Chapter 2
Notes on the Control of the Liouville Equation

Roger Brockett

Abstract In these notes we motive the study of Liouville equations having control terms using examples from problem areas as diverse as atomic physics (NMR), biological motion control and minimum attention control. On one hand, the Liouville model is interpreted as applying to multiple trials involving a single system and on the other, as applying to the control of many identical copies of a single system; e.g., control of a flock. We illustrate the important role the Liouville formulation has in distinguishing between open loop and feedback control. Mathematical results involving controllability and optimization are discussed along with a theorem establishing the controllability of multiple moments associated with linear models. The methods used succeed by relating the behavior of the solutions of the Liouville equation to the behavior of the underlying ordinary differential equation, the related stochastic differential equation, and the consideration of the related moment equations.

2.1 Introduction

In these notes we describe a number of problems in automatic control related to the Liouville equation and various approximations of it. Some of these problems can be cast either in terms of designing a single feedback controller which effectively controls a particular system over repeated trials corresponding to different initial conditions or, alternatively, using a broadcast signal to simultaneously control many copies of a particular system. Sometimes these different points of view lead to problems that are identical from the mathematical point of view. In many cases a certain continuum limit can be formulated, either by considering an infinity of trials

R. Brockett (✉)
Division of Engineering and Applied Sciences, Harvard University, 345 Maxwell Dworkin, 33 Oxford Street, Cambridge, MA 02138, USA
e-mail: brockett@seas.harvard.edu

F. Alabau-Boussouira et al., *Control of Partial Differential Equations*, Lecture Notes in Mathematics 2048, DOI 10.1007/978-3-642-27893-8_2,

or an infinity of copies. In this situation we are often led to problems involving the control of an associated Liouville equation.

The use of feedback as part of a regulatory mechanism is a standard idea in engineering, biology, and even economics. This stands in contrast to the many other uses of feedback in communication, adaptive sensing, learning algorithms and, more typically in engineering, tracking problems where it is used to improve the speed and accuracy of the response of servomechanisms. Its main virtue is that it is a single mechanism capable of dealing with a great variety of disturbances.

Before introducing the controlled Liouville equation and some mathematical problems that go along with it, we will discuss some additional motivation.

2.2 Some Limitations on Optimal Control Theory

An optimal control problem, as usually formulated, assumes that one has exact knowledge of the equations of evolution. The problem is posed as that of finding a control that transfers the state of the system from a given initial condition to a final one, or possibly a manifold of final states, while minimizing some performance measure. This formulation fits well a number of real-world problems, such as finding the minimum fuel trajectory for getting a payload from the earth to Mars. On the other hand it is less useful as a tool for designing feedback compensators for tracking servomechanisms, a typical problem in robotics, and other path following problems. In these situations there is no fixed initial state and no fixed final state. We do not know what the initial condition will be at a particular time; it is as if the system needs to be ready for a wide variety of challenges.

The development of the various least squares methods for linear systems has led to tools that address more directly the issues raised by such tracking problems. By exploiting the linear structure and by assuming that the desired end state is the point 0, least squares theory produces a feedback control rule that is simultaneously optimal for all initial conditions. Of course the fact that the control can be expressed in feedback form is the key to the invariance with respect to initial conditions. However, the assumptions include the fact that there is a fixed desired steady state and this is a strong limitation.

Moreover, and here we are beginning to discuss a second major point, there are a great many applications in which the payoff for implementing a linear relationship between sensed signals and control variables does not justify the cost of the equipment needed to achieve it. For example, in high volume consumer goods, such as dish washers and clothes dryers, it is inexpensive to sense the temperature of the water or air but the benefits associated with implementing a linear relationship between the temperature of the mixed water and the flow from the hot and cold water lines do not justify the cost. Acceptable performance is obtainable using on-off control which can be implemented much more cheaply. Even in the case of audio equipment, where there is a significant payoff for building systems that are very close to linear, the benefits of linearity are confined to finite range of amplitudes and a subset of frequencies. Standard optimal control theory provides no mechanisms to incorporate implementation costs. This is a major reason why we can not consider the usual optimal control formulation, even when robustness is taken into account, to be completely satisfactory.

