} = Cx - y$ to state estimate $\hat{x}$. Though simple, this approach can lead to very conservative results, since it cannot exploit the time-varying nature of the problem.

Worst-case measurement disturbances. The measurement condition $||Cx - y|| \leq a$ can also be used in a time varying worst-case analysis, or a mixed worst-case/stochastic analysis if stochastic disturbances are kept in the model. Estimation with worst-case process and measurement disturbances is treated in [Bertsekas and Rhodes, 1971; Durieu *et al.*, 2001]. Mixed estimation is explored in [Hanebeck and Horn, 2001]. [Morrell and Stirling, 1988] treats the mixed case when only the initial conditions have a worst-case component.

Use the information in events only. Given that events provide linear measurements, ignoring the absence of an event makes the conditional state distribution become Gaussian. The estimator then reduces to a time varying Kalman filter, see [Kalman, 1960; Kalman and Bucy, 1961]. This approach is also taken in [Sandee *et al.*, 2007], where the control law is sampled as a function of distance rather than time.

Grid filtering. The state x can be discretized onto a finite grid, turning the stochastic state process into a Markov chain, another case where the state of the Bayesian estimator is finite dimensional. The estimator's state dimension grows exponentially with that of the process, however, making this approach infeasible except for low-order systems.

Particle filtering. In a particle filter, the conditional state density is approximated by a cloud of point densities, see [Arulampalam *et al.*, 2002]. This approach is perhaps even more generally applicable than grid filtering, and more efficient for higher order system.

JMAP estimation. Instead of trying to keep track of the conditional density $f_{X|Y}$, a simpler estimation problem is to find the most likely state trajectory given the measurements, and use the end point as a state estimate. This is the *joint maximum a posteriori* (JMAP) formulation, see [Cox, 1964]. With log-concave noise densities, the estimation problem becomes a convex optimization problem, see [Schön *et al.*, 2003]. With Gaussian disturbances and quantized measurements, the estimation problem becomes a quadratic program, suitable to use with e.g. moving horizon estimation [Rawlings and Bakshi, 2006].

5.4 Estimation Problems in Papers VI and VII

Paper VI studies the Bayesian estimator for discrete-time linear systems with Gaussian process noise and log-concave measurement likelihoods, a generalization of (5.2) when X is convex, that can include stochastic measurement uncertainty. A number of properties are shown for the Bayesian estimator in this case e.g. that it preserves log-concavity of the conditional density $f_{X|Y}$, which implies among other things that the density is unimodal. A heuristically designed Kalman filter is shown to work well in the example.

Paper VII considers an estimation problem with mixed stochastic and convex set-bounded process and measurement noise. An estimator is proposed with state dimension roughly double that of a time-varying Kalman filter, employing ellipsoidal over-bounding and online convex optimization. The estimator is shown to come quite close to the performance of a grid filter approximation to the Bayesian estimator in a double integrator simulation example.

6

Outlook

Though event-based control has received increasing attention over the last decade, the field is still very far from maturity. This chapter outlines some interesting directions for future work, both in general, and related to the contributions of the thesis.

Some broad topics that deserve further investigation are:

Fundamentals. There remains a need to seek simplicity in order to handle the complexities of event-based control, to understand which aspects of a control problem actually do require to develop new event-based approaches, and to separate the investigation of analysis techniques from architectures for event-based control.

- Which kinds of tools and properties do easily carry over from linear control? Which questions do these tools already answer?
- Which aspects of a problem are intrinsically event-based?
- Is there a core set of concepts to describe such event-based aspects of a given problem?
- How can specific model assumptions be separated from modeling, analysis, and design techniques, in ways that simplify the treatment instead of complicating it?

Finding new problem formulations. The success of classic linear control relies on asking questions that are both relevant and tractable. The best questions in event-based control may not yet have been found.

- What are the control problems where event-based control has the most impact?
- What are the practical problems encountered when trying to use event-based control?
- What relevant problem formulations are easiest to solve? Are there formulations that are simplified by introducing events?

Quantitative comparison. Heuristic elements are probably necessary to address most event-based control design problems. Quantitative comparison between proposed approaches is needed, as well as tools to facilitate the comparison.

Some important aspects that deserve more attention are:

Robust event-based control. In many cases of event-based control, robust stability can likely be resolved using classic linear methods. Still, some interesting questions remain:

- Can event-based control be more robust than periodic control?
- Can a robust event-based controller guarantee a lowered event rate, given that the controlled system is close to the nominal one?

Robustness to mismodeled disturbance intensity. Unlike in linear control, some methods for event-based control depend on the scale of disturbances specified in the model.

- How do event-based controllers compare when the disturbance intensity is not nominal?
- How can a controller be made to adapt to changing ins-tensity?

Control over networks. Some interesting topics that seem to have received little attention are

- event-based control to obviate the need for synchronized clocks.
- bandwidth allocation between event-based control nodes as a resource control problem. Wireless networks are particularly challenging, since throughput decreases with over-utilization.
- event-based control in the face of unacknowledged packet losses.

A promising case to find well-performing and robust event-based controllers is for multiple control loops closed over a shared medium, with discrete-time slots afforded to the sender that asserts the highest priority, such as is possible with the CAN bus.

Path constraints are introduced in Paper I. Some extensions that should be rather straightforward are:

- Making the framework applicable to general polynomial stochastic hybrid systems [Hespanha, 2005].
- Using path constraints to find bounds on induced $\mathcal{L}_2$ gain from disturbance inputs to performance outputs.

Other interesting directions for path constraints include

- to deepen the theoretical foundations.
- to improve the numerics. Monomials are a notoriously ill-conditioned basis. Proper variable scaling is necessary with any basis, either for conditioning or to be able derive useful bounds. Ideas include
  - Progressive variable scaling: Solve a problem with low polynomial order, extract scales of the variables, and use when solving the same problem using a higher polynomial order.
  - Orthogonal polynomials as basis functions, perhaps Hermite polynomials.
  - Other basis functions with better conditioning. They must still be compatible with the operators used in a model.
- to use a basis for the value function that goes to zero for large state norm $||x||$. With a constant value function term given by e.g. a periodic LQ solution, this should allow the value function derived from lower bounding the control objective to be used for controller design.
- to implement local optimization to improve an initial controller with respect to upper bound on cost.

Some ideas for applications to investigate with path constraints include

- multiple identical loops closed over a shared network. Symmetries with respect to the state should allow to use a finite number of basis functions independent of the number of loops, but growing with polynomial order.
- systems that contain inherently event-based dynamics. An example is server systems, with queuing dynamics, actuation to turn on or off machines, etc.

Possible directions to investigate for the sensing constrained problems such as in Papers VI and VII include

- Implement an event-based observer using joint maximum a posteriori (JMAP) estimation and a fast gradient optimization method [Nesterov, 2003].
- Find a simple event-based state estimator and show that it works well to close the loop.

References

Anta, A. and P. Tabuada (2010): “To sample or not to sample: Self-triggered control for nonlinear systems.” *IEEE Transactions on Automatic Control*, **55:9**, pp. 2030–2042.

Antunes, D. J., W. P. M. H. Heemels, J. P. Hespanha, and C. J. Silvestre (2012a): “Scheduling measurements and controls over networks - Part I: Rollout strategies for protocol design.” In *Proceedings of the 2012 American Control Conference*.

Antunes, D. J., W. P. M. H. Heemels, J. P. Hespanha, and C. J. Silvestre (2012b): “Scheduling measurements and controls over networks - Part II: Rollout strategies for simultaneous protocol and controller design.” In *Proceedings of the 2012 American Control Conference*.

Arulampalam, S., S. Maskell, and N. Gordon (2002): “A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking.” *IEEE Transactions on Signal Processing*, **50**, pp. 174–188.

Åström, K. J. (2006): *Introduction to Stochastic Control Theory*. Dover, New York. Reprint. Originally published by Academic Press 1970.

Åström, K. J. (2007): “Event based control.” In *Analysis and Design of Nonlinear Control Systems: In Honor of Alberto Isidori*. Springer Verlag.

Åström, K. J. and B. Bernhardsson (2002): “Comparison of Riemann and Lebesque sampling for first order stochastic systems.” In *Proceedings of the 41st IEEE Conference on Decision and Control*.

Bertsekas, D. and I. Rhodes (1971): “Recursive state estimation for a set-membership description of uncertainty.” *IEEE Transactions on Automatic Control*, **16:2**, pp. 117–128.

Blind, R. and F. Allgöwer (2011a): “Analysis of networked event-based control with a shared communication medium: Part I - pure ALOHA.” In *Proceedings of the 2011 IFAC World Congress*.

Blind, R. and F. Allgöwer (2011b): "Analysis of networked event-based control with a shared communication medium: Part II - slotted ALOHA." In *Proceedings of the 2011 IFAC World Congress*.

Bommannavar, P. and T. Basar (2008): "Optimal control with limited control actions and lossy transmissions." In *Proceedings of the 47th IEEE Conference on Decision and Control*, pp. 2032–2037.

Cogill, R. (2009): "Event-based control using quadratic approximate value functions." In *Proceedings of the 48th IEEE Conference on Decision and Control, combined with the 28th Chinese Control Conference, Shanghai, China*, pp. 5883–5888.

Cogill, R., S. Lall, and J. P. Hespanha (2007): "A constant factor approximation algorithm for optimal estimation subject to communication costs." In *Proceedings of the 2007 American Control Conference*.

Cox, H. (1964): "On the estimation of state variables and parameters for noisy dynamic systems." *IEEE Transactions on Automatic Control*, **9:1**, pp. 5–12.

Curry, R. E. (1970): *Estimation and Control with Quantized Measurements*. Research Monograph No. 60. M.I.T. Press.

Dimarogonas, D. V. and K. H. Johansson (2009): "Event-triggered control for multi-agent systems." In *Proceedings of the 2009 IEEE Conference on Decision and Control*.

Dodds, S. J. (1981): "Adaptive, high precision, satellite attitude control for microprocessor implementation." *Automatica*, **17:4**, pp. 563–573.

Donkers, M. C. F. and W. P. M. H. Heemels (2012): "Output-based event-triggered control with guaranteed $\mathcal{L}_\infty$-gain and improved and decentralized event-triggering." *IEEE Transactions on Automatic Control*, **57:6**, pp. 1362–1376.

Donkers, M. C. F. and W. P. M. H. Heemels (2013): "Model-based periodic event-triggered control for linear systems." *Automatica*. To appear.

Donkers, M. C. F., P. Tabuada, and W. P. M. H. Heemels (2011): "On the minimum attention control problem for linear systems: A linear programming approach." In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 4717–4722.

Durand, S., J. F. Guerrero Castellanos, and N. Marchand (2011): "Simple Lyapunov sampling for event-driven control." In *Proceedings of the 18th IFAC World Congress*.

Durieu, C., E. Walter, and B. Polyak (2001): "Multi-input multi-output ellipsoidal state bounding." *Journal of Optimization Theory and Applications*, **111**, November, pp. 273–303(31).

Feldbaum, A. A. (1960–1961): "Dual control theory. I–IV." *Automation and Remote Control*, **21, 22**, pp. 874–880, 1033–1039, 1–12, 109–121.

Grüne, L., S. Jerg, O. Junge, D. Lehmann, J. Lunze, F. Müller, and M. Post (2010): "Two complementary approaches to event-based control." *Automatisierungstechnik*, **58:4**, pp. 173–182.

Guinaldo, M., D. Lehmann, J. S. Moreno, S. Dormido, and K. H. Johansson (2012): "Distributed event-triggered control with network delays and packet losses." In *Proceedings of the 51st IEEE Conference on Decision and Control*. Maui, Hawaii, USA. To appear.

Hanebeck, U. and J. Horn (2001): "New estimators for mixed stochastic and set theoretic uncertainty models: the general case." *Proceedings of the 2001 American Control Conference*, **6**, pp. 5040–5045 vol.6.

Heemels, W., R. Gorter, A. van Zijl, P. van den Bosch, S. Weiland, W. Hendrix, and M. Vonder (1999): "Asynchronous measurement and control: a case study on motor synchronization." *Control Engineering Practice*, **7:12**, pp. 1467–1482.

Heemels, W., K. Johansson, and P. Tabuada (2012): "An introduction to event-triggered and self-triggered control." In *Proceedings of the 51st IEEE Conference on Decision and Control*. Maui, Hawaii, USA. To appear.

Heemels, W. and J. Sandee (2006): "Practical stability of perturbed event-driven controlled linear systems." In *Proceedings of the 2006 American Control Conference*.

Heemels, W. P. M. H., M. C. F. Donkers, and A. R. Teel (2011): "Periodic event-triggered control based on state feedback." In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 2571–2576.

Heemels, W. P. M. H., J. H. Sandee, and P. P. J. Van Den Bosch (2008): "Analysis of event-driven controllers for linear systems." *International Journal of Control*, **81:4**, pp. 571–590.

Hespanha, J. P. (2005): "Polynomial stochastic hybrid systems." In *Proceedings of the 8th International Conference on Hybrid Systems: Computation and Control*, number 3414, pp. 322–338. Springer-Verlag, Berlin.

Imer, O. and T. Basar (2005): “Optimal estimation with limited measurements.” In *Proceedings of the 44th IEEE Conference on Decision and Control and European Control Conference*, pp. 1029–1034.

Imer, O. and T. Basar (2006a): “Optimal control with limited controls.” In *Proceedings of the 2006 American Control Conference*. Minneapolis, MN.

Imer, O. and T. Basar (2006b): “To measure or to control: optimal control with scheduled measurements and controls.” In *Proceedings of the 2006 American Control Conference*.

Imer, O. C. and T. Basar (2010): “Optimal estimation with limited measurements.” *International Journal of Systems, Control and Communications*, **2:1/2/3**, pp. 5–29.

Kalman, R. E. (1960): “A new approach to linear filtering and prediction problems.” *Transactions of the ASME–Journal of Basic Engineering*, **82**, pp. 35–45.

Kalman, R. E. and R. S. Bucy (1961): “New results in linear filtering and prediction theory.” *Transactions of the ASME - Journal of Basic Engineering*, **83**, pp. 95–107.

Lehmann, D. and J. Lunze (2011a): “Event-based output-feedback control.” In *Proceedings of the 19th Mediterranean Conference on Control and Automation*. Corfu, Greece.

Lehmann, D. and J. Lunze (2011b): “Extension and experimental evaluation of an event-based state-feedback approach.” *Control Engineering Practice*, **19:2**, pp. 101 – 112.

Lehmann, D. and J. Lunze (2012): “Event-based control with communication delays and packet losses.” *International Journal of Control*, **85:5**, pp. 563–577.

Lehmann, D., J. Lunze, and K. Johansson (2012): “Comparison between sampled-data control, deadband control and model-based event-triggered control.” In *Proceedings of the 4th Conference on Analysis and Design of Hybrid Systems*. Eindhoven, Netherlands.

Lemmon, M., T. Chantem, X. S. Hu, and M. Zyskowski (2007): “On self-triggered full-information H-infinity controllers.” In *Proceedings of the 10th International Conference on Hybrid Systems: Computation and Control*.

Li, L. and M. Lemmon (2011): “Performance and average sampling period of sub-optimal triggering event in event triggered state estimation.” In *Proceedings of the 2011 IEEE Conference on Decision and Control*.

Li, L., M. Lemmon, and X. Wang (2010): “Event-triggered state estimation in vector linear processes.” In *Proceedings of the 2010 American Control Conference*, pp. 2138–2143.

Lipsa, G. M. and N. C. Martins (2009): “Optimal state estimation in the presence of communication costs and packet drops.” In *Proceedings of the 47th annual Allerton conference on Communication, control, and computing*, Allerton’09, pp. 160–169.

Lipsa, G. M. and N. C. Martins (2011): “Remote state estimation with communication costs for first-order LTI systems.” *IEEE Transactions on Automatic Control*, **56:9**, pp. 2013–2025.

Lunze, J. and D. Lehmann (2010): “A state-feedback approach to event-based control.” *Automatica*, **46:1**, pp. 211–215.

Marchand, N. (2008): “Stabilization of Lebesgue sampled systems with bounded controls: the chain of integrators case.” In *Proceedings of the 17th IFAC World Congress*.

Mazo, M., J. and P. Tabuada (2011): “Decentralized event-triggered control over wireless sensor/actuator networks.” *IEEE Transactions on Automatic Control*, **56:10**, pp. 2456–2461.

Molin, A. and S. Hirche (2009): “On LQG joint optimal scheduling and control under communication constraints.” In *Proceedings of the 48th IEEE Conference on Decision and Control, combined with the 28th Chinese Control Conference, Shanghai, China*, pp. 5832–5838.

Molin, A. and S. Hirche (2010a): “Optimal event-triggered control under costly observations.” In *Proceedings of the 19th International Symposium on Mathematical Theory of Networks and Systems*.

Molin, A. and S. Hirche (2010b): “Structural characterization of optimal event-based controllers for linear stochastic systems.” In *Proceedings of the 49th IEEE Conference on Decision and Control, Atlanta, Georgia, USA*, pp. 3227–3233.

Molin, A. and S. Hirche (2010c): “Suboptimal event-based control of linear systems over lossy channels.” In *Proceedings of 1st IFAC Workshop on Distributed Estimation and Control in Networked Systems*. Annecy, France.

Molin, A. and S. Hirche (2011a): “Optimal design of decentralized event-triggered controllers for large-scale systems with contention-based communication.” In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference, Orlando, FL, USA*, pp. 4710–4716.

Molin, A. and S. Hirche (2011b): "Order reduction in optimal event-triggered control design for linear stochastic systems." In *Proceedings of the 2011 American Control Conference*, pp. 1642–1649.

Molin, A. and S. Hirche (2012a): "An iterative algorithm for optimal event-triggered estimation." In *Proceedings of the 4th IFAC Conference on analysis and design of hybrid systems, Eindhoven, the Netherlands*.

Molin, A. and S. Hirche (2012b): "On the optimality of certainty equivalence for event-triggered control systems." *IEEE Transactions on Automatic Control*. To appear.

Morrell, D. and W. Stirling (1988): "Set-valued filtering and smoothing." *Twenty-Second Asilomar Conference on Signals, Systems and Computers, 1988.*, pp. 236–240.

Nesterov, Y. (2003): *Introductory Lectures on Convex Optimization: A Basic Course*. Applied Optimization. Springer.

Neugebauer, M. and K. Kabitzsch (2004): "A new protocol for a low power sensor network." In *Proceedings of the 2004 IEEE International Conference on Performance, Computing, and Communications*, pp. 393–399.

Rabi, M. (2006): *Packet Based Inference and Control*. University of Maryland, College Park.

Rabi, M. and J. Baras (2007): "Level-triggered control of a scalar linear system." In *Proceedings of the 2007 Mediterranean Conference on Control Automation*.

Rabi, M., J. Baras, and G. Moustakides (2006): "Efficient sampling for keeping track of an Ornstein-Uhlenbeck process." In *Proceedings of the 14th Mediterranean Conference on Control and Automation*, pp. 1–6.

Rabi, M., K. Johansson, and M. Johansson (2008): "Optimal stopping for event-triggered sensing and actuation." In *Proceedings of the 47th IEEE Conference on Decision and Control*.

Rabi, M. and K. H. Johansson (2009a): "Optimal stopping for updating controls." In *Proceedings of the second International workshop on sequential methods*. UTT, Troyes, France.

Rabi, M. and K. H. Johansson (2009b): "Scheduling packets for event-triggered control." In *Proceedings of the 2009 European Control Conference*.

Rabi, M., G. Moustakides, and J. Baras (2006): "Multiple sampling for estimation on a finite horizon." In *Proceedings of the 45th IEEE Conference on Decision and Control*. USA.

Rabi, M., G. Moustakides, and J. Baras (2012): "Adaptive sampling for linear state estimation." *SIAM Journal on Control and Optimization*, **50:2**, pp. 672–702.

Ramesh, C., H. S, and K. H. Johansson (2009): "LQG and medium access control." In *Proceedings of the 1st IFAC Workshop on Estimation and Control of Networked Systems*.

Ramesh, C., H. Sandberg, L. Bao, and K. Johansson (2011a): "On the dual effect in state-based scheduling of networked control systems." In *Proceedings of the 2011 American Control Conference*, pp. 2216–2221.

Ramesh, C., H. Sandberg, and K. H. Johansson (2011b): "Steady state performance analysis of multiple state-based schedulers with CSMA." In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference, Orlando, Forida*, pp. 4729–4734.

Ramesh, C., H. Sandberg, and K. H. Johansson (2012): "Stability analysis of multiple state-based schedulers with CSMA." In *Proceedings of the 51st IEEE Conference on Decision and Control*. Maui, Hawaii, USA. To appear.

Rantzer, A. (2006): "A separation principle for distributed control." In *Proceedings of the 45th IEEE Conference on Decision and Control*. San Diego, CA.

Rawlings, J. and B. Bakshi (2006): "Particle filtering and moving horizon estimation." *Computers and Chemical Engineering*, **30:10-12**, pp. 1529–1541.

Sandee, J., W. Heemels, S. Hulsenboom, and P. van den Bosch (2007a): "Analysis and experimental validation of a sensor-based event-driven controller." In *Proceedings of the 2007 American Control Conference*, pp. 2867–2874.

Sandee, J. H., W. P. M. H. Heemels, and P. P. J. van den Bosch (2007b): "Case studies in event-driven control." In *Proceedings of the 10th International Conference on Hybrid Systems: Computation and Control*, pp. 762–765.

Schön, T., F. Gustafsson, and A. Hansson (2003): "A note on state estimation as a convex optimization problem." In *Proceedings of the 2003 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 6, pp. VI–61–4.

Seyboth, G., D. V. Dimarogonas, and K. H. Johansson (2011): "Control of multi-agent systems via event-based communication." In *Proceedings of the 18th IFAC World Congress*, pp. 10086–10091. Milan, Italy.

Sijs, J. and M. Lazar (2009): "On event based state estimation." In *Proceedings of 12th International Conference on Hybrid Systems: Computation and Control, San Francisco, CA, USA*, pp. 336–350.

Sijs, J., M. Lazar, and W. P. M. H. Heemels (2010): "On integration of event-based estimation and robust MPC in a feedback loop." In *Proceedings of the 13th ACM International Conference on Hybrid Systems: Computation and Control, Stockholm, Sweden*, pp. 31–40.

Speranzon, A., C. Fischione, B. Johansson, and K. Johansson (2007): "Adaptive distributed estimation over wireless sensor networks with packet losses." In *Proceedings of the 46th IEEE Conference on Decision and Control*, pp. 5472–5477.

Stoecker, C. and J. Lunze (2011): "Event-based control of nonlinear systems: An input-output linearization approach." In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 2541–2546. Orlando, FL, USA.

Tabuada, P. (2007a): "Event-triggered real-time scheduling of stabilizing control tasks." *IEEE Transactions on Automatic Control*, **52:9**, pp. 1680–1685.

Tabuada, P. (2007b): "Event-triggered real-time scheduling of stabilizing control tasks." *IEEE Transactions on Automatic Control*, **52:9**, pp. 1680–1685.

Velasco, M., P. Martí, and E. Bini (2009): "Equilibrium sampling interval sequences for event-driven controllers." In *Proceedings of the 2009 European Control Conference*.

Wan, P. and M. D. Lemmon (2009): "Event-triggered distributed optimization in sensor networks." In *Proceedings of the 8th International Conference on Information Processing in Sensor Networks, San Francisco, California, USA*, pp. 49–60.

Wang, X. and N. Hovakimyan (2010): "Performance prediction in uncertain networked control systems using $\mathcal{L}_1$-adaptation-based distributed event-triggering." In *Proceedings of the 49th IEEE Conference on Decision and Control*, pp. 7570–7575.

Wang, X. and M. Lemmon (2010): "Asymptotic stability in distributed event-triggered networked control systems with delays." In *Proceedings of the 2010 American Control Conference*, pp. 1362–1367.

Wang, X. and M. Lemmon (2011a): "Attentively efficient controllers for event-triggered feedback systems." In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 4698–4703.

Wang, X. and M. Lemmon (2011b): “Event-triggering in distributed networked control systems.” *IEEE Transactions on Automatic Control*, **56:3**, pp. 586–601.

Wang, X. and M. Lemmon (2011c): “On event design in event-triggered feedback systems.” *Automatica*, **47:10**, pp. 2319–2322.

Wang, X. and M. D. Lemmon (2008): “Event design in event-triggered feedback control systems.” In *Proceedings of the 47th IEEE Conference on Decision and Control, Cancún, México*, pp. 2105–2110.

Weimer, J., J. Araujo, and K. H. Johansson (2012): “Distributed event-triggered estimation in networked systems.” In *Proceedings of the 4th IFAC Conference on Analysis and Design of Hybrid Systems*. Eindoven, the Netherlands.

Witsenhausen, H. S. (1968): “A counterexample in stochastic optimum control.” *SIAM Journal on Control*, **6:1**, pp. 131–147.

Xu, Y. and J. Hespanha (2004a): “Optimal communication logics in networked control systems.” In *Proceedings of the 43rd IEEE Conference on Decision and Control*, vol. 4, pp. 3527–3532.

Xu, Y. and J. P. Hespanha (2004b): “Communication logics for networked control systems.” In *Proceedings of the 2004 American Control Conference*, pp. 572–577.

Xu, Y. and J. P. Hespanha (2005): “Estimation under uncontrolled and controlled communications in networked control systems.” In *Proceedings of the 44th Conference on Decision and Control*.

Xu, Y. and J. P. Hespanha (2006): “Communication logic design and analysis for networked control systems.” In *Current trends in nonlinear systems and control*. Birkhäuser, Boston.

Yu, H. and P. J. Antsaklis (2011): “Event-triggered real-time scheduling for stabilization of passive and output feedback passive systems.” In *Proceedings of the 2011 American Control Conference*, pp. 1674–1679.

Paper I

Sporadic Event-Based Control Using Path Constraints and Moments

Toivo Henningsson

Abstract

Control is traditionally applied using periodic sensing and actuation. In some applications, it is beneficial to use instead event based control, to communicate or make a change only when necessary. There are no known general closed form solutions to such event based control problems. We consider stationary event-based control problems with mixed continuous/discrete time dynamics and stochastic disturbances. The system is modelled by a set of path constraints, which are converted into constraints on trajectories' moments up to some order N; upper and lower bounds on the control objective for any system that meets the constraints are derived using sum-of-squares techniques and convex semidefinite programming. Joint optimization of upper bound and controller parameters is non-convex in general; approaches to such controller optimization are investigated, including local optimization using bilinear matrix inequalities. Examples show that the bounds are significantly tighter than earlier results obtained using quadratic value functions.

1. Introduction

Digital control is traditionally carried out using periodic sampling and actuation. Sometimes, however, there is a bottleneck in the control loop. There may be a fixed cost or a minimum time between *events* such as to transmit a state estimate or change a control signal. In *event-based control*, the decision when to generate an event is taken dynamically, rather than to pick a fixed sample rate a priori.

Event-based control can mean many different things. It can be phrased in a stochastic, deterministic, or worst-case setting, with linear or non-linear dynamics, in continuous or discrete time, with the aim to reduce computation, communication or actuation. In a non-stochastic setting, some authors predict the next event time in advance, see, e.g. [Tabuada, 2007; Wang and Lemmon, 2009].

This paper considers systems with linear dynamics and stochastic disturbances, and the objective to reduce communication or actuation. Both continuous time (CT) and discrete time (DT) settings will be considered; in fact, trajectories may switch back and forth between flow (CT) and jump (DT), see Figure 1.

One way to approach the class of problems considered in this paper is to discretize the system into a Markov chain, and then solve the optimal control problem using dynamic programming [Arapostathis *et al.*, 1993]; this is applied to single state plants in [Henningsson *et al.*, 2008]. This method has exponential complexity in the number of state variables. To deal with more than a few states, we will consider instead value functions up to some fixed polynomial degree N, which gives polynomial complexity.

In [Åström and Bernhardsson, 1999], impulse control of a continuous time (CT) integrator plant with a white noise disturbance was considered. It was shown that the mean event frequency can be reduced to a

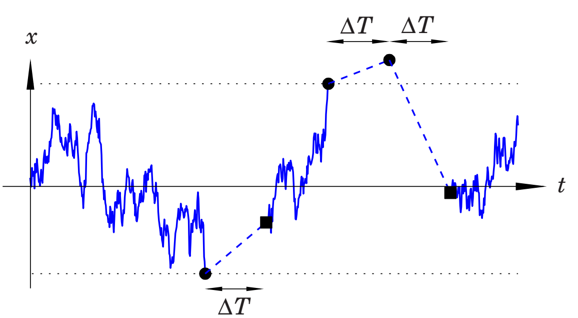

Figure 1. Example of a mixed flow/jump trajectory. When entering a jump (dots), the system jumps to a new state x_+ and time $t_+ = t + \Delta T$ (squares).

third by using a threshold based event triggering strategy instead of periodic events, for the same state variance. However, such a control policy is *aperiodic*; the time between two events may be arbitrarily short, making it hard to implement in practice. Several other authors have also investigated aperiodic CT problems, e.g. [Rabi, 2006], [Rabi *et al.*, 2008], [Hristu-Varsakelis and Kumar, 2002]. To get an implementable control law, some authors, e.g. [Cogill *et al.*, 2007], [Cogill, 2009], [Imer and Basar, 2010], [Sandee *et al.*, 2007] have considered event-based control in discrete time (DT), with a cost term for each sample with an event.

We are interested in the slightly broader class of *sporadic* controllers [Henningsson *et al.*, 2008], with a guaranteed waiting time between any two events. After this period of *inactive state*, the controller may begin to monitor the plant state continuously, or at some sample rate. CT and DT sporadic control is also considered in [Heemels *et al.*, 2008], (where sporadic CT is called non-uniform control) under the objective of ultimate boundedness.

In the last decade, moment relaxations (see, e.g. [Savorgnan *et al.*, 2009]), and their dual, sum-of-squares (SOS) restrictions (see, e.g. [Prajna *et al.*, 2004], [Prajna *et al.*, 2005]), have gained popularity to approximate nonlinear optimal control problems without closed form solutions. Typically, lower bounds on achievable cost are found, which improve as the problem size grows with relaxation order. This paper is an adaptation of such techniques to event-based optimal control problems. By including the controller in the model, we can also find and optimize upper bounds on cost.

One motivating example that can be (approximately) solved with the methods in this paper is the following sporadic control problem: a classic linear quadratic (LQ) problem with the added constraints that 1) the control signal is zero except for *control events*, when it may be a (vector) Dirac impulse, 2) there is a minimum time ΔT between control events. A fixed cost per control event may be added, and a filter on the plant input to shape the control waveform. A jump transition is created by sampling the system for a time ΔT after each control event (see Figure 1), which recasts the sporadic control problem into a mode switching control problem (see Figure 2). The mode without control may be CT or DT (possibly with a time step $\neq \Delta T$).

The paper is outlined as follows: After preliminaries in Section 2, the event-based control problem is formulated in Section 3. Path constraints to model the system are described in Section 4, and combined in Section 5 using convex optimization to show bounds on cost for any system that meets them; these problems are cast as semidefinite programs (SDP:s) in Section 6 to facilitate efficient solution.

For lower bound problems, the degrees of freedom of the controller can

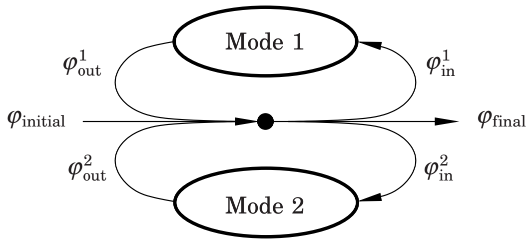

Figure 2. General mode switching model: Each time the trajectory leaves a mode, the controller decides to enter either Mode 1, Mode 2, or terminate.

be left unconstrained; the bound will hold for any controller, including the optimal. For upper bound problems, the controller must be included as a constraint. Section 7 considers approaches to joint optimization of controller parameters and upper bounds, which is in general non-convex. Results are presented in Section 8 and conclusions are given in Section 9.

The source code for the toolbox used to produce the numerical results in this paper is available online [Henningsson, 2011].

2. Preliminaries

For matrices A, B, let $A \succeq B$ denote that $A - B$ is positive semidefinite. Given that $\mathcal{X} = \mathbb{R}^n$: Let $\mathcal{V}(\mathcal{X})$ be a space of test functions (typically polynomials) $V : \mathcal{X} \mapsto \mathbb{R}$. For $f, g \in \mathcal{V}(\mathcal{X})$, let $f \geq g$ denote pointwise inequality: $f(x) - g(x) \geq 0, \forall x \in \mathcal{X}$. Let $\mathcal{V}_+(\mathcal{X}) \subset \mathcal{V}(\mathcal{X})$ be the convex cone of (pointwise) positive functions $V \geq 0, V \in \mathcal{V}(\mathcal{X})$.

Let $\mathcal{V}_N(\mathcal{X})$ be the space of (multivariate) polynomials over $\mathcal{X}$ of degree $\leq N$. Let $\Sigma_N(\mathcal{X}) \subset \mathcal{V}_N(\mathcal{X})$ be the convex cone of sum-of-squares polynomials of degree $\leq N$, i.e. the convex closure of $V_{N/2}(\mathcal{X}) \cdot V_{N/2}(\mathcal{X})$. Given a basis $\psi(x)$ for $V_{N/2}(\mathcal{X})$, it is well known that $\lambda \in \mathcal{V}_N(\mathcal{X})$ is also $\in \Sigma_N(\mathcal{X})$ iff there is a matrix $\Lambda \succeq 0$ such that $\lambda(x) = \psi(x)^T \Lambda \psi(x)$.

3. Problem Formulation

Consider a system that can switch between two *modes* $m \in \mathcal{M} = \{\text{flow}, \text{jump}\}$, with different dynamics for the *state* $x \in \mathcal{X} = \mathbb{R}^{n_x}$. A *trajectory* (or path) consists of *parts* $k \in \mathcal{K} = \{1, 2, 3, \dots\}$, each within one mode $m_k \in \mathcal{M}$. The controller may switch modes freely between parts, see Figure 2. The trajectory begins at time $k = 0, t = t^{\text{in}}_k = t_{\text{initial}}$ and state $x = x^{\text{in}}_k = x_{\text{initial}}$.

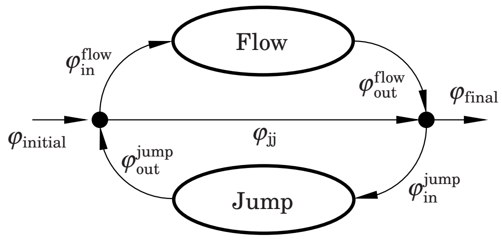

Figure 3. Flow–jump mode switching model: When the controller decides to exit flow, it must take a jump. After a jump, it may decide either way.

