

Reduced Complexity Control Systems

Roger W. Brockett

** School of Engineering and Applied Sciences, Harvard University,
Cambridge MA, e-mail: brockett@deas.harvard.edu*

Abstract: This paper discusses methods for the analysis and design of systems controlled, at least in part, by finite state controllers, such as might be found in adaptive or highly autonomous robotic systems. We focus on the role that feedback can play in simplifying the characterization of trajectories and, in particular, the extent to which elementary feedback rules based on finite state automata can be used to reduce the complexity of both the controller and the analysis. Finally, we introduce a new control paradigm based on randomized finite state controllers and present an analysis of a class of such systems.

1. INTRODUCTION

Over the last 50 years the most useful mathematical models for thinking about the analysis and design of control systems have been based on relatively uncomplicated abstractions such as linear systems, asymptotic stability, Gaussian random variables, etc. These abstractions provide an imperfect representation of reality; in the real world there are bounds on the inputs and the state variables, nonlinearities, limitations on the size of random variables, etc. Often the inaccuracies inherent in these abstractions are of little importance because they can be ignored or treated in an ad hoc way. However, in certain areas of growing importance, such as highly autonomous robotic systems, the familiar abstractions are less effective because they do not provide an efficient characterization of a large enough part of the problem and/or its solution. This may come about because there are multiple modes of operation, hard bounds on state variables, large and irregularly spaced discontinuities, etc. For such problems there is a need for efficient methods for piecing together localized descriptions in something like the way locally defined splines are pieced together in numerical analysis.

This process of analyzing complex systems by localizing and piecing together can take several forms. In the area of numerically controlled machine tools one pervasive idea involves language driven machines. In this case the set of possible paths is extremely large but complexity is managed by dividing the possible motions into easily characterized paths such as straight lines, circles, etc. These elementary paths are treated as words in a formal language which forms the instruction set for the machine. In this case simplification is achieved by localizing in time and at the expense of some loss in “expressiveness” in that one can not generate completely arbitrary paths. In other situations the localization occurs in the state space and is achieved by local high gain feedback loops which are used to force trajectories onto specific submanifolds or to “sandwich” trajectories between narrowly separated

submanifolds. In particular, in some aspects of systems biology one sees a large number of high gain feedback loops, essentially operating in on-off modes to keep variables within bounds. In this way systems with a large number of variables are kept within desired regions of the state space even when disturbances are present.

We are concerned here with systems whose description may involve: i) observations that are discontinuous (e.g., quantized data) and/or differential equations with discontinuous right-hand sides. ii) hard limits on the state variables and/or the controls. iii) communication constraints limiting the type of feedback signals available, iv) computational constraints restricting the complexity of the mapping from observations to the control values.

Our main points are captured with a few examples which are intended to illustrate the possibilities for progress on some of these problems. Specifically, we focus on three questions.

- How can we best harness the capability of feedback to simplify the implementation of control systems?
- How can we use finite automata to model and implement control strategies for continuous systems?
- In seeking to simplify the implementation of control laws, is there a role for randomized control laws?

2. REDUCING COMPLEXITY BY FEEDBACK

In spite of its central role in the subject of control, the concept of feedback and the reasons for using it prove to be surprisingly difficult to catalogue. In an introductory control course the scope may sufficiently narrow so that this not an important issue but it becomes important in the larger context where discussions frequently involve biology, economics, or social dynamics. The goal of the work reported in Egerstedt-Brockett [2003] was to show in a quantitative way that the use of feedback can simplify the set of instructions required to define a trajectory leading from point A to point B. The model used is a finite state machine and the feedback is based on the set of observations postulated to exist at each point in the state space, i.e., whatever observations are available at a particular time and at a specific state lying on the

* This work was supported in part by the Army Research Office, under the aMURI award N0. N00014-01-1-0859. The views expressed in this paper are those of the author alone and not of the sponsor.

path. The complexity of the most succinct open loop description of the path was compared with the complexity of descriptions which make use of the observations. It is not a surprise that the use of feedback can shorten the description of a provably correct navigation scheme, but the quantitative analysis serves to focus attention on which aspect of feedback are responsible for the improvement.