Finally, we might ask why optimal control theory has not been more useful in understanding the control mechanisms found in biology. The questions there range from understanding control of the operation of an individual cell to the motor control of the complete organism. In particular, given that evolution has had as long as it has to optimize the neuroanatomy and the muscle/skeletal structures, why is that we don't find optimal control theory to be more effective in explaining these structures?

2.3 Measuring Implementation Cost

The expense required to implement a control policy in an industrial setting where each control signal is generated by a box requiring both a capital investment and continuing maintenance cost, can be accounted in a straightforward way. Unfortunately, such costs are strongly dependent of the technology being used. We wish to focus instead on measures which are intrinsic in the sense that they might apply, at least to some degree, to a range of situations including both those found in engineering and those found in biology. Some considerations that are relevant here have been discussed in our paper [1] which we now paraphrase.

Our point of view is that the easiest control law to implement is a constant input. Anything else requires some attention. The more frequently the control changes, the more effort it takes to implement it. Because the control law will depend on the state x and the time t, it can be argued that the cost of implementation is linked to the rate at which the control changes with changing values of x and t. This rate of change may also affect the effort required to compute the desired control or some suitable approximation to it. In any case, solutions that require less frequent adjustments as x and t change are to be preferred over those that require more frequent adjustments. From the point of view of an animal controlling its body, or a systems engineer allocating the cpu cycles of a computer controlling a machine tool, control laws with small values of $\|\partial u/\partial t\|$ and $\|\partial u/\partial x\|$ require less frequent updating and will be more robust with respect to small changes in the data. These considerations suggest that a suitable quantification of what is meant by "attention" might include a measure of the size of the partial derivatives, $\partial u/\partial x$ and $\partial u/\partial x$. For example, the numerical measure of the attention of a given control law might be might be a weighted Sobolev norm of $u(t, x)$.

This reasoning suggests a class of optimization problems associated with selecting the architecture of a control system. The general structure of the optimization problem will involve minimizing functionals of the form

$$\eta_a = \int_{\Omega} \phi\left(x, t, \frac{\partial u}{\partial x}, \frac{\partial u}{\partial t}\right) dx dt$$

subject to constraints on u such as will insure that the performance is adequate for the task. We can think of η as an *attention functional* and use it as a guide to suggest which control laws might be more or less expensive to implement. To this we may add the observation that although textbooks on control often discuss the difference between open-loop and closed-loop control, the distinction is either vague or applicable only in highly restrictive situations. In many cases, e.g., fixed end-point linear-quadratic optimal control on finite time intervals, it is unclear what might be meant by a closed-loop solution. This makes it difficult for researchers in other fields to discuss the distinction in a precise way. At an intuitive level, it seems that biological motor control involves not only "pure" open-loop control but also a gradation of modalities spanning a range between open-loop and closed-loop operation. Intuitively, one thinks that large values of $\|\partial u/\partial x\|$ indicate closed-loop control and that large values of $\|\partial u/\partial t\|$ indicate open-loop control. By modifying the attention functional we can change the ratio of the penalty put on the closed-loop $\|\partial u/\partial x\|$ terms relative to the penalty put on the open-loop $\|\partial u/\partial t\|$ terms. In this way we create a continuum and arrive at a characterization which makes possible a quantitative study of the trade-offs between open-loop and closed-loop control.

Example 2.3.1. To give some indication about where these ideas can lead it may be helpful to have an example. Consider the scalar control problem $\dot{x} = u$ with the distribution of initial conditions given by a density $\rho_0(x)$. Our goal is to minimize

$$\eta = \int_0^\infty \int_{\mathbb{R}} \rho(t,x) a x^2 dx dt + \int_{\mathbb{R}} \left(\frac{\partial u}{\partial x}\right)^2 dx$$

where $\rho(t,x)$ denotes the density at time t. To avoid complication, we constrain u to be a function of x alone. The calculations leading to a characterization of u now follow.