Entering the *flow mode* at time $t = t_k^{\text{in}}$ and state $x_k(t_k^{\text{in}}) = x_k^{\text{in}}$, the state x evolves until $t = t_k^{\text{out}}, x_k(t_k^{\text{out}}) = x_k^{\text{out}}$, by the (stochastic differential equation) dynamics

$$dx_k = Ax_k dt + Bu_k^{\text{flow}} dt + dw, \quad \mathrm{E}(dw dw^T) = R dt, \tag{1}$$

where $u_k^{\text{flow}}(t) \in \mathcal{U}_{\text{flow}} = \mathbb{R}^{n_{u_{\text{flow}}}}$ is the control signal, w is a Wiener Process, (independent of the past trajectory), and $R \succeq 0, A, B$ are model matrices of appropriate dimensions. The controller may decide to exit the flow mode at any time.

Entering the *jump mode* at $t = t_k^{\text{in}}$ causes a jump that ends at $t = t_k^{\text{out}} = t_k^{\text{in}} + \Delta T, \Delta T \geq 0$ and state

$$x_k^{\text{out}} = \Phi x_k^{\text{in}} + \Gamma u_k^{\text{jump}} + w_k, \quad w_k \in \mathcal{N}(0, P_{\text{jump}}), \tag{2}$$

where $u_k^{\text{jump}} \in \mathcal{U}_{\text{jump}} = \mathbb{R}^{n_{u_{\text{jump}}}}$ is the control signal, the Gaussian disturbance w_k is independent of the past trajectory, and $P_{\text{jump}} \succeq 0, \Phi, \Gamma$ are model matrices of appropriate dimensions. The jump time ΔT is also a model parameter.

Remark 1

For brevity, we describe only the case with one flow and one jump mode. The switching model of Figure 3 is appropriate in this case, since it disallows consecutive flow parts; we will still use Figure 2 in calculations for brevity. The methods in this paper apply also in the case of two jump modes, possibly with different time steps ΔT_i. □

The expected cost over trajectories j_{acc} is a sum of integrals over each flow interval and a term for each jump:

$$j_{\text{acc}} = \mathrm{E}\left[\sum_{k \in \mathcal{K}_{\text{flow}}} \int_{T_k} c_{\text{flow}}(z_k(t))\, dt + \sum_{k \in \mathcal{K}_{\text{jump}}} c_{\text{jump}}(z_k)\right], \tag{3}$$

where the index sets $\mathcal{K}_m$ and part intervals T_k are given by

$$\mathcal{K}_m = \{k \in \mathcal{K}; m_k = m\}, \qquad T_k = [t_k^{\text{in}}, t_k^{\text{out}}],$$

the *extended state* z in flow and jump respectively by

$$z_k(t) = \begin{pmatrix} x_k(t) \\ u_k^{\text{flow}}(t) \end{pmatrix} \in \mathcal{Z}_{\text{flow}}, \quad z_k = \begin{pmatrix} x_k^{\text{in}} \\ u_k^{\text{jump}} \end{pmatrix} \in \mathcal{Z}_{\text{jump}},$$

$\mathcal{Z}_m = \mathcal{X} \times \mathcal{U}_m$, and the cost functions $c_m \in \mathcal{V}_+(\mathcal{Z}_m)$.

REMARK 2
The function $c_{\text{flow}}(z)$ is the cost *per time unit* in flow mode, while $c_{\text{jump}}(z)$ is the cost *per jump*. □

The controller consists of two parts:

- A switching law $\theta(x)$ to choose mode $m = \text{flow}$ when $\theta(x) \geq 0$, and mode $m = \text{jump}$ otherwise.
- Modal control laws $u_m = f_m(x), m \in \mathcal{M}$.

The control objective is to minimize the *average cost*

$$\mathcal{J} = \bar{\mathcal{R}}(j_{\text{acc}}) = \limsup_{t_{\text{spent}} \to \infty} \frac{1}{t_{\text{spent}}} j_{\text{acc}}, \tag{4}$$

where the trajectory duration t_{spent} is given by

$$t_{\text{spent}} = t_{\text{final}} - t_{\text{initial}} = \sum_{k \in \mathcal{K}} t_k^{\text{out}} - t_k^{\text{in}}. \tag{5}$$

4. Path Constraints

We will now list a number of path constraints to model the considered system. In order to show bounds on path integrals such as the cost (3) in the next section, nonnegative path integrals are derived from the constraints. We first introduce a compact notation for path integrals using measures.

4.1 Path Measures

Define the *occupation measure* μ and *jump event measure* φ, with arguments $f \in \mathcal{V}(Z_{\text{flow}}), f \in \mathcal{V}(\mathcal{Z}_{\text{jump}})$ respectively:

$$\mu(f) = \mathrm{E} \sum_{k \in \mathcal{K}_{\text{flow}}} \int_{T_k} f\big(z_k(t)\big) dt, \quad \varphi(f) = \mathrm{E} \sum_{k \in \mathcal{K}_{\text{jump}}} f(z_k).$$

Given a function $f(z)$ of the extended state $z = (x, u^m)$, $\mu(f)$ can be thought of as an accumulator that integrates $f(z)dt$ along the parts of the trajectory in flow, and $\varphi(f)$ as one that adds up $f(z)$ for each jump.

Using μ and φ, the accumulated cost (3) and trajectory duration (5) can be expressed more compactly as

$$j_{\text{acc}} \qquad = \mu(c_{\text{flow}}) + \varphi(c_{\text{jump}}), \tag{6}$$

$$t_{\text{spent}} = \sum_{k \in \mathcal{K}_{\text{flow}}} \int_{T_k} dt + \sum_{k \in \mathcal{K}_{\text{jump}}} \Delta T \qquad = \mu(1) + \varphi(\Delta T), \tag{7}$$

where 1 in $\mu(1)$ means the constant function $f(z) = 1$, and in the same way for $\varphi(\Delta T)$.

To describe mode switching such as in Figs. 2 and 3, we define, for the initiation and termination events, measures

$$\varphi_{\text{initial}}(f) = \mathrm{E}\, f(x_{\text{initial}}), \qquad \varphi_{\text{final}}(f) = \mathrm{E}\, f(x_{\text{final}}),$$

and, accumulating mode entry and exit events, measures

$$\varphi^m_{\text{dir}}(f) = \mathrm{E} \sum_{k \in \mathcal{K}_m} f(x_k^{\text{dir}}), \qquad m \in \mathcal{M}, \text{dir} \in \{\text{in}, \text{out}\}.$$

Note that the jump event measure φ and jump entry measure $\varphi_{\text{in}}^{\text{jump}}$ are not the same, since φ is defined over the extended state z_{jump}, and $\varphi_{\text{in}}^{\text{jump}}$ over the state x only. However, they coincide for u_{jump}-independent test functions:

$$\varphi(V) = \varphi_{\text{in}}^{\text{jump}}(V), \quad \forall V \in \mathcal{V}(\mathcal{X}). \tag{8}$$

Having defined the path measures, we will now use them to formulate path constraints and nonnegative path integrals.

4.2 Pointwise Path Constraints

The simplest form of path constraints express feasible regions of the (extended) state space. (Such algebraic equations can be used for differential-algebraic equation (DAE) systems modelling.) Consider the constraint that $f(z_{\text{flow}}) = 0$ when the trajectory is in flow mode, for some given function $f(z)$. Then also $f(z_{\text{flow}})V(z_{\text{flow}}) = 0$ for any function $V \in \mathcal{V}(\mathcal{Z}_{\text{flow}})$, as is the path integral

$$\mu(fV) = 0, \qquad \forall V \in \mathcal{V}(\mathcal{Z}_{\text{flow}}).$$

The same can be done for event measures, e.g. $\varphi(fV) = 0, \forall V \in \mathcal{V}(\mathcal{Z}_{\text{jump}})$ if $f(z_{\text{jump}}) = 0$ for all jumps.

Now consider the inequality constraint that $f(z_{\text{flow}}) \geq 0$ when in flow mode. Then also $f(z_{\text{flow}})\lambda(z_{\text{flow}}) \geq 0$ for any nonnegative function λ, as is the path integral

$$\mu(f\lambda) \geq 0, \qquad \forall \lambda \in \mathcal{V}_+(\mathcal{Z}_{\text{flow}}).$$

The constraint $f(z) = 1 \geq 0$ apparently holds in any mode, and will be used since it establishes positivity of the path measures.

4.3 Control Laws

Control laws can be expressed as path constraints; deterministic ones usually as pointwise ones. Examples:

- A switching law such that $\theta(x) \geq 0$ in flow and $\theta(x) \leq 0$ in jump.
- A control law $u_{\text{jump}} = f_{\text{jump}}(x)$ is equivalent to the constraint that $g(z_{\text{jump}}) = u_{\text{jump}} - f_{\text{jump}}(x) = 0$ in jumps.
- A random switching law, causing Poisson jumps in flow with a state-dependent intensity such that $n_{\text{jump}}(x)$ jumps are expected per $t_{\text{flow}}(x)$ time in flow, where $n_{\text{jump}}, t_{\text{flow}} \in \mathcal{V}_+(\mathcal{X})$. Then

  $$\mu(\theta n_{\text{jump}}) - \varphi(\theta t_{\text{flow}}) = 0, \qquad \forall \theta \in \mathcal{V}(\mathcal{X}). \tag{9}$$

  This is not a pointwise constraint since the control law is random, but it holds in expectation, which is what we need.

4.4 Dynamics Constraints

Dynamics constraints express how the trajectory may evolve from one instant to another.

Mode switching. The mode switching dynamics of the model in Figure 2 are contained in the center point. Since each trajectory initiation and mode exit event is paired with exactly one termination or mode entry event, with the state x preserved across transitions, the switching constraint

$$\varphi_{\text{initial}} + \varphi_{\text{out}}^{\text{flow}} + \varphi_{\text{out}}^{\text{jump}} - (\varphi_{\text{final}} + \varphi_{\text{in}}^{\text{flow}} + \varphi_{\text{in}}^{\text{jump}}) = 0 \tag{10}$$

holds, where the argument $V \in \mathcal{V}(\mathcal{X})$ to each measure has been suppressed for brevity. For the mode switching dynamics of Figure 3 we have two switching points; they are modelled in the same way by pairing inflow and outflow,

$$\begin{aligned} \varphi_{\text{initial}} + \varphi_{\text{out}}^{\text{jump}} - (\varphi_{\text{in}}^{\text{flow}} + \varphi_{\text{jj}}) = 0, \\ \varphi_{\text{out}}^{\text{flow}} + \varphi_{\text{jj}} - (\varphi_{\text{in}}^{\text{jump}} + \varphi_{\text{final}}) = 0, \end{aligned} \tag{11}$$

again with the common argument $V \in \mathcal{V}(\mathcal{X})$ in either equation suppressed. We see that the sum of these two equations is (10), thus (11) is a stronger constraint than (10).

Flow dynamics. Consider the flow dynamics (1). Given a (twice differentiable) function $V \in \mathcal{V}(\mathcal{X})$, the expected change in $V(x)$ by the dynamics, conditioned on the extended state z, is (using Itō's Lemma)

$$\begin{aligned} \mathrm{E}(dV|z) &= \mathrm{E}(dx^T)\nabla V(x) + \frac{1}{2}\mathrm{tr}\Big(\mathrm{E}(dx dx^T)\nabla^2 V(x)\Big) \\ &= \Big((Ax+Bu)^T \nabla V(x) + \frac{1}{2}\mathrm{tr}\big(R\nabla^2 V(x)\big)\Big)dt \\ &= (\mathcal{A}^*_{\text{flow}} V)(z)dt; \end{aligned}$$

this defines the *backwards flow dynamics operator* $\mathcal{A}^*_{\text{flow}}$, a Kolmogorov backwards operator. Equating the expectations of the left and right hand sides over the time spent in flow gives the *flow dynamics constraint*

$$\begin{aligned} 0 = \mathrm{E}\sum_{k\in\mathcal{K}_{\text{flow}}}\int_{T_k}\Big((\mathcal{A}^* V)(z)dt - dV\Big) &= \mu(\mathcal{A}^*_{\text{flow}} V) - \Big[V\Big]_{x_k^{\text{in}}}^{x_k^{\text{out}}} \\ = \mu(\mathcal{A}^*_{\text{flow}} V) + \varphi_{\text{in}}^{\text{flow}}(V) - \varphi_{\text{out}}^{\text{flow}}(V), &\qquad \forall V \in \mathcal{V}(\mathcal{X}). \end{aligned} \tag{12}$$

Jump dynamics. Consider the jump dynamics (2). Given a function $V \in \mathcal{V}(\mathcal{X})$, the expected value of $V(x^{\text{out}})$ after a jump, conditioned on $z = (x, u)$ before the jump, is

$$\begin{aligned} \mathrm{E}\Big(V(x^{\text{out}})\Big| z\Big) &= \mathrm{E}\Big(V(\Phi x + \Gamma u + w)\Big| z\Big) \\ &= (\phi * V)(\Phi x + \Gamma u), \\ &= (\mathcal{H}^* V)(z), \end{aligned}$$

where the probability density ϕ is Gaussian $\sim \mathcal{N}(0, P_{\text{jump}})$; this defines the *backwards single jump operator* $\mathcal{H}^*$. Summing over all events gives the *jump dynamics constraint*

$$\begin{aligned} &\mathrm{E}\sum_{k\in\mathcal{K}_{\text{jump}}} V(x_k^{\text{out}}) = \mathrm{E}\sum_{k\in\mathcal{K}_{\text{jump}}}(\mathcal{H}^* V)(z_k) \\ &\implies \varphi_{\text{out}}^{\text{jump}}(V) = \varphi(\mathcal{H}^* V), \qquad \forall V \in \mathcal{V}(\mathcal{X}). \end{aligned} \tag{13}$$

5. Bounds on Cost by Convex Optimization

To show bounds $\underline{J} \leq J \leq \bar{J}$ on the average cost (4) of a system, we will show positivity of path integrals such as $l = j_{\text{acc}} - \underline{J} t_{\text{spent}}$ and $l = \bar{J} t_{\text{spent}} - j_{\text{acc}}$, by expressing them as a sum of nonnegative path integrals. In practice, it is sufficient to show that

$$l + \varphi_{\text{initial}}(V) - \varphi_{\text{final}}(V) \geq 0, \tag{14}$$

for some *value function* $V \in \mathcal{V}(\mathcal{X})$ such that $\varphi_{\text{final}}(V)$ is uniformly bounded from below as $t_{\text{spent}} \to \infty$. This boundedness can be established in many ways:

- For a lower bound, it may be sufficient that the bound holds for solutions with bounded moments of x_{final}; then $\varphi_{\text{final}}(V)$ will be bounded as well, for polynomial V.
- x_{final} will have bounded moments if the flow region $\{x \in \mathcal{X}; \theta(x) \geq 0\}$ is bounded and the jump dynamics (2) are exponentially stable.
- $\varphi_{\text{final}}(V)$ is uniformly bounded from below if V is.

5.1 Lower Bound

To show the lower bound $\underline{J} \leq J$, we want to show that

$$j_{\text{acc}} + \varphi_{\text{final}}(V) - \varphi_{\text{initial}}(V) \geq \underline{J} t_{\text{spent}}. \tag{15}$$

Note that the sign of V has been chosen opposite from (14). Using first (10), and then (8), (12) and (13), we see that

$$\begin{aligned} &\varphi_{\text{final}}(V) - \varphi_{\text{initial}}(V) \\ =&\varphi^{\text{flow}}_{\text{out}}(V) - \varphi^{\text{flow}}_{\text{in}}(V) + \varphi^{\text{jump}}_{\text{out}}(V) - \varphi^{\text{jump}}_{\text{in}}(V) \\ =&\mu(\mathcal{A}^*_{\text{flow}} V) + \varphi(\mathcal{H}^* V) - \varphi(V) \end{aligned} \tag{16}$$

The inequality (15) is then implied by

$$\begin{aligned} &j_{\text{acc}} + \mu(\mathcal{A}^*_{\text{flow}} V) + \varphi(\mathcal{H}^* V) - \varphi(V) \\ =&\underline{J} t_{\text{spent}} + \mu(\lambda_{\text{flow}}) + \varphi(\lambda_{\text{jump}}) \geq \underline{J} t_{\text{spent}}, \quad \lambda_m \in \mathcal{V}_+(\mathcal{Z}_m), \end{aligned}$$

where we have used (16) and $\mu, \varphi \geq 0$. Collecting terms inside μ and φ, this condition is in turn implied by

$$\begin{aligned} c_{\text{flow}} + \mathcal{A}^*_{\text{flow}} V &= \underline{J} + \lambda_{\text{flow}}, \\ c_{\text{jump}} + \mathcal{H}^* V - V &= \underline{J} \Delta T + \lambda_{\text{jump}}, \end{aligned} \tag{17}$$

for some $\lambda_{\text{flow}} \in \mathcal{V}_+(\mathcal{Z}_{\text{flow}}), \lambda_{\text{jump}} \in \mathcal{V}_+(\mathcal{Z}_{\text{jump}})$.

5.2 Lower Bound with Controller

To add a switching law such that $\theta(x) \geq 0$ in flow, and $-\theta(x) \geq 0$ in jump, we use

$$\mu(\theta\nu_{\text{flow}}) - \varphi(\theta\nu_{\text{jump}}) \geq 0, \quad \forall \nu_m \in \mathcal{V}_+(\mathcal{Z}_m). \tag{18}$$

The control law $u_{\text{jump}} = f_{\text{jump}}(x)$ is incorporated by adding

$$\varphi(gW) = 0, \quad \forall W \in \mathcal{V}(\mathcal{Z}_{\text{jump}}),$$

to the left hand side of (15), where $g(z_{\text{jump}}) = u_{\text{jump}} - f_{\text{jump}}(x)$. With these control laws, (17) is strengthened into

$$\begin{aligned} c_{\text{flow}} + \mathcal{A}^*_{\text{flow}} V &= \underline{J} + \lambda_{\text{flow}} + \theta\nu_{\text{flow}}, \\ c_{\text{jump}} + \mathcal{H}^* V - V + gW &= \underline{J}\Delta T + \lambda_{\text{jump}} - \theta\nu_{\text{jump}}. \end{aligned} \tag{19}$$

5.3 Upper Bound with Controller

To show the upper bound $J \leq \bar{J}$, we want to show that

$$j_{\text{acc}} + \varphi_{\text{final}}(V) - \varphi_{\text{initial}}(V) \leq \bar{J} t_{\text{spent}}.$$

We proceed as before, but now all inequality terms have to be introduced with opposite sign. With controller constraints, the conditions (19) are turned into

$$\begin{aligned} c_{\text{flow}} + \mathcal{A}^*_{\text{flow}} V &= \bar{J} - \lambda_{\text{flow}} - \theta\nu_{\text{flow}}, \\ c_{\text{jump}} + \mathcal{H}^* V - V + gW &= \bar{J}\Delta T - \lambda_{\text{jump}} + \theta\nu_{\text{jump}}. \end{aligned} \tag{20}$$

We see that the bound conditions above are convex, since they are linear with convex constraints on $\{\lambda_m\}, \{\nu_m\}$. Thus maximization of $\underline{J}$ subject to (17) or (19) is a convex problem, as is minimization of $\bar{J}$ subject to (20).

6. Practical Optimization

To get problems that can be solved by a convex programming solver, we must choose some finite basis for the test functions $V, \{\lambda_m\}, \{\nu_m\}$ and W. We will use polynomials up to some degree N of trajectory moments. A sum-of squares restriction yields semidefinite programs (SDP:s).

We let the terms in (17), (19), and (20) be polynomials of degree $\leq N$. Since it is in general hard to determine the global positivity of a polynomial, we use $\Sigma_N \subset \mathcal{V}_{N,+}$ to assure positivity; this can be expressed as a linear matrix inequality (LMI). The optimal bound can only improve with increasing N, as a solution to the bounds with lower N is still valid with higher N.

Making sure that no term in (17), (19), and (20) has higher degree than N, we can optimize over $\underline{J} \in \mathbb{R}$ or $\bar{J} \in \mathbb{R}$, and

$$V \in \mathcal{V}_N(\mathcal{X}), \qquad \lambda_m \in \Sigma_N(\mathcal{Z}_m),$$
$$\nu_m \in \Sigma_{N-\deg\theta}(\mathcal{Z}_m), \quad W \in \mathcal{V}_{N-\deg g}(Z_{\text{jump}}).$$

The conditions (20) still give an SDP if we fix ν_m and W, and include instead as optimization variables

$$\theta \in \mathcal{V}_{N-\max_{m\in\mathcal{M}} \deg \nu_m}(\mathcal{X}), \quad g \in \mathbb{G} \subseteq \mathcal{V}_{N-\deg W}(\mathcal{Z}),$$

where the space $\mathbb{G}$ is chosen to give a desirable form for the u_{jump} controller, e.g. linear feedback.

7. Controller Optimization

Now that we can model a system and derive upper and lower bounds $\underline{J} \leq J \leq \bar{J}$ on the average cost J, how can we optimize for good controllers? We would like to prescribe a form for the switching law and modal controllers such as $\theta \in \mathcal{V}_{N_\theta}(\mathcal{X}), \{f_m \in \mathcal{V}_{n_f}(\mathcal{Z}_m)\}_{m\in\mathcal{M}}$, and then find the controller parameters that give the lowest cost.

Since the actual cost J is unknown, we have to content with minimizing an upper bound $\bar{J}$ instead. Unfortunately, joint optimization of upper bound and controller is generally non-convex because of the product terms between controller parameters and dual variables that appear in controller constraints, such as $\theta\nu_{\text{flow}}$ in (20).

These product terms make the controller optimization into a bilinear matrix inequality (BMI) problem; we can still optimize locally given an initial guess. The formulation also allows various structural constraints on the controller such as limited polynomial degrees of θ and $\{f_m\}$, or sparsity constraints, e.g. limiting the set of states that a control signal may depend on.

The controller optimization problem becomes convex if we fix enough decision variables so that no product terms with free variables remain. It is then possible to do global optimization by gridding over remaining variables. By making the problem simple, with low relaxation order N and few constraints, few parameters have to be scanned.

We next give some results relating tightness of the upper bound $\bar{J}$ and problem complexity, and consider especially the case when global optimization can be done by scanning over a single real parameter.

7.1 Mixing Controllers

Consider a deterministic switching controller modelled by

$$\theta(x) \geq 0, \quad \text{in flow}, \quad -\theta(x) \geq 0, \quad \text{in jump}, \tag{21}$$

and a controller stochastically mixing time in flow : jump as $t_{\text{flow}}(x)$: $n_{\text{jump}}(x)$. By section 4.3, the positive path integral given by the former is

$$\mu(\theta \nu_{\text{flow}}) - \varphi(\theta \nu_{\text{jump}}) \geq 0, \qquad \forall \{\nu_m \in \mathcal{V}_+(Z_m)\}_{m \in \mathcal{M}}.$$

This is exactly the same term as (9), if we identify $n_{\text{jump}} = \nu_{\text{flow}}, t_{\text{flow}} = \nu_{\text{jump}}$. The bound derived from the deterministic switching constraint (21) can thus be achieved by a stochastically mixing controller with $n_{\text{jump}} = \nu_{\text{flow}}, t_{\text{flow}} = \nu_{\text{jump}}$! Since we expect the optimal switching law to be deterministic, this gives a hint of how tight the upper bound can be as function of the polynomial order $\deg \nu_m \leq N - \deg \theta$.

The result does not hold in general if we introduce more constraints for the deterministic switching law, such as

$$\theta \varphi_{\text{out}}^{\text{flow}} = 0, \quad \theta \varphi_{\text{in}}^{\text{flow}} \geq 0, \quad \theta \varphi_{\text{jj}} \leq 0,$$

where $\theta \varphi_{\text{out}}^{\text{flow}} = 0$ holds only in the mixed flow/jump setting. These tighter constraints have been used to produce the upper bounds in the results, except for when the equivalence to random switching has been exploited.

7.2 Single Parameter Sweep: Poisson Controller

We will now describe a case when global optimization can be performed by scanning over a single real variable. Consider the upper bound problem with constraints (20), no modal control law u_m (i.e. $gW = 0$), $N = 2$ and quadratic threshold $\theta \in \mathcal{V}_2(\mathcal{X})$ to be optimized. Since $\deg \nu_m \leq N - \deg_\theta = 0$, the polynomials ν_m are constants, e.g. $\nu_m \in \mathbb{R}$. The problem can thus be solved globally by sweeping the ratio $\nu_{\text{flow}} : \nu_{\text{jump}}$ (a common scaling can be accommodated in θ). This is the procedure outlined in [Cogill, 2009] for the case of two jump modes with the same ΔT.

Since $\deg \nu_m = 0$, the upper bound $\bar{J}$ optimized in this formulation can be achieved by a *Poisson controller*; a random switching controller with state independent switching ratio! Still, the derived threshold θ may realize a better cost than the Poisson controller, and can be used as an initial guess for local optimization.

[Cogill, 2009] considers also the case with modal control law $u_{\text{jump}} = -Kx$. This can be accommodated in our formulation by solving the lower bound problem with Poisson switching constraint, since the solution turns out to be exact for $N = 2$ in this case.

8. Results

Consider an integrator process (with state $x \in \mathcal{X} = \mathbb{R}$)

$$dx = udt + dw, \quad \mathrm{E}(dw^2) = dt, \tag{22}$$

where w is a Wiener Process. The control input u is a train of Dirac pulses with minimum time between them $\Delta T = 1$,

$$u(t) = \sum_{i=1}^{n_{\text{events}}} u_i \delta(t - t_i), \quad t_{i+1} - t_i \geq \Delta T.$$

We let $u_i = -x(t_i - 0)$ to immediately reset the state at any control event t_i. The cost function is

$$j_{\text{acc}} = \int_T x(t)^2 dt + \rho n_{\text{events}},$$

where T is the interval of time spent in the system. We want to find an event triggering strategy to minimize $J = \bar{\mathcal{R}}(j_{\text{acc}})$, the average cost as $t_{\text{spent}} = |T| \to \infty$.

To achieve the minimum inter-event time ΔT, the jump mode is constructed as an immediate reset to $x = 0$, followed by the dynamics (22) sampled for time ΔT. The flow mode is just (22) with control input $u = 0$.

Figure 4 shows the optimal average cost J as a function of event cost ρ (calculated in [Henningsson *et al.*, 2008] for this problem), the cost of periodic control with optimal period $h \geq \Delta T$, and lower and upper bounds, which fit quite tightly around the optimum. The upper bound $\bar{J}_{N=4}$ was found by BMI optimization using the solver PENBMI [Kočvara and Stingl, 2006]. The curve $\bar{J}_{N=6}, \theta_{\text{BMI}_4}$, calculated with the same thresholds, show that they are in fact almost optimal. The cost of optimal Poisson Sampling lies far above the other bounds, almost coinciding with the cost of periodic control. (In fact, they both choose periodic sampling with $h = \Delta T$ when $\rho \leq 0.5$) The upper bound $\bar{J}_{N=6}, \theta_{\text{Poisson}}$ shows that the thresholds from Poisson control are considerably better than the bound.

Now consider a double integrator process (with state $x \in \mathcal{X} = \mathbb{R}^2$)

$$dx_1 = x_2 dt, \quad dx_2 = udt + dw, \quad \mathrm{E}(dw^2) = dt,$$

with immediate reset to $x = 0$ at events, minimum time between them ΔT, and the cost function

$$j_{\text{acc}} = \int_T x_1(t)^2 dt + \rho n_{\text{events}}.$$

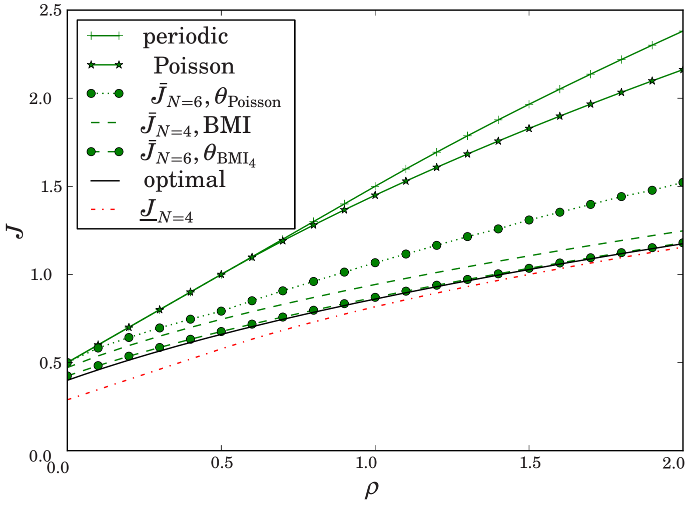

Figure 4. Cost J as a function of event cost ρ for the integrator with $\Delta T = 1$.

Figure 5 shows upper and lower bounds for the cost J as a function of event cost ρ. All upper bounds were found using thresholds from Poisson control. We see that Poisson control and periodic control are comparable, but that the Poisson thresholds perform distinctly better. Still, the gap between upper and lower bounds suggests that there is room to realize a lower cost with better thresholds.

9. Conclusions and Future Work

We have modelled a broad class of event based optimal control problems using path constraints, and shown how to derive interval bounds on the control objective from these using convex semidefinite programming. Joint optimization of upper bound and controller parameters is non-convex in general; approaches to it using global and local optimization have been investigated. The examples show that the bounds are significantly tighter than previous results using quadratic value functions; they also clearly demonstrate that event-based control is superior to periodic control in the examples.

Interesting directions for future work include further case studies and extension to other kinds of stochastic hybrid control problems, improved controller optimization and numerical conditioning.

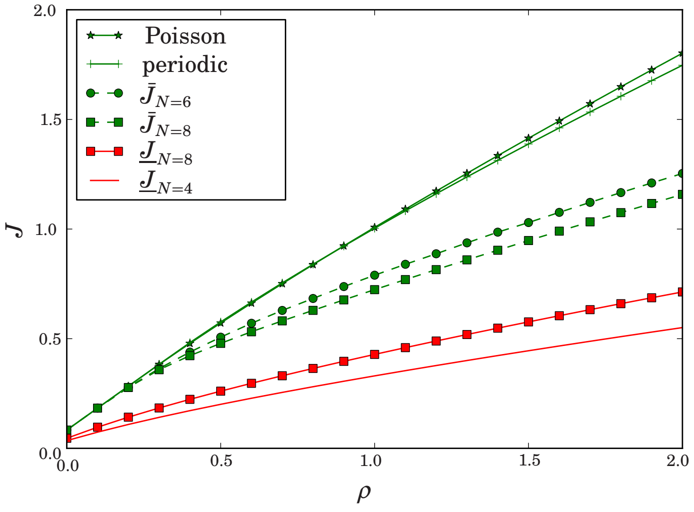

Figure 5. Cost J as a function of event cost ρ for the double integrator with $\Delta T = 1$. Thresholds for upper bounds are from Poisson control.

Acknowledgment

The author would like to thank A. Cervin, B. Bernhardsson and A. Ghulchak for fruitful discussions on the topic of event-based control, C. Savorgnan for a helpful introduction to polynomial and moment approximations, and M. Kočvara and M. Stingl for letting me use their solver PENBMI.

References

Arapostathis, A., V. S. Borkar, E. Fernández-Gaucherand, M. K. Ghosh, and S. I. Marcus (1993): "Discrete-time controlled Markov processes with average cost criterion: a survey." *SIAM Journal on Control and Optimization*, **31**, pp. 282–344.

Åström, K. J. and B. Bernhardsson (1999): "Comparison of periodic and event based sampling for first-order stochastic systems." In *Preprints of the 14th World Congress of IFAC*, Beijing, P.R. China.

Cogill, R. (2009): "Event-based control using quadratic approximate value functions." In *Proceedings of the 48th IEEE Conference on Decision and Control*, Shanghai, P.R. China.

Cogill, R., S. Lall, and J. P. Hespanha (2007): "A constant factor approximation algorithm for optimal estimation subject to communication costs." In *Proceedings of the 2007 American Control Conference*, New York, New York.

Heemels, W. P. M. H., J. H. Sandee, and P. P. J. V. D. Bosch (2008): "Analysis of event-driven controllers for linear systems." *International Journal of Control*, **81**, pp. 571–590.

Henningsson, T. (2011): "EventMually toolbox." `http://www.control.lth.se/Research/tools/eventmually.html`.

Henningsson, T., E. Johannesson, and A. Cervin (2008): "Sporadic event-based control of first-order linear stochastic systems." *Automatica*, **44:11**, pp. 2890–2895.

Hristu-Varsakelis, D. and P. Kumar (2002): "Interrupt-based feedback control over a shared communication medium." In *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vegas, Nevada.

Imer, O. C. and T. Basar (2010): "Optimal estimation with limited measurements." *International Journal of Systems, Control and Communications*, **2**, pp. 5–29.

Kočvara, M. and M. Stingl (2006): "PENBMI." `http://www.penopt.com/penbmi.html`.

Prajna, S., A. Papachristodoulou, P. Seiler, and P. A. Parrilo (2005): "Sostools and its control applications." In Henrion and Garulli, Eds., *Positive Polynomials in Control*, vol. 312 of *Lecture Notes in Control and Information Sciences*, pp. 580–580. Springer Berlin/Heidelberg.

Prajna, S., P. Parrilo, and A. Rantzer (2004): "Nonlinear control synthesis by convex optimization." *IEEE Transactions on Automatic Control*, **49:2**, pp. 310–314.

Rabi, M. (2006): *Packet based Inference and Control*. PhD thesis, Institute for Systems Research, University of Maryland. `http://hdl.handle.net/1903/6612`.

Rabi, M., K. Johansson, and M. Johansson (2008): "Optimal stopping for event-triggered sensing and actuation." In *Proceedings of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico.

Sandee, J. H., W. P. M. H. Heemels, and P. P. J. van den Bosch (2007): "Case studies in event-driven control." In *Proceedings of the 10th International Conference on Hybrid Systems: Computation and Control*, Pisa, Italy.

Savorgnan, C., J. B. Lasserre, and M. Diehl (2009): "Discrete-time stochastic optimal control via occupation measures and moment relaxations." In *Proceedings of the 48th IEEE Conference on Decision and Control*, Shanghai, P.R. China.

Tabuada, P. (2007): "Event-triggered real-time scheduling of stabilizing control tasks." *IEEE Transactions on Automatic Control*, **52:9**, pp. 1680–1685.

Wang, X. and M. Lemmon (2009): "Self-triggered feedback control systems with finite-gain L_2 stability." *IEEE Transactions on Automatic Control*, **54:3**, pp. 452–467.

Paper II

A Riccati-like Equation for Finding Optimal Elliptical Triggering Rules

Bo Bernhardsson   Toivo Henningsson

Abstract

The paper gives explicit solutions to two event-based optimal control problems. Though being "toy examples" we believe they are interesting and can serve as helpful test examples, e.g. for fine tuning numerical optimization based approaches. The setup is a traditional linear quadratic control problem with white Gaussian noise (LQG) where there is an added extra penalty on actuation events. We use the results in [Henningsson, 2011] where the optimal state distribution is determined by a convex primal optimization problem and the optimal value function from a dual problem. We thereby derive optimal control laws for two problem setups: One problem is a multi-dimensional integrator system with non-diagonal matrices Q and R describing the state cost and noise covariance respectively. It is shown that the optimal threshold is described by an ellipsoid $x^T Px \leq 2\sqrt{\rho}$ in the state space, where ρ is the cost per control event, and the matrix P can be determined from a novel Riccati-like equation $PRP + \frac{1}{2}\mathrm{tr}(RP)P = Q$. We show that this equation has a unique non-negative definite solution P when $R > 0$ and $Q \geq 0$. The other problem includes a first-order system with a non-zero inter-event time $\Delta T > 0$, describing, e.g., a minimum time between control events. This problem has previously been studied by simulations, and by a numerical convex optimization approach. We show that the optimal controller is threshold based, and obtain analytical expressions for the threshold.