Feedback can simplify in a second sense. In biology and economics one sees extremely complicated systems operating in a decentralized ways, accomplishing very complex tasks. In these settings the feedback mechanisms at work appear to be rather simple rules based on locally observed prices, concentrations, weather, etc. In this context it is not hard to argue that any open loop description would be overwhelmingly complex. However, when one looks at the elemental processes in detail there are many, if not most, that involve discrete transactions, a cell divides or it does not, a pound of coffee is purchased or not, etc. Moreover, some of these decisions are frequently based on history; e.g., I bought coffee yesterday so I won't buy it today.

Although there are formal definitions of complexity that have proven to be useful in information theory and theoretical computer science (see, e.g., Cover [1991]), they are based on counting arguments and have a less compelling interpretation in settings where real numbers are involved. We feel that it is more promising to measure complexity in terms of the topology of the set of acceptable, or expected, trajectories.

3. DISCRETE CONTROL AND OBSERVATION

As a first step toward formalizing the ideas to be developed we define a type input-output models suitable for use when the system itself is described by ordinary differential equations but the inputs and outputs are limited to finite sets. Given a "physical" system modeled as

$$\dot{x} = f(x, t, u)$$

we consider two rather different ways in which the input set can be constrained. On one hand there is the idea of simply specifying a list of possible inputs $U = \{u_1, u_2, \dots, u_k\}$ and imposing the condition that these are the only admissible values for u but that one can switch from one of these values to another at any time. We will call this quantized control and call the set of values the control set. A second model, one which provides a better fit in some circumstances is to say that there is a set $U = \{u_1, u_2, \dots, u_k\}$ and at any time it is possible to apply an impulse of strength u_i whose effect is to displace the state vector by a certain amount. If the model is

$$\dot{x} = f(x) + g(x)u$$

and a unit strength impulse applied at $t = t_i$ is interpreted as having the effect of displacing x according to

$$x(t_i^+) = x(t_i^-) + g(x(t_i^-))$$

That we consider $x : [0, t] \rightarrow \mathbb{R}^n$ to be continuous from the left and consider the value of x to be "frozen" at he previous value until the impulse is fully delivered. We will refer to this as *quantized impulse control*.

An aspect of this interpretation that will be used below is that for special choices of the functions and initial

conditions these equations can provide a realization of a finite state machine. For example, if $u(t)$ is constrained to be a sequence of unit strength impulses and if $z(0) \in \{-1, 1\}$ then the solution of

$$\dot{z} = -2u(t)z(t)$$

will take on values in $\{-1, 1\}$ for all time.

Turning now to models for observation, Standard, uniformly spaced quantizers can be thought of as rescaled models of $q : \mathbb{R} \rightarrow \mathbb{Z}$ whereby a scalar x is mapped to an integer according to a rule such as $x \mapsto \lfloor x + 1/2 \rfloor$. This is idealized in that it may only be possible to accomodate values of x in a certain range and the system may require a certain time to "settle" to the right value. In cases where the quantization levels are widely separated it is sometimes necessary to look more carefully at the errors that can arise in the quantization process. As a rule, quantizers are subject to inaccuracy when the levels are not crossed cleanly because the derivative of the signal has a small absolute value when the value of the signal is near a threshold. This can be illustrated using a phase plot which shows regions in (\dot{x}, x) -space where the output of the quantizer is unreliable.




Fig. 1. Illustrating the regions of potential inaccuracy for a scalar quantizer in terms of subsets of phase space.