We rewrite the equation of motion as $dx/u = dt$. Using this we see that if $\int_0^\infty x^2 dt$ is finite then

$$\int_0^\infty x^2 dt = \int_0^{x(0)} \frac{-x^2}{u(x)} dx \quad \text{and} \quad \int_0^\infty u^2 dt = \int_0^{x(0)} -u(x) dx$$

Thus

$$\int_0^\infty \int_{\mathbb{R}} \rho(t,x) a x^2 dx dt = \int_{\mathbb{R}} \rho_0(x_0) \left(\int_0^{x(0)} a \frac{x^2 dx}{u(x)} + b u(x) \right) dx_0$$

and the functional to be minimized can be written as

$$\eta = \int_{\mathbb{R}} \rho_0(x) \left(\int_0^x a \frac{w^2}{u(w)} dw \right) dx + \int_{\mathbb{R}} \left(\frac{\partial u}{\partial x}\right)^2 dx$$

If ρ_0 is a delta function this expression can be further simplified. Let $\rho(0,x)$ be a delta function centered at $x_0 > 0$. In this case the term involving ρ_0 can be simplified giving

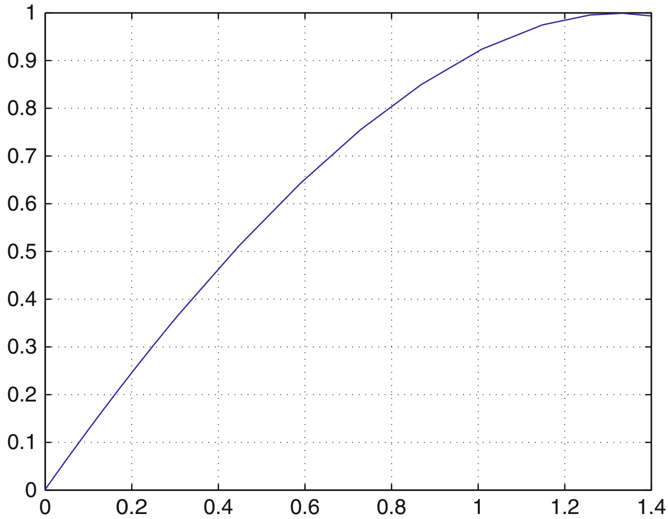

Fig. 2.1 The graph of the optimal gain function for $x > 0$

$$\eta = \int_0^{x_0} \left(a \frac{x^2}{u(w)} \right) dx + \int_{\mathbb{R}} \left(\frac{\partial u}{\partial x} \right)^2 dx$$

The indicated partial derivative is actually a total derivative and so an application of the Euler–Lagrange operator gives

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

Because the solution should remain at zero when reaching zero, it is necessary that $u(0) = 0$. Because the support of the density will be confined to the interval $[0, x_0]$, and because we are minimizing the square of du/dx, the optimal u will have $du/dx = 0$ at $x = x_0$. A numerical solution of these equations corresponding to $a = 1$ and $x_0 = 1$ is shown in Fig. 2.1.

This control law is nearly linear near 0 and approaches saturation as x approaches 1, reflecting the fact that we are putting a penalty on the derivative.

2.4 Ensemble Control

There are several areas of work that have been called "ensemble control" but generally this term applies to problems involving a large number of more or less identical subsystems which are being manipulated by a single source of command signals. (See [2–4].)

A finite number of copies of a system, controlled by an t-dependent function of time can be investigated both as an approximation to the Liouville equation and as something of interest in its own right. Such collections are of interest as models for flocks, swarms, and ensembles of various kinds. Their study gives rise a number of interesting questions, centering around the topics of controllability and stabilizability, but also the control of various averages.

Special aspects of the replicated systems include the degeneracy that will occur when two or more elemental systems are in the same state. In addition, a direct application of Lie algebraic controllability conditions, while in principle quite routine, can be tedious because of the large number of subsystems.