Manuscript in preparation.

1. Introduction

Instead of sampling signals equidistantly in time it is well known that more efficient control might in some situations be achieved by event-based sampling, see, e.g. [Åström and Bernhardsson, 1999]. The main drawback with the approach is that the analysis and controller design becomes much more cumbersome, since no general solutions are known. Only a very limited number of problems have hitherto allowed for analytic expressions of the optimal event-mechanism.

This paper considers a standard LQG setup with linear dynamics and stochastic disturbances, with the additional objective of reducing the instances of controller action events. The notation and analysis method is taken from [Henningsson, 2011], to which we refer the reader for more background to the theory. To describe the system and controller dynamics a mixture of continuous (CT or "flow") and discrete time (DT, or "jump") is considered.

In [Åström and Bernhardsson, 1999], it was shown that the average event frequency f_u can be cut to a third compared to periodic control for the same state variance, for an integrator plant disturbed by white noise using impulse control. The event-based policy was *aperiodic:* events may occur arbitrarily close in time. Other authors have also investigated related aperiodic setups, e.g. [Rabi, 2006], [Rabi *et al.*, 2008], [Hristu-Varsakelis and Kumar, 2002]. To get an implementable control law, some authors have investigated *periodic* discrete time formulations, with a penalty for each sample with an event, e.g. [Cogill *et al.*, 2007], [Cogill, 2009], [Imer and Basar, 2010], [Sandee *et al.*, 2007], [Heemels *et al.*, 2011].

In [Henningsson *et al.*, 2008], *sporadic* control policies are investigated, where events may occur at any time as long as there is a minimum time ΔT between any two events. The paper contains numerical results and a sketch of analytical results; the latter are extended this paper. In [Henningsson, 2011], methods to derive bounds on costs in event-based control problems is presented. The current work is based on these methods, illustrating them by providing explicit solutions in the case of (multidimensional) integrator plants.

In Section 2 we describe the problem setup and how to find the optimal state distribution (primal problem) and the optimal value function (dual problem) using the results from [Henningsson, 2011]. Sections 3 and 4 then describe explicit solutions to two illustrative examples. The first is a multidimensional integrator system with quadratic cost matrix Q and noise covariance matrix R. Similar setups have been studied previously in [Cogill *et al.*, 2007; Cogill, 2009]. It turns out that the optimal controller is an event-based controller which actuates when the state leaves an ellipsoidal domain $\{x \mid x^T P x \leq 2\sqrt{\rho}\}$ where P is determined by a Riccati-like equation from system data, and where ρ is the cost of nonzero controller actuation. The second example is a first-order system with a nonzero inter-event time ΔT.

2. The Event-Based Control Problem

2.1 Dynamics

Consider a system with state $x \in \mathbb{R}^n$. The state normally evolves in the *flow mode*, with dynamics here assumed given by a multidimensional integrator system driven by white noise,

$$dx = dw, \quad \mathrm{E}(dwdw^T) = Rdt,$$

where w is a Wiener Process in $\mathbb{R}^n$ and $R \in \mathbb{R}^{n\times n} > 0$. At any time, the controller may initiate a *jump transition*

$$\begin{aligned} x^+ &= x + u + w_{\text{jump}}, \quad w_{\text{jump}} \sim \mathcal{N}(0, P_{\text{jump}}), \\ t^+ &= t + \Delta T, \end{aligned}$$

where $P_{\text{jump}} \in \mathbb{R}^{n\times n} \geq 0$ is the covariance of the Gaussian noise w_{jump}. The disturbances dw and w_{jump} are white and independent of the future trajectory. We will in this article assume that the control actions reset the state using $u = -x$ and that further control actions are prohibited for a period ΔT of time. This corresponds to $u = -x$, $P_{\text{jump}} = R\Delta T$.

2.2 The Optimal Control Problem

Let the system trajectory go from initial time $t = 0$ and state $x = x_i$ to final time $t = t_f$ and state x_f (x_i, x_f and t_f may be random variables). Define the finite time *path measures*

$$\begin{aligned} \langle \mu_{\text{ft}}, f\rangle &= \mathrm{E}\int_{t\in\mathcal{T}_{\text{flow}}} f(x)dx, \\ \langle \varphi_{\text{ft}}, f\rangle &= \mathrm{E}\sum_{k\in\mathcal{K}_{\text{jump}}} f(x_k), \end{aligned} \tag{1}$$

where $\mathcal{T}_{\text{flow}}$ is the time spent in flow, $\mathcal{K}_{\text{jump}}$ indexes the jumps, and x_k is the state upon entry of jump k. Let the control objective be

$$J_{\text{avg}} = \limsup_{t_{\text{spent}}\to\infty} \frac{j}{t_{\text{spent}}}, \quad j = \langle \mu_{\text{ft}}, c_{\text{flow}}\rangle + \langle \varphi_{\text{ft}}, c_{\text{jump}}\rangle,$$

where $c_{\text{flow}}(x) = x^T Qx$ and $c_{\text{jump}}(x) = \rho + \text{tr}(QR\Delta T)$. Here ρ is the cost of each control event, and the second term is the expected cost during the successive period ΔT of open loop operation. We want to find a controller to minimize J, by choosing when to jump (=actuate) assuming noise-free knowledge of x.

The search for the resulting optimal state distributions can be written as the following convex optimization problem: Find $\mu(x)$ and $\varphi(x)$ such that

$$\begin{aligned}\text{Primal:}\quad & \inf J = \langle \mu, c_{\text{flow}}\rangle + \langle \varphi, c_{\text{jump}}\rangle \\ & \text{s.t.}\quad \mathcal{A}_{\text{flow}}\mu + \mathcal{A}_{\text{jump}}\varphi = 0 \\ & \qquad\ \ \langle \mu, 1\rangle + \langle \varphi, \Delta T\rangle = 1 \\ & \qquad\ \ \mu \geq 0,\ \ \varphi \geq 0\end{aligned}$$

where

$$\begin{aligned}(\mathcal{A}_{\text{flow}}\mu)(x) &= \frac{1}{2}\text{tr}(R\nabla^2\mu), \\ (\mathcal{A}_{\text{jump}}\varphi)(x) &= \phi_{R\Delta T}(x)\int \varphi dx - \varphi(x),\end{aligned}$$

are the Kolmogorov forward operators for the flow and jump modes of the system. Here $\phi_\Sigma = \det(2\pi\Sigma)\exp\left(-\frac{1}{2}x^T\Sigma^{-1}x\right)$ if $\Delta T > 0$ and $\phi_0 = \delta(x)$. For details see [Henningsson, 2011]. The last operator $\mathcal{A}_{\text{jump}}$ describes the effect of the system operating in open loop for time ΔT after a control event. The dual Kolmogorov backward operators give the expected change of a C^2 function $V(x)$ given the flow and jump dynamics respectively:

$$\begin{aligned}\text{E}(dV|x) \quad &= \mathcal{A}^*_{\text{flow}}V = \frac{1}{2}\text{tr}(R\nabla^2 V), \\ \text{E}\Big(V(x^+) - V(x)|x\Big) \quad &= \mathcal{A}^*_{\text{jump}}V = [\phi_{R\Delta T}\star V](0) - V(x).\end{aligned}$$

The calculation of optimal controllers can be based on the following result based on [Henningsson, 2011].

THEOREM 1
Suppose a bounded function $V(x)$ and constant J are found satisfying

$$c_{\text{flow}}(x) + (\mathcal{A}^*_{\text{flow}}V)(x) \geq J, \qquad \forall x \in \mathbb{R}^n, \tag{2}$$

$$c_{\text{jump}}(x) + (\mathcal{A}^*_{\text{jump}}V)(x) \geq J\Delta T, \qquad \forall x \in \mathbb{R}^n, \tag{3}$$

where for each x equality is achieved in either (2) or (3). Then the optimal cost is J and it is optimal to actuate when equality is achieved in

(3). Furthermore, for finite time path distributions, the accumulated cost satisfies

$$\frac{1}{T}j + \mathrm{E}(V(x_f) - V(x_i)) \geq J, \tag{4}$$

where equality is achieved by the optimal controller given above.

Proof: See Appendix A. □

3. Example: The Elliptic Integrator Case

3.1 The System

We will consider an n-dimensional integrator system of the form above with $\Delta T = 0$. The control signal may hence at any time reset the state x to the origin, incurring a cost

$$c_{\text{jump}} = \rho$$

per such event. We are interested in keeping the mean event rate f_u of control actions rare while minimizing an average state cost $c_{\text{flow}}(x) = x^T Q x$, i.e. minimizing

$$J_x + \rho f_u = \limsup_{T\to\infty} \frac{1}{T}\left(\int_0^T x^T Q x\, dt + \rho\,\{\#\text{ events up to } T\}\right) \tag{5}$$

Note that when $\Delta T = 0$, the backward operators are

$$\mathcal{A}^*_{\text{flow}} V = \frac{1}{2}\mathrm{tr}(R\nabla^2 V),$$
$$\mathcal{A}^*_{\text{jump}} V = V(0) - V(x).$$

3.2 Optimal Sporadic Controller

Theorem 2
The controller minimizing the expected average cost $J := J_x + \rho f_u$ uses an ellipsoidal threshold and resets the state when $g(x) \leq 0$ where

$$g(x) = g_0 - x^T P x, \tag{6}$$

with $g_0 = 2\sqrt{\rho}$ and where P is determined by the Riccati-like design equation

$$PRP + \frac{1}{2}\mathrm{tr}(RP)P = Q. \tag{7}$$

This equation has a unique solution $P \geq 0$ for any $R > 0$ and $Q \geq 0$. The optimal cost is

$$J = \sqrt{\rho}\,\mathrm{tr}(RP),$$

the average state cost is $J_x = \frac{1}{2}\sqrt{\rho}\ \mathrm{tr}(RP)$ and the mean event rate is $f_u = \frac{1}{2\sqrt{\rho}}\mathrm{tr}(RP)$.

Proof: We will show that with g given by (6) and P given by (7) the function [1]

$$V(x) = \begin{cases} -\dfrac{1}{4}g(x)^2, & g(x) \geq 0, \\ 0 & \text{otherwise}, \end{cases}$$

is a value function for the dynamics satisfying the conditions in Theorem 1 (with $\Delta T = 0$).

The jump dynamics constraint (3) for the value function is satisfied since resets are instant and go from $g(x) \leq 0$ to $x = 0$ and the reset cost satisfies

$$\mathcal{A}^*_{\text{jump}}V + c_{\text{jump}} = V(0) - V(x) + c_{\text{jump}} = -\frac{1}{4}g_0^2 + \rho = 0$$

Looking at the interior ($g > 0$) we have that

$$\begin{aligned} \nabla g &= -2Px, & \nabla^2 g &= -2P, \\ \nabla V &= -\frac{1}{2}g\nabla g = gPx, & \nabla^2 V &= gP - 2Pxx^T P. \end{aligned}$$

The expected change in V per time unit is

$$\mathcal{A}^*_{\text{flow}}V = \frac{1}{2}g(x)\,\mathrm{tr}(RP) - x^T PRPx.$$

We hence have

$$\begin{aligned} &c_{\text{flow}} + \mathcal{A}^*_{\text{flow}}V = \\ &= x^T\Big(Q - \frac{1}{2}\mathrm{tr}(RP)P - PRP\Big)x + \frac{1}{2}g_0\mathrm{tr}(RP). \end{aligned}$$

This becomes constant equal to J precisely when (7) is satisfied.

We now show that there is exactly one symmetric $P \geq 0$ that satisfies (7). To construct such P, consider the equations

$$PRP + tP = Q.$$

[1] V is not $\mathcal{C}^2$ at the trigger threshold, but this is easy to handle.

for a, yet to be determined, scalar $t \geq 0$. We can after an invertible coordinate transformation assume that $R = I$ and Q is diagonal, i.e. the equation becomes $P^2 + tP = Q$. Since P and P^2 commute, they must commute also with $Q = P^2 + tP$. Thus there must exist a basis of common eigenvectors for P and Q, in which both will be diagonal. Assuming such an eigenvalue decomposition $Q = U\Lambda_Q U^T$ we see that $P = U\Lambda_P U^T$ solves the equation iff

$$\lambda_p^2 + t\lambda_p = \lambda_q.$$

This equation has a unique nonnegative solution λ_p for any $t \geq 0$ and $\lambda_q \geq 0$. The solution $\lambda_p(t)$ is strictly decreasing in t (unless $\lambda_q = 0$ when the nonnegative solution is $\lambda_p(t) = 0$ for any t). When $t = 0$, the solution is $P = Q^{\frac{1}{2}}$. When $t \to \infty$, the solution goes to $P = 0$. Thus, unless $Q = 0$ there must be a unique value $t^* \in (0, \infty)$ such that $t^* = \frac{1}{2}\mathrm{tr}(RP(t^*))$. The unique solution to (7) is then $P = P(t^*)$.

To calculate the mean event rate we note that the expected time to next actuation event, is easily determined by the fact that $\mathcal{A}^*_{\text{flow}} g = -\mathrm{tr}(RP)$ is constant. The mean exit time starting at state x is hence

$$\theta(x) = \frac{1}{\mathrm{tr}(RP)} g(x),$$

which gives the mean inter-event time

$$T_{\text{inter-event}} = \theta(0) = \frac{g_0}{\mathrm{tr}(RP)},$$

and thus the mean event rate

$$f_u = \frac{1}{T_{\text{inter-event}}} = \frac{1}{g_0}\mathrm{tr}(RP) = \frac{1}{2\sqrt{\rho}}\mathrm{tr}(RP).$$

The average state cost becomes

$$J_x = \mathrm{E}\, c_{\text{flow}}(x) = J - \rho f_u = \frac{1}{2}\sqrt{\rho}\,\mathrm{tr}(RP) = \frac{1}{2}J.$$

□

3.3 Optimal State Distribution in $\mathbb{R}^2$

In the elliptic case (R and Q not proportional), we have only been able to calculate explicit expressions for the resulting optimal state distribution for $n = 2$. It is given by fundamental solutions to the Laplace operator $\Delta f = \delta(x)$ satisfying $f = 0$ on the boundary of the elliptic domain obtained above. To describe it we use complex parametrization $w = x_1 + ix_2$. The following function satisfies $\Delta f(w) = \delta(w)$ and $f(w) = 0$ on the ellipse $(x_1/\cosh(\xi))^2 + (x_2/\sinh(\xi))^2 = 1$.

$$f(w) = -\frac{1}{2\pi}\log\left|\sqrt{k}\,\mathrm{sn}\left(\frac{2K(k)}{\pi}\arcsin(w), k\right)\right|,$$

where "sn(z,k)" denotes the Jacobi elliptic function and where $K(k)$ is given by the elliptic complete integral of the first kind,

$$K(k) = \int_0^{\pi/2} \frac{d\theta}{\sqrt{1-k^2\sin^2\theta}}.$$

The value $0 < k < 1$ is related to the shape of the ellipse and is determined implicitly by

$$\xi = \frac{\pi}{4}\frac{K(\sqrt{1-k^2})}{K(k)}.$$

3.4 Optimal State Distribution in $\mathbb{R}^n$—the Spherical Case

For the case with spherical symmetry, where $Q = qI$ and $R = rI$, explicit expressions for the optimal state distributions can be given as fundamental solutions to the Laplace operator in a sphere. The optimal state distribution is the minimum for the following convex optimization problem in $\mathbb{R}^n$, a variant of the primal problem:

$$\inf \int |\Delta f|\,dx$$
$$f \geq 0, \quad \int f\,dx = 1, \quad \int |x|^2 f(x)\,dx = \sigma^2.$$

The minimum equals $\frac{2n^2}{(n+2)}\sigma^{-2}$ and is attained for[2]

$$\begin{aligned}
f(x) &= \frac{1}{6\sigma^2}(R-|x|)_+, && n=1\\
f(x) &= \frac{1}{2\pi\sigma^2}\log_+\left(\frac{R}{|x|}\right), && n=2\\
f(x) &= \frac{n}{c_n(n^2-4)\sigma^2}(|x|^{2-n} - \mathbb{R}^{2-n})_+, && n\geq 3
\end{aligned}$$

where $R = (2(n+2)/n)^{1/2}\sigma$ and where $c_n = \pi^{n/2}/\Gamma(\frac{n}{2}+1)$ denotes the volume of the unit ball in $\mathbb{R}^n$.

[2] the notation $f_+ := \max(f, 0)$ is used

To see this, we first notice that the minimization can be reduced to rotationally symmetric functions: Given a function $f(x)$ define the rotational symmetrization as the following Haar integral

$$\tilde{f}(|x|) = \int_{U \in SO(n)} f(Ux)\,dU = \int_{|\omega|=1} f(\omega|x|)\,d\omega \Big/ \int_{|\omega|=1} d\omega,$$

where the integral is with respect to the unit measure in $SO(n)$. It is clear that $f(Ux)$, and hence $\tilde{f}$, satisfies the same linear conditions as f. From Jensen's inequality it follows that the convex and rotationally invariant map $f \to \int |\Delta f| =: T(f)$ is smaller for $\tilde{f}$ than for f:

$$T(\tilde{f}) = T\left(\int_{SO(n)} f(Ux)\,dU\right) \leq \int_{SO(n)} T(f(Ux))\,dU = T(f).$$

Now using the following facts it is easy to calculate the costs, resulting in the same result as in the previous section,

$$\begin{aligned}
\Delta(x^{2-n}) &= -n(n-2)c_n\delta(x), && n \geq 3,\\
\Delta(\log(x)) &= -2\pi\delta(x), && n = 2,\\
s_{n-1} &= nc_n, \quad \text{area of unit sphere in } \mathbb{R}^n, &&\\
\Delta(f) &= \frac{\partial^2 f}{\partial r^2} + \frac{N-1}{r}\frac{\partial f}{\partial r} + \frac{1}{r^2}\Delta_{S^{N-1}} f, &&
\end{aligned}$$

where the Laplace-Beltrami operator $\Delta_{S^{N-1}}$ disappears for rotationally symmetric functions.

4. Example: First-Order System with $\Delta T > 0$

As a second example we will give analytic formulas for the optimal sporadic controller for a first-order integrator example with minimum inter-event time of $\Delta T > 0$. Each control action moves the state to the origin, but there is then a period of time ΔT where no new control action is possible, and where the system state will hence evolve as Brownian motion. We will first solve the dual problem and then the primal problem.

To aid the presentation, we introduce some notation. Hence, let $\phi(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2}$ and $\Phi(x) = \int_{-\infty}^{x} \phi(t)dt$; note that $\phi'(x) = -x\phi(x)$, $\Phi(-\infty) = 0$, $\Phi(0) = 1/2$, and $\Phi(\infty) = 1$. Also note that the normal distribution $N(0, \sigma^2)$ is given by $\frac{1}{\sigma}\phi(x/\sigma)$. To compactify notation put $\widetilde{\Phi}(x) = \Phi(x) - \Phi(0)$. Introduce also the even function

$$\Psi(x) := \int_0^x \widetilde{\Phi}(x)dt = x\widetilde{\Phi}(x) + \phi(x) - \phi(0).$$

It is easy to see that $\Psi''(x) = \widetilde{\Phi}'(x) = \phi(x)$ and that hence

$$\frac{1}{2}f''(x) = k\frac{1}{\sigma}\phi(x/\sigma) \Leftrightarrow f(x) = 2k\Psi(x/\sigma) + xf'(0) + f(0).$$

The operators are now

$$\begin{aligned} \mathcal{A}^*_{\text{flow}}V &= \frac{1}{2}V''_{xx}, \\ \mathcal{A}^*_{\text{jump}}V &= [\phi_\sigma \star V](0) - V(x). \end{aligned}$$

4.1 Dual: Value Function When $\Delta T > 0$

We will show that a value function of the form

$$V(x) = \begin{cases} x^2(J - x^2/6), & |x| \le d, \\ V(d) & \text{otherwise.} \end{cases}$$

satisfies the condition in Theorem 1, for suitable values of J and d and that (2) is satisfied with equality in the flow region $|x| \le d$ and (3) in the jump region $|x| \ge d$.

Direct verification shows that this V satisfies (2) with equality when $|x| < d$. Some thoughts reveal that to satisfy the conditions in Theorem 1, we must have $V'_x(d) = 0$, which is equivalent to

$$d^2/3 = J.$$

A smaller value of d would give $V'_x(d-) > 0$ and $V'_x(d+) = 0$ giving a negative Dirac distribution in V''_{xx} violating (2) at $|x| = d$ and a larger value of d would give a decreasing function V in the region $\sqrt{3J} \le |x| \le d$ violating (3) in either $|x| = \sqrt{3J}$ or $|x| \ge d$.

To have equality in (3) we should have

$$J\sigma^2 = \rho + \frac{\sigma^4}{2} - V(d) + [\phi_\sigma \star V](0), \tag{8}$$

where we have used the notation $\sigma^2 = \Delta T$. This gives an implicit equation for d as a function of ΔT and ρ which we now show how to simplify.

To calculate the convolution of V with the normal distribution $N(0, \sigma)$ we need the following integrals:

$$I_k(d) = \int_0^d x^k \phi(x) dx$$

Since $\phi'(x) = -x\phi(x)$ partial integration shows

$$I_{k+2}(d) = -d^{k+1}\phi(d) + (k+1)I_k(d),$$
$$I_0(d) = \widetilde{\Phi}(d),$$
$$I_1(d) = \phi(0) - \phi(d).$$

This gives

$$\int_0^d V(x)\phi_\sigma(x)dx = \int_0^d (Jx^2 - x^4/6)\phi_\sigma(x)dx =$$
$$= J\sigma^2 I_2(d/\sigma) - \frac{\sigma^4}{6} I_4(d/\sigma) =$$
$$= J\sigma^2\left(-D\phi(D) + \widetilde{\Phi}(D)\right) -$$
$$\frac{\sigma^4}{6}\left(-D^3\phi(D) - 3D\phi(D) + 3\widetilde{\Phi}(D)\right).$$

We also have

$$\int_d^\infty V(x)\phi_\sigma(x)dx = V(d)\left(\frac{1}{2} - \widetilde{\Phi}(D)\right),$$

where we have introduced the notation $D := d/\sigma$. This gives

$$J\sigma^2 = \rho + \frac{\sigma^4}{2} - V(d) + 2J\sigma^2\left(-D\phi(D) + \widetilde{\Phi}(D)\right) -$$
$$\frac{\sigma^4}{3}\left(-D^3\phi(D) - 3D\phi(D) + 3\tilde{\Phi}(D)\right) + V(d)(1 - 2\widetilde{\Phi}(D))$$

Using $J = d^2/3$ and $V(d) = d^4/6$ we get the equation

$$\rho = \sigma^4\left(-\frac{1}{2} + \frac{D^2}{3}\left(1 + D\phi(D) - 2\widetilde{\Phi}(D) + D^2\widetilde{\Phi}(D)\right) - D\phi(D) + \widetilde{\Phi}(D)\right) \tag{9}$$

The right hand side can with some work be shown to be a monotonously increasing function in D, this means that there is a unique D for any choice of ρ and $\Delta T = \sigma^2$.

Figure 1 shows the value function $V(x)$ for some different values of $\Delta T = 0, 0.2, 1$ and $\rho = 1/6$, note that the threshold increases when ΔT increases.

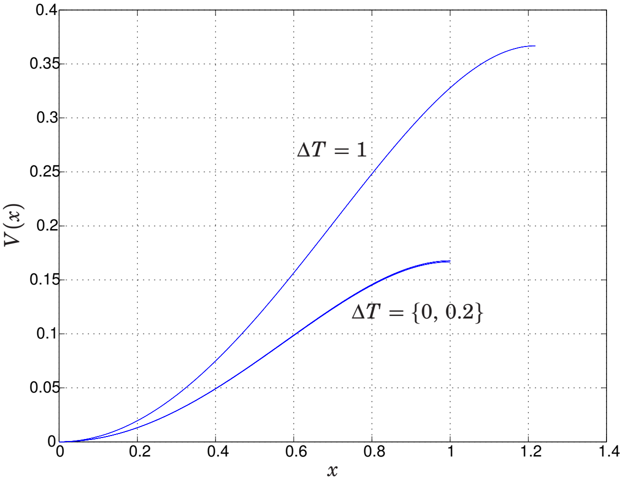

Figure 1. Value function for $\Delta T = 0$, 0.2, 1, with $\rho = 1/6$. The curves for $\Delta T = 0$ and 0.2 almost overlap. The thresholds are $d = 1, 1.01$ and 1.22 respectively. The value function is constant equal to $V(d)$ for $x > d$

.

4.2 Primal: Optimal State Distribution When $\Delta T > 0$

We will now show that for each threshold d and for suitable constants k_i we have

$$\mu(x) = 2k_1\sigma(\Psi(d/\sigma) - \Psi(x/\sigma))1_{|x|\le d},$$
$$\varphi(x) = \frac{k_2}{\sigma}\phi(x/\sigma)1_{d\le|x|} + k_3(\delta(x+d) + \delta(x-d)),$$

where we still use the notation $\sigma = \sqrt{\Delta T}$.

Using the equations for stationary dynamics

$$\mathcal{A}_{\text{flow}}\mu + \mathcal{A}_{\text{jump}}\varphi = 0$$

in the regions $|x| < d$ and $|x| > d$ respectively, and since we have $\mathcal{A}_{\text{flow}}\mu = \frac{1}{2}\mu''$ and $\mathcal{A}_{\text{jump}}\varphi = \frac{1}{\sigma}\phi(x/\sigma)\langle\varphi, 1\rangle - \varphi$. and the we see that the constants k_1, k_2, k_3 satisfy

$$\langle\varphi, 1\rangle = k_1 = k_2, \tag{10}$$

whereas the unit mass condition is

$$\langle\mu, 1\rangle + \Delta T\,\langle\varphi, 1\rangle = 1 \tag{11}$$

Equation (10) gives

$$2k_2(1 - \Phi(d/\sigma)) + 2k_3 = k_2 \tag{12}$$

and we can obtain k_1 from (11)

$$\begin{aligned} 1 &= k_1 \cdot \left(4\sigma \int_0^d (\Psi(d/\sigma) - \Psi(x/\sigma))dx + \sigma^2 \right) \\ &= k_1 \cdot \left(4\sigma \left(d\Psi(d/\sigma) - \sigma \int_0^{d/\sigma} \Psi(x)dx \right) + \sigma^2 \right) \end{aligned}$$

The integral can be evaluated using partial integration

$$\begin{aligned} &\int_0^x \Psi(y)dy = \int_0^x \left[y\widetilde{\Phi}(y) + \phi(y) - \phi(0)\right] dy \\ &= \frac{x^2}{2}\widetilde{\Phi}(x) - \int_0^x \frac{y^2}{2}\phi(y)dy + \Phi(x) - \Phi(0) - x\phi(0) \\ &= \frac{x^2+1}{2}\widetilde{\Phi}(x) + \frac{x}{2}\phi(x) - x\phi(0), \end{aligned}$$

where we have used

$$\int_0^x y^2\phi(y)dy = -x\phi(x) + \widetilde{\Phi}(x).$$

We hence get

$$\frac{1}{k_1} = \sigma^2 \left(1 + 2(D^2 - 1)\widetilde{\Phi}(D) + 2D\phi(D) \right),$$

where $D = d/\sigma$. From (10) and (12) we then get k_2 and k_3.

For a given $\Delta T > 0$ we can now evaluate the stationary cost as a function of the control event cost ρ:

$$J(\rho) = \int_{-d}^{d} x^2\mu(x)dx + \left(\rho + \frac{(\Delta T)^2}{2} \right) k_1.$$

Here again the integral can be evaluated using repeated partial integrations. We have

$$\int_{-d}^{d} x^2\mu(x)dx = \frac{\sigma^4 k_1}{3} \left((D^4 - 3)\widetilde{\Phi}(D) + (D^3 + 3D)\phi(D) \right)$$

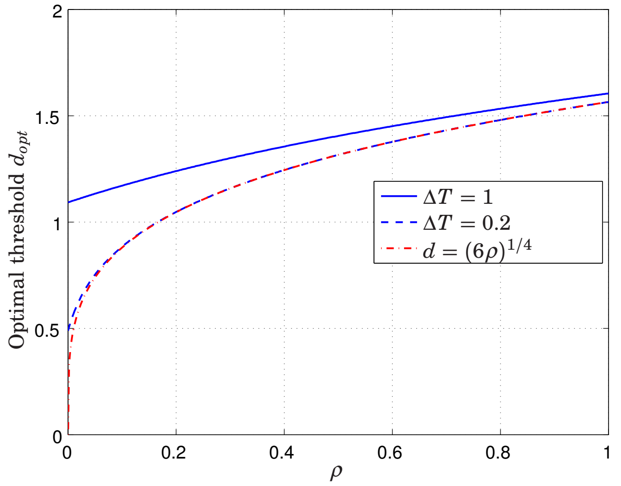

Figure 2. Optimal threshold as a function of control event cost ρ.

which means that the cost equals

$$J(\rho) = \frac{\frac{\sigma^4}{3}\Big((D^4-3)\widetilde{\Phi}(D) + (D^3+3D)\phi(D)\Big) + \Big(\rho + \frac{\sigma^4}{2}\Big)}{\sigma^2\Big(1 + 2(D^2-1)\widetilde{\Phi}(D) + 2D\phi(D)\Big)},$$

with $\sigma^2 = \Delta T$ and $D = d/\sigma$. We now obtain the optimal threshold D either from (9) or by directly optimizing over the threshold d. See Figures 2 and 3. (It is comforting to notice that numerical calculations verify that both methods give the same result.) Numerical experiments also verify that this D gives $J = d^2/3$ as before.

When $\sigma \to 0$ it is easy to see that we get

$$\begin{aligned} J(\rho, d) &= \frac{d^2}{6} + \frac{\rho}{d^2} \\ d_{opt} &= (6\rho)^{1/4} \\ J_{opt} &= 2\left(\frac{\rho}{6}\right)^{1/2} \\ (k_1, k_2, k_3) &\to \left(\frac{1}{d^2}, \frac{1}{d^2}, \frac{1}{2d^2}\right), \end{aligned}$$

so we recover the familiar triangular distribution for the case $\Delta T = 0$ in [Åström and Bernhardsson, 1999].

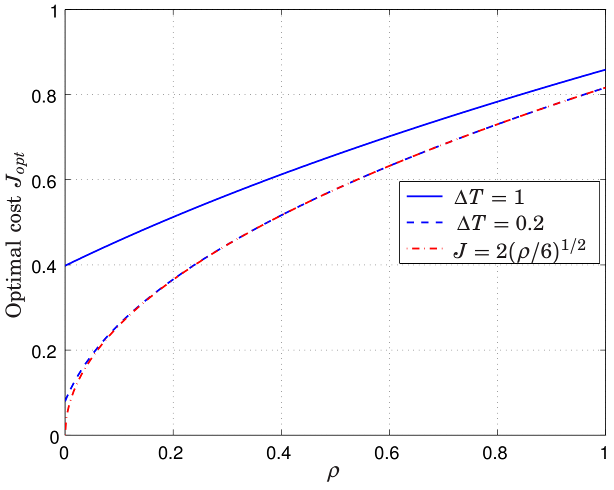

Figure 3. Optimal cost as a function of control event cost ρ.

5. Summary

We have managed to obtain more or less explicit solutions to the sporadic optimal control problems in the case with quadratic cost Q and general covariance matrix R for the multi-dimensional integrator case. We have not been able to obtain similar nice explicit results to the general case $\dot{x} = Ax + Bu$. Numerical optimization indicates that the "flow" domain is still a connected region containing the origin, but it no longer seems to be ellipsoidal.

Similar analysis as in Section 4 shows that the formula (4) also holds when $\Delta T > 0$, but with another, larger, constant g_0. Formulas for this constant can be obtained with similar techniques as in Section 5.

References

Åström, K. J. and B. Bernhardsson (1999): "Comparison of periodic and event based sampling for first-order stochastic systems." In *Preprints of the 14th World Congress of IFAC*, Beijing, P.R. China.

Cogill, R. (2009): "Event-based control using quadratic approximate value functions." In *Proceedings of the 48th IEEE Conference on Decision and Control*, Shanghai, P.R. China.

Cogill, R., S. Lall, and J. P. Hespanha (2007): "A constant factor approximation algorithm for optimal estimation subject to communication costs." In *Proceedings of the 2007 American Control Conference*, New York, New York.

Heemels, W. P. M. H., M. C. F. Donkers, and A. R. Teel (2011): "Periodic event-triggered control based on state feedback." In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 2571–2576. IEEE.

Henningsson, T. (2011): "Sporadic event-based control using path constraints and moments." In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, Orlando, Florida.

Henningsson, T., E. Johannesson, and A. Cervin (2008): "Sporadic event-based control of first-order linear stochastic systems." *Automatica*, **44:11**, pp. 2890–2895.

Hristu-Varsakelis, D. and P. Kumar (2002): "Interrupt-based feedback control over a shared communication medium." In *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vegas, Nevada.

Imer, O. C. and T. Basar (2010): "Optimal estimation with limited measurements." *International Journal of Systems, Control and Communications*, **2**, pp. 5–29.

Rabi, M. (2006): *Packet based Inference and Control*. PhD thesis, Institute for Systems Research, University of Maryland. `http://hdl.handle.net/1903/6612`.

Rabi, M., K. Johansson, and M. Johansson (2008): "Optimal stopping for event-triggered sensing and actuation." In *Proceedings of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico.

Sandee, J. H., W. P. M. H. Heemels, and P. P. J. van den Bosch (2007): "Case studies in event-driven control." In *Proceedings of the 10th International Conference on Hybrid Systems: Computation and Control*, Pisa, Italy.