If quantization errors can not be ignored, it may be appropriate to treat them stochastically. One approach is to introduce a stochastic variable $z \in \{-1/2, 1/2\}$ and model the observation as

$$y(t) = q(x(t)) + z(t) ; dz = -2z dN$$

with N being a Poisson counter whose rate is given by $\phi(\dot{x}, x)$, with ϕ having the general features of the function

$$\phi = \sum_{i=-\infty}^{\infty} e^{-(\dot{x}^2 + (x-i)^2)/2\sigma}$$

with σ small.

We mention a second option. Astrom and Bernhardsson [2003] discuss the use of what they call *Lebesgue sampling*. It is defined in a stochastic control context, and its performance is compared with ordinary quantization. Here we make use of something very similar. Suppose that $x : [t_1, t_2] \rightarrow \mathbb{R}$ is a differentiable function. If a is a real number then there may or may not be solutions of the equation $x(t) = a$. Denote the set of all solutions by $x^{-1}(a_i)$, with the understanding that this set may be empty. As in fixed point theory, (see Milnor [1965]) we denote the number of points in this set by $\#x^{-1}(a_i)$. In some cases it happens that instead of observing a quantity $cx(t)$ or a noisy version, $cx(t) + n(t)$, we may only be able to observe the level crossings of $cx(t)$. In more generality, we may postulate a number of levels $\{a_1, a_2, \dots, a_k\}$ and assume that one can observe

$$y_i = \#(cx)^{-1}(a_i)$$

for y defined on $[0, t]$. We will return to this observation model in section 5.

4. STATIC APPROXIMATIONS FOR FEEDBACK

If we are to limit the input values to a finite set the values should be chosen to approximate, in some sense, the optimal feedback control law. There are many versions of the problem of approximating a continuous function on a compact interval by a piecewise constant function with a specified number of discontinuities. Suppose that the measure of fit is integral squared error, that the function to be approximated is $f(\cdot)$, and that the interval is $[x_0, x_{n+1}]$. This problem is then solved by minimizing with respect to a sequence of domain values $x_1 < x_2 < \dots < x_n$ and a sequence of range values $g_0, g_1, g_2, \dots, g_n$ the quantity

$$\eta = \sum_{i=0}^n \int_{x_i}^{x_{i+1}} (f(x) - g_i)^2 dx$$

Of course best values for the g_i are the average values f over the corresponding interval. The quality of the approximation depends on the the integral of the square of the deviation from these. The Wirtinger inequality asserts that if $f(0) = 0$ then

$$\int_0^a f^2(x) dx \leq \frac{4\pi^2}{a^2} \int_0^a \left(\frac{df}{dt}\right)^2 dt$$

and when applied to the problem at hand we get an inequality that shows that the integral of the square of the derivative of f controls the quality of the fit. If the approximation is assigned a complexity that grows with n then we see that the quality of the approximation improves with increasing n and that it improves with decreasing values of

$$\int \left(\frac{df}{dx}\right)^2 dx$$

Constraining the input set can be expected to harm performance, but will the degradation be significant? How much degradation in performance occurs if an optimal control law is approximated by a piecewise constant one? As a preliminary, and more tractable version of this problem, consider the following optimization problem which includes a penalty on the the "sensitivity" of the control with respect to the state. More precisely, we want a feedback control $u = f(x)$ but include in the performance measure a term of the form

$$\eta_1 = \int_R \left(\frac{\partial f}{\partial x}\right)^2 dx$$

which, of course, would be zero if f were constant. The following formulation addresses two issues: the fact that the control law only needs to work over a specific range of values and that it should have low sensitivity.