Example 2.4.1. Consider k copies of the scalar system $\dot{x} = -x^3 + u$ and note that the lie bracket of two power law vector fields is given by

$$[x^m \frac{\partial}{\partial x}, x^n \frac{\partial}{\partial x}] = (n-m)x^{m+n-1} \frac{\partial}{\partial x}$$

Thus the Lie algebra generated by the drift vector field and the control vector field is infinite dimensional and contains all polynomial vector fields. To investigate the controllability of k copies of the scalar system we need to look at the distribution generated by bracketing

$$[\sum x_i^3 \frac{\partial}{\partial x_i}, \sum \frac{\partial}{\partial x_i}] = -3 \sum x_i^2 \frac{\partial}{\partial x_i}$$

In this case the Lie algebra contains all the vector fields of the form

$$L_i^j = \sum x^j \frac{\partial}{\partial x_i}; \quad j = 0, 1, 2, \ldots, k$$

The distribution generated by these vector fields at the point $x_1 = p_1, x_2 = p_2, \ldots,$ $x_k = p_k$ is the range space of the Vandermonde matrix

$$V = \begin{bmatrix} 1 & p_1 & p_1^2 & \cdots & p_1^{k-1} \\ 1 & p_2 & p_2^2 & \cdots & p_2^{k-1} \\ 1 & p_3 & p_3^2 & \cdots & p_3^{k-1} \\ \cdots & \cdots & \cdots & & \cdots \\ 1 & p_k & p_k^2 & \cdots & p_k^{k-1} \end{bmatrix}$$

Because the Vandermonde matrix is nonsingular if and only if the p_i are distinct, we see that this distribution spans $\mathbb{R}^k$ at all points corresponding to unrepeated values of the p_i.

Any given ordering of the x_i, for example $x_1 < x_2 < \cdots < x_k$, defines a connected, open subset of $\mathbb{R}^k$ in which the Vandermonde matrix does not vanish. Each of the $n-1$ co-dimension one planar subsets of the boundary is an integral manifold for both the control vector field and the drift and so these can not be crossed. Natural questions then arise about the reachable set in such a cone. In particular, can any point in the cone be reached from any other point in it, independent of the value of k?

Example 2.4.2. Consider a system consisting of a pair of identical second order systems with nonlinear restoring forces.

$$\ddot{x} + x + x^2 = u; \quad \ddot{y} + y + y^2 = u$$

Clearly the set $\{(x, y, \dot{x}, \dot{y}) \mid x = y \ ; \ \dot{x} = \dot{y}\}$ is an invariant set. We may ask if there are other invariant sets and, if not, can we drive an arbitrary initial state to this invariant set. Rewriting the system as

$$\frac{d}{dt}\begin{bmatrix} x \\ \dot{x} \\ y \\ \dot{y} \end{bmatrix} = \begin{bmatrix} \dot{x} \\ -x - x^2 \\ \dot{y} \\ -y - y^2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \\ 1 \end{bmatrix} u$$

we see that in addition to the vector fields defined by f and g the Lie algebra contains

$$[f, g] = \begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \end{bmatrix}; \quad [f, [f, g]] = \begin{bmatrix} 0 \\ -1 - 2x \\ 0 \\ -1 - 2y \end{bmatrix}; \quad [f, [f, [f, g]]] = \begin{bmatrix} -1 - 2x \\ 0 \\ -1 - 2y \\ 0 \end{bmatrix}$$

Thus the distribution associated with the vector fields in the Lie algebra generated by drift and the control vector field spans a four-dimensional space at all points except those on the co-dimension one hyperplane defined by $x = y$.

Our language will be to call the overall system "the system" and to refer to the individual subsystems as being "the elemental systems". Examples of what we have in mind can be found in the literature the on the following topics:

A. *Classical thermodynamics* deals with the control of ensembles, usually modeled as collections of identical particles. Viewed as a control problem, the conversion of heat into work concerns the control of various averages such as temperature and pressure (the intensive variables) using heat flow and adjustable volume. Here the elemental systems consist, for example, of gas molecules; the over all system would be described by a combination of intensive and extensive variables. One might take the controls to be heat flow and volume. Formulated as a control problem, a possible goal is to extract as much mechanical work as possible given constraints on the path. In elementary thermodynamics the system is described in terms of the thermodynamic "state". It is typical to assume that the controls are applied in such a way as to keep the system in thermodynamic equilibrium; which is to say, all paths are adiabatic.