A. Proof of Theorem 1

Lemma 1

Consider a distribution of finite time system trajectories. Let $V : \mathbb{R}^n \mapsto \mathbb{R}$ be a $\mathcal{C}^2$ function. Then the expected change in $V(x)$ from entry to exit is given by

$$\mathrm{E}(\Delta V) = \mathrm{E}(V(x_f) - V(x_i)) = \langle \mu_{\mathrm{ft}}, \mathcal{A}^*_{\mathrm{flow}} V \rangle + \left\langle \varphi_{\mathrm{ft}}, \mathcal{A}^*_{\mathrm{jump}} V \right\rangle.$$

□

Proof: Splitting the difference into accumulated differences and using the definitions of $\mathcal{A}_{\mathrm{flow}}, \mathcal{A}_{\mathrm{jump}}$, we get

$$\begin{aligned}
\mathrm{E}(\Delta V) &= \mathrm{E}\int_{\mathcal{T}_{\mathrm{flow}}} dV + \mathrm{E}\sum_{k\in\mathcal{K}_{\mathrm{jump}}} V(x_k^+) - V(x_k) \\
&= \mathrm{E}\int_{t\in\mathcal{T}_{\mathrm{flow}}} (\mathcal{A}^*_{\mathrm{flow}} V)(x)dt + \mathrm{E}\sum_{k\in\mathcal{K}_{\mathrm{jump}}} (\mathcal{A}^*_{\mathrm{jump}} V)(x_k) \\
&= \langle \mu_{\mathrm{ft}}, \mathcal{A}^*_{\mathrm{flow}} V \rangle + \left\langle \varphi_{\mathrm{ft}}, \mathcal{A}^*_{\mathrm{jump}} V \right\rangle.
\end{aligned}$$

□

Lemma 2

Suppose there exist a $\mathcal{C}^2$ function $V : \mathbb{R}^n \mapsto \mathbb{R}$ and $\underline{J} \in \mathbb{R}$ such that

$$\begin{aligned}
c_{\mathrm{flow}}(x) + (\mathcal{A}^*_{\mathrm{flow}} V)(x) &\geq \underline{J}, & \forall x \in \mathcal{X}_{\mathrm{flow}}, \\
c_{\mathrm{jump}}(x) + (\mathcal{A}^*_{\mathrm{jump}} V)(x) &\geq \underline{J}\Delta T, & \forall x \in \mathcal{X}_{\mathrm{jump}}.
\end{aligned}$$

Then for any distribution of finite time paths consistent with the system dynamics, such that $x \in \mathcal{X}_{\mathrm{flow}}$ when flowing, and $x \in \mathcal{X}_{\mathrm{jump}}$ at the beginning of each jump, the accumulated cost is lower bounded by

$$j = \langle \mu_{\mathrm{ft}}, c_{\mathrm{flow}} \rangle + \langle \varphi_{\mathrm{ft}}, c_{\mathrm{jump}} \rangle \geq \underline{J}T - \mathrm{E}(\Delta V).$$

Upper bound: Reversing the inequalities in the assumptions above gives the converse inequality in the conclusion. □

Proof: Using Lemma 1, we see that

$$\begin{aligned}
j + \mathrm{E}(\Delta V) &= \langle \mu_{\mathrm{ft}}, c_{\mathrm{flow}} + \mathcal{A}^*_{\mathrm{flow}} V \rangle + \left\langle \varphi_{\mathrm{ft}}, c_{\mathrm{jump}} + \mathcal{A}^*_{\mathrm{jump}} V \right\rangle \\
&\geq \langle \mu_{\mathrm{ft}}, \underline{J} \rangle + \langle \varphi_{\mathrm{ft}}, \underline{J}\Delta T \rangle = \underline{J}T,
\end{aligned}$$

where the inequality follows since $\langle \mu_{\mathrm{ft}}, f \rangle$ and $\langle \varphi_{\mathrm{ft}}, f \rangle$ are increasing in f (see (1)), and independent of the values of $f(x)$ that the trajectories do not visit, and the last equality follows from

$$T = \mathrm{E}(t_f) = \mathrm{E}\int_{\mathcal{T}_{\mathrm{flow}}} dt + \mathrm{E}\sum_{\mathcal{K}_{\mathrm{jump}}} \Delta T = \mu_{\mathrm{ft}}(1) + \varphi_{\mathrm{ft}}(\Delta T).$$

□

Proof of Theorem 1

The inequality (4) follows directly from Lemma 2 with $\underline{J} = J$. We can let $\mathcal{X}_{\text{flow}}$ and $\mathcal{X}_{\text{jump}}$ be the sets where equality is achieved in (2) and (3) respectively; from the preconditions we have $\mathcal{X}_{\text{flow}} \cup \mathcal{X}_{\text{jump}} = \mathbb{R}^n$. Using these sets with Lemma 2, this time using $c_{\text{flow}}(x) + (\mathcal{A}^*_{\text{flow}}V)(x) \leq \bar{J}$, $c_{\text{jump}}(x) + (\mathcal{A}^*_{\text{jump}}V)(x) \leq \bar{J}\Delta T$, letting $\bar{J} = J$, we see that $j + \mathrm{E}(\Delta V) \leq JT$ with the switching sets $(\mathcal{X}_{\text{flow}}, X_{\text{jump}})$.

Since V is bounded, (4) gives that

$$J_{\text{avg}} = \lim_{T\to\infty} \frac{j}{T} \geq \lim_{T\to\infty} \frac{JT - \mathrm{E}(\Delta V)}{T} = J,$$

with equality using the prescribed controller. □

Paper III

Sporadic Event-Based Control of First-Order Linear Stochastic Systems

Toivo Henningsson   Erik Johannesson
Anton Cervin

Abstract

The standard approach in computer-controlled systems is to sample and control periodically. In certain applications, such as networked control systems or energy-constrained systems, it could be advantageous to instead use event-based control schemes. Aperiodic event-based control of first-order stochastic systems has been investigated in previous work. In any real implementation, however, it is necessary to have a well-defined minimum inter-event time. In this paper, we explore two such sporadic control schemes for first-order linear stochastic systems and compare the achievable performance to both periodic and aperiodic control. The results show that sporadic control can give better performance than periodic control in terms of both reduced process state variance and reduced control action frequency.

1. Introduction

Digital feedback controllers are most often implemented using periodic sampling, computation, and actuation. This approach enables the control designer to utilize standard sampled-data system theory or to discretize a continuous-time controller assuming a fixed sampling rate and constant hold intervals [Åström and Wittenmark, 1997].

For some applications, however, event-based control schemes may have an advantage over periodic schemes. In networked control applications, it could make sense to only transmit information when something significant has occurred in the system, in order to save bandwidth. In embedded applications, it may be essential to minimize the number of control actions in order to save energy. In the application of inventory control it seems rational to replenish stock only when it is low rather than on a periodic basis, if there is a fixed transportation cost. Some sensors such as rotary motion encoders only give new measurements at ahead-of-time unknown events.

Event-based control as a technology is of course not new. Mostly, however, it has been applied in an ad-hoc way. This can be attributed to the lack of a comprehensive theory, which in turn can be explained by the mathematical difficulties involved. A discrete-time formulation can sometimes make it slightly easier to obtain a solution. Some recent papers have thus solved optimal discrete-time estimation problems, with limited [Imer and Basar, 2005] or event-triggered [Cogill *et al.*, 2006] measurements.

From a control-theoretic point of view, event-based control systems can be viewed as hybrid systems. In this paper, we consider first-order linear stochastic systems, where an exogenous random disturbance (modelled as white noise) causes the process state to drift. The control law generates discrete events when the state crosses certain boundaries. Hence, our system falls into the category of stochastic hybrid systems as defined in [Hu *et al.*, 2000].

Event-based control of first-order linear stochastic systems was studied in [Åström and Bernhardsson, 1999]. It was shown that, compared to periodic control, the output variance could be significantly reduced assuming the same mean time between events. The control was realized by applying an impulse action whenever the magnitude of the system state exceeded a certain threshold. This work was elaborated in [Rabi, 2006], which explores, among other things, event-based control with piecewise constant control signals and level-triggered sampling.

From a real-time systems point of view, however, tasks triggered by asynchronously generated events cannot be guaranteed service unless there is a well-defined minimum inter-arrival time. For the controller presented in [Åström and Bernhardsson, 1999] there was no such mini-

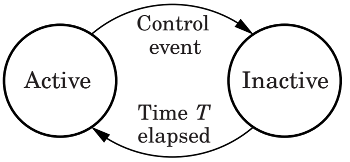

Figure 1. Controller state transitions. Control events may only be generated in the active state.

mum inter-arrival time. In accordance with real-time systems terminology [Buttazzo, 1997], we will refer to such a control policy as *aperiodic*.

In this paper, we explore the class of *sporadic* event-based controllers for first-order linear stochastic systems. With a minimum inter-arrival time T between events, such a controller can be guaranteed not to consume more than a certain network bandwidth or CPU utilization. Two sporadic controllers will be studied. The first controller measures the process state continuously and can take control actions at any time, but no more often than every T seconds. The second controller measures the process state every T_s seconds until a control action is applied, and resumes measurements T seconds after the last control action.

2. Problem Formulation

The process to be controlled is given by the linear stochastic differential equation

$$dx = axdt + udt + \sigma dw, \qquad x(0) = 0, \tag{1}$$

where x is the state, u the control signal, w is a Wiener process with unit incremental variance, a is the pole of the process, and $\sigma > 0$ is the intensity of the process noise. The control signal is zero except at *events* t_k, when it is allowed to be a Dirac pulse of magnitude u_k:

$$u(t) = \sum_{k=0}^{\infty} \delta(t - t_k) u_k. \tag{2}$$

The controller chooses when to generate an event based on the state of the system. After each event there is a period of *inactive state* of duration T, when no new events can be generated, see Figure 1.

The performance is measured by the stationary state cost,

$$J_x = \limsup_{t\to\infty} \mathrm{E}\, \frac{1}{t} \int_0^t x^2 ds,$$

and by the average control rate (or control cost),

$$J_u = \limsup_{t\to\infty} \mathrm{E}\, \frac{1}{t} N_u(0, t),$$

where $N_u(t_1, t_2)$ is the number of control actions in the interval (t_1, t_2). The total cost to be minimized is

$$J = J_x + \rho J_u, \tag{3}$$

where $\rho \geq 0$ is the relative cost of control actions.

Normalized Formulation

To reduce the number of free parameters we can use coordinate scaling to fix $\sigma = T = 1$. The parameters that remain, a and ρ, suffice to specify the problem up to coordinate scaling, and the original variables can be retrieved from inverse scaling. The parameters σ and T will be kept in the presentation when they add insight.

Let the transformed variables be described by

$$dt = T d\tau, \qquad dw = \sqrt{T} dv, \qquad x = \sigma\sqrt{T} x'$$

The dynamics become

$$dx' = a'x'd\tau + u'd\tau + dv,$$

where $u' = \sqrt{T}\sigma^{-1}u$, and $a' = aT$ is the relevant measure of process speed. The original costs are retrieved as

$$J_x = \sigma^2 T J'_x, \qquad J_u = T^{-1} J'_u,$$

so $\rho' = \frac{\rho}{\sigma^2 T^2}$ is the proper weighting after normalization. The normalized problem is described by the parameters

$$a' = aT, \qquad \rho' = \frac{\rho}{\sigma^2 T^2}.$$

3. Sporadic Control

3.1 General Observations

A sporadic controller is defined by two properties: when it generates an event and what control signal is used at the event. It is easy to see that an optimal controller for the problem above must satisfy the following:

- At any event t_k, the control signal u_k is chosen to bring x to the origin, i.e. $u_k = -x(t_k - 0)$.
- When an event is permitted, the decision of whether to generate one is a function only of the state x, and due to symmetry only of the absolute value $|x|$.
- If an event should be generated when $|x| = r$, one should also be generated whenever $|x| \geq r$.

Thus, the optimal control policy is a threshold policy where an event is triggered to bring x to zero whenever permitted and $|x| \geq r$.

To find the optimal threshold r, we evaluate J as a function of r in the closed loop system. To facilitate this, we first consider what happens between events.

3.2 Evolution Between Events

Between events the control signal is known, and the system evolves as a linear stochastic process. Assume that an event occurs at time $t_k = 0$, and that we want to predict the evolution from that time, from the state prior to the event $x_0 = x(0-)$. Let

$$m(t) = \mathrm{E}\big(x(t)\big), \quad P(t) = \mathrm{E}\big(x(t)^2\big) - m(t)^2$$

be the expected state trajectory and the expected state variance due to process noise entering after the event respectively, with initial conditions

$$m(0) = x_0 + u_k, \qquad P(0) = 0.$$

The distribution of $x(t)$ will be Gaussian with mean $m(t)$ and variance $P(t)$.

The expected state cost during the interval $(0,t)$ can be expressed as the sum of one contribution $V_P(t)$ from P and one $V_m(t)$ from m according to

$$\int_0^t \mathrm{E}\big(x(s)^2\big)ds = \int_0^t \big(P(\tau) + m(\tau)^2\big)d\tau = V_P(t) + V_m(t).$$

Since there is no feedback between events, u will enter the evolution only through $m(t)$. We find that

$$\begin{aligned}
\mathrm{E}(dP) &= \mathrm{E}(2xdx + dx^2) - 2m\,\mathrm{E}(dm) \\
&= \mathrm{E}(2x(axdt + udt) + \sigma^2 dw^2) - 2m(am + u) \\
&= (2aP + \sigma^2)dt.
\end{aligned}$$

Starting from $P(0) = 0$, the solution is

$$P(t) = \begin{cases} (1 - e^{2at})\dfrac{\sigma^2}{-2a}, & a \neq 0, \\ \sigma^2 t, & a = 0. \end{cases} \tag{4}$$

Integrating, the process noise contribution to the state cost during the interval $(0, t)$ is

$$V_P(t) = \begin{cases} \dfrac{\sigma^2}{-2a}\left(t - \dfrac{e^{2at} - 1}{2a}\right), & a \neq 0, \\ \dfrac{1}{2}\sigma^2 t^2, & a = 0. \end{cases} \tag{5}$$

The expected trajectory evolves according to $\mathrm{E}(dm) = \mathrm{E}(dx) = (am + u)dt$, giving the prediction

$$m(t) = e^{at}m(0) + \int_{0+}^{t} e^{a(t-\tau)}u(\tau)d\tau.$$

With no control during the interval $(0, t)$, the cost is

$$V_m(t) = \int_0^t m(s)^2 ds = Q(t)m(0)^2,$$

$$Q(t) = \begin{cases} \dfrac{e^{2at} - 1}{2a}, & a \neq 0, \\ t, & a = 0. \end{cases} \tag{6}$$

3.3 Sporadic Control with Continuous Measurements

We assume that the process state is measured continuously in the active state. As soon as the state leaves the region $|x| < r$, an event is generated and the controller is put in the inactive state for an interval of length T.

Since the system is reset to the same state at each event, the expected cost and time from one event to the next are enough to find the stationary costs, as

$$J_x = \frac{V_{\text{active}} + V_{\text{inactive}}}{T_{\text{active}} + T}, \qquad J_u = \frac{1}{T_{\text{active}} + T},$$

where V_{active} and T_{active} are the expected state costs and dwell times during one period of active state, and $V_{\text{inactive}} = V_P(T)$. We will characterize the behavior between two events by modifying the system so that it starts at one event and is stopped at the next.

The expected cost and dwell time during one period of active state can be found as

$$V_{\text{active}} = \int x^2 F(x)dx, \qquad T_{\text{active}} = \int F(x)dx, \tag{7}$$

where $F(x) = \int_0^\infty f(x,t)dt$ is the accumulated state density of the density $f(x,t)$ in the active state.

The system enters the active state as

$$f(x, t=0) = \begin{cases} \varphi(x), & |x| < r \\ 0, & |x| \geq r \end{cases}$$

where $\varphi(x)$ is Gaussian with zero mean and variance $P(T)$. The time evolution is given by the Fokker-Planck equation (see, e.g. [Åström, 1970], [Feller, 1971]) (with $\sigma = 1$):

$$\frac{\partial f}{\partial t} = \frac{1}{2}\frac{\partial^2}{\partial x^2}(f\sigma^2) - \frac{\partial}{\partial x}(fax) = \frac{1}{2}\frac{\partial^2 f}{\partial x^2} - ax\frac{\partial f}{\partial x} - af,$$

with absorbing boundary conditions $f(\pm r, t) = 0$. Since $f(x,t) \to 0$ as $t \to \infty$ we can integrate over $t \in [0, \infty)$ to find a differential equation for $F(x)$:

$$-\varphi(x) = \int_0^\infty \frac{\partial f}{\partial t} dt = \frac{1}{2}F''(x) - axF'(x) - aF(x),$$

with boundary conditions $F(\pm r) = 0$. The solution exists as long as $\varphi(x)$ does, and can be found numerically with a linear ODE Boundary Value Problem (BVP) solver or analytically as

$$F(x) = 2\int_{y=-r}^{x} e^{a(x^2-y^2)} \int_{z=y}^{0} \varphi(z)dzdy, \qquad |x| \leq r. \tag{8}$$

Figure 2 shows the costs as a function of r for the case $T = \sigma = 1$ and $a \in \{-0.5, 0, 0.5\}$. Other cases can be reconstructed by scaling as explained in Section 2. We see an initial decrease in the state cost as the threshold is increased, so the optimal threshold is non-zero even when $\rho = 0$. We also see that both costs decrease as a decreases, since the system becomes easier to control.

The cost functions can alternatively be found from

$$V_{\text{active}} = \int \varphi(x)V_x(x)dx, \quad T_{\text{active}} = \int \varphi(x)\theta(x)dx, \tag{9}$$

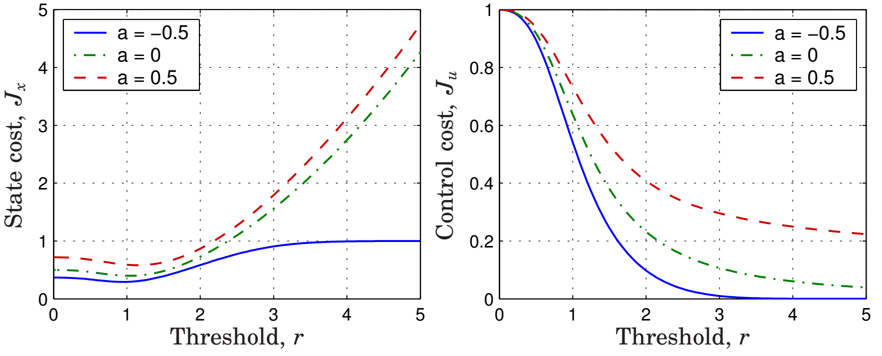

Figure 2. Cost functions for sporadic control with continuous-time measurements assuming $\sigma = T = 1$. Top: State cost J_x as a function of threshold r. Bottom: Control cost J_u as a function of threshold r. Both functions are plotted for systems with $a = -0.5$, $a = 0$ and $a = 0.5$ respectively.

where $V_x(x)$ is the expected state cost until the next event starting in the active state at x, and $\theta(x)$ is the corresponding expected dwell time (or first passage time, (see [Feller, 1971]). The value function $V(x) = V_x(x) - J\theta(x)$ can be used for dynamic programming.

When $x = \pm r$, $V_x(x) = \theta(x) = 0$, and when $|x| < r$,

$$\mathrm{E}\big(dV_x(x)\big) = -x^2 dt, \quad \mathrm{E}\big(d\theta(x)\big) = -1dt$$

which together with the dynamics (1) gives

$$-x^2 dt = \mathrm{E}\left(V_x' dx + \frac{1}{2} V_x'' dx^2\right) = \left(axV_x' + \frac{1}{2}V''\right)dt,$$

for $V_x(x)$, and similarly for $\theta(x)$. The solutions can be found numerically with an ODE BVP solver, or as

$$\begin{pmatrix} V_x(x) \\ \theta(x) \end{pmatrix} = 2\int_{y=x}^{r}\int_{z=0}^{y} e^{a(z^2-y^2)} \begin{pmatrix} z^2 \\ 1 \end{pmatrix} dzdy. \tag{10}$$

We note that problem can be extended in a few ways that fit well with our solution methods. Behavior in the inactive state only affects the solution through T_{inactive}, V_{inactive}, and the state density when entering the active state $\varphi(x)$. Possible extensions include a delay $\tau \leq T$ from the issue of an event to the actuation of the control impulse, and a stochastically varying inactive time T.

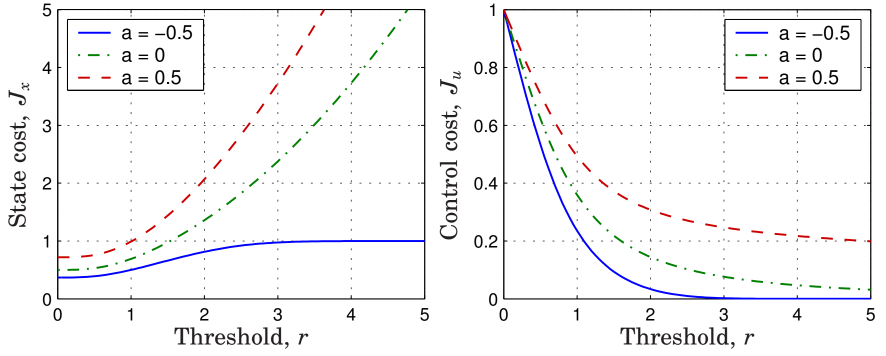

Figure 3. Cost functions for sporadic control with discrete-time measurements assuming $\sigma = T = T_s = 1$. Top: State cost J_x as a function of threshold r. Bottom: Control cost J_u as a function of threshold r. Both functions are plotted for systems with $a = -0.5$, $a = 0$ and $a = 0.5$ respectively.

3.4 Sporadic Control with Discrete Measurements

We now assume that the process is sampled with the interval $T_s \le T$ in the active state. Any deviations of the state outside the threshold between samples will go unnoticed. As before, when a deviation is detected at time t_k, the controller issues a control event and enters the inactive state, where it stays for T seconds. We now let $\{t_k\}$ denote all sampling instants, which progress as

$$t_{k+1} = \begin{cases} t_k + T_s, & |x_k| < r, \\ t_k + T, & |x_k| \ge r. \end{cases}$$

To find the optimal threshold r, the cost is characterized as a function of r. To this end, we compute the stationary state distribution (see [Feller, 1971]) at the sampling instants. Between sampling instants, the state evolves as

$$x_{k+1} = \begin{cases} e^{aT_s} x_k + w_k(T_s), & |x_k| < r, \\ w_k(T), & |x_k| \ge r, \end{cases} \tag{11}$$

where $w_k(t)$ is a Gaussian random variable with zero mean and variance $P(t)$. The stationary density always exists since there is a positive probability to go from any state x to any state interval (x_1, x_2) in one step, and for $|x| \ge r$ the density after any time step falls off as a Gaussian with variance $P(T)$. The accumulated state cost from time t_k to time t_{k+1} is given by

$$V_{\text{stay}} = Q(T_s)\ \mathrm{E}\left\{x_k^2 \,\middle|\, |x_k| < r\right\} + V_P(T_s) \tag{12}$$

if the controller stays in the active state and by

$$V_{\text{exit}} = V_P(T) \tag{13}$$

if the controller enters the inactive state.

Finally, assuming stationarity, the costs become

$$J_x = \frac{p_{\text{stay}} V_{\text{stay}} + p_{\text{exit}} V_{\text{exit}}}{p_{\text{stay}} T_s + p_{\text{exit}} T}, \quad J_u = \frac{p_{\text{exit}}}{p_{\text{stay}} T_s + p_{\text{exit}} T}$$

where

$$p_{\text{stay}} = \text{Prob}\,\{|x_k| < r\} = 1 - p_{\text{exit}}.$$

The stationary distribution of x_k can be found numerically by discretizing the state space and then iterating the distribution according to (11) until convergence.

Figure 3 shows the costs as a function of r for the case $T = T_s = \sigma = 1$ and $a \in \{-0.5, 0, 0.5\}$. Here, the state cost increases monotonically with r. With $T_s < T$ we would have an initial decrease, approaching the behavior for continous measurements as $T_s/T \to 0$. The control action frequency J_u falls off faster with increasing threshold than for the continous measurement case, since x is checked against the threshold less often with discrete measurements. As expected, both costs decrease with a.

Alternatively, the expected state cost $V_x(x)$ and dwell time $\theta(x)$ until the next event starting from state x can be iterated until convergence. As in the continuous case, we could extend the problem formulation with actuation delay and stochastically varying inactive time.

4. Comparison of Control Schemes

Sporadic control with continuous and discrete measurements (with $T_s = T$) will now be compared to periodic and aperiodic control. We first discuss how to make the comparison.

4.1 Periodic and Aperiodic Control

An aperiodic controller sets the process state x to zero whenever $|x| \geq r$ using an impulse control action [Åström and Bernhardsson, 1999]. The cost functions can be found by letting $\varphi(x)$ approach a unity Dirac pulse in (8) or (9), yielding

$$J_x = V_{\text{active}}/T_{\text{active}}, \qquad J_u = T_{\text{active}}^{-1}.$$

We assume that periodic control is also implemented with impulse control action, such that x is periodically reset to zero. The sampling interval is restricted to be no shorter than for the sporadic schemes. The costs become

$$J_x = V_P(T)/T, \qquad J_u = T^{-1}. \tag{14}$$

4.2 Preliminaries

For the sporadic controllers, minimization of the loss function J for a given ρ determines an optimal threshold r, which maps to an optimal average event rate J_u. The same holds for aperiodic control. In periodic control, however, there is no threshold. Instead, ρ determines the optimal sampling interval. Hence, we can parametrize controllers from all four classes by average event rate.

The four controllers differ by the constraints on when they can generate control events. A scheme with fewer restrictions will be harder to implement but give a lower cost J. As $p_{\text{inactive}} = J_u T \to 0$ and events become rare, sporadic control should approach aperiodic since T becomes negligible. When $\rho \to 0$, sporadic control with discrete measurements and $T_s = T$ will approach periodic since there remains no incentive to omit an event.

When $a < 0$, J_x and therefore J is bounded by the variance achieved in open loop. As ρ increases, all controllers will generate fewer events so that $J_u \to 0$, and ultimately J_x will approach a maximum. The limit can be found from (14), where $J_x \to -1/2a$ as $T \to \infty$.

4.3 Comparison

The trade-off between state variance and average event frequency is made explicit in Figure 4, where J_x is plotted against J_u for the four controllers. The results for $\sigma \neq 1$ are found by scaling J_x by σ^2. It is seen that the controllers are strictly ranked in performance by how much freedom they have to generate events, and that the sporadic controller with discrete measurements always outperforms the periodic one.

Figure 4 also shows what we consider the main advantage of event-based control: fewer events are needed for the same state cost. With periodic control, the variance increases quite rapidly with lower sampling rate. However, with sporadic control the average control rate can be reduced much further without the same penalty. For example, when $a = 0.5$ the average control rate may be decreased by about 40 % for only slightly more variance, using sporadic control with discrete measurements.

A notable result is that for sporadic control with continuous measurements, J_x can be made somewhat smaller *with fewer events*. This is also seen in the upper plot of Figure 2, where J_x attains a minimum for $r > 0$. Apparently, there is a hidden cost in issuing a control event, due to the

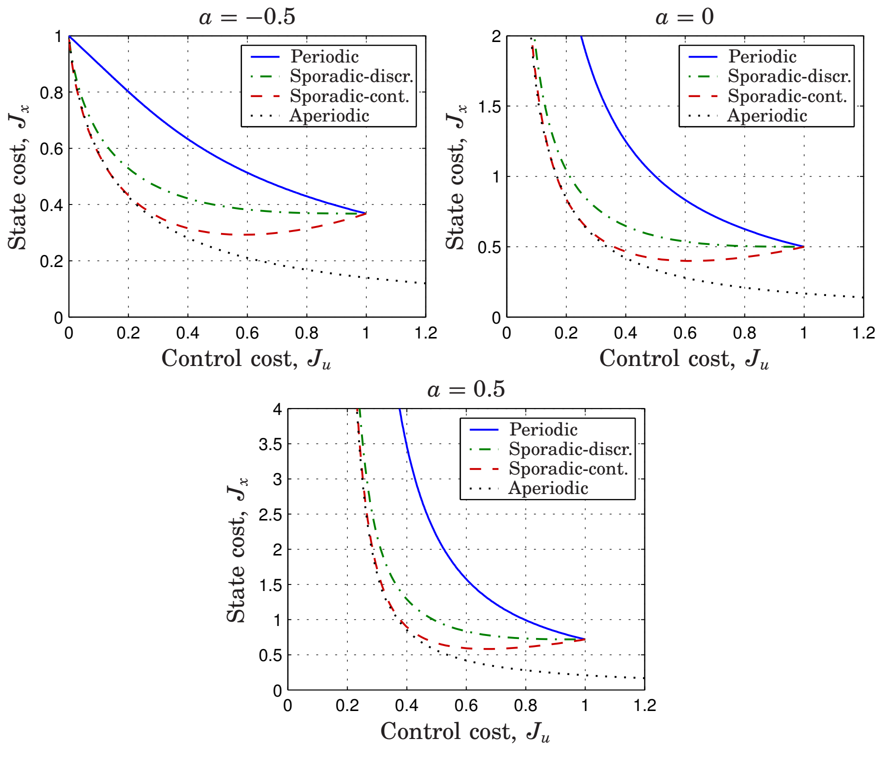

Figure 4. Trade-off between state cost J_x and control cost J_u for the four classes of controllers. Note the different vertical scales.

risk that large state errors will arise while in the inactive state. This phenomenon is absent for discrete measurements and $T_s = T$ since in this case events are generated independent of past actions.

Figure 5 shows the optimal achievable cost J^* for the four controllers. It is notable that for the stable system $a = -0.5$ the optimal periodic controller chooses to never sample when $\rho \geq 1$, while the sporadic controllers just raise their thresholds and remain ready to deal with large disturbances.

5. Higher-Order Systems

So far, we have only considered first-order systems. When raising the state dimension, there are many different generalizations worthy of study, depending on which is the constraining resource that motivates using

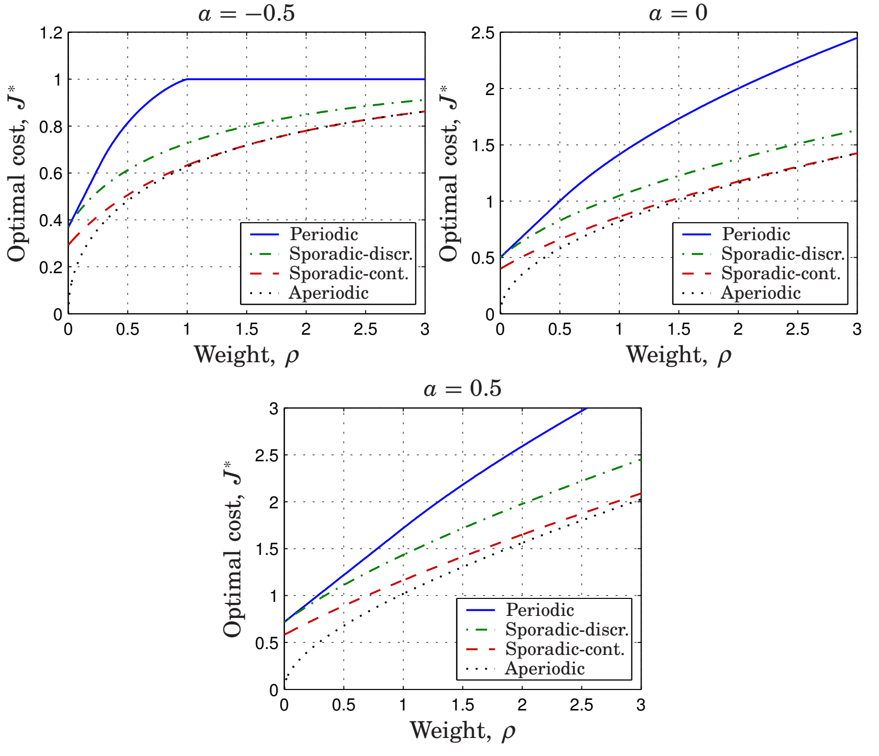

Figure 5. Optimal achievable cost J^* for the four classes of controllers. Note the different vertical scales.

event based control. We will briefly discuss some possibilities.

5.1 Formulations

The dynamics and cost J_x are naturally extended to

$$dx = Axdt + Budt + dw,$$
$$J_x = \limsup_{t\to\infty} \mathrm{E}\, \frac{1}{t}\int_0^t x^T Q x ds.$$

where now x, w and possibly u are vectors. One natural generalization of the measurement equation is

$$dy = Cxdt + dw_m,$$

where dw_m is measurement noise.

The possible forms of the controller, actuators, and sensors are more varied. Some scenarios are:

Communication constraints. Events are packets sent over a communication channel, from an observer at the sensor to a controller at the actuator. The observer decides when to send a state estimate to the controller, which predicts the plant state in open loop in between. Each event resets the prediction error.

Actuator constraints. The actuator only generates pulses of certain shapes, with some cost per pulse. The controller plans for an optimal and possibly long sequence of pulses, which is sensitive to timing.

Sensor constraints. The sensor only gives measurements under some conditions, e.g. at or beyond some thresholds. The control problem becomes a state estimation problem with nonstandard measurement information, for which the Kalman Filter is not optimal.

Processing power constraints. A simple control law is needed. The best bet is probably to postulate one and optimize over a few parameters.

We can consider a single control loop, or multiple loops sharing the same limited resource. The loops can be independent, or cooperate to control a single plant. It seems unreasonable that the controllers should know each others' state, especially with communication constraints.

5.2 Methods

The discretizations applied in this paper can be generalized to higher state dimensions, but become impractical beyond a few states due to the curse of dimensionality. Sometimes the dimension can be reduced somewhat; e.g. if the state is estimated with a stationary Kalman Filter, the distribution of the actual state is known conditioned on the estimate. Otherwise, nonlinear process dynamics come at a modest additional cost. Optimal stochastic control is in principle applicable to both the communication and actuation constrained scenarios.

Beyond a few states, simpler formulations are necessary for a solvable problem. This may include reducing the amount of uncertainty in the problem. A communication constrained problem easily becomes pointless with too little uncertainty, while an actuator constrained problem may still preserve its major features.

6. Conclusions

In some applications there is a cost related to the execution of a control signal, regardless of the magnitude of that signal. If that cost is included in the performance objective of the controller, it will be meaningful to reduce the frequency of control actions. This may be accomplished with a periodic controller by lengthening the sampling interval. However, the penalty in terms of increased process state variance is significant. Trying to improve the tradeoff by not acting on small state errors naturally leads to the notion of event-based control.

In this paper, we have shown that sporadic control can provide a better tradeoff between control objectives as well as better overall control performance than periodic control, when there is a fixed cost of control actions. It is noted that the average frequency of control events can be reduced with only a small increase in state variance. Moreover, we show that sporadic control can actually reduce both the average frequency of control events and the state variance simultaneously. When the objective is to reduce the frequency of events as well as the state variance, the sporadic control schemes presented here even perform almost as well as aperiodic control, while respecting a prespecified shortest inter-event time.

Event-based control has an additional threshold parameter that should scale with the size of disturbances. If they are bigger than expected, the control approaches periodic control. If they are smaller, the theshold will act as a tolerable margin of error. Both responses are reasonable in the face of a mismatched disturbance intensity.

Obviously, to implement sporadic control where periodic control is currently used requires some changes. Unless the hardware supports continuous measurements, discrete measurements are an easier option and approach the continuous performance quite fast if one can measure more often than control. The change from periodic to sporadic control with the same measurement and control interval should require minimal modifications.

References

Åström, K. J. (1970): *Introduction to Stochastic Control Theory*. Academic Press, New York.

Åström, K. J. and B. Bernhardsson (1999): "Comparison of periodic and event based sampling for first-order stochastic systems." In *Preprints of the 14th World Congress of IFAC*, Beijing, P.R. China.