Example 1: Find f as a function of x such that for $u = f(x)$ and x governed by the first order system

$$\dot{x} = u ; x(0) = \pm a$$

we minimize

$$\eta = \int_0^\infty x^2 dt + \int_{\mathbb{R}} \left(\frac{\partial f}{\partial x}\right)^2 dx$$

Solution: From the equation $\dot{x} = f(x)$ we have

$$\frac{\dot{x}x^2}{f(x)} = x^2$$

Notice that for any feedback control f that drives x to zero we have

$$\int_0^\infty x^2 dt = \int_a^0 \frac{x^2}{f} dx$$

The corresponding Euler-Lagrange equation for the sum of this integral and the integral of the square of the derivative of f is

$$\frac{d^2 f}{dx^2} + \frac{x^2}{f^2} = 0$$

Clearly we have the boundary conditions $f(0) = 0$ and $df/dx|_a = 0$. The condition on the derivative at a is a transversality condition whose intuitive explanation is that the square of the derivative enters the integral and at the end point there is no benefit associated with it being nonzero. The equation can be solved numerically, sweeping out a range of initial conditions on the derivative at zero to find a solution that has $\partial f/\partial x|_a = 0$. The numerical solution is shown in the left-hand panel of Figure 2.




Fig. 2. Illustrating the best feedback law with a derivative penalty and a piecewise constant approximation.

Problem 1: Given an unstable system $\dot{x} = Ax + Bu$ with observations $y_i = \#(cx)^{-1}(a_i)$, find conditions on A, b, c and the m -vector a such that there exists a control law $u(y(t))$ having the property that solutions of $\dot{x} = Ax + \sum Bu(y_i)$ starting close to $x = 0$ are contained in a ball of a given radius. Compare with Brockett-Liberson [2000].

5. AUTOMATA AS OBSERVERS

In analyzing the properties of geometrical objects it is sometimes helpful to approximate the object of interest with a collection of piecewise flat structures, say plane triangles, and compute with, or reason about, the original object using such an approximation. Analogous thinking, applied to differential equations, suggests that it may be possible to replace dynamic compensation such as a lead/lag filter by an automaton which provides similar dynamic compensation. In this way one can hope to achieve a reduction in the complexity and perhaps improve reliability.

In some cases we can think of associating to a differential equation an automaton which will provide a rough model for its trajectories. For example, easy to see that if $x(k)$ takes on integer values then a clocked automaton of the simple form

$$x(k+1) = x(k) + 1 ; x \in \mathbb{Z}$$

has much in common with the integrator

$$\dot{x} = 1 ; x \in \mathbb{R}$$

In the same vein, a clocked automaton of the form

$$x(k+1) = x(k) + 1 \pmod{p}$$

models the angle $\tan^{-1}(\dot{x}/x)$ associated with the harmonic oscillator

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

when the harmonic oscillator is described in polar coordinates.

The modes associated with linear systems without repeated roots are either spirals, if the eigenvalue is complex or real exponentials. If we want track the qualitative aspects of the trajectories of such modes we need to know which type of behavior to look for. We begin with a variation on one of the classic problems in automatic control, an example usually described in terms of hysteresis. In the usual setting the hysteresis function has an internal state variable that takes on a continuum of values. What we now describe is simpler, although in the normal operating mode the trajectories are nearly the same. This solution illustrates a solution to a regulator problem using a three state automaton.

Example 2: Let c and d be real numbers $c > d > 0$ and consider the first order scalar system with observations

$$\dot{x} = -x + u$$

$$y_1 = \# \hat{x}^{-1}(c) ; ; y_1 = \hat{x}^{-1}(d)$$

with x being regarded as a function on $[0, t)$. Consider a three state automaton taking on the values 1, 0, -1 and driven by the observations in accordance with

$$\dot{z} = -(2z + 1) \frac{d}{dt} (\#x^{-1}(d)) + (-2z + 1) \frac{d}{dt} (\#x^{-1}(c))$$

With this definition of z , the control law

$$u(t) = az(k) ; a > 1$$

results in a trajectory that ultimately lies in the range

$$|x(t) - (c + d)/2| \leq r$$

for suitable r .

This solution is robust in the sense that if we replace the level crossing observation by the stochastic version of section 3 the resulting stochastic equation has similar properties.




Fig. 3. Typical open loop trajectory and automaton for closed loop control

The next example addresses a similar problem but now with a system whose response is oscillatory. In complete generality, the problem of determining information about $x(0)$ from an observation of the level crossings of $ce^{At}x(0)$ ranges from the well understood to the unrewarding, depending on how densely the level crossings are spaced and the degree of observability of the pair (A, c) . If the levels are narrowly spaced, as in a conventional quantizer, this is almost the same as recovering $x(0)$ from $ce^{At}x(0)$; however the situation becomes less clear if the eigenvalues of A are close to each other. The following second order case avoids this issue and allows us to focus on a situation with few levels.

Example 3: Consider the system

$$\ddot{x} + 2\xi\dot{x} + x = u$$

with $|\xi| < 1$. Of course if $u = 0$ this system has solutions that are exponentially weighted sinusoids

$$x(t) = e^{\xi t} \sin(\omega t + \phi)$$

Suppose now that $\xi < 0$ so that the uncontrolled system is unstable and that we wish to provide a feedback control that will keep the solution within some bound. For this purpose there are available observations in the form of the level crossings or coincidence detectors as discussed above.

$$y_1 = \#x^{-1}(1) ; y_0 = \#x^{-1}(0) ; y_{-1} = \#x^{-1}(-1)$$

with x being regarded as a function on $[0, t)$. As in the previous example, we want to construct an automaton driven by these observations and with a small number of states such that there exists a choice of u depending on the state of the automaton which accomplishes this task.




Fig. 4. The state diagram of an automata for controlling the system of the example.

The state of the automaton will be denoted by z with the labels for the states being such that $z(t) \in \{1, 2, 3, 4, 5\}$. Roughly speaking, the states will be identified with regions of the phase space in accordance with

- $z = 1$ if $x \in (0, 1) ; \dot{x} > 0$
- $z = 2$ if $x > 1$
- $z = 3$ if $|x| < 1 ; \dot{x} < 0$
- $z = 4$ if $x < -1$
- $z = 5$ if $x \in (-1, 0] ; \dot{x} > 0$

but the precise definition of the automaton is as follows. Let $I_k(z)$ be the indicator function for state k , i.e., $I_k(z)$ is one if z is in state k and zero otherwise. The evolution equation for the state is

$$\dot{z} = (I_1(z) + I_2 \frac{d}{dt} y_1 - 4I_0(z) \frac{d}{dt} y_0 + (I_3(z) + I_4(z)) \frac{d}{dt} y_{-1}$$

The level lines, together with the $\dot{x} = 0$ axis divide (\dot{x}, x) -space into 6 sectors but we merge two of these, $\dot{x} < 0, 1 > x > 0$ and $\dot{x} < 0, 0 > x > -1$ as indicated by the listing of the states. Thus if the trajectory x starts in sector i and $z(0) = i$ Z and x will stay synchronized. But if $i \neq j$ then a complete rotation will result in synchronization.

The control law itself is to be selected so that it is only nonzero when the automaton is in state 1. Observe that the uncontrolled system has period $2\pi/\sqrt{1 - \xi^2}$ and thus over one period the the uncontrolled dynamics would satisfy

$$\frac{d}{dt} (\dot{x}^2 + x^2) = -2\xi\dot{x}^2 > 0$$

When in state one the control should provide a pulse of size u_0 scaled so that $u_0\dot{x}$ removes somewhat more energy than the natural dynamics provides over one cycle. This need not be closely calibrated; it is only important that u_0 be scaled so that its strength is enough to more than offset the growth in $\dot{x}^2 + x^2$ that occurs as a result of the natural dynamics.

Problem 2: Given the unstable linear time invariant system

$$\dot{x} = Ax + bu ; ; y = c^T x$$




Fig. 5. A typical trajectory and the level crossing pattern.