Åström, K. J. and B. Wittenmark (1997): *Computer-Controlled Systems*. Prentice Hall.

Buttazzo, G. C. (1997): *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. Kluwer Academic Publishers.

Cogill, R., S. Lall, and J. P. Hespanha (2006): "A constant factor approximation algorithm for event-based sampling." In *Proceedings of the 44th Annual Allerton Conference on Communication, Control, And Computing*, Monticello, Illinois.

Feller, W. (1971): *An Introduction to Probability Theory and Its Applications*, vol. 2. John Wiley & Sons.

Hu, J., J. Lygeros, and S. Sastry (2000): "Towards a theory of stochastic hybrid systems." In *Proceedings of the Third International Workshop on Hybrid Systems: Computation and Control*, Pittsburgh, PA.

Imer, O. C. and T. Basar (2005): "Optimal estimation with limited measurements." In *Proceedings of the 44th IEEE Conference on Decision and Control*, Seville, Spain.

Rabi, M. (2006): *Packet based Inference and Control*. PhD thesis, Institute for Systems Research, University of Maryland. `http://hdl.handle.net/1903/3970`.

Paper IV

Scheduling of Event-Triggered Controllers on a Shared Network

Anton Cervin   Toivo Henningsson

Abstract

We consider a system where a number of independent, time-triggered or event-triggered control loops are closed over a shared communication network. Each plant is described by a first-order linear stochastic system. In the event-triggered case, a sensor at each plant frequently samples the output but attempts to communicate only when the magnitude of the output is above a threshold. Once access to the network has been gained, the network is busy for T seconds (corresponding to the communication delay from sensor to actuator), after which the control action is applied to the plant. Using numerical methods, we compute the minimum-variance control performance under various common MAC-protocols, including TDMA, FDMA, and CSMA (with random, dynamic-priority, or static-priority access). The results show that event-triggered control under CSMA gives the best performance throughout.

1. Introduction

Networked feedback control systems are normally implemented using periodic sampling at the sensor nodes, combined with either time-triggered or event-triggered communication between the sensor, controller, and actuator nodes. Periodic sampling allows for standard sampled-data control theory (e.g. [Åström and Wittenmark, 1997]) to be used, although network-induced delay and jitter may limit the performance [Cervin *et al.*, 2003].

In recent work [Åström and Bernhardsson, 1999; Hristu-Varsakelis and Kumar, 2002; Rabi, 2006; Johannesson *et al.*, 2007], event-triggered sampling has been proposed as a means for more efficient resource usage in networked control. The basic idea is to sample, communicate, and control only when something significant has occurred in the system. For first-order stochastic systems, it has been shown that event-based sampling can significantly reduce the output variance and/or the average control rate compared to periodic sampling [Åström and Bernhardsson, 1999]. A similar idea is to introduce a deadband in the sensor. The trade-off between network traffic and control performance for higher-order control loops with deadband sampling was studied in [Otanez *et al.*, 2002].

When multiple control loops are closed over a shared medium (like a communication bus or a wireless local-area network), a multiple access method such as TDMA (time division multiple access), FDMA (frequency division multiple access), or CSMA (carrier sense multiple access) is needed to multiplex the data streams. It is clear that the choice of access method can have a great impact on the control performance. Intuitively, TDMA should be suitable for time-triggered control loops, while CSMA, being a random-access method, would seem to be well suited for event-based control. FDMA provides a way to share the bandwidth without regard to synchronization among the loops, which could potentially be beneficial for both time-triggered and event-triggered control. At the same time, less bandwidth per control loop means longer transmission times and hence longer feedback delays.

Multi-loop networked control systems—taking into account issues such as clock synchronization, medium access, communication protocols, imperfect transmissions, delay and jitter, and event-triggered sampling, as well as the control algorithms themselves—are very complex systems. To facilitate analysis, great simplifications are needed. In this paper, we study a scenario where a number of independent control loops are closed over a shared network (see Fig. 1). Using very simple models for the plants, controllers, and network arbitration, we are able to numerically compute and compare the minimum-variance control performance under the various medium access protocols. In particular, we apply recent results in

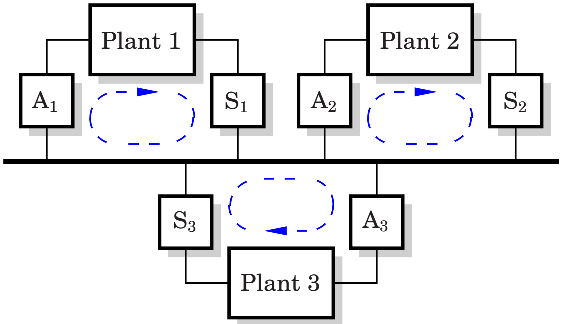

Figure 1. Multiple control loops are closed over a shared communication medium. The controller in each loop may be co-located with either the sensor (S) or the actuator (A).

sporadic event-based control of first-order systems [Johannesson *et al.*, 2007; Cervin and Johannesson, 2008] to model and analyze the interaction between control loops and medium-access schemes. Although far from an exhaustive study, the results offer some interesting insight into the suitability of the studied MAC-protocols for networked control.

The remainder of this paper is outlined as follows. In Section II, the system description is given. Section III reviews how to calculate the stationary variance under time-triggered and event-triggered sampling. In Section IV, we model the medium access schemes and describe the co-design problem associated with each scheme. Section V reports numerical results for symmetrical integrator plants. In Section VI, we digress and compare the achievable performance under global vs local scheduling decisions. Section VII contains a case study with three asymmetric plants. Finally, the conclusions are given in Section VIII.

2. System Description

We consider a system where N control loops are closed over a shared network. Each plant $i \in 1 \ldots N$ is described by a first-order stochastic differential equation

$$dx_i(t) = a_i x_i(t)dt + u_i(t)dt + \sigma_i dw_i(t), \quad x_i(0) = 0, \tag{1}$$

where x_i is the state, a_i is the process pole, u_i is the control signal, w_i is a Wiener process with unit incremental variance, and $\sigma_i > 0$ is the intensity of the noise. All noise processes are assumed independent.

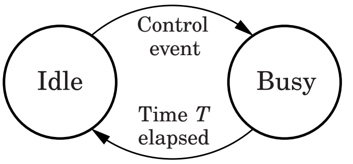

Figure 2. Network state transitions. Control events may only be generated in the idle state.

A sensor located at each plant i takes samples of the plant state at certain discrete time instants $\{t_{i,k}\}_{k=0}^{\infty}$:

$$x_{i,k} = x_i(t_{i,k}). \tag{2}$$

The sampling can be either time-triggered or event-triggered, depending on the medium access scheme. After obtaining a sample, the sensor tries to initiate a control event by transmitting the value to the actuator. The network is however a shared resource that only one control loop may access at a time[3]. If two or more sensors attempt to transmit at the exact same time, a resolution mechanism determines who will gain access to the network. (The other nodes will simply discard their samples.) Once access has been gained, the network stays occupied for T seconds, corresponding to the transmission delay from sensor to actuator. During this interval, no new control events may be generated (see Fig. 2).

The controller in each loop may be co-located with either the sensor or the actuator; the network delay is assumed constant and known, so it does not matter which. The overall goal is to minimize the total cost

$$J = \sum_{i=1}^{N} J_i, \tag{3}$$

where the performance of loop i is measured by the stationary state variance

$$J_i = \lim_{t\to\infty} \frac{1}{t} \mathrm{E} \int_0^t (x_i(s))^2 ds. \tag{4}$$

In response to a sample taken at time $t_{i,k}$, the actuator is allowed to emit a Dirac pulse of size $u_{i,k}$. It is clear (see [Cervin and Johannesson, 2008]) that minimum variance is achieved by driving the expected value of the state at time $t_{i,k} + T$ to zero, implying the deadbeat control law

$$u_{i,k} = -e^{a_i T} x_{i,k}. \tag{5}$$

[3]This is not true under FDMA. Under FDMA, we rather assume that each control loop has access to its own private network with lower bandwidth.

The control signal generated by actuator i is hence given by the pulse train

$$u_i(t) = \sum_{k=0}^{\infty} \delta\big(t - t_{i,k} - T\big) u_{i,k}. \tag{6}$$

While it may seem unrealistic to allow Dirac controls, it allows for a fair and straightforward comparison between time-triggered and event-triggered control. The Dirac pulse may be replaced by an arbitrary pulse shape of length no longer than T at the expense of slightly more complicated cost calculations.

3. Evaluation of Cost

We here briefly review how to compute the cost (4) under time-triggered and event-triggered sampling with a delay and minimum inter-event interval T. For more details, see [Åström, 1970; Johannesson *et al.*, 2007; Cervin and Johannesson, 2008]. For clarity, we here drop the plant index i.

3.1 Time-Triggered Sampling

Under time-triggered sampling, the stationary variance (4) can be calculated analytically. The sampling instants t_k are known a-priori and do not depend on the plant state, which will be normal distributed at all times. The (possibly irregularly) sampled closed-loop system becomes

$$x_{k+1} = w_k, \tag{7}$$

where $\{w_k\}_{k=0}^{\infty}$ are independent, zero-mean Gaussian variables with variance $P(t_{k+1} - t_k)$, where

$$P(t) = \begin{cases} \sigma^2 \dfrac{e^{2at} - 1}{2a}, & a \neq 0, \\ \sigma^2 t, & a = 0. \end{cases} \tag{8}$$

(Note that the delay does not affect the state distribution at the sampling instants.) Sampling the cost function gives

$$\mathrm{E} \int_{t_k}^{t_{k+1}} x^2 ds = Q(T)\, \mathrm{E}(x_k)^2 + J_v(t_{k+1} - t_k), \tag{9}$$

where

$$Q(T) = \begin{cases} \frac{e^{2aT} - 1}{2a}, & a \neq 0, \\ T, & a = 0 \end{cases} \tag{10}$$

is the state weight due to delay, while

$$J_v(t) = \begin{cases} \frac{e^{2at}-2at-1}{4a^2}, & a \neq 0, \\ \frac{t^2}{2}, & a = 0 \end{cases} \tag{11}$$

accounts for the inter-sample noise (see, e.g. [Åström, 1970]). Finally, we know that $\mathrm{E}\, x^2(t_k) = P(t_k - t_{k-1})$. Using the expressions above, it is straightforward to evaluate the cost under any static cyclic schedule.

3.2 Event-Triggered Sampling

Under event-triggered sampling, control events may only be generated when the network is idle and $|x(t)| \geq r$, where r is the event detection threshold. The state will no longer be Gaussian, which complicates the calculation of $\mathrm{E}\, x^2(t_k)$. A useful and realistic approximation is to assume that the sensor does not measure x continuously, but rather uses fast sampling with the interval $T_s \ll T$. The (irregularly) sampled closed-loop system then becomes

$$x_{k+1} = \begin{cases} e^{aT_s} x_k + w_k(T_s), & |x_k| < r \\ w_k(T), & |x_k| \geq r \ \ \&\ \text{won} \\ e^{aT} x_k + w_k(T), & |x_k| \geq r \ \ \&\ \text{lost} \end{cases} \tag{12}$$

where $\{w_k(t)\}_{k=0}^{\infty}$ is a sequence of independent, zero-mean Gaussian variables with variance $P(t)$; "won" means that the sensor node won the network arbitration, while "lost" means the opposite. Letting the system run in open loop between the fast samples, the expressions (8)–(11) for the sampled cost are still valid.

The update equation (12) is useful both for calculation of the state distribution and for Monte Carlo simulations. Because of the shared medium, the stationary probability distributions of $x_1, \ldots, x_N$ are not independent. To evaluate the cost using the first approach, it is hence necessary to find the multi-dimensional probability distribution $f(x_1, \ldots, x_N)$. This can in theory be done by gridding the state space and then iterating the distribution according to (12) until convergence. In practice, this can be done for a few dimensions, forcing us to rely on Monte Carlo simulations for $N \geq 3$ in this paper.

4. Medium Access Schemes and Control Policies

In this section, we present simple scheduling and control models for three medium access schemes and discuss how to derive optimal schedules and control policies.

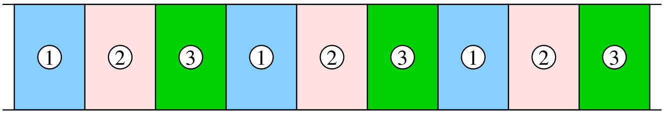

Figure 3. Time division multiple access (TDMA). A static cyclic schedule determines which sensor node samples and transmits in which time slot.

4.1 TDMA (Time Division Multiple Access)

In TDMA (see Fig. 3), a cyclic access schedule is determined off-line. In each slot in the schedule, one control loop has access to the network for T seconds. Since there is no cost associated with using the network in our problem formulation, it is obvious that no slot should be left empty, and that the sensor should always sample and transmit in its slot. Hence, the optimal control scheme associated with TDMA will be a pure time-triggered scheme.

For symmetric plants (with $a_i = a$, $\sigma_i = \sigma$, $\forall i$), a simple round-robin schedule is optimal. For asymmetric plants, an optimal schedule of length n can be found by evaluating the resulting cost for each possible schedule. (The search for an optimal schedule can be done more efficiently. The LQ-optimal cyclic scheduling and control problem for multiple higher-order plants is treated in [Rehbinder and Sanfridson, 2004].)

4.2 FDMA (Frequency Division Multiple Access)

In FDMA (see Fig. 4), the communication bandwidth is divided between the nodes, such that each loop receives a fixed fraction U_i of the total capacity $\sum_{i=1}^{N} U_i = 1$. Accounting for the lower transmission rate, the delay from sensor i to actuator i is now T/U_i.

It is previously known [Johannesson *et al.*, 2007] that event-triggered sampling with a minimum inter-event interval T is superior to time-triggered sampling with the interval T, also when there is delay in the system. Hence, event-triggered control is the better choice for FDMA. The optimal event detection threshold and the associated optimal cost can be

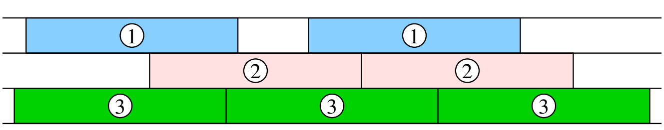

Figure 4. Frequency division multiple access (FDMA). The bandwidth is divided into fixed shares, giving each loop a dedicated channel. Within each share, an event-triggered control loop is implemented.

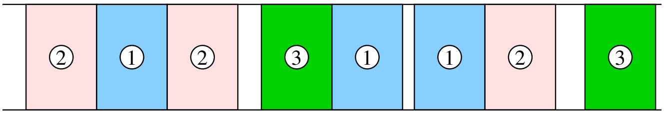

Figure 5. Carrier sense multiple access (CSMA). Each loop is event-triggered. A static, dynamic, or random priority function determines who will transmit if many nodes try to access the network at the same time.

found numerically by sweeping r and computing the cost for each value.

For symmetric plants, an even division of the bandwidth is optimal. For asymmetric plants, the shares U_i can be found using optimization. Since the cost functions $J_i(U_i)$ are smooth and strictly decreasing, it is feasible to use standard nonlinear optimization tools to find the shares.

4.3 CSMA (Carrier Sense Multiple Access)

In CSMA (see Fig. 5), any node may try to access the network as soon as it becomes idle, making it suitable for event-triggered control loops. If many nodes want to transmit at the same time, some resolution mechanism must be used. In shared-medium Ethernet for instance, the collision detection and random back-off strategy will grant a random node access to the network (after some delay). In the Controller Area Network (CAN) on the other hand, access can be resolved based on either fixed (node) priorities or dynamic (message) priorities.

We will consider three different resolution mechanisms:

Random (CSMA-rand). As in Ethernet or WLAN, a random node will eventually win the contention. For simplicity, it is assumed that the resolution time is very small compared to the transmission time so that it can be neglected. The overall performance is optimized by selecting suitable event detection thresholds for the control loops. This is done by sweeping r_i and computing the cost for each value.

Static priority (CSMA-statprio). Each sensor node is assigned a static priority, which determines who will win the arbitration. Such a scheme can be useful for asymmetric plants where it is known that some plants are more sensitive to long access delays than others.

Dynamic priority (CSMA-dynprio). For symmetric first-order plants, it can make sense to use the control error as a dynamic priority. (This idea was put forth in [Walsh *et al.*, 1999], where it was called the Maximum-Error-First (MEF) scheduling technique.) It is assumed that the network interface provides a mechanism (such as message priorities in CAN) so that priority access can be given to the node with the largest control error. It is obvious that this scheme will be better than random priorities. Again, the overall performance is optimized by selecting event thresholds for the loops.

5. Results for Symmetric Integrator Plants

We here present numerical results for N symmetric integrator plants with $a_i = 0$ and $\sigma_i = 1$. We assume that the network bandwidth scales in proportion to the number of plants, such that the transmission delay from sensor to actuator is $T = 1/N$ when the full bandwidth is utilized. For the numerical computations, we assume fast sampling with $T_s = T/100$.

Under TDMA, the optimal cyclic transmission schedule is $\{1, 2, \ldots, N\}$. The sampling period of each loop is 1 and the delay is $T = 1/N$, giving the following exact value for the cost per loop:

$$J_i = \left(J_v(T) + Q(T)\,\mathrm{E}\,x^2(t_k)\right)/T = \frac{1}{2} + \frac{1}{N}. \tag{13}$$

Under FDMA, each loop receives a share $U_i = 1/N$ of the bandwidth, implying the same performance regardless of the number of nodes. Computing the stationary state distribution under event-triggered sampling for different values of r, we find the optimal threshold $r = 1.06$, yielding the cost

$$J_i = 1.40. \tag{14}$$

For the CSMA case, we use Monte Carlo simulations to find the stationary variance of the plants under random or dynamic priority access. For each N, we sweep r to find the optimal threshold and the corresponding optimal cost. Each configuration was simulated for 10^8 time steps, corresponding to in the order of 10^6 simulated seconds. (The simulation time was around $15n$ seconds for each configuration on an Intel Core 2 CPU @1.83 GHz.)

The optimal costs under the various policies described above for $N = 1 \ldots 10$ nodes are reported in Fig. 6, and the optimal thresholds under CSMA are shown in Figs. 7. It is seen that TDMA outperforms FDMA, except for $N = 1$ where sporadic event-based control has the edge over periodic control. In turn, both variants of CSMA outperform TDMA, CSMA with dynamic priorities performing slightly better than CSMA with random access. The results are not surprising, since CSMA with event-triggered sampling dynamically allocates the bandwidth to the loop(s) most in need. A higher event threshold is needed for the random priority scheme in order to be more selective about which plant to control.

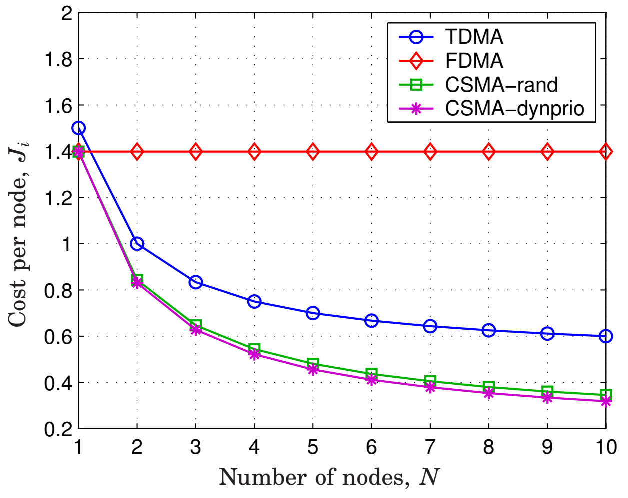

Figure 6. Optimal cost per node vs number of nodes when controlling symmetric integrator plants.

It is possible to reason about what happens when $N \rightarrow \infty$ under the various access schemes. Under TDMA, the performance approaches $J_i = 1/2$, while under FDMA, the performance is unaffected by N and is constant $J_i = 1.40$. CSMA approaches aperiodic event-based control [Åström and Bernhardsson, 1999] when $N \rightarrow \infty$, regardless of the priority scheme used. For integrator plants, the optimal cost per plant approaches $J_i = 1/6$. Hence, CSMA asymptotically gives 67% lower cost than TDMA and 88% lower cost than FDMA when the number of control loops increases. Equivalently, one can reason about the network capacity needed to maintain the same performance as the number of integrator plants grows. Here, again, CSMA will asymptotically require 67% less bandwidth than TDMA and 88% less bandwidth than FDMA to achieve the same cost per loop.

6. Local vs Global Knowledge

One important assumption in our model is that the decisions as to whether to transmit or not are taken locally at each sensor node. It was seen above that event-triggered control under CSMA with dynamic priority access gave the lowest cost among all the considered schemes. It is interesting to compare the performance to a controller with global knowledge of the plant states. Such a controller would of course not be implementable in a

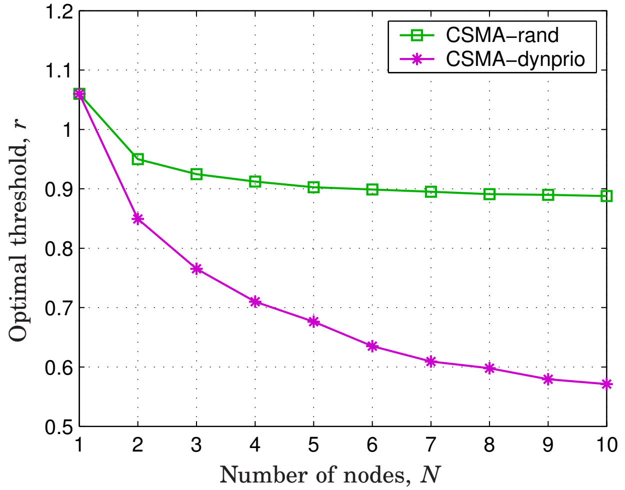

Figure 7. Optimal threshold vs number of nodes for CSMA with random or dynamic priority access when controlling symmetric integrator plants.

networked setting but can provide a lower bound on the achievable cost.

We consider the special case of $N = 2$ symmetric integrator plants with the minimum inter-control interval and delay $T = 1/2$. The optimal local scheme under CSMA with dynamic priorities was computed above, giving the optimal cost $J_i = 0.834$ for the threshold $r = 0.85$. For the global scheme, we gridded the plant state space in the two dimensions and applied dynamic programming to derive the optimal control policy. For each state (x_1, x_2), the controller has the choice to control to the first plant, the second plant, or to idle. The resulting optimal global control policy is shown in Fig. 8, together with the local CSMA policy with dynamic priorities. It is seen that the control policies are quite similar. One difference is that the global controller will idle if both plants have about the same error magnitude, waiting to see where the processes will go next. The resulting cost under the global policy is found to be $J_i = 0.828$, which is only one percent lower than the cost for the optimal local scheme.

7. Results for Three Asymmetric Plants

As a final numerical example, we consider a case where three asymmetric first-order systems should be controlled: one asymptotically stable plant, one integrator, and one unstable plant. The plant parameters are $\sigma_i = 1$

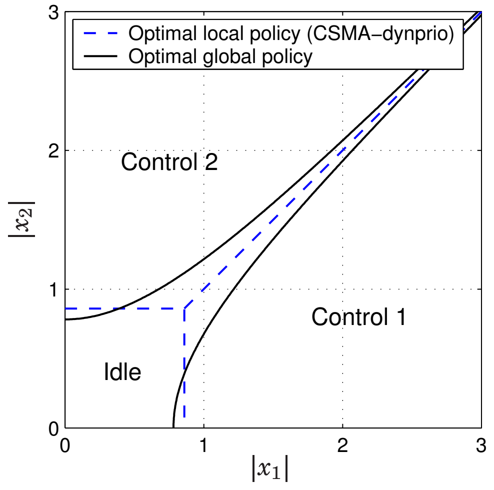

Figure 8. Event-triggered control of two integrators: optimal global and local policies.

and

$$a_1 = -0.5, \qquad a_2 = 0, \qquad a_3 = 0.5.$$

Further, we let $T = 1/3$. Here, intuition tells us that more resources should be allocated to the unstable plant (Plant 3) while the stable plant (Plant 1) can manage with less resources.

For TDMA, the total cost is computed for all possible cyclic schedules of length $n = 2, \ldots, 12$. Since the unstable plant must be controlled at least once per cycle, we fix the first entry in the schedule to 3, leaving about 3^{n-1} schedules to test per value of n (including "necklace duplicates"). The optimal schedule for each value of n is reported in Table 1. It is seen that the best schedule is of length 6: $\{3, 2, 3, 2, 3, 1\}$, giving a total cost of $J = 2.56$. In the optimal schedule, the stable plant is controlled once per cycle, the integrator twice, and the unstable plant three times per cycle.

For FDMA, we optimize over the bandwidths U_1, U_2, U_3 to find the lowest total cost. For each plant, we first approximate the cost function $J_i(U)$ by sweeping r for each value of U. We then apply nonlinear optimization to find the optimal shares, yielding $U_1 = 0$, $U_2 = 0.397$, $U_3 = 0.603$ and the total cost $J = 3.49$. It is interesting to note that the long delay associated with FDMA apparently makes it pointless to control the stable plant.

For CSMA, we consider two arbitration mechanisms: random access and static priorities. For the random access scheme, we sweep the three thresholds to find the minimum cost, giving $r_1 = 1.12$, $r_2 = 0.92$, $r_3 = 0.77$,

Table 1. Optimal cyclic schedules for the three asymmetric plants.

Length n	Cyclic schedule	Total cost J
2	$\{3,2\}$	2.651
3	$\{3,3,2\}$	2.708
4	$\{3,2,3,1\}$	2.588
5	$\{3,2,3,2,1\}$	2.650
6	$\{3,2,3,2,3,1\}$	2.563
7	$\{3,2,3,3,2,3,1\}$	2.589
8	$\{3,2,3,2,3,2,3,1\}$	2.567
9	$\{3,2,3,3,2,3,2,3,1\}$	2.591
10	$\{3,2,3,2,3,1,3,2,3,1\}$	2.573
11	$\{3,2,3,3,2,3,1,3,2,3,1\}$	2.588
12	$\{3,2,3,2,3,1,3,2,3,2,3,1\}$	2.563

Table 2. Optimal costs for the three asymmetric plants under the various medium access schemes.

Scheme	J_1	J_2	J_3	$J = \sum J_i$
TDMA	0.690	0.889	0.984	2.56
FDMA	1.000	1.177	1.319	3.49
CSMA-rand	0.554	0.618	0.772	1.94
CSMA-statprio	0.562	0.641	0.723	1.92

and the total cost $J = 1.96$. The three loops occupy the network on average 14%, 22%, and 38% of the time, while it is idle 26% of the time. The relative shares for the loops are not that different from the ones generated by the optimal cyclic schedule.

For the static priority CSMA case, we assume that the unstable plant has the highest priority, the integrator has medium priority, while the stable plant has the lowest priority. Again sweeping the three thresholds and evaluating the costs gives the optimal thresholds $r_1 = 0.95$, $r_2 = 0.87$, $r_3 = 0.77$, and the total cost $J = 1.94$. The priorities allow for tighter thresholds to be utilized. The three loops occupy the network on average 15%, 25%, and 38% of the time, while it is now idle 22% of the time.

The results under the various access schemes are summarized in Table 2. We can again conclude that CSMA can provide better control performance than both TDMA and FDMA. For this example, CSMA gives 23% percent lower total cost than TDMA and 44% lower cost than FDMA.

We further note that there is only a very modest improvement by using priorities, which is good news for wireless systems where random access schemes may be the only realistic choice for the implementation.

8. Discussion and Conclusion

This paper has studied a prototypical networked control co-design problem, where both the control policy and network scheduling policy have been taken into account. Although very simple mathematical models were used, some interesting conclusions regarding the various medium access schemes could be drawn. CSMA with event-triggered sampling was the superior scheme in all presented examples, while FDMA performed poorly due to the long transmission delay.

The simulation-based design approach adopted in this paper is conceptually easy to extend to higher-order plants and controllers. We have noted that the simulation time required to evaluate the cost with a given accuracy grows slower than the number of states in the system. Rather, the main problem with more realistic systems is the number of controller parameters that need to be optimized. For higher-order systems, it is probably necessary to impose restrictions on the controller structure and only optimize over a small subset of the parameters.

Another interesting approach would be to develop a way to characterize the performance of an event-triggered control loop as a function of its network resource usage pattern. Integrating several control loops, it should be possible to provide guarantees on the worst-case performance of each controller. Apart from higher-order plants and controllers, several other extensions to the work in this paper are possible to imagine, including

- having the controller located in a separate node, meaning that both the transmission from sensor to controller and from controller to actuator need to be scheduled.

- having more detailed models of real network protocols, including, e.g., the random back-offs in CSMA/CD.

- allowing MIMO systems, where each sensor and actuator may reside on a different node in the network.

- modeling measurement noise, variable transmission times, and lost packets.

References

Åström, K. J. (1970): *Introduction to Stochastic Control Theory*. Academic Press, New York.

Åström, K. J. and B. Bernhardsson (1999): “Comparison of periodic and event based sampling for first-order stochastic systems.” In *Preprints of the 14th World Congress of IFAC*, Beijing, P.R. China.

Åström, K. J. and B. Wittenmark (1997): *Computer-Controlled Systems*. Prentice Hall.

Cervin, A., D. Henriksson, B. Lincoln, J. Eker, and K.-E. Årzén (2003): “How does control timing affect performance?” *IEEE Control Systems Magazine*, **23:3**, pp. 16–30.

Cervin, A. and E. Johannesson (2008): “Sporadic control of scalar systems with delay, jitter and measurement noise.” In *Proceedings of the 17th IFAC World Congress*. Seoul, Korea.

Hristu-Varsakelis, D. and P. Kumar (2002): “Interrupt-based feedback control over a shared communication medium.” In *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vegas, Nevada.

Johannesson, E., T. Henningsson, and A. Cervin (2007): “Sporadic control of first-order linear stochastic systems.” In *Proceedings of the 10th International Conference on Hybrid Systems: Computation and Control*, Pisa, Italy.

Otanez, P. G., J. R. Moyne, and D. M. Tilbury (2002): “Using deadbands to reduce communication in networked control systems.” In *Proceedings of the 2002 American Control Conference*, Anchorage, Alaska.

Rabi, M. (2006): *Packet based Inference and Control*. PhD thesis, Institute for Systems Research, University of Maryland. `http://hdl.handle.net/1903/6612`.

Rehbinder, H. and M. Sanfridson (2004): “Scheduling of a limited communication channel for optimal control.” *Automatica*, **30:3**, pp. 491–500.

Walsh, G. C., H. Ye, and L. Bushnell (1999): “Stability analysis of networked control systems.” In *Proceedings of the 1999 American Control Conference*, San Diego, California.

Paper V

A Simple Model for the Interference Between Event-Based Control Loops Using a Shared Medium

Toivo Henningsson   Anton Cervin

Abstract

Traditionally, control loops are closed using periodic sensing and actuation. When communication resources are scarce, much may be gained by instead transmitting only when something important has happened in the loop. However, there are no known closed form solutions to this kind of control problem. This paper presents a simple model of the interference between event-based control loops caused by sharing a common medium, based on approximating the behavior of all loops except one foreground loop. The stationary state distribution can be computed at low computational cost using mostly standard linear time-invariant system theory (applied in the spatial dimension). Control laws are optimized to minimize state variance using the limited communication resources. Comparison to Monte Carlo simulations of a full model shows the simple model to be remarkably accurate. The model is applied to investigate how the performance of N control loops sharing a common Carrier Sense Multiple Access channel approaches the ideal case of aperiodic control as the number of loops grows.

1. Introduction

Event-based control holds the promise of better control performance and lower resource consumption compared to standard sampled-data control. In their seminal paper on event-based control, Åström and Bernhardsson (Å&B) [Åström and Bernhardsson, 1999] showed that, for an integrator process driven by white noise, event-triggered control requires only one third of the number of samples to achieve the same output variance as periodic, time-triggered control. However, the reduction comes at the cost of more irregular events. In fact, Å&B's *aperiodic* controller would require infinite bandwidth to be implemented in a networked control setting.

Aiming for implementable controllers, we have previously proposed the concept of *sporadic control* [Henningsson *et al.*, 2008], where a minimum inter-event time T is enforced. The parameter T can be used to model that the communication channel stays busy for some time when a packet is transmitted.

This paper explores how the ideal performance of Å&B can be approached by letting N sporadic control loops share a communication medium with limited bandwidth. As the number of loops grows, it is expected that the medium can be used more and more efficiently, essentially transforming the sporadic constraint into a constraint on the average communication rate.

A major theoretical challenge in event-based control of stochastic systems is to find the stationary probability distribution of the state. For low-order systems, gridding of the state-space may be used, but this quickly becomes infeasible for higher-order systems. In a previous study [Cervin and Henningsson, 2008], we used Monte Carlo simulations to numerically evaluate the performance for a modest number of sporadic controllers using a shared medium.

By contrast, in the current paper we derive a simple model that allows rapid evaluation of the performance, even as the number of loops becomes large. The control loops are partitioned into one *foreground loop*, and $N-1$ *background loops*. Only the state probability distribution of the foreground loop and the discrete state of the medium are modelled in detail. When the resulting problem has been solved, key model parameters are matched to make the behavior of the foreground and background loops appear equal.

Comparing with results from Monte Carlo simulations of the full system model, it is seen that the simple model is remarkably accurate in predicting the behavior of the full system, given its simplicity. The agreement becomes better as N grows, and, as N becomes very large, the ideal performance of Å&B is approached.

The medium access policy assumed in the paper is Carrier Sense Multiple Access (CSMA) with random, delay-free arbitration if several nodes attempt to transmit at the same time. The network is hence collision free—this is a key assumption for our results to hold. It is well known that the performance of ALOHA/CSMA with collisions and retransmissions breaks down when utilization approaches 1, e.g., [Tobagi and Hunt, 1980; Jelenković and Tan, 2007]. Collisions are highly relevant for wireless sensor/actuator networks and will be studied in future papers.

Related Work

Rabi and Johansson [Rabi and Johansson, 2009] analyzed how packet loss impacts the performance of event-triggered control loops. The analysis was done for a single networked loop, assuming that the loss rate was independent of the event-triggering threshold of the controller.

Event-triggered state estimation or control of higher-order systems has been studied by Cogill [Cogill *et al.*, 2006; Cogill, 2009], however not in the context of several competing control loops.

In all the works mentioned so far, the events are triggered by fixed thresholds in the state space. Alternative ways to trigger event-based controllers are proposed in [Tabuada, 2007; Wang and Lemmon, 2009].

Outline

The rest of the paper is laid out as follows: Section 2 presents the system model and the optimal control problem. The equations for the stationary state distribution are derived in Section 3 and solved in Section 4, by reducing the linear problem from infinite dimension to a small finite dimension that can be solved efficiently. Model parameters are adjusted in Section 5 to match key parameters between the foreground loop and the background loops. The influence of the packet length distribution is investigated in Section 6. Results and comparison to other models are presented in Section 7 and possible directions for future work are discussed in Section 8. Conclusions are given in Section 9.

2. Problem Formulation

This section presents the system model, which includes the foreground loop and medium state, the conditions to match the foreground and background loops, and the optimal control problem that we want to solve.