Suppose that there exists no static feedback control law $u = f(y)$ that makes the null solution asymptotically stable. Under what additional assumption will it be possible to find a quantization scheme and an automaton with N or fewer states such that when it is used as part of a feedback system which is otherwise memoryless there is a ball, $S = \{x \mid \|x\| \leq a\}$, such that any $x(0) \in S$ generates a bounded trajectory.

6. RANDOMIZED FINITE STATE CONTROL

In this section we give the outline of a new and promising approach to finite state control. It is based on a particular kind of randomization involving finite state Markov processes whose transition rates are adjustable.

The subject of continuous time Markov chains is concerned with stochastic processes which take on a finite set of values and which jump discontinuously between these values at rates identified with the corresponding infinitesimal generator. That is, the evolution of the vector of probabilities whose i^{th} component is the probability that the system is in state i at time t takes the form $\dot{p} = Ap$ with A being a, possibly time varying, infinitesimally stochastic matrix. One general method for realizing a sample path description of such processes uses counting processes $N(t)$. These are random processes taking on values in the nonnegative integers, monotone increasing, and having a counting rate such that

$$\mathcal{E}(N(t) - N(\tau)) = \int_{\tau}^t \lambda(\sigma) d\sigma$$

Such counters can be used to generate sample paths using stochastic equations of the Itô type

$$dz = \sum_{i=1}^m f_i(z, t) dN_i \quad ; \quad z(t) \in S = \{s_1, s_2, \dots, s_k\}$$

provided that the f_i are chosen appropriately. If the counting rates are allowed to depend on other variables, it is possible to exercise some control over the transitions and there is a well developed stochastic calculus for computing statistical properties of the solutions. We will show that if we allow the counting rates to depend on the state of the system to be controlled then such systems can create a containment region in the state space of the system to be controlled, even though the controller is finite state and the feedback control takes on only a finite set of values. (Compare with Wong-Brockett [1999].) The size of the containment region and the size of the steady state error can be estimated from the properties of an associated variance equation. For suitable choices of the rate dependencies, this equation is linear if the system to be controlled is linear, even though the controller is a finite state system. In this way the model provides guidance as to how the the gains affect the containment region.

The general theory behind this circle of ideas will be given elsewhere but many of the basic ideas are present in the following example, which involves controlling a neutrally stable harmonic oscillator.

Example 4: Consider the system

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ b \end{bmatrix} z$$

$$dz = -2z dN \quad ; \quad z(0) \in \{\pm 1\}$$

with the rate of the counter N being a suitably chosen function of x . It is convenient to relabel z as x_3 and to describe the system and controller using a single vector equation expressed in Itô notation

$$\begin{bmatrix} dx_1 \\ dx_2 \\ dx_3 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} dt + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} dN$$

If we take expectations we can replace this by a deterministic equation, however first we must decide on how the counting rate is to depend on the state. This is a critical choice because it is the mechanism by which control is exercised. For reasons that will become clear we chose the form $\lambda + cx_2z$, noting that this is meaningful only for values of x such that $\lambda + cx_1z \geq 0$. There are two points to be made immediately. The first is that the dependence of the rate on z is actually just a notational convenience and will not be the source of technical difficulties. The second is that the bound $\lambda + cx_1z \geq 0$ has important implications for the outer limit of the domain of confinement as will be seen. With the given choice for the rate, the expected value of x satisfies

$$\frac{d}{dt} \mathcal{E} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & b \\ 2c & 0 & -2\lambda \end{bmatrix} \mathcal{E} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$$

The relevant characteristic equation is $s^3 + 2\lambda s^2 + s + 2(\lambda - bc) = 0$ so that if $0 < bc < \lambda$ then the eigenvalues have negative real parts. The corresponding equation for the matrix of second moments

$$\Sigma = \mathcal{E}xx^T = \mathcal{E} \begin{bmatrix} x_1^2 & x_1x_2 & x_1x_3 \\ x_1x_2 & x_2^2 & x_2x_3 \\ x_1x_3 & x_2x_3 & x_3^2 \end{bmatrix}$$

is given by

$$\dot{\Sigma} = F\Sigma + \Sigma F^T + G$$

with

$$F = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & b \\ 2c & 0 & -2\lambda \end{bmatrix} \quad ; \quad G = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & b \\ 0 & b & 0 \end{bmatrix}$$

Thus we have

$$\frac{d}{dt} \mathcal{E}x_1^2 = 2\mathcal{E}x_1x_2 \quad ; \quad \frac{d}{dt} \mathcal{E}x_1x_2 = 2c\mathcal{E}(x_2^2 - x_1^2 + bx_1x_3)$$

$$\frac{d}{dt} \mathcal{E}x_1x_3 = \mathcal{E}(x_2x_3 + 2cx_1^2 - 2\lambda x_1x_3)$$

$$\frac{d}{dt} \mathcal{E}x_2^2 = \mathcal{E}(-2x_1x_2) + 2b$$

$$\frac{d}{dt} \mathcal{E}x_2x_3 = 2c\mathcal{E}(x_1x_2 - x_1x_3 - 2x_2x_3) + 2b$$

Solving for the steady state values for these variances we have

$$\mathcal{E}x_1^2 = \frac{\lambda b}{\lambda + bc} \quad ; \quad \mathcal{E}x_1x_3 = \frac{bc}{\lambda + bc} \quad ; \quad \mathcal{E}x_2^2 = b \quad ; \quad \mathcal{E}x_3^2 = 1$$

and the remaining values are 0.

As noted, the model will only be meaningful in that region of (x_1, x_2) -space having the property that $\lambda + cx_1z \geq 0$. Because z can change to -1 at any time this means $\lambda > cx_1$. Given the expression for the variance, this suggests we make λ a multiple of the root mean square value of x_1 as computed above. Calling this multiple m we impose

$$\lambda > mc\sqrt{\frac{\lambda b}{\lambda + bc}}$$

With such a choice the containment region will be approximated by an annulus with inner radius $c\sqrt{\frac{\lambda b}{\lambda + bc}}$ and outer radius $mc\sqrt{\frac{\lambda b}{\lambda + bc}}$.

What if the system to be controlled had been exponentially unstable? To see the implications we need to examine anew the characteristic equation that lead to the stability condition. Suppose we have

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} \xi & 1 \\ -1 & \xi \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ b \end{bmatrix} z$$

$$dz = -2z dN \quad ; \quad z(0) \in \{\pm 1\}$$

the analysis proceeds as above yielding, eventually, the equation for the expected values

$$\frac{d}{dt} \mathcal{E} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} \xi & 1 & 0 \\ -1 & \xi & b \\ 2c & 0 & -2\lambda \end{bmatrix} \mathcal{E} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$$

The relevant characteristic equation is now more complex,

$$s^3 + 2(\lambda - \xi)s^2 + (1 - 4\xi\lambda + \xi^2)s + 2(\lambda + \lambda\xi^2 - bc) = 0$$

The presence of ξ takes away the possibility of achieving stability by letting λ become large and points to the need for controllers with more states.

We only have room to sketch the more general theory in which the finite state controller has an arbitrary number of states. In that setting it is convenient to represent the states as the n standard basis vectors in \mathbb{R}^n , i.e., the state space is $\{e_1, e_2, \dots, e_n\}$, as in Brockett [2008]. The sample path description is constructed using Poisson counters N_1, N_2, \dots, N_m in an Itô equation of the form

$$dx = \sum_{i=1}^m G_i x dN_i$$

with the rates of the counters being dependent on the state. That is, the transition probabilities associated with a Markov chain can be functions of x . The matrices G_i are chosen to have entries either zero or ± 1 . The off-diagonals are nonnegative and the columns sum to zero. The resulting Itô equation generates a Markov process whose transition probabilities are related to the rates of the Poisson counters in accordance with

$$A(t) = \sum_{i=1}^m G_i \lambda_i(t)$$

We refer this representation of the sample paths of a Markov process as a *unit vector representation*. Insofar as finite state Markov processes are concerned, this representation is completely general. Note that $\mathcal{E}x(t) = p(t)$.