2.1 System Model

The process controlled in the foreground loop is modelled as an integrator disturbed by white noise according to

$$dy = udt + \sigma dw, \tag{1}$$

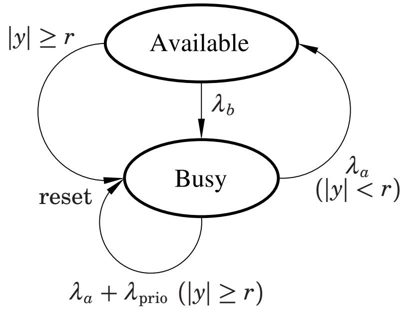

Figure 1. Continuous-time Markov chain model of the channel state dynamics under CSMA. The dynamics depends on the integrator state y: the reset transitions are triggered only when $|y| \geq r$; when $|y| < r$ there is instead a transition to the Available state.

where w is a Wiener process with unit incremental variance. The control signal u will be zero except when the foreground loop is able to generate a control event, at which time it will contain a Dirac impulse.

The control law is simple: whenever $|y| \geq r$ for some threshold r, the controller tries to take the channel and transmit a new packet. If it succeeds, the state y is reset to zero immediately[4].

The CSMA channel model is summarized in Figure 1:

- Whenever the channel is Available, anyone may transmit, causing a transition into the Busy state. The foreground loop will do so when the integrator state reaches the threshold $|y| \geq r$; the background loops will do so at an expected rate λ_b.
- The transmission of a packet takes an expected time $T = \lambda_p^{-1}$. A new packet may then be initiated immediately, or the channel will become available, which will happen at a rate $\lambda_a \leq \lambda_p$ when $|y| < r$.
- If several transmitters are contending for the channel when it is released, one is picked at random. The foreground loop will be able to gain the channel in the Busy state at a rate $\lambda_{\text{prio}} \leq \lambda_p - \lambda_a$ when it wants to (i.e., when $|y| \geq r$).

The state of the system is thus $(y, k) \in \mathbb{R} \times \mathbb{Z}_n$ ($n = 2$ in this case), where k is the state of the channel, so the stationary probability density function (pdf) $f(y)$ and probability flow in the y direction $\varphi(y)$ are

$$f(y) = \begin{pmatrix} f_a(y) \\ f_b(y) \end{pmatrix}, \qquad \varphi(y) = \begin{pmatrix} \varphi_a(y) \\ \varphi_b(y) \end{pmatrix}.$$

[4]It is trivial to account for a fixed delay in the control action, see [Cervin and Johannesson, 2008].

The control law imposes the linear constraints

$$f_a(y) = 0 \ \forall\, |y| \geq r, \qquad \varphi(0) = \begin{pmatrix} 0 \\ \varphi_{\text{fg}} \end{pmatrix},$$

where φ_{fg} is the rate of reset, since the channel never stays available when $|y| \geq r$, and all resets cause inflow into the Busy state at $y = 0$.

When $|y| < r$, the channel state will evolve according to

$$\begin{pmatrix} \dot{p}_a \\ \dot{p}_b \end{pmatrix} = \begin{pmatrix} -\lambda_b & \lambda_a \\ \lambda_b & -\lambda_a \end{pmatrix} \begin{pmatrix} p_a \\ p_b \end{pmatrix},$$

where p_a and p_b are the probabilities to be in the Available and Busy states respectively. When $|y| \geq r$, the channel state will evolve according to

$$\dot{p}_b = -(\lambda_a + \lambda_{\text{prio}})p_b;$$

the leakage corresponds to the rate of reset triggered from the Busy state.

2.2 Matching Conditions

We must require the foreground packet rate φ_{fg} to be a proportional share of the total packet rate $Nf_u = \lambda_p p_{\text{b}}$, i.e.

$$\lambda_p p_{\text{b}} = N\varphi_{\text{fg}}. \tag{2}$$

To model equal prioritization of all loops, we demand that in the Busy state with $|y| \geq r$, the intensity to gain the channel is the total packet rate λ_p divided by the expected number of loops waiting to gain the channel, i.e.

$$\lambda_{\text{outside}} = \lambda_a + \lambda_{\text{prio}} = \frac{\lambda_p}{1 + (N-1)p_{\text{outside|b}}}. \tag{3}$$

Finally, we choose the clearing rate λ_a according to

$$\lambda_a = \lambda_p - \lambda_b, \tag{4}$$

meaning that the average rate of background packets is λ_b in the Busy state as well as in the Available state.

2.3 Optimal Control

The design objective is to choose an optimal threshold $r = r^*$ to minimize the state variance

$$J_y = \mathrm{E}(y^2)$$

for each loop. We may also want to trade some increase in J_y to decrease the average event rate f_u of each loop. In the Å&B case, the tradeoff is given by (see [Åström and Bernhardsson, 1999])

$$J_y = \frac{1}{6}\sigma^2 f_u^{-1} = \frac{1}{6}r^2. \tag{5}$$

3. The Stationary Distribution Problem

We consider now the general case of simultaneous evolution of an integrator state according to (1) (with $u = 0$) and a continuous time Markov chain. The state of the system is $(y, k) \in \mathbb{R} \times \mathbb{Z}_n$, where k is the state of the Markov chain. The pdf over the state is $f : \mathbb{R} \mapsto \mathbb{R}^n$.

3.1 The Spatial Dynamics

The Fokker-Planck equation for the pdf f under the Brownian motion (1) with $u = 0$ is

$$\dot{f}(y) = \frac{1}{2}\sigma^2 f''(y) = -\varphi'(y),$$

where $\dot{f}$ and f' are the derivatives with respect to time and integrator state y respectively, and $\varphi(y)$ is the flow of probability in the y direction. We may thus take $\varphi(y)$ to be $\varphi(y) = -\frac{1}{2}\sigma^2 f'(y)$. Simultaneously, the Markov chain state k evolves by the transition intensity matrix A_m as

$$\dot{f} = A_m f.$$

Summing the probability flows from the integrator drift and the Markov chain gives the combined dynamics

$$\dot{f}(y) = \frac{1}{2}\sigma^2 f''(y) + A_m f.$$

Assuming stationarity, $\dot{f} = 0$ now gives the spatial dynamics

$$f''(y) + \frac{2}{\sigma^2}A_m f = \mathcal{D} f = 0, \tag{6}$$

which can also be expressed in state space form as

$$x'(y) = \begin{pmatrix} 0 & -\frac{2}{\sigma^2}A_m \\ I & 0 \end{pmatrix} x(y), \qquad x(y) = \begin{pmatrix} f'(y) \\ f(y) \end{pmatrix}.$$

3.2 Moments

Aside from the actual stationary pdf $f(y)$, we will also need moments such as marginal probability over the Markov chain states $F(0)$ and state variance V, according to

$$F(0) = \int_0^\infty f(y)dy, \quad V = \int_0^\infty y^2 \mathbf{1}^T f(y)dy.$$

Since the model is symmetric with respect to the origin, we consider f to be defined only over $y \geq 0$; the extension to the non-symmetric case is straightforward. The easiest way to find these moments is to work them into the spatial dynamics.

The moments can be computed from

$$\begin{pmatrix} F(y) \\ F_1(y) \\ F_2(y) \end{pmatrix} = \int_y^\infty \begin{pmatrix} f(y) \\ \mathbf{1}^T F(y) \\ F_1(y) \end{pmatrix} dy, \tag{7}$$

where we collect individual integrals for the zeroth moment, but sum over the Markov chain states for higher moments to reduce the computational complexity. The variance is found as $V = 2F_2(0)$ by partial integration twice to eliminate the factor y^2 in the integrand.

Combining (6) and (7) gives the extended dynamics

$$x_e'(y) = \underbrace{\begin{pmatrix} 0 & -\frac{2}{\sigma^2}A_m & & & \\ I & 0 & & & \\ & -I & 0 & & \\ & & -\mathbf{1}^T & 0 & \\ & & & -1 & 0 \end{pmatrix}}_{A_{\text{full}}} \underbrace{\begin{pmatrix} f'(y) \\ f(y) \\ F(y) \\ F_1(y) \\ F_2(y) \end{pmatrix}}_{x_e(y)}, \tag{8}$$

with boundary conditions $x \to 0$ as $|y| \to \infty$. We see that (8) has a triangular structure such that the the integrated densities F_i depend on f and f' but not vice versa.

3.3 Causal/Anticausal Decomposition

The operator $\mathcal{D}$ of (6) can be factored into a causal part $\mathcal{D}_-$ and an anticausal part $\mathcal{D}_+$ according to

$$\mathcal{D} = \frac{d^2}{dy^2} + \frac{2}{\sigma^2}A_m f = \mathcal{D}_+\mathcal{D}_- = \mathcal{D}_-\mathcal{D}_+,$$

$$\mathcal{D}_+ = \frac{d}{dy} - A_+, \quad \mathcal{D}_- = \frac{d}{dy} + A_+, \quad A_+ = \sqrt{-\frac{2}{\sigma^2}A_m}.$$

The square root exists since A_m must have a full zero eigenspace, or the temporal dynamics $\dot{p} = A_m p$ would have an unbounded solution. Since A_m has all eigenvalues in the left half plane, A_+ will have eigenvalues λ in the sector $\Re(\lambda) \geq |\Im(\lambda)|$. The poles of the spatial dynamics (6) will be the eigenvalues of $\pm A_+$, thus satisfying $|\Re(\lambda)| \geq |\Im(\lambda)|$.

4. Solving the Spatial Dynamics

The stationary density $f(y)$ is the solution to a set of linear equations composed of the spatial dynamics (8) which is piecewise constant in y, additional linear conditions specified in the model, and the normalization condition $\int f(y)dy = 1$. By solving (8) over an interval, we can eliminate the interior values of $x_e(y)$, leaving a low-dimensional set of linear equations to solve for $f(y)$. The moments and interior values of $f(y)$ can then be reconstructed.

The extended dynamics (8) can be solved using standard linear time invariant system theory. We must be careful since the dynamics is reversible, consisting of matching stable and antistable, or rather causal and anticausal parts.

4.1 Bounded Intervals

We want to solve (8) over an interval $[y_0, y_1]$. LTI system theory gives that this relation can be expressed as

$$x_e(y_1) = e^{A_{\text{full}}\Delta y} x_e(y_0), \qquad \Delta y = y_1 - y_0. \tag{9}$$

This formulation may, however, be very ill conditioned.

Premultiplying (9) by the scaling matrix $(e^{A_{\text{full}}\Delta y} + I)^{-1}$ gives the equivalent relation

$$h\big(-\frac{1}{2}A_{\text{full}}\Delta y\big)x_e(y_1) = h\big(\frac{1}{2}A_{\text{full}}\Delta y\big)x_e(y_0), \tag{10}$$

where the analytical function h is given by

$$h(At) = \left(e^{At} + e^{-At}\right)^{-1} e^{At} = \frac{1}{2}\big(I + \tanh(At)\big).$$

This formulation will in general be much better conditioned than (9). [5]

[5] Half of the eigenvalues of $h(At)$ will satisfy $|\lambda_+| \in [0.5, 1.07]$, each corresponding to an eigenvalue $|\lambda_-| \in [0, 0.5]$ of $h(-At)$ with the same eigenspace, and vice versa for the other half of the eigenvalues. (This can be seen by looking at the behavior of $\log(|h(at)|)$ along the boundary line $at = (1 \pm i)t$); $\log(|h(at)|)$ is a harmonic function on $|\Re(at)| \geq |\Im(at)|$).

The function $h(\pm At)$ can be evaluated efficiently through the doubling recursion

$$h(2At) = \left(h(At)^2 + h(-At)^2\right)^{-1} h(At)^2;$$

$h(\pm 2^{-n}At)$ is first evaluated for some suitable n such that $e^{\pm 2^{-n}At}$ is reasonably conditioned, $h(\pm At)$ is then found in a modest number of steps. If A is first put on Schur form $h(At)$ will be triangular, and all necessary matrix inversions can be done efficiently through back substitution.

4.2 Semiinfinite Intervals

When solving (6) over a semiinfinite interval $[y_0, \infty)$ we must insist that A_m has some leakage (thus having eigenvalues strictly in the left half plane) to be able to satisfy the boundary conditions that $f \to 0$ as $y \to \infty$. The solution of (6) is then

$$f(y) = e^{-A_+(y-y_0)} f(y_0), \tag{11}$$

implying the equivalent boundary conditions at $y = y_0$

$$f'(y_0) + A_+ f(y_0) = 0.$$

The integrals F_i can be found by direct integration of (11).

4.3 Interpolating the Probability Density Function

Given the boundary conditions $x(y_0), x(y_1)$, the pdf $f(y)$ can be interpolated according to

$$\mathcal{D}_- \mathcal{D}_+ f = \mathcal{D}_- g = 0, \quad \mathcal{D}_+ f = g.$$

First, g is solved for in the causal direction, with initial conditions given by

$$g(y_0) = (\mathcal{D}_+ f)(y_0) = f'(y_0) - A f(y_0).$$

Then, f is solved for in the anticausal direction, with initial conditions $f(y_1)$ and input $g(y)$. With semiinfinite intervals, it is enough to solve outwards from the finite endpoint.

5. Matching the Loops

Now that we can solve a problem instance given specific model parameters, we want to adjust the model parameters λ_a, λ_b and λ_{prio} so that all N loops behave the same, i.e. the matching conditions (2), (3) and (4) are fulfilled. The condition (4) is simply realized by parametrizing λ_a in λ_b, which implies the restriction $\lambda_b \in [0, \lambda_p]$.

Suppose that we know λ_b, λ_a and want to find λ_{prio} to satisfy (2). The spatial dynamics inside the threshold is then known, and the corresponding linear relations can be used. The spatial dynamics for the outside gives that

$$p_{\text{b,outside}} = \frac{f_b(r)^2}{-f_b'(r)},$$

since $f_b(y)$ is a decaying exponential function for $y \geq r$. Inserting this expression into (2) gives the relation

$$\frac{f_b(r)^2}{-f_b'(r)} + \underbrace{p_{\text{b,inside}} - \frac{N}{\lambda_p}\varphi_{\text{fg}}}_{\text{linear in } (f_b(r), f_b'(r))} = 0.$$

Fixing e.g. $f_b'(r)$ we get a quadratic equation in $f_b(r)$, with one positive solution. We can now solve for λ_{outside} from

$$f_b'(r) = -\sqrt{\frac{2}{\sigma^2}\lambda_{\text{outside}}}\, f_b(r).$$

When all parameters are known, the solution is finally normalized to unit total probability. We can normalize afterwards, since all other conditions on $f(y)$ are purely linear.

To match also the priority according to (3), we can use i.e. secant search over λ_b. We know that $\lambda_{\text{outside}} \in [\lambda_a, \lambda_p]$; a secant search over λ_b to satisfy (2) for each of these fixed endpoints gives a suitable starting interval for the priority matching. Sometimes there is no $\lambda_b \in [0, \lambda_p]$ that satisfies (2) with $\lambda_{\text{outside}} = \lambda_p$; we then use $\lambda_b = \lambda_p$ as the right endpoint of the search instead.

6. Correction for the Waiting Time Distribution

So far, we have assumed that each packet occupies the medium for an exponentially distributed waiting time, which allows to use a simple Markov chain model for the channel state. We will now investigate the influence of the waiting time distribution on the state variance $V_y(t) = \mathrm{E}(y(t)^2)$ to derive a first order correction to the state cost J_y.

By the dynamics (1), $V_y(t)$ evolves between events as

$$V_y(t) = \sigma^2 t + V_y(0);$$

during a waiting time τ this gives the accumulated variance

$$V_{\mathrm{acc}}(\tau) = \int_0^\tau V_y(t)dt = \frac{1}{2}\sigma^2\tau^2 + \tau V_y(0).$$

The waiting time of one packet is $\tau, \mathrm{E}(\tau) = T$, which gives the expected final and accumulated variances

$$\mathrm{E}\big(V_y(\tau)\big) = \sigma^2 T + V_y(0),$$
$$\mathrm{E}\big(V_{\mathrm{acc}}(\tau)\big) = \frac{1}{2}\sigma^2\big(T^2 + \mathrm{V}(\tau)\big) + TV_y(0).$$

Keeping T fixed, the final state variance when the packet is completed stays fixed as well. The accumulated variance, however, depends also on the variance $\mathrm{V}(\tau)$. Compared to a fixed waiting time, an exponential waiting time will make V_{acc} bigger by the term

$$\Delta V_{\mathrm{acc}} = \frac{1}{2}\sigma^2\,\mathrm{V}(\tau) = \frac{1}{2}\sigma^2 T^2.$$

Since the total packet rate is $Nf_u = \lambda_p p_{\mathrm{b}}, \lambda_p = T^{-1}$, the state cost J_y with exponential waiting times is bigger by

$$\Delta J_y = Nf_u \Delta V_{\mathrm{acc}} = \frac{1}{2}\sigma^2 p_b \lambda_p^{-1}. \tag{12}$$

7. Results

All results are derived with the parameters $\sigma = 1$ and $\lambda_p = N$. The latter means that the network bandwidth scales in proportion to the number of loops. This scaling is convenient since it makes the Å&B performance independent of N. The relative state variance and packet rate of the control schemes are unaffected by the choice of σ and λ_p since the integrator has no time constant.

We will compare the following models:

- Å&B's ideal, aperiodic controller, with $r = f_u = 1$ and $J_y = \frac{1}{6}$ [Åström and Bernhardsson, 1999]. This is a lower bound on the achievable performance.
- A Monte Carlo simulation model with N sporadic control loops competing for the network under CSMA with random arbitration [Cervin and Henningsson, 2008]. The transmission time of each packet is assumed fixed and equal to $T = N^{-1}$. Via bisection search over r, the

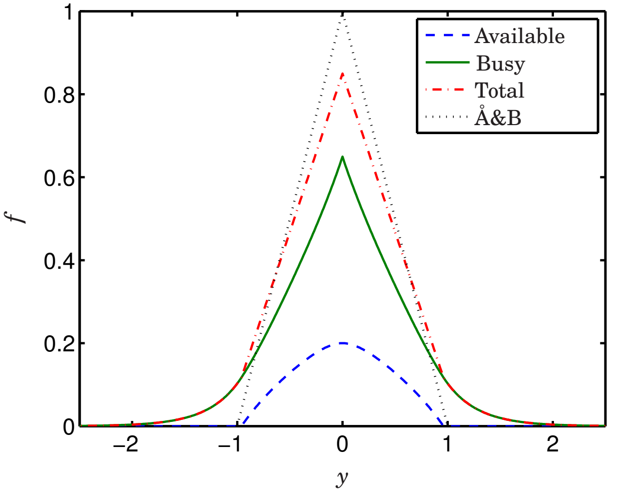

Figure 2. The stationary pdf $f(y)$ according to the simple model with $N = 10$ loops and optimal threshold $r = 0.96$. The Å&B pdf is shown for comparison.

optimal performance has been found for $N \in \{1, 2, 3, 5, 10, 20, 45, 100, 200\}$. The time granularity of the model was $N^{-1}10^{-3}$ and each simulation ran for 10^8 time steps.

- The simple model proposed in this paper. In the model, the transmission times are exponentially distributed with mean $T = N^{-1} = \lambda_p^{-1}$; the cost correction (12) has been subtracted from J_y to predict the state variance with fixed rather than exponential packet times.

7.1 Probability Densities

Figure 2 shows the stationary pdf $f(y)$ obtained by the simple model with $N = 10$ loops. The optimal threshold $r = r^*$ is chosen to minimize state variance J_y. The pdf of Å&B is shown for reference.

We see that, just as in the Å&B case, the pdf of the integrator state y varies linearly inside the threshold, with a break in the origin because of inflow from reset. In the simple model however, the inflow goes into the Busy state, while the outflow from reset is divided between the Available state at the threshold and the Busy state outside the threshold.

The resulting difference between the pdf:s of the Å&B case and the simple model is that the latter includes an exponential tail outside the threshold r, while the loop is waiting to gain access to the channel.

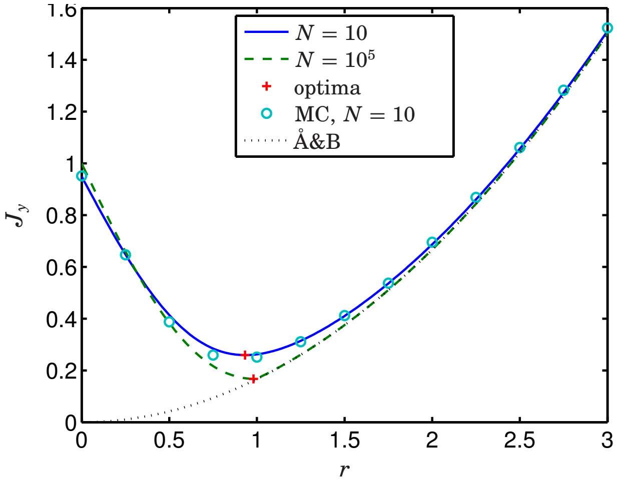

Figure 3. State cost J_y as a function of threshold r according to the simple model for different number of loops N, along with the Å&B case and Monte Carlo results for the full model. The Å&B case serves as a lower bound for both $N = 10$ and $N = 10^5$, with the two essentially coinciding in the latter case beyond the optimum $r = r^*$.

7.2 Threshold Dependence

Figures 3 and 4 show the dependence of state cost J_y and packet rate f_u on threshold r, with $N = 10$ and $N = 10^5$ loops. Monte Carlo (MC) results with $N = 10$ loops and Å&B results are shown for comparison.

We see that the state cost J_y has a minimum around $r = 1$, while f_u drops monotonically with the threshold. By using a threshold $r > r^*$, it is possible to trade a decreased packet rate f_u for an increased state cost J_y; nothing is gained by using $r < r^*$. The simple model is remarkably accurate in predicting the Monte Carlo results for the full model, given the radically lower computation complexity.

As the threshold r grows, all curves tend to the Å&B case, which serves as a lower bound on J_y and an upper bound on f_u. For $N = 10$ loops, the convergence is gradual. For $N = 10^5$, we can distinguish two domains: when $r \geq r^*$, J_y and f_u essentially coincide with Å&B; when $r \leq r^*$, f_u lies at the channel capacity while the state cost J_y deteriorates as the threshold r decreases.

It thus seems that for large N, performance follows the Å&B case as soon as the necessary average packet rate f_u lies within the peak packet rate of the channel; the minimal state cost J_y^* is achieved at this break point.

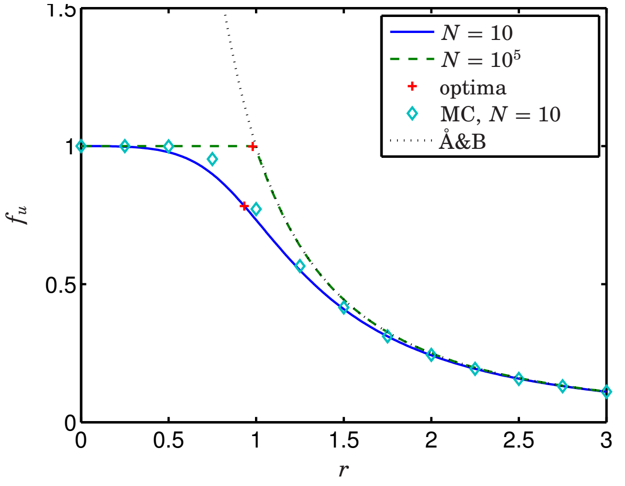

Figure 4. Packet rate f_u as a function of threshold r according to the simple model for different number of loops N, along with the Å&B case and Monte Carlo results for the full model. The Å&B case serves as an upper bound for both $N = 10$ and $N = 10^5$, with the two essentially coinciding in the latter case beyond the optimum $r = r^*$.

7.3 Dependence on the Number of Loops

The threshold r was optimized using golden section search to minimize state cost J_y as a function of N. Figure 5 shows optimal state cost J_y, packet rate f_u and threshold r as N varies from 1 to 10^5 for the simple model, along with full model Monte Carlo and Å&B results for comparison.

We see that as N increases, J_y drops from about 0.4 towards the Å&B case of $J_y = \frac{1}{6}$. The simple model is slightly optimistic about J_y for $N = 1$, and slightly pessimistic for $N \geq 2$. Convergence to the Å&B case is very close at $N = 10^3$.

The rise in packet rate f_u with N shows a similar pattern to the drop in J_y; this rise is in fact necessary to bring J_y down since the Å&B performance (5) lower bounds $J_y(f_u)$. In accordance, the slightly lower state cost J_y of the full model Monte Carlo results is accompanied by a higher packet rate f_u, realized by a quicker drop in threshold r.

8. Future Work

A main direction for continued research is to apply the same kind of modelling developed in this paper to other channel models, such as with partial and full collisions. It is then probably better to trigger control events at

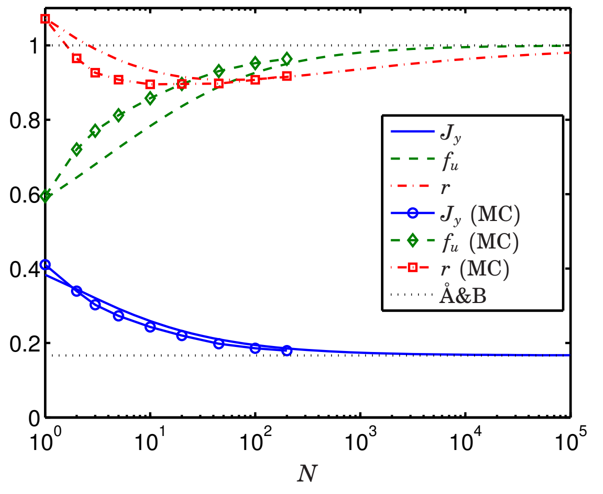

Figure 5. Optimal state cost J_y, packet rate f_u, and threshold r as a function of number of loops N. Results are shown for the simple model, full model Monte Carlo simulations, and the aperiodic Å&B case (for which $J_y = \frac{1}{6}$, $f_u = r = 1$).

a bounded intensity once the threshold is crossed. Points that deserve further study include

- How should the shared medium best be utilized when collisions lead to packet drop?
- How does a partial collisions model transition from CSMA behavior to full collision behavior as time to detect that the channel is busy changes?
- Can the model be extended to general first order process dynamics, or to processes with higher state dimension?

9. Conclusion

Event based control offers the promise to better utilize limited communication resources than traditional periodic control. The potential benefit increases with the number of control loops sharing the same medium, as they may be able to trade use of the communication channel between each other to be able to gain access when it is most needed. The potential for interference, however, also increases; the question is which of the effects dominates.

This paper presents a simple model for the interaction between N event based control loops sharing a common medium in the form of a CSMA channel. The model includes the medium and a single foreground loop, approximating the behavior of the other (background) loops. The resulting model can be evaluated with little computational resources, independent of N.

Comparison to Monte Carlo simulations of the simultaneous evolution of all loops shows that the simple model predicts the behavior of the full model with remarkable accuracy, given its simplicity. As the number of loops increases, they are able to share the medium more and more efficiently, making the sporadic channel constraint appear more and more like an average capacity constraint, resulting in significantly improved performance .

Acknowledgment

The authors would like to thank M. Rabi for fruitful discussions leading up to the model used in this paper. This work has been supported by the Swedish Research Council, ELLIIT, and EU/FP7/ArtistDesign.

References

Åström, K. J. and B. Bernhardsson (1999): "Comparison of periodic and event based sampling for first-order stochastic systems." In *Preprints of the 14th World Congress of IFAC*. Beijing, P.R. China.

Cervin, A. and T. Henningsson (2008): "Scheduling of event-triggered controllers on a shared network." In *Proceedings of the 47th IEEE Conference on Decision and Control*. Cancun, Mexico.

Cervin, A. and E. Johannesson (2008): "Sporadic control of scalar systems with delay, jitter and measurement noise." In *Proceedings of the 17th IFAC World Congress*. Seoul, Korea.

Cogill, R. (2009): "Event-based control using quadratic approximate value functions." In *Proceedings of the IEEE Conference on Decision and Control*. Shanghai, P.R. China.

Cogill, R., S. Lall, and J. P. Hespanha (2006): "A constant factor approximation algorithm for event-based sampling." In *Proceedings of the 44th Annual Allerton Conference on Communication, Control, And Computing*, Monticello, Illinois.

Henningsson, T., E. Johannesson, and A. Cervin (2008): "Sporadic event-based control of first-order linear stochastic systems." *Automatica*, **44:11**, pp. 2890–2895.

Jelenković, P. R. and J. Tan (2007): "Is ALOHA causing power law delays?" In *Managing Traffic Performance in Converged Networks*, vol. 4516 of *Lecture Notes in Computer Science*. Springer.

Rabi, M. and K. H. Johansson (2009): "Scheduling packets for event-triggered control." In *Proceedings of the 2009 European Control Conference*. Budapest, Hungary.

Tabuada, P. (2007): "Event-triggered real-time scheduling of stabilizing control tasks." *IEEE Transactions on Automatic Control*, **52:9**, pp. 1680–1685.

Tobagi, F. A. and V. B. Hunt (1980): "Performance analysis of carrier sense multiple access with collision detection." *Computer Networks*, **4:5**, pp. 245–259.

Wang, X. and M. Lemmon (2009): "Self-triggered feedback control systems with finite-gain $\mathcal{L}_2$ stability." *IEEE Transactions on Automatic Control*, **45:3**, pp. 452–467.

Paper VI

Log-concave Observers

Toivo Henningsson Karl Johan Åström

Abstract

The Kalman filter is the optimal state observer in the case of linear dynamics and Gaussian noise. In this paper, the observer problem is studied when process noise and measurements are generalized from Gaussian to log-concave. This generalization is of interest for example in the case where observations only give information that the signal is in a given range. It turns out that the optimal observer preserves log-concavity. The concept of strong log-concavity is introduced and two new theorems are derived to compute upper bounds on optimal observer covariance in the log-concave case. The theory is applied to a system with threshold based measurements, which are log-concave but far from Gaussian.

Reprinted from *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems*. Kyoto, Japan, 2006.

1. Introduction

The Kalman filter (see [Kalman, 1960], [Kalman and Bucy, 1961]) is one of the most widely used schemes for state estimation from noisy measurements. It is optimal for linear measurements and Gaussian noise, but it is often applied in a more general setting. Although the Extended Kalman filter (see [Gelb and Corporation., 1974]) often works well in practice, sometimes it does not, and it is in general not easy to see how altered conditions change the observer problem.

In this paper, a particular generalization is investigated where measurements and noise are allowed to be log-concave (see [Prékopa, 1971], [Prékopa, 1973], [Bagnoli and Bergstrom, 1989], [An, 1996]). The model of log-concave measurements is applicable in many instances where the assumption of independent additive measurement noise is too limited, for instance with heavy quantization, or with the problem of event based sampling discussed in [Åström and Bernhardsson, 2002].

Within this framework, the problem of moving horizon ML/MAP estimation becomes a convex optimization problem, see [Schön *et al.*, 2003]. This paper will however focus on the covariance of the Bayesian Observer, which is investigated and compared with the Kalman filter.

Strongly log-concave functions are introduced as a means to quantify observer properties. Two new theorems are applied to derive upper bounds on optimal observer covariance.

It turns out that the observer problem is still quite well behaved so that, especially with some insight gained in the analysis, a Kalman filter might often be usable for this more general measurement setting. For a more thorough treatment, see [Henningsson, 2005].

The paper is organized as follows. A motivating example is presented in Section 2. The notion of log-concavity is introduced in Section 3, where we state the main results as Theorems 1 and 2. In Section 4 we treat the observer problem. The results in section 3 are used to investigate the observer properties. Finally in Section 5 the results are applied to the example.

2. Example: A MEMS Accelerometer

Consider an accelerometer based on the following design. A test mass is suspended to move freely in one dimension and is affected by an external acceleration. Sensors detect deviations from the origin exceeding a detection threshold and report the sign of the deviation. An input signal is available to accelerate the test mass so as to keep it close to the origin. The aim of the design is to estimate the external acceleration as accurately as possible.

The discrete time dynamics are given by

$$x(k) = \begin{pmatrix} 1 & h \\ 0 & 1 \end{pmatrix} x(k-1) + \begin{pmatrix} \frac{1}{2}h^2 \\ h \end{pmatrix} u(k-1) + v(k-1),$$

where x is the state, u the input signal, v the external acceleration and h the sampling period. The state consists of position x_1 and velocity x_2. With the external acceleration as a white noise disturbance, sampling yields v to be Gaussian white noise with covariance

$$P_N = \sigma^2 \begin{pmatrix} \frac{1}{3}h^3 & \frac{1}{2}h^2 \\ \frac{1}{2}h^2 & h \end{pmatrix},$$

where σ^2 is the process noise intensity.

The measurements are given by

$$y(k) = \begin{cases} \operatorname{sign}\big(x_1(k)\big), & |x_1(k)| \geq 1 \\ 0, & \text{otherwise,} \end{cases}$$

which is the only non-classical assumption used in the model. The output $y(k)$ is not readily described as a linear combination of state and uncorrelated measurement noise, making a straightforward application of Kalman filter theory difficult.

In fact, it is not at all obvious what properties to expect for this observer problem; will the observer error remain bounded, how large will it be, how does it depend on the measurement sequence, how complex observer is necessary, and so on. To answer questions about the observer problem, the Bayesian observer for the system will be analyzed. Other examples where similar measurement conditions apply are when measurements are coarsely quantized or come in the form of level triggered events.

3. Log-concavity

Many results are available on general log-concavity, see, e.g. [Prékopa, 1971], [Prékopa, 1973], [Bagnoli and Bergstrom, 1989], and [An, 1996]. The book [Boyd and Vandenberghe, 2004] contains much material on convex functions that can easily be transfered to the log-concave case. Here, only the properties that are most relevant in the context of this paper will be stated.

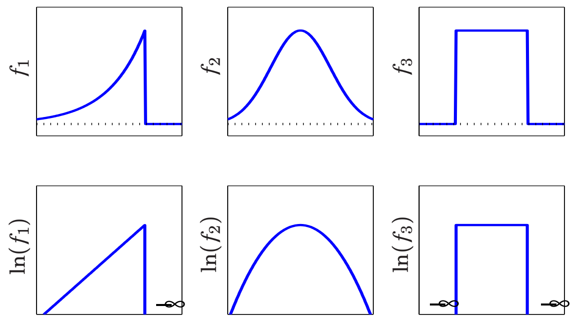

Figure 1. Some examples of log-concave functions in one variable; the function is plotted above and its logarithm below. The dotted line is $f = 0$, and $\ln(0)$ is taken to be $-\infty$. f_1: Truncated exponential function, f_2: Gaussian function, f_3: rectangular window.

A log-concave function is a function with concave logarithm. Log-concave functions are well suited for applying convexity theory to probability densities; many common densities are log-concave and several useful operations preserve log-concavity. In contrast, no probability density on $\mathbb{R}^n$ is either convex or concave since probability densities have a finite integral while convex and concave functions on $\mathbb{R}^n$ do not.

Definition 1—Log-concave Function
A function $f : \mathbb{R}^n \longrightarrow \mathbb{R}$ is logarithmic concave or *log-concave*, iff $f(x) \geq 0$, f has convex support and $\ln\big(f(x)\big)$ is concave on this support. □

For some simple examples of log-concave functions see Figure 1, and for some counterexamples Figure 2. Among common log-concave densities are Gaussian and exponential densities.

Log-concave functions are unimodal, meaning that the superlevel sets $\{x; f(x) \geq a\}, a \in \mathbb{R}$ are convex. Many attractive properties of log-concave functions are analogous to those for concave functions. A useful fact is that multiplication takes the place of addition so that the product of two log-concave functions is log-concave. Another very useful result derived by Prékopa is

Proposition 1—Prékopa
Let $f(x, y)$ be jointly log-concave in $x \in \mathbb{R}^m, y \in \mathbb{R}^n$. Then the integral

$$g(x) = \int f(x, y)dy$$

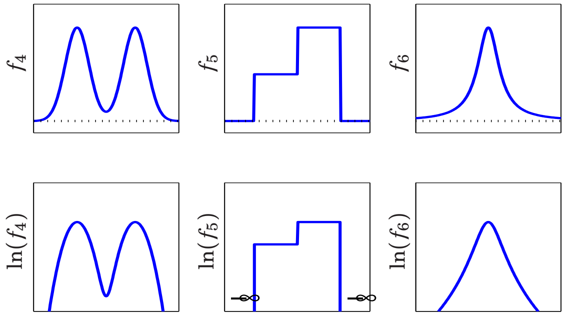

Figure 2. Some examples of functions in one variable that are *not log-concave*; the function is plotted above and its logarithm below. f_4: Not unimodal, f_5: Discontinuous on interior of support, $f_6(x) = \frac{1}{1+x^2}$: sub-exponential decay.

is a log-concave function of x.

Proof: See [Prékopa, 1971] and [Prékopa, 1973]. □

This theorem implies for instance that the marginal densities of log-concave densities are log-concave, and that the convolutions of log-concave functions are log-concave. It will be central in the proof of Theorems 1 and 2.

3.1 Strong Log-concavity

Log-concavity is in its nature only a qualitative property. To allow for quantitative statements, the following class of functions is introduced.

Definition 2—Strongly Log-concave Function
Let $P \in \mathbb{R}^{n\times n}$ be positive definite and define the set

$$\mathcal{LC}(P^{-1}) = \left\{ f;\, f_0(x) = \frac{f(x)}{e^{-\frac{1}{2}x^T P^{-1} x}} \text{ log-concave} \right\}.$$

The function f is *strongly log-concave* of strength P^{-1} iff $f \in \mathcal{LC}(P^{-1})$. □

All strongly log-concave functions are log-concave, bounded and go to zero as $|x| \to \infty$ at least as fast as a Gaussian function.

Membership in $\mathcal{LC}(P^{-1})$ can be seen as an inequality, in the sense that

$$\begin{aligned} & f \in \mathcal{LC}(P^{-1}), \quad P \le R \\ \implies & f \in \mathcal{LC}(R^{-1}). \end{aligned}$$

The inclusion $f \in \mathcal{LC}(P^{-1})$ is *tight* iff $\mathcal{LC}(P^{-1})$ is a subset of all $\mathcal{LC}(R^{-1})$ that contain f.

A Gaussian density with covariance P is tightly in $\mathcal{LC}(P^{-1})$, and can be seen as the *corresponding Gaussian* to this class. The definition states that any strongly log-concave function can be written as the product of a log-concave function and a corresponding Gaussian. Also, the following properties hold:

THEOREM 1—ENCAPSULATION PROPERTY
If $f \in \mathcal{LC}(F^{-1})$ and $g \in \mathcal{LC}(G^{-1})$ then

$$\begin{aligned} f(Ax+b) &\in \mathcal{LC}(A^T F^{-1} A) \\ (f * g)(x) &\in \mathcal{LC}\big((F+G)^{-1}\big) \\ f(x)\cdot g(x) &\in \mathcal{LC}(F^{-1}+G^{-1}), \end{aligned}$$

where $x, b \in \mathbb{R}^n, A \in \mathbb{R}^{n\times n}$ and $f * g$ is the convolution of f and g.

Proof: See Appendix A. □

The inclusions are as narrow as the premises allow, being tight when f and g are the corresponding Gaussians.

THEOREM 2—COVARIANCE BOUND
If $f \in \mathcal{LC}(P^{-1})$ is a probability density then

$$V = \int (x - m_x)(x - m_x)^T f(x)dx \le P,$$

where $m_x = \int x f(x)dx$. The bound is tight for the corresponding Gaussian.

Proof: See Appendix B. □

The matrix expressions for strength of log-concavity correspond exactly to the way that the operations propagate inverse covariances for Gaussian functions. By the latter theorem, the inverse strength of log-concavity is an upper bound on the covariance.

The theorems form a chain of inequalities that can be used to propagate upper bounds on covariance under the operations of affine transformation, convolution and multiplication. For more properties of strongly log-concave functions, see [Henningsson, 2005].

4. Log-concave Observers

The observer problem that will be considered is for processes with linear dynamics and log-concave noise and measurements, as defined below.

The dynamics are given by

$$x(k) = Ax(k-1) + Bu(k-1) + v(k-1),$$

where x is the state, u the input and v the process noise. The noise has log-concave probability density f_N. The matrices A and B, as well as f_N may be time-varying.

The measurements are described by the stochastic variables $Y(k)$,

$$f_{Y(k)|X(k)}\big(y|x(k)\big) = f_M\big(y, x(k)\big),$$

where the measurement function f_M is log-concave in x for each y and may be time-varying.

4.1 The Bayesian Observer

As a basis for the analysis, the online Bayesian observer for estimation of $x(k)$ from the history of y and u will be considered. The state of the observer at any time is fully described by the function

$$f_k(x) = f_{X_k|y_{1:k}, f_{X_0}}(x),$$

where $y_{1:k}$ is the measurement history and f_{X_0} is the assumed initial density.

The observer update from f_{k-1} to f_k is best described in three steps taking into account dynamics, process noise, and measurements:

$$f_k^{\mathrm{d}}(x) \;\propto\; f_{k-1}\big(A^{-1}x - A^{-1}Bu(k-1)\big), \tag{1}$$

$$f_k^{\mathrm{dn}}(x) \;=\; (f_N * f_k^{\mathrm{d}})(x), \tag{2}$$

$$f_k(x) \;\propto\; f_M(y(k), x) \cdot f_k^{\mathrm{dn}}(x), \tag{3}$$

where $\propto$ denotes proportionality and A is assumed to be invertible. For the derivation, see [Henningsson, 2005]. The dynamics update corresponds to an affine transformation, the noise update to a convolution with f_N, and the measurement update to a multiplication with $f_M(y(k), \cdot)$. For an illustration, see Figures 3, 4 and 5. If f_{X_0}, f_N and $f_M(y, \cdot)$ are Gaussian, the observer updates (1)-(3) reduce to a Kalman filter.

4.2 Properties

Since log-concavity is preserved under affine parameter transformation, convolution, and multiplication, all f_k are log-concave if f_{X_0} is log-concave.

Theorem 1 can be used to propagate upper bounds on observer covariance. This approach can be used to asses the merits of a particular sensor setup, or together with some information about the localization of f_k to give state estimates with error bounds. The computations of covariance propagation have the structure of a Kalman filter applied to corresponding Gaussians.

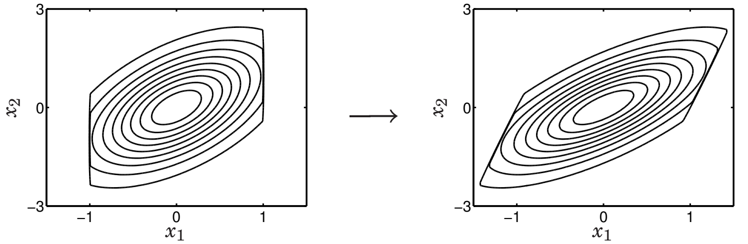

Figure 3. Illustration of the dynamics update for the MEMS accelerometer observer. The transformation amounts to a shear in this case.

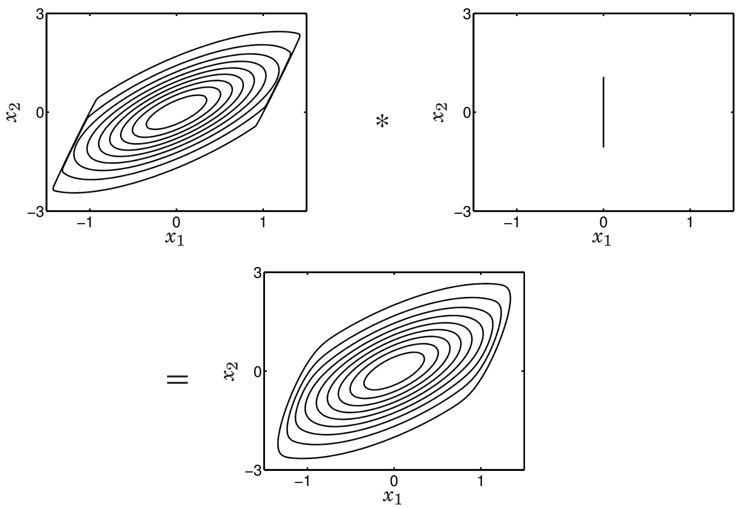

Figure 4. Illustration of the process noise update for the MEMS accelerometer observer. The Gaussian noise enters almost exclusively in the x_2 direction.

5. An Application

The MEMS accelerometer will now be used to illustrate how the theory can be applied in the analysis of a concrete observer problem.

5.1 Analytical Covariance Bounds

The accelerometer has linear dynamics and log-concave noise and measurements. The process noise density f_N is Gaussian with covariance P_N, so that $f_N \in \mathcal{LC}(P_N^{-1})$. The measurement function $f_M(y, x)$ is log-concave in x for all y, see Figure 6.

Applying Theorem 1 directly leads in this case to a highly conservative

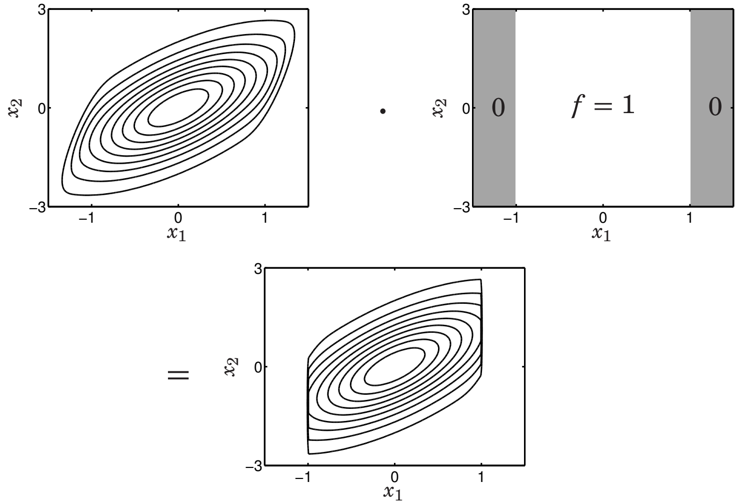

Figure 5. Illustration of the measurement update for the MEMS accelerometer observer. The measurement is $y = 0$.

covariance bound, achieved when completely ignoring the measurements. The bound grows cubically with time. Grid based finite difference simulations of the Bayesian observer do however indicate that the covariance is bounded, and if the output changes frequently, small.

The reason why the bound is so conservative is that f_M is not strongly log-concave for any y; strength of log-concavity being the only measure of information that the theorem considers. In lack of stronger proven results, a slight approximation will allow to account for the major source of state information.

The most important source of state information under normal conditions is the events when y goes from being 0 to ± 1, at which time x_1 is known to be almost exactly equal to y. This can be modeled as a Gaussian measurement of x_1 with expectation y and variance σ_M^2.

The variance σ_M^2 will depend on the process noise and uncertainty in velocity, but will be small when h is small. The modified measurement function $\hat{f}_M$ can be seen in Figure 7. For events, $\hat{f}_M(\pm 1, \cdot) \in \mathcal{LC}(Q_M)$ where

$$Q_M = \begin{pmatrix} \sigma_M^{-2} & 0 \\ 0 & 0 \end{pmatrix},$$

and otherwise $\hat{f}_M(0, \cdot) \in \mathcal{LC}(0)$.

Under this approximation, the variance of the optimal estimate $\hat{x}_2$ of

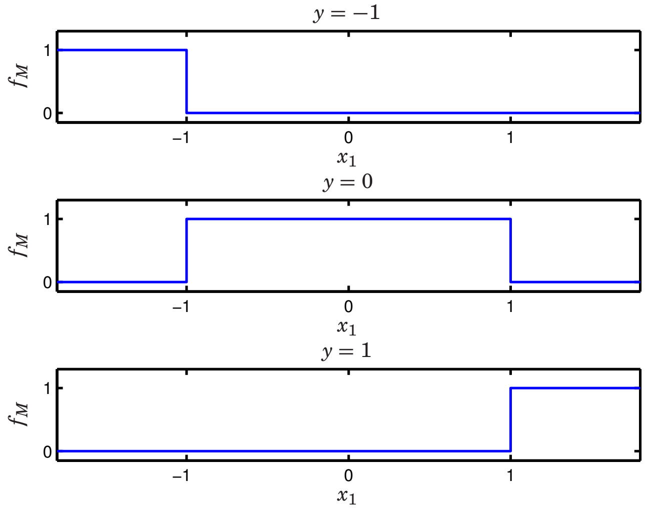

Figure 6. The measurement function $f_M(y, x)$ for the MEMS accelerometer describing the relative probability of state x when $y = -1, 0, -1$. The function is independent of x_2.

x_2 right after an event can be shown to satisfy

$$V(\hat{x}_2) \leq \frac{1}{3}\sigma^2 t + 2\sigma_M^2 t^{-2},$$

where t is the time since the last event. For the derivation, see [Henningsson, 2005].

The bound illustrates that the accuracy of the accelerometer depends strongly on the rate of events. If the objective of control is good measurements, the controller should keep the rate above a certain level, for instance sending the test mass bouncing in a ping pong fashion between the detection boundaries.

5.2 Kalman Filter Approximation

A Kalman filter was tuned to give a reasonable approximation of the Bayesian observer. The crucial issue was to assign the covariance of the measurement $y = 0$. While a single measurement $y = 0$ predicts x_1 to have expectation zero with variance $\sigma^2 = \frac{1}{3}$, there is much less additional information in the measurement $y = 0$ at the next time step.

In this case it is reasonable to design the Kalman filter by choosing the stationary variance p_{11}^{stat} of x_1 when $y = 0$. The variance would typically be $p_{11}^{\text{stat}} = \frac{1}{\sqrt{3}}$ (rectangular distribution) or a little less. From solving the

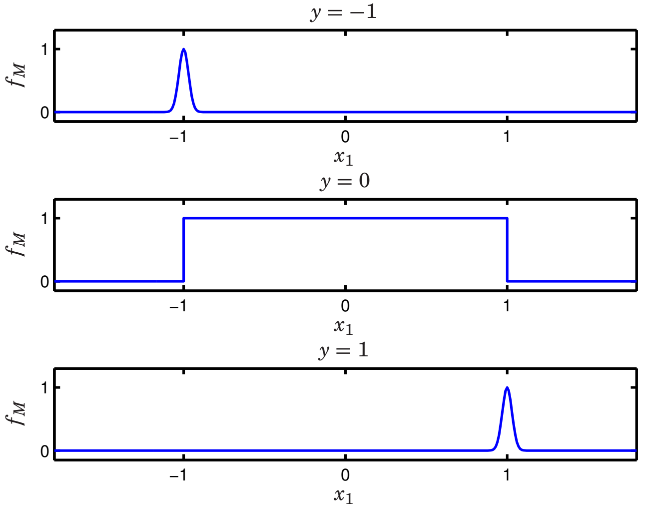

Figure 7. The modified measurement function $\hat{f}_M(y, x)$ for the MEMS accelerometer when $y = -1, 0, -1$. For $y = \pm 1$, the function has been changed to a narrow Gaussian centered on the detection threshold.

Riccati equation, it is found that the measurement variance $\sigma_o^2 h^{-1}$ for $y = 0$ must be chosen according to

$$\sigma_o = 2^{-1/3}(p_{11}^{\text{stat}})^{2/3}\sigma^{-1/3},$$

where σ^2 is the process noise intensity.

5.3 Simulation

Figure 8 shows a comparison of actual and predicted variances for a simulation of the Bayesian observer. The variance σ_M^2 was chosen so that the approximate upper bound would always be conservative. The upper bound is quite tight some time after each event, but then diverges. The variance of the tuned Kalman filter appears to be an only mildly conservative approximation of the actual variance. As long as the rate of events is reasonably high, the approximate upper bound is very tight.

A simple control law was devised to control the rate of events, and simulations were run for different rates to compare observer performance for the grid filter and the tuned Kalman filter. Figure 9 shows the observer error as a function of mean time between events t_{mean}. The grid filter is slightly better than the tuned Kalman filter and considerably better than the approximate covariance bound down to values of t_{mean} around 0.4.

For lower t_{mean} it seems that the grid filter scheme encounters discretization issues. At the same time, the tuned Kalman filter comes very

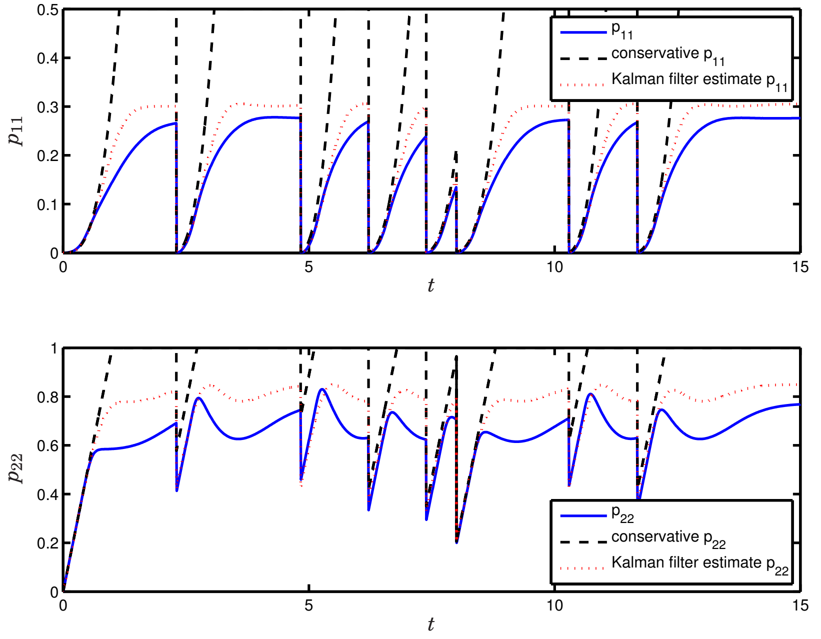

Figure 8. Observer covariances during a simulation for the MEMS accelerometer: grid filter, approximate upper bound and tuned Kalman filter.

close to the approximate upper bound which appears to be very tight in this region, indicating that the observer problem is very similar to the Kalman filter case for high rates. This similarity is not surprising since when the covariance is small, the bulk of probability mass is concentrated in a small region which is only seldom affected by the non Gaussian measurements.

Thus it is seen that the upper bound derived from the theory is quite tight when the rate of events is high and that if the inherent correlation in the non Gaussian measurements is considered, a Kalman filter can be applied as a close to optimal observer.

Example 1—Quantized Measurements

In the previous example it was necessary to rely on approximations because the measurement functions were not strongly log-concave. If the measurement function can be chosen freely, much stronger results are possible.

Consider the general problem of estimating a scalar variable from a series of independent identically distributed quantized measurements. The objective is to find a conditional measurement distribution, or mea-

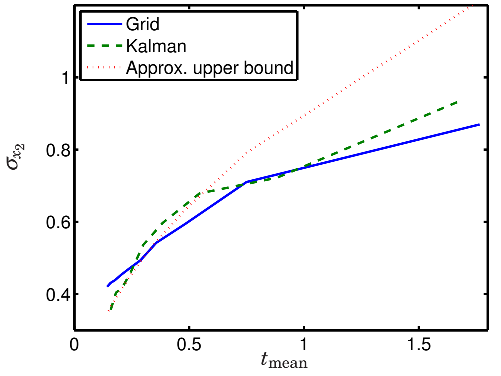

Figure 9. RMS x_2 estimation error as a function of mean time between events: grid filter estimation error, tuned Kalman filter estimation error, and approximate upper bound. For too high event rates, the grid filter suffers from discretization problems.

surement function, that is in some sense optimal. Using strength of log-concavity as an optimality criterion one can formulate the following problem:

Let the independent measurements y be distributed according to

$$f_{Y|x}(y|x) = f(x-y), \quad y \in \mathcal{Z},$$

where x is the variable to be estimated. Find a function $f \in \mathcal{LC}(p^{-1})$, where $p > 0$ is as small as possible, such that

$$f(x) \geq 0,$$
$$f(-x) = f(x),$$
$$\sum_{k=-\infty}^{\infty} f(x-k) = 1.$$

The solution is given by the function

$$f(x) = \begin{cases} 2^{-4|x|^2}, & |x| \leq \frac{1}{2}, \\ 1 - 2^{-4(1-|x|)^2}, & \frac{1}{2} < |x| \leq 1, \\ 0, & \text{otherwise}, \end{cases}$$

satisfying $f(x) \in \mathcal{LC}\left(8\ln(2)\right)$. The function is plotted in Figure 10, and in log-scale in Figure 11. A series of n measurements with f as measurement

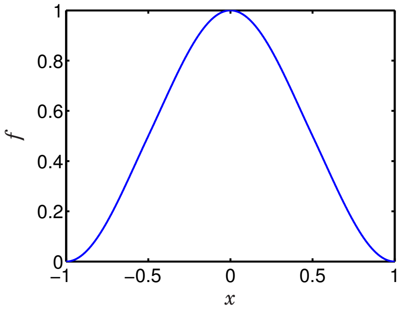

Figure 10. The measurement function in Example 1. The function is Gaussian when $|x| \leq \frac{1}{2}$ and zero when $|x| \geq 1$.

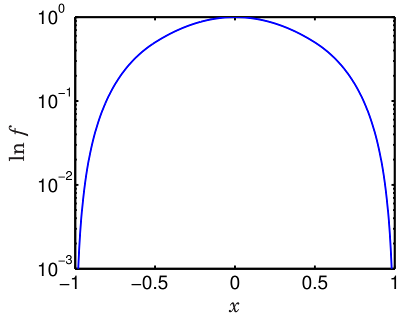

Figure 11. The measurement function in Example 1 in log-scale. The logarithm is clearly concave, being quadratic when $|x| \leq \frac{1}{2}$.

function is guaranteed to yield a probability density in $\mathcal{LC}\big(n \cdot 8\ln(2)\big)$ and therefore a variance satisfying $\sigma^2 \leq \frac{1}{n \cdot 8\ln(2)}$.

$\square$

6. Conclusion

Log-concavity is a powerful tool when dealing with probability densities. The generalization to allow log-concave densities in the observer widens the range of application considerably compared to the Kalman filter. Although no closed form solution exists in the general case, the observer problem is still very accessible to mathematical treatment.

Regarding observability and observer performance, strongly log-concave functions together with Theorems 1 and 2 can be applied to derive simple bounds on achievable observer covariance.

An in-depth treatment of the log-concave case gives a greater understanding of the performance of an Extended Kalman filter. In design of instrumentation, striving for log-concave measurement functions can facilitate the observer problem.

Acknowledgement

The authors would like to thank Anders Rantzer, Maria Karlsson, Peter Alriksson and Ather Gattami for helpful comments.

References

An, M. (1996): “Log-concave probability distributions: Theory and statistical testing.” Papers 96-01. Centre for Labour Market and Social Research, Denmark.

Åström, K. J. and B. Bernhardsson (2002): “Comparison of Riemann and Lebesque sampling for first order stochastic systems.” In *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vegas, Nevada.

Bagnoli, M. and T. Bergstrom (1989): “Log-concave probability and its applications.” Papers 89-23. Michigan Center for Research on Economic & Social Theory.

Boyd, S. and L. Vandenberghe (2004): *Convex Optimization*. Cambridge University Press, New York, NY.

Gelb, A. and A. S. Corporation. (1974): *Applied optimal estimation*. M.I.T. Press, Cambridge, Mass.

Henningsson, T. (2005): “Logarithmic concave observers.” Master’s Thesis ISRN LUTFD2/TFRT--5747--SE. Department of Automatic Control, Lund University, Sweden.

Kalman, R. E. (1960): “A new approach to linear filtering and prediction problems.” *Transactions of the ASME–Journal of Basic Engineering*, **82**, Series D, pp. 35–45.

Kalman, R. E. and R. S. Bucy (1961): “New results in linear filtering and prediction theory.” *Transactions of the ASME—Journal of Basic Engineering*, **83**, pp. 95–107.

Prékopa, A. (1971): “Logarithmic concave measures with application to stochastic programming.” *Acta Sci. Math.*, **32**, pp. 301–316.

Prékopa, A. (1973): "On logarithmic concave measures and functions." *Acta Sci. Math.*, **34**, pp. 335–343.

Schön, T., F. Gustafsson, and A. Hansson (2003): "A note on state estimation as a convex optimization problem." In *Proceedings of the 2003 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Hong Kong.

A. Proof of Theorem 1

The proofs are based on the fact that a function f is in $\mathcal{LC}(F^{-1})$ iff it can be factored as

$$f(x) = e^{-\frac{1}{2}x^T F^{-1} x} f_0(x), \tag{4}$$

where $f_0(x)$ is log-concave. This follows from the definition.

The proofs for affine transformation and multiplication are straightforward, while the proof for convolution is a little more involved.

A.1 Affine Transformation

Let $f \in \mathcal{LC}(F^{-1})$, $A \in \mathbb{R}^{n\times n}$, $b \in \mathbb{R}^n$ and $y = Ax + b$. Then

$$\begin{aligned} g(x) &= f(Ax+b) \\ &= e^{-\frac{1}{2}(Ax+b)^T F^{-1}(Ax+b)} \cdot f_0(y) \\ &= e^{-\frac{1}{2}(x^T A^T F^{-1} Ax + 2b^T F^{-1} Ax + b^T F^{-1} b)} \cdot f_0(y) \\ &= e^{-\frac{1}{2}x^T (A^T F^{-1} A)x} \cdot \underbrace{\left(e^{-\frac{1}{2}b^T F^{-1} b} e^{-(A^T F^{-1} b)^T x} f_0(y)\right)}_{g_0(x)}. \end{aligned}$$

We see that g_0 is the product of a constant, an exponential function and a log-concave function, since log-concavity is preserved under affine parameter transformation. Then g_0 is log-concave because each of the factors is log-concave. Thus $g \in \mathcal{LC}(A^T F^{-1} A)$.

A.2 Convolution

For the proof we need the following matrix identity. Let A, B and C be positive definite matrices such that $C^{-1} = A^{-1} + B^{-1}$, or $C = A(A+B)^{-1}B$. Let x, y and $z = y - (A+B)^{-1}Bx$ be vectors. Then

$$z^T(A+B)z = y^T(A+B)y - 2x^T By + x^T B(A+B)^{-1}Bx$$

and

$$\begin{aligned} x^T Cx + z^T (A + B)z &= x^T A(A + B)^{-1} Bx + z^T (A + B)z \\ &= x^T Bx + y^T (A + B)y - 2x^T By \\ &= y^T Ay + (x - y)^T B(x - y), \end{aligned}$$

that is,

$$y^T Ay + (x - y)^T B(x - y) = x^T Cx + z^T (A + B)z, \tag{5}$$

which can be seen as completion of squares in x.

Let $f \in \mathcal{LC}(F^{-1})$ and $g \in \mathcal{LC}(G^{-1})$. Then

$$\begin{aligned} h(x) &= (f * g)(x) \\ &= \int f(y)g(x - y)dy \\ &= \int e^{-\frac{1}{2}y^T F^{-1} y} e^{-\frac{1}{2}(x-y)^T G^{-1}(x-y)} \cdot f_0(y)g_0(x - y)dy \\ &= \int e^{-\frac{1}{2}\left(y^T F^{-1} y + (x-y)^T G^{-1}(x-y)\right)} \cdot f_0(y)g_0(x - y)dy. \end{aligned}$$

Applying (5) with $A = F^{-1}, B = G^{-1}$ and $C = H^{-1}$ yields $H^{-1} = (F+G)^{-1}$ and

$$\begin{aligned} h(x) &= \int e^{-\frac{1}{2}\left(x^T H^{-1} x + z^T (F^{-1}+G^{-1})z\right)} \cdot f_0(y)g_0(x - y)dy \\ &= e^{-\frac{1}{2}x^T H^{-1} x} \underbrace{\int e^{-\frac{1}{2}z^T (F^{-1}+G^{-1})z} \cdot f_0(y)g_0(x - y)dy}_{h_0(x)}. \end{aligned}$$

The integrand is log-concave since it is the product of a Gaussian function and two log-concave functions, and thus h_0 is log-concave according to Theorem 1. This proves that $h \in \mathcal{LC}(H^{-1}) = \mathcal{LC}\big((F + G)^{-1}\big)$.

A.3 Multiplication

Let $f \in \mathcal{LC}(F^{-1})$ and $g \in \mathcal{LC}(G^{-1})$. Then

$$\begin{aligned} h(x) &= f(x)g(x) \\ &= e^{-\frac{1}{2}x^T F^{-1} x} e^{-\frac{1}{2}x^T G^{-1} x} \cdot f_0(x)g_0(x) \\ &= e^{-\frac{1}{2}x^T (F^{-1}+G^{-1})x} \cdot h_0(x), \end{aligned}$$

where f_0 and g_0 are log-concave and $h_0(x) = f_0(x)g_0(x)$. Thus h_0 is log-concave and so $h \in \mathcal{LC}(F^{-1} + G^{-1})$.

B. Proof of Theorem 2

The factorization (4) will be central also in this proof. Consider first the theorem in one dimension. Let $f \in \mathcal{LC}(p^{-1}), p > 0$ be a nonincreasing probability density defined for $x \geq 0$. Then f can be factored as

$$f(x) = e^{-\frac{1}{2}p^{-1}x^2} f_0(x),$$

where $f_0(x), x \geq 0$ is log-concave. The right derivative $f'(0)$ exists since any convex function is almost everywhere differentiable which transfers trivially to log-concave functions. Furthermore $f_0'(0) = f'(0) \leq 0$, and since f_0 is log-concave it is nonincreasing for all $x \geq 0$.

Let $C > 0$ be defined such that

$$\int_0^\infty Ce^{-\frac{1}{2}p^{-1}x^2} dx = \int_0^\infty \underbrace{e^{-\frac{1}{2}p^{-1}x^2} f_0(x)}_{f(x)} dx = 1.$$

Then, since $f_0(x)$ is nonincreasing, there must exist some $x_0 > 0$ such that

$$\begin{aligned} f_0(x) &\geq C, \quad x < x_0 \\ f_0(x) &\leq C, \quad x > x_0. \end{aligned}$$

The second moment of f is

$$\begin{aligned} \int_0^\infty x^2 f(x) dx &= \int_0^\infty x^2 Ce^{-\frac{1}{2}p^{-1}x^2} dx + \int_0^\infty x^2 \Big(f(x) - Ce^{-\frac{1}{2}p^{-1}x^2} \Big) dx \\ &= p + \int_0^\infty x^2 e^{-\frac{1}{2}p^{-1}x^2} \Big(f_0(x) - C \Big) dx \\ &= p + \int_0^\infty e^{-\frac{1}{2}p^{-1}x^2} \Big(x_0^2 + (x^2 - x_0^2) \Big) \Big(f_0(x) - C \Big) dx \\ &\leq p + x_0^2 \int_0^\infty e^{-\frac{1}{2}p^{-1}x^2} \Big(f_0(x) - C \Big) dx \\ &= p, \end{aligned}$$

where we have used that $(x^2 - x_0^2)\big(f_0(x) - C\big) \leq 0$. Thus the second moment of f around $x = 0$ is $\leq p$.

Now assume that $f(x) \in \mathcal{LC}(p^{-1})$ is an arbitrary strongly log-concave function in one dimension that assumes its maximum value at $x = M_x$. All strongly log-concave functions are bounded and go to zero as $|x| \to \infty$, so if f does not assume its maximum it can be made to do so by changing the value in one point, which does not affect integrals of f and preserves strong log-concavity. Then $g(x) = f(x - M_x)$ can be written as a convex combination of two probability densities in $\mathcal{LC}(p^{-1})$ such that one has support on $x < 0$ and is nondecreasing and one has support on $x \geq 0$ and is nonincreasing. The second moment of g around 0 is a convex combination of the moments of the two densities, and so

$$\int (x - M_x)^2 f(x) dx \leq p.$$

Since the covariance of the density f is the minimum of the second moment around any point,

$$\int (x - m_x)^2 f(x) dx = \min_y \int (x - y)^2 f(x) dx \leq p,$$

where m_x is expectation of the density. This proves the theorem in one dimension.

For the proof in R^n we shall need another matrix inequality. In the Cauchy-Schwartz inequality $(u^T v)^2 \leq (u^T u)(v^T v)$, let $u = P^{-\frac{1}{2}} x$ and $v = P^{\frac{1}{2}} e_z$, where $P > 0$, $||e_z|| = 1$. This yields

$$\begin{aligned} (x^T e_z)^2 &\leq (x^T P^{-1} x)(e_z^T P e_z) \\ \Longrightarrow x^T e_z (e_z^T P e_z)^{-1} e_z^T x &\leq x^T P^{-1} x \\ \Longrightarrow Q_r = e_z (e_z^T P e_z)^{-1} e_z^T &\leq P^{-1}. \end{aligned} \tag{6}$$

Now consider a density $f \in \mathcal{LC}(P^{-1}), P > 0$. Without loss of generality, assume the expectation to be zero. The covariance is then

$$V = \int x x^T f(x) dx,$$

and for a given unit vector e_z,

$$e_z^T V e_z = \int (e_z^T x)^2 f(x) dx = \int_{t \in \mathbb{R}} t^2 \underbrace{\int_{y \perp e_z} f(t e_z + y) dy}_{g(t)} dt,$$

where $x = t e_z + y$ and $g(t)$ is the marginal density of f in the e_z direction, having zero expectation. We see that

$$\begin{aligned} g(t) &= \int_{y \perp e_z} e^{-\frac{1}{2}x^T P^{-1} x} f_0(x) dy \\ &= \int_{y \perp e_z} e^{-\frac{1}{2}x^T Q_r x} e^{-\frac{1}{2}x^T (P^{-1} - Q_r) x} f_0(x) dy \\ &= e^{-\frac{1}{2}(e_z^T P e_z)^{-1} t^2} \underbrace{\int_{y \perp e_z} e^{-\frac{1}{2}x^T (P^{-1} - Q_r) x} f_0(x) dy,}_{g_0(t)} \end{aligned}$$

since $y^T e_z = 0$ so that $x^T Q_r x = t e_z^T Q_r e_z t = (e_z^T P e_z)^{-1} t^2$. From (6) $P^{-1} - Q_r \geq 0$ so that g_0 is log-concave. Thus $g \in \mathcal{LC}\big((e_z^T P e_z)^{-1}\big)$ so that

$$e_z^T V e_z \leq e_z^T P e_z \Longrightarrow V \leq P,$$

which proves the theorem. The bound is tight for the corresponding Gaussian by definition.