If we combine this with the state evolution of a linear system we have

$$\begin{bmatrix} dx \\ dz \end{bmatrix} = \begin{bmatrix} A & B \\ C & 0 \end{bmatrix} \begin{bmatrix} x \\ z \end{bmatrix} dt + \sum \begin{bmatrix} 0 & 0 \\ 0 & G_i \end{bmatrix} \begin{bmatrix} x \\ z \end{bmatrix} dN_i$$

Now after a suitable x -dependent choice of the rates of the counters and after taking expectations we get a coupled set of equations of the form

$$\frac{d}{dt} \mathcal{E} \begin{bmatrix} x \\ z \end{bmatrix} = \begin{bmatrix} A & B \\ C & G \end{bmatrix} \mathcal{E} \begin{bmatrix} x \\ z \end{bmatrix}$$

The matrix B can be chosen freely but because the matrix C comes from an x -dependent choice of the counting rates it must have columns that sum to zero. This then provides a setting in which to study the stabilization and control of linear systems using randomized finite state machines.

7. CONCLUSIONS

At one point in time computer control was thought to mean sampled data control of linear systems; now it is more typically thought to encompass various aspects robotic control, vision guided control, etc. These new problems have led to an expanded view of what control should achieve, especially when it comes to problems involving many variables, high levels of autonomy, and distributed solutions. The subject of hybrid control attempts to bring many of these ideas together and is perhaps the most aggressive view of nonlinear control currently being studied. Our hope is that the ideas presented here will help inspire further progress.

REFERENCES

- Magnus Egerstedt and Roger Brockett, "Feedback can Reduce the Specification Complexity of Motor Programs," *IEEE Trans. on Automatic Control*, vol. 48, (Feb, 2003).
- John Milnor, *Topology from a differentiable viewpoint*, University Press at Virgina, Charlottesville, VA, 1965.
- Karl Astrom and Bo Bernhardsson, Systems with Lebesgue Sampling," in *Directions in Mathematical Systems Theory and Optimization*, LNCIS, (Anders Rantzer and Christopher Byrnes, Eds.) Springer-Verlag, 2003.
- Wing Wong and R. W. Brockett, "Systems with finite bandwidth constraints I: State Estimation Problems," *IEEE Trans. Automatic Control*, vol. 42, pp. 1294-1298, (1997).
- Wing Wong and R. W. Brockett, "Systems with Finite Communication Bandwidth Constraints II: Stabilization with Limited Information Feedback," *IEEE Trans. Automatic Control*, Vol. 44, pp. 1049-1053, (1999)
- Roger W. Brockett, "Optimal Control of the Liouville Equation", *Proceedings of the International Conference on Complex Geometry and Related Fields*, (Zhijie Chen et al. Eds.) American Mathematical Society, 2007 pp. 23-35.
- Roger Brockett and Daniel Liberzon, "Quantized feedback stabilization of linear systems," *IEEE Transactions on Automatic Control*, vol. 45,(2000) pp. 1279-1289.
- Thomas M. Cover, Joy A. Thomas. /Elements of information theory/, 1st Edition. New York: Wiley-Interscience, 1991. ISBN 0-471-06259-6
- Pascal Grieder, Michal Kvasnica, Mato Baotic' and Manfred Morari, "Stabilizing low complexity feedback control of constrained piecewise affine systems" *Automatica*, Vol.41, pp.1683-1694, 2005
- Roger Brockett, "Optimal Control of Observable Continuous Time Markov Chains," (submitted, 2008)