Paper VII

Recursive State Estimation for Linear Systems with Mixed Stochastic and Set-Bounded Disturbances

Toivo Henningsson

Abstract

Recursive state estimation is considered for discrete time linear systems with mixed process and measurement disturbances that have stochastic and (convex) set-bounded terms. The state estimate is formed as a linear combination of initial guess and measurements, giving an estimation error of the same mixed type (and causing minimal interference between the two kinds of error). An ellipsoidal over-approximation to the set-bounded estimation error term allows to formulate a linear matrix inequality (LMI) for optimization of the filter gain, considering both parts of the estimation error in the objective. With purely stochastic disturbances, the standard Kalman Filter is recovered. The state estimator is shown to work well for an event based estimation example, where measurements are very coarsely quantized.

1. Introduction

In many control systems, there exist some disturbances that are best modelled as stochastic, and other disturbances that are better modelled as set-bounded uncertainties. The classical approach to state estimation in such cases is to approximate the set-bounded uncertainties by stochastic ones, allowing to use a standard Kalman Filter. Another approach is to approximate the stochastic disturbances by set-bounded ones, and use a state estimator for set-bounded uncertainty.

It is, however, not straightforward to translate between stochastic and set-bounded disturbances, since they do not combine in the same way. Two measurements of the same variable with independent identically distributed (I.I.D.) stochastic noise combine to form an estimate with only half the error variance. Two measurements with set-bounded uncertainty $y_i = x + z_i, |z_i| \leq 1$ may on the other hand be little better than just one if $y_1 \approx y_2$, not uncommon of situations where this kind of disturbance model is applied.

Thus, it is useful to be able to deal with both kinds of disturbances at the same time. The contribution of this paper is the formulation of an estimator that can deal with general state estimation problems with mixed disturbances. The optimization of the filter gain required in each step is expressed as an LMI. Since the basic structure is that of a Kalman Filter, the estimator reduces to a Kalman Filter in the case of purely stochastic disturbances.

There is much previous work for the cases of only stochastic or only set-bounded disturbances, and also some variations on mixing the two. With only stochastic disturbances, the optimal solution is the classical Kalman Filter (see [Kalman, 1960], [Kalman and Bucy, 1961]). State estimation with set bounded disturbances is considered in [Bertsekas and Rhodes, 1971] and [Durieu *et al.*, 2001]. Kalman Filtering with a set-bounded initial expectation in the prior is treated in [Morrell and Stirling, 1988]. For a different approach to mixed disturbance estimation, see [Hanebeck and Horn, 2001] and references therein.

When dealing with set-bounded disturbances, there is the issue of how to represent the uncertainty sets that arise as data is combined. Unlike Gaussian noise, there is no general exact closed form representation of limited complexity. We first present the general equations, which can be used with polytopic uncertainty sets. These will however grow quickly in complexity. We will thus focus on the ellipsoidal approximation of uncertainty sets; together with a recursive formulation of the estimator this gives a fixed complexity for the estimator operations.

The rest of the paper is laid out as follows. The mixed state estimation problem to be solved is stated in section 2, including the basic estimator structure. Section 3 covers some preliminaries used in the solution. The first step of the solution is taken in section 4, which shows how to decompose the problem into the stochastic part, treated in section 5, and the set-bounded part, treated in section 6. The latter section contains the central theorem to express the set-bounded part of the filter's optimization criterion for a combination of polytopic and ellipsoidal uncertainties, which is proved in the appendix. Section 7 compares the proposed estimator with a grid based Bayesian estimator and a Kalman Filter for an example problem. Conclusions are given in section 8.

2. Problem Formulation

The objective is to perform recursive state estimation for discrete time dynamic systems modelled by

$$x_k = Ax_{k-1} + u_{k-1} + e^{\text{proc.}}_{k-1} \tag{1}$$

$$y_k = Cx_k + e^{\text{meas.}}_k \tag{2}$$

where A and C are the dynamics and measurements matrices, and the state x_k, the known control input u_k, the measurements y_k, the process disturbance $e^{\text{proc.}}_k$, and the measurement disturbance $e^{\text{meas.}}_k$ are vectors. Also A and C may be time dependent.

All error terms e^i are the sum of a stochastic term w^i and a set-bounded term δ^i,

$$\begin{aligned} e^i &= w^i + \delta^i \\ \mathrm{E}(w^i) &= 0, \qquad \mathrm{E}(w^i(w^i)^T) = R^i \\ \delta^i &\in \Delta^i \end{aligned}$$

for some positive semidefinite covariance matrix R^i and convex uncertainty set Δ^i. The stochastic terms of the process and measurement disturbance $w^{\text{proc.}}_k$ and $w^{\text{meas.}}_k$ for all times are assumed mutually uncorrelated.

Given the system above and an initial state estimate $\hat{x}_0$ with mixed error

$$e_0 = x_0 - \hat{x}_0$$

we want to form a running state estimate as a linear combination of the initial state and the measurements. The dynamics (1) are used to form the *predicted* estimate $\hat{x}_{k|k-1}$ from the previous *filtered* estimate $\hat{x}_{k-1|k-1}$:

$$\hat{x}_{k|k-1} = A\hat{x}_{k-1|k-1} + u_{k-1}. \tag{3}$$

The measurement y_k is then used to form the current filtered estimate

$$\begin{aligned}\hat{x}_{k|k} &= \hat{x}_{k|k-1} + L_k\Big(y_k - C\hat{x}_{k|k-1}\Big) \\ &= \underbrace{\begin{pmatrix} I - L_kC & L_k \end{pmatrix}}_{X_k} \begin{pmatrix} \hat{x}_{k|k-1} \\ y_k \end{pmatrix}\end{aligned} \tag{4}$$

using some suitable filter gain L_k. We wish to choose L_k to minimize the estimation error in some appropriate sense. The matrix X_k specifies how to weigh together the predicted state estimate and the current measurement, and represents the action of the filtering step.

3. Notation and Preliminaries

The Minkowski sum of two sets X_k and Y is defined as

$$X + Y = \{x + y; x \in X, y \in Y\}.$$

Similarly, we will let the sum $X + y$ of a set X and a vector y be the translation $X + \{y\}$. The product of a set X and a matrix A will be interpreted as the element-wise product

$$AX = \{Ax; x \in X\}.$$

We will also use the product of two sets X, Y as the stacked Cartesian product

$$X \times Y = \left\{\begin{pmatrix} x \\ y \end{pmatrix}; x \in X, y \in Y\right\}.$$

For a matrix A, we denote by $A > 0$ ($A \geq 0$) that A is positive (semi-)definite. For a block matrix

$$M = \begin{pmatrix} A & B \\ B^T & D \end{pmatrix}$$

with $D > 0$, the conditions that $M \geq 0$ and that the Schur Complement (see [Boyd *et al.*, 1994, ch. 2.1, pp. 7-8]) of D in M

$$\Delta = A - BD^{-1}B^T$$

is positive semidefinite, $\Delta \geq 0$, are equivalent.

4. Problem Decomposition

We begin by decomposing the problem into a stochastic and a set-bounded part. The dynamics (1) combined with the prediction (3) gives the next prediction error

$$e_{k|k-1} = Ae_{k-1|k-1} + e_{k-1}^{\text{proc.}} \tag{5}$$

while the measurement equation (2) combined with the filtering step (4) gives the next filtered error

$$e_{k|k} = X_k \begin{pmatrix} e_{k|k-1} \\ e_k^{\text{meas.}} \end{pmatrix}. \tag{6}$$

The minimization of the expected/worst-case estimation error will guide the selection of the filter gain L_k, which will then be used to update the point estimate according to (4). L_k can be optimized online, or, since it is independent of the point estimate, it can be calculated ahead of time if the disturbance characteristics are known, e.g. if they are periodic or stationary.

The estimation errors $e_{k|k-1}$ and $e_{k|k}$ are composed of a stochastic and a set-bounded part, and are formed by forming each part separately. The two parts will be coupled only in the search for the optimal filter gain L_k in the filtering step, which we find by minimizing the cost function

$$V(L) = \text{tr } W\big(R_{k|k}(L) + \alpha r(L)^2 P(L)\big) \tag{7}$$

where $W > 0$ is a weight on the estimation error for different states, $\alpha > 0$ is the relative penalty on set-bounded error, $R_{k|k}(L)$ is the filtered error covariance, and $P_k(L)$ and $r(L)$ bound the set-bounded error after filtering $\delta_{k|k} \in \Delta_{k|k}(L)$ inside an ellipsoid:

$$\delta_{k|k}^T P(L)^{-1} \delta_{k|k} \le r(L)^2 \qquad \forall\, \delta_{k|k} \in \Delta_{k|k}(L). \tag{8}$$

Either P or r can be fixed for the optimization step, depending on whether we want to prespecify the shape of the ellipsoid circumscribed around $\Delta_{k|k}(L)$.

To carry out the minimization, we take the following steps:

- Form LMI conditions linear in L for
  - the stochastic part: $R \ge R_{k|k}(L)$
  - the set-bounded part: (P, r) satisfying (8)
- Minimize

  $$\bar{V} = \text{tr } W(R + \alpha r^2 P)$$

  under these LMI conditions.

When we introduce ellipsoidal approximation of the set-bounded error $\Delta_{k|k}$, we will merge the prediction and filtering steps for this part to reduce conservatism.

5. Stochastic Part

We consider the update and optimization of the stochastic estimation error terms. The prediction and filtering steps (5) and (6) give the stochastic error covariances

$$R_{k|k-1} = AR_{k-1|k-1}A^T + R_{k-1}^{\text{proc.}} \tag{9}$$

$$R_{k|k} = X_k \underbrace{\begin{pmatrix} R_{k|k-1} & 0 \\ 0 & R_k^{\text{meas.}} \end{pmatrix}}_{R_k^{\text{pm}}} X_k^T \tag{10}$$

for $w_{k|k-1}$ and $w_{k|k}$ respectively, since if $\mathrm{E}(ww^T) = R$,

$$\mathrm{E}\Big((Aw)(Aw)^T\Big) = A\,\mathrm{E}(ww^T)A^T = ARA^T.$$

The prediction step (9) is straightforward. To form an LMI for the filtering step (10), we first factor R_k^{pm} as

$$R_k^{\text{pm}} = SR^{\text{pm0}}S^T, \qquad R^{\text{pm0}} > 0.$$

By the Schur Complement, the condition $R \geq R_{k|k}$ or

$$R - X_k S R^{\text{pm0}} S^T X_k^T \geq 0$$

is then equivalent (since $R^{\text{pm0}} > 0$) to the LMI

$$\begin{pmatrix} R & X_k S \\ S^T X_k^T & (R^{\text{pm0}})^{-1} \end{pmatrix} \geq 0,$$

which is linear in L and R.

6. Set-Bounded Part

We now consider the update and optimization of the set-bounded estimation error terms. The operations are first formulated for general uncertainty sets, and then the case of ellipsoidal over-approximation is treated.

6.1 General Uncertainty Sets

From the prediction step (5), we must have $\delta_{k|k-1} \in \Delta_{k|k-1}$,

$$\Delta_{k|k-1} = A\Delta_{k-1|k-1} + \Delta_{k-1}^{\text{proc.}}.$$

If $\Delta_{k-1|k-1}$ and $\Delta_{k-1}^{\text{proc.}}$ are polytopes, so is $\Delta_{k|k-1}$.

For the filtering step, we have

$$\delta_{k|k} = X_k \underbrace{\begin{pmatrix} \delta_{k|k-1} \\ \delta_k^{\text{meas.}} \end{pmatrix}}_{\delta_k^{\text{pm}}}.$$

The constraint (8) can be expressed for any $\delta_k^{\text{pm}} \in \Delta_k^{\text{pm}} = \Delta_{k|k-1} \times \Delta_k^{\text{meas.}}$ as a second order cone constraint when P is fixed:

$$r \geq ||P^{-\frac{1}{2}}\delta_{k|k}|| = ||P^{-\frac{1}{2}}X_k\delta_k^{\text{pm}}||$$

or in general by the Schur Complement (since $P > 0$) as an LMI

$$\begin{aligned} & r^2 - (\delta_k^{\text{pm}})^T X_k^T P^{-1} X_k \delta_k^{\text{pm}} \geq 0 \\ \iff & \begin{pmatrix} P & X_k\delta_k^{\text{pm}} \\ (\delta_k^{\text{pm}})^T X_k^T & r^2 \end{pmatrix} \geq 0. \end{aligned}$$

If Δ_k^{pm} is a polytope, it is enough to consider the constraint at the vertices, since an ellipsoid contains a set of vertices iff it contains the convex hull of those vertices (the polytope).

6.2 Ellipsoidal Uncertainty Sets

Now suppose that the filtered set-bounded error from the previous step $\Delta_{k-1|k-1}$, and possibly the process or measurement disturbance parts $\Delta_{k-1}^{\text{proc.}}$ and $\Delta_k^{\text{meas.}}$, are described by ellipsoids. In this case we can use the ellipsoid (8) to find an ellipsoidal over-approximation for $\Delta_{k|k}$ to use in the next step. To formulate (8) as an LMI in this case, we need the following theorem.

Theorem 1—Ellipsoid Bounding Weighted Ellipsoid Sum
Given a number of ellipsoids $\mathcal{E}_i, i = 1 \ldots n$:

$$z_i \in \mathcal{E}_i \iff \begin{cases} z_i = G_i x_i + b_i \\ x_i^T Q_i x_i \leq r_i^2 \end{cases}$$

the weighted Minkowski sum

$$\mathcal{A} = X \sum_i \mathcal{E}_i = \left\{ x = Xz; z = \sum_i z_i, z_i \in \mathcal{E}_i \; \forall i \right\}$$

can be proved by the S-procedure (see [Boyd *et al.*, 1994, ch. 2.6.3, pp. 23-24]) to be contained in the centered target ellipsoid $\mathcal{E}$,

$$x \in \mathcal{E} \iff x^T P^{-1} x \le r^2 \tag{11}$$

iff the LMI condition

$$\begin{pmatrix} P & XG & Xb \\ G^T X^T & Q\tau & \\ b^T X^T & & r^2 - \sum_i \tau_i r_i^2 \end{pmatrix} \ge 0 \tag{12}$$

is satisfied for some scalars $\tau_i \ge 0$, where $b = \sum_i b_i$, and

$$G = \begin{pmatrix} G_1 & G_2 & \dots & G_n \end{pmatrix}, \qquad Q_\tau = \mathrm{diag}\big(\{\tau_i Q_i\}_i\big).$$

If $n = 1$ and $r_1 > 0$, the condition (12) is also necessary for $\mathcal{A} \subseteq \mathcal{E}$. □

Proof: See the appendix.

Using the theorem. We let $P = P$ and $z = \delta_k^{\mathrm{pm}}$, where Δ_k^{pm} is a sum of ellipsoids. With one centered ellipsoid ($b_i = 0$) containing each of the previous filtered error, the process and measurement disturbances:

$$\Delta_{k-1|k-1} \subseteq \mathcal{E}_1, \qquad \Delta_{k-1}^{\mathrm{proc.}} \subseteq \mathcal{E}_2, \qquad \Delta_k^{\mathrm{meas.}} \subseteq \mathcal{E}_3$$

the set-bounded part gets the prediction step $\Delta_{k|k-1} \subseteq A\mathcal{E}_1 + \mathcal{E}_2$ and the filtering step

$$\Delta_{k|k} \subseteq X_k(\Delta_{k|k-1} \times \mathcal{E}_3) \subseteq X_k\big((A\mathcal{E}_1 + \mathcal{E}_2) \times \mathcal{E}_3\big).$$

The ellipsoid sum for $\Delta_{k|k}$ can thus be expressed with the theorem, plugging in the ellipsoids $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$, and

$$G_1 = \begin{pmatrix} A \\ 0 \end{pmatrix}, \quad G_2 = \begin{pmatrix} I \\ 0 \end{pmatrix}, \quad G_3 = \begin{pmatrix} 0 \\ I \end{pmatrix}.$$

Thus we can use the LMI condition (12) to circumscribe an ellipsoid around $\Delta_{k|k}$.

Variations. We can use more or fewer ellipsoidal terms for the uncertainty sets Δ_i, and also polytopic terms. For polytopic terms, the sum $\mathcal{P}$ of all such terms is first formed. As in the case with only polytopic terms, the LMI must be written once for each vertex of $\mathcal{P}$. If $\mathcal{P}$ is symmetric, we need only write half as many LMIs since the centered target ellipsoid $\mathcal{E}$

sees no difference between the vertices v and $-v$. A polytope vertex can be represented by a zero-dimensional ellipsoid with $b_i \neq 0$.

A polytope that is the sum of one-dimensional polytopes (line segments) may expressed more economically as a sum of one-dimensional ellipsoids. However, the result may be more conservative since forming the sum of ellipsoids relies on the S-procedure.

The use of both P and r as variables in the condition (11) for the target ellipsoid may seem redundant, but it allows to state a possibly simpler optimization problem if the shape of the target ellipsoid is fixed. (I.e. to some shape desired in a stationary situation.) It is of course possible to constrain P to other spaces than to be fully free or with a prespecified shape. Another use for r could be to improve the numerical conditioning of the optimization problem by guessing the size of the resulting ellipsoid before optimizing for P.

7. Simulations

7.1 Example System

Consider a double integrator process with dynamics

$$x_{k+1} = \underbrace{\begin{pmatrix} 1 & h \\ 0 & 1 \end{pmatrix}}_{A} x_k + \underbrace{\begin{pmatrix} \frac{1}{2}h^2 \\ h \end{pmatrix}}_{B} u_k + w_k^{\text{proc.}}$$

$$\mathrm{E}\big(w_k^{\text{proc.}}\big) = 0, \quad \mathrm{E}\Big(\big(w_k^{\text{proc.}} w_k^{\text{proc.}}\big)^T\Big) = \underbrace{\frac{1}{4}\begin{pmatrix} \frac{1}{3}h^3 & \frac{1}{2}h^2 \\ \frac{1}{2}h^2 & h \end{pmatrix}}_{R_k^{\text{proc.}}}$$

where $h = 0.1$ is the sample time, $(x_k)_1$ is the position and $(x_k)_2$ the velocity. White process noise enters along with the control acceleration u_k.

The measurements are coarsely quantized:

$$y_k = \mathrm{round}\big(Cx_k\big), \quad C = \begin{pmatrix} 1 & 0 \end{pmatrix},$$

where round(x) rounds x to the nearest integer. Using the current framework, we can model the measurement by

$$y_k = Cx_k + \delta_k^{\text{meas.}}, \quad \delta_k^{\text{meas.}} \in \Delta_k^{\text{meas.}} = [-\tfrac{1}{2}, \tfrac{1}{2}].$$

With the sampling time h small enough, we may consider $(x_k)_1$ to be almost completely known at all *events*, when y_k changes value. This measurement may be modelled as

$$
\begin{aligned}
\frac{1}{2}\big(y_k + y_{k-1}\big) &= Cx_k + w_k^{\text{meas.}}, \\
\mathrm{E}\big(w_k^{\text{meas.}}\big) = 0, \quad \mathrm{E}\Big(&(w_k^{\text{meas.}} w_k^{\text{meas.}})^T\Big) = R_k^{\text{meas.}},
\end{aligned}
\tag{13}
$$

where $R_k^{\text{meas.}}$ gives a suitable approximation of the error in the guess $Cx_k \approx \frac{1}{2}\big(y_k + y_{k-1}\big)$. We take $R_k^{\text{meas.}} = (R_k^{\text{proc.}})_{11}$.

Since the system is unstable, we stabilize it with the control law

$$
u_k = -\begin{pmatrix} 1 & 2 \end{pmatrix} \hat{x}_k,
$$

which places the poles in approximately $z = e^{-h}$. The state estimate $\hat{x}_k$ is taken from a simple heuristic state estimator that:

- runs in open loop between events
- updates at events:

$$
\begin{aligned}
(\hat{x}_k)_1 &= \frac{1}{2}\Big(y_k + y_{k-1}\Big) \\
(\hat{x}_k)_2 &= \frac{(\hat{x}_k)_1 - (\hat{x}_{k_{\text{last}}})_1}{h(k - k_{\text{last}})}
\end{aligned}
$$

  where k_{last} is the time index of the last event or known initial state.

The process was simulated with the heuristic controller to produce the test sequence u_k, y_k in Fig. 1. The corresponding state sequence x_k can seen in Fig. 2 (together with state estimates from different estimators).

7.2 Estimator Implementation For The Example

In this example, the process noise is purely stochastic, and the set-bounded measurement error $\Delta_k^{\text{meas.}}$ can be represented as an interval symmetric around the origin, so the target ellipsoid $\mathcal{E} \supseteq \Delta_{k|k}$ should enclose the sum of an ellipsoid for $\Delta_{k|k-1}$ and the polytope for $\Delta_k^{\text{meas.}}$. Since we have only one ellipsoid in the sum, (12) is both necessary and sufficient for the target ellipsoid $\mathcal{E}$ to enclose it. Since the polytope $\Delta_k^{\text{meas.}}$ is symmetric with two vertices, we need only one instance of the LMI condition (12).

7.3 Performance Comparison

Three filters were compared on the test sequence:

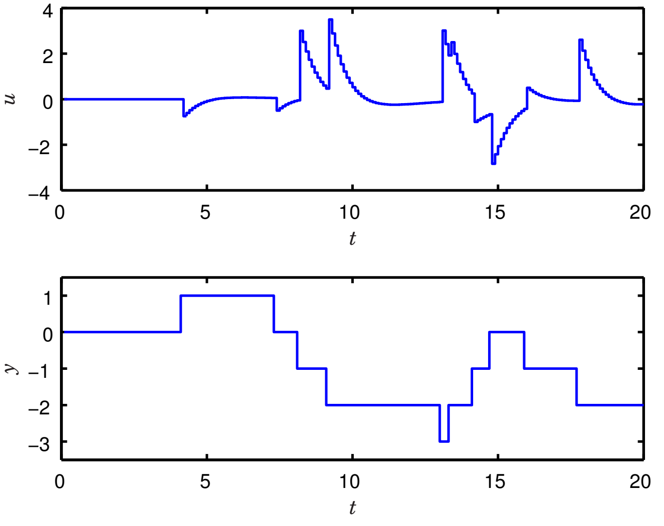

Figure 1. Test sequence for the observers

Table 1. Mean quadratic errors over a 10^5 time step test sequence.

E_{Mixed}	E_{Grid}	E_{Kalman}
$\begin{pmatrix} 0.054 & 0.063 \\ 0.063 & 0.180 \end{pmatrix}$	$\begin{pmatrix} 0.045 & 0.053 \\ 0.053 & 0.157 \end{pmatrix}$	$\begin{pmatrix} 0.444 & 0.223 \\ 0.223 & 0.285 \end{pmatrix}$

- The Mixed Estimator proposed in this paper using ellipsoidal overbounding of $\Delta_{k|k}$ in each step, with

$$\alpha = 1, \qquad W = \begin{pmatrix} 1 & -0.3 \\ -0.3 & 0.4 \end{pmatrix}.$$

  The weight matrix W was chosen by letting W^{-1} be roughly proportional to the error covariance of the Grid Filter (see below) a long time after an event.

- A *Grid Filter*; a discretization of the Bayesian Estimator for the system (with approximately 32 000 states). See [Henningsson and Åström, 2006] for more about the Bayesian Estimator for this system.

- A Kalman Filter that uses only the measurements (13) at events, and runs in open loop in between.

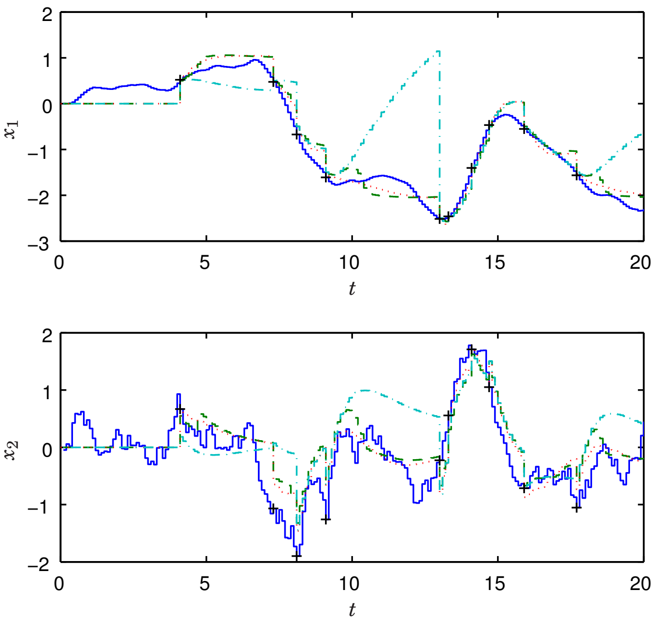

Figure 2. Actual states and state estimates generated by the observers. Actual states (solid), Mixed Estimator (dashed), Grid Filter (dotted), Kalman Filter (dash-dotted). Events are marked with a + sign.

Table 1 shows the average estimation error of the filters over a test sequence of 10^5 time steps, evaluated as

$$E = \frac{1}{N}\sum_{k=1}^{N}(x_k - \hat{x}_k)(x_k - \hat{x}_k)^T.$$

The Mixed Estimator is seen to come quite close to the Grid Filter performance, but the Kalman Filter is far behind. Fig. 2 shows actual state trajectories together with the estimates. Events are marked with + signs. When events are frequent, all estimators seem to follow the state trajectories reasonably well, especially for the position x_1. When there is longer time between events, the Kalman Filter seems to lose track. The Mixed Filter is much better at following the Bayesian estimate. The strategy it uses seems to be something like:

- At an event, update the state estimate.
- Continue by open loop predictions some time after each event, while the prediction error is small.

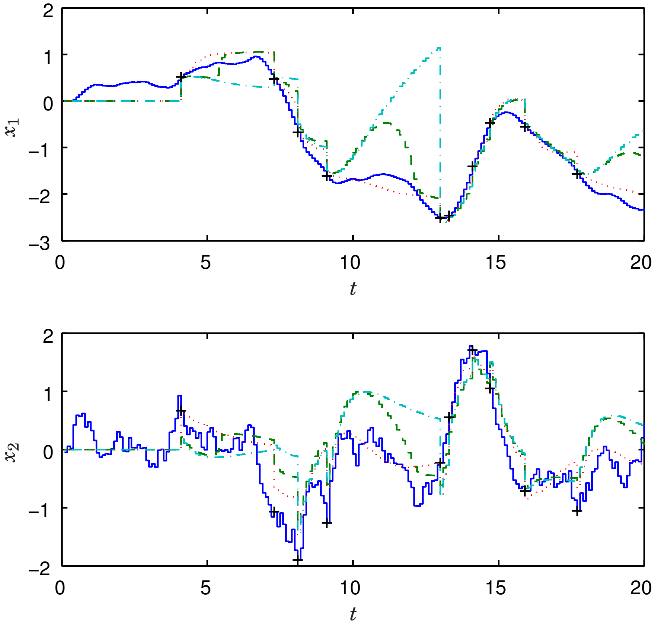

Figure 3. Actual states and state estimates, with $\alpha = 10$ for the Mixed Estimator, which makes it follow the Kalman Filter for too long.

- When the prediction error becomes too large, start to incorporate the imprecise measurements available.

Fig. 3 shows the same simulation with $\alpha = 10$ for the Mixed Estimator. The weight α adjusts the tradeoff between stochastic and set-bounded estimation error. With higher α it is seen that the Mixed Filter waits longer to incorporate the uncertain measurements after each events. The value $\alpha = 1$ used in Fig. 2 seems to give a more reasonable tradeoff.

The uncertainty set $\Delta_{k|k}$ (a polytope in this example) and the recursive ellipsoidal over-approximation $\hat{\Delta}_{k|k}$ used by the mixed filter can be seen in Fig. 4, just prior to the event at $t = 13$. The actual set takes up perhaps $\frac{2}{3}$ of the ellipsoid's volume, and that they more or less touch at the sharpest corners of the polytope.

8. Conclusion

This paper describes the design of a state estimator for linear systems with process and measurement disturbances containing both stochastic

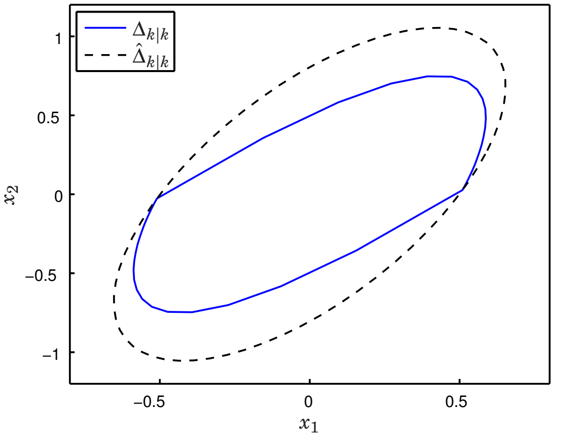

Figure 4. Actual set bounded error and ellipsoidal approximation used by the Mixed Filter at $t = 12.9$, just before an event.

and set bounded terms. The estimator structure that is borrowed from the Kalman Filter is optimal for purely stochastic disturbances, and allows the two parts of the estimation error to be treated efficiently and almost independently. The filter gain is optimized by solving a Linear Matrix Inequality (LMI) problem.

The estimator can value the usefulness of measurements corrupted by different amounts of stochastic and set bounded disturbances, with a parameter α that can be used to tune the tradeoff between the two kinds of error. An example shows that the estimator performs quite close to an optimal Bayesian Estimator, and that α can be used to adjust how long to wait after receiving a good measurement before incorporating measurements with interval uncertainty.

The estimator reproduces the behavior of the Kalman Filter with set-bounded initial expectation in [Morrell and Stirling, 1988] under the circumstances assumed in that work, when the weight α goes to zero. When α is nonzero, the estimator applies a higher filter gain to eliminate the set-bounded uncertainty faster.

An open issue is how to choose the state weighting matrix W in a systematic fashion.

Acknowledgments

The author would like to thank A. Cervin for valuable feedback, and A. Gattami for helpful tips on LMIs.

References

Bertsekas, D. and I. Rhodes (1971): "Recursive state estimation for a set-membership description of uncertainty." *IEEE Transactions on Automatic Control*, **16:2**, pp. 117–128.

Boyd, S., L. El Ghaoui, E. Feron, and V. Balakrishnan (1994): *Linear Matrix Inequalities in System and Control Theory*, vol. 15 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA.

Durieu, C., E. Walter, and B. Polyak (2001): "Multi-input multi-output ellipsoidal state bounding." *Journal of Optimization Theory and Applications*, **111**, pp. 273–303(31).

Hanebeck, U. and J. Horn (2001): "New estimators for mixed stochastic and set theoretic uncertainty models: the general case." *Proceedings of the 2001 American Control Conference*, Arlington, Virginia.

Henningsson, T. and K. J. Åström (2006): "Log-concave observers." In *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems*. Kyoto, Japan.

Kalman, R. E. (1960): "A new approach to linear filtering and prediction problems." *Transactions of the ASME—Journal of Basic Engineering*, **82**, Series D, pp. 35–45.

Kalman, R. E. and R. S. Bucy (1961): "New results in linear filtering and prediction theory." *Transactions of the ASME—Journal of Basic Engineering*, **83**, pp. 95–107.

Morrell, D. and W. Stirling (1988): "Set-valued filtering and smoothing." *Twenty-Second Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, California.

A. Proof of Theorem 1

This development is based on [Boyd *et al.*, 1994, ch. 3.7.4, pp. 46-47]. The construction is extended to be linear in the transformation X, to handle ellipsoids that are flat in some dimensions, and to specify the centers b_i separately, but is reduced in that we are only interested in centered target ellipsoids $\mathcal{E}$.

To handle the Minkowski sum of ellipsoids, we need a condition for when one ellipsoid contains the intersection of a number of ellipsoids. Given a set of quadratic functions $\{f_i(x)\}_i, i = 1 \dots n$, one sufficient condition to verify that a quadratic function $f(x) \geq 0$ whenever all $f_i(x) \geq 0$

is given by the S-procedure:

$$\exists\, \tau_i \geq 0, i = 1 \dots n: \quad f(x) \geq \sum_i \tau_i f_i(x) \quad \forall x.$$

The condition is also necessary, e.g., when $n = 1$ and $f_1(x) > 0$ for some x, see [Boyd *et al.*, 1994, ch. 2.6.3, pp. 23-24].

The condition (12) which we seek to derive is formed by first constructing an extended space where each term of the ellipsoid sum has its own coordinates, and forming the set where all coordinates are within their respective ellipsoids, which is the intersection of ellipsoidal cylinders. We then used the S-procedure to circumscribe an ellipsoidal cylinder parametrized in the sum coordinates.

Let

$$x^T = \begin{pmatrix} x_1^T & x_2^T & \dots & x_n^T \end{pmatrix}, \qquad z = \sum_i z_i.$$

Then, according to the definitions in the theorem,

$$z = Gx + b = \underbrace{\begin{pmatrix} G & b \end{pmatrix}}_{G_e} \underbrace{\begin{pmatrix} x \\ 1 \end{pmatrix}}_{x_e} = G_e x_e.$$

We take the first step of the S-procedure (using $\tau_i \geq 0\, \forall i$) by forming the condition

$$\sum_i \tau_i (r_i^2 - x_i^T Q_i x_i) = \left(\sum_i \tau_i r_i^2 \right) - x^T Q_\tau x \geq 0 \tag{14}$$

which will always be fulfilled when $z_i \in \mathcal{E}_i\, \forall i$.

The condition for the target ellipsoid, $x \in \mathcal{E}$, $x = Xz = XG_e x_e$ is equivalent to

$$r^2 - x_e^T G_e^T X^T P^{-1} X G x_e \geq 0. \tag{15}$$

Subtracting (14) from (15), we form our S-procedure condition, which can clearly only be fulfilled for all x if (15) is fulfilled whenever (14) is:

$$x_e^T \left(\underbrace{\begin{pmatrix} Q_\tau & \\ & r^2 - \sum_i \tau_i r_i^2 \end{pmatrix}}_{Q_e} - G_e^T X^T P^{-1} X G_e \right) x_e \geq 0.$$

As we assume x to be arbitrary, we might as well assume x_e to be arbitrary since scaling of x_e with a nonzero constant does not affect whether the condition holds. The case when the last entry of x_e is zero is approached when $||x|| \to \infty$. Thus we can equivalently consider positive semidefiniteness of the matrix that stands between x_e^T and x_e above.

By the Schur Complement, since $P^{-1} > 0$, this condition is equivalent to

$$\begin{pmatrix} P & XG_e \\ G_e^T X^T & Q_e \end{pmatrix} \geq 0,$$

which is exactly (12).