B. *Quantum control of ensembles of identical, weakly interacting, particles.* This arises in the model used in many discussions of nuclear magnetic resonance (NMR) problems. The control is an electromagnetic field consisting of short bursts, or pulses, consisting of different frequencies of controlled duration. The goal is usually to manipulate the orientation of a collection of quantum mechanical spins, say those of the hydrogen ions coming from water molecules, such that the majority of the elemental systems align in a particular nonequilibrium configuration.
C. *Quantum control of a parameterized family* of nearly identical systems using a common control. Here again, a well studied model comes from NMR. Because of slight variations in the magnetic field the resonant frequencies of the individual hydrogen ions differ over the ensemble. Because of this, the control has a different effect on the various elemental systems. Consequently, even if the elemental systems were to start from the same state it requires great care to steer the largest possible fraction of them to a desired end state.
D. *Control of flocks*: It is of interest to understand the extent to which a leader can shape and stabilize the motion of the elemental systems comprising a flock using a broadcast signal. A natural constraint would be to ask that any feedback signal be based on a symmetric function of the states of the elemental systems, for example on the average velocity of the elemental systems.

These applications have in common the goal of controlling a (large) number of weakly interacting individual systems with a single, or perhaps small number, of control inputs. In some quantum mechanical applications the Liouville–von Neumann density equation is appropriate to describe the situation; in other situations the Fokker–Planck equation, or even many copies of a finite state model may serve better.

2.5 The Liouville Equation

Given an ordinary differential equation, $\dot{x}(t) = f(x(t))$ defined on a manifold X, and having the property that there exists a unique solution through each point, there is an associated partial differential equation which describes the evolution of an initial density of points. Let $\rho(0,\cdot)$ be the initial density, thought of as a probability density for $x(0)$. As such it is nonnegative and normalized

$$\int_X \rho(0,x)\,dx = 1$$

Let ψ be a smooth function $\psi : X \to \mathbb{R}^+$ having compact support. The expected value of ψ at some future time is

$$\mathscr{E}\psi(x(t)) = \int_X \psi(x)\rho(t,x)\,dx$$

The time derivative of this expression can be expressed in two ways. It is expressible in terms of $\partial\rho/\partial t$ but also in terms of $\langle \partial\psi/\partial x, f(x)\rangle\rho(t,x)$. These possibilities are the basis for the expression

$$\int_X \psi(x)\frac{\partial\rho}{\partial t}\, dx = \int_X \langle \partial\psi/\partial x, f(x)\rangle\rho(t,x)\, dx$$

Integrating the right-hand side by parts we have

$$\int_X \psi(x)\frac{\partial\rho}{\partial t}\, dx = -\int_X \psi(x)\langle \partial/\partial x, f(x)\rho(t,x)\rangle\, dx$$

Because ψ is arbitrary this implies, subject to mild additional assumptions, that

$$\frac{\partial\rho(t,x)}{\partial t} = -\left\langle \frac{\partial}{\partial x},\ f(x)\rho(t,x)\right\rangle$$

We can think of this as a Cauchy problem to be solved, subject to an initial condition $\rho(0,x) = \rho_0(x)$. It describes how the density evolves in time under the flow defined by the given deterministic equation. It is easy to verify that if $\rho(t,x)$ is nonnegative and satisfies this equation then

$$\frac{d}{dt}\int_{\mathbb{R}^n} \rho(t,x)dx = 0$$

The solution of the Liouville equation can be expressed in terms of the general solution of $\dot{x} = f(x,t)$. If the solution of $\dot{x} = f(x,t)$ is such that the initial value x_0 goes to $\phi(t,x_0)$ at time t then the solution of the Liouville equation is

$$\rho(t,x) = \rho_0(\phi^{-1}(t,x))/\det J_\phi(x)$$

where ϕ^{-1} denotes the result of solving $x = \phi(t,x_0)$ for x_0 and J is the Jacobian of this map; its determinant is necessarily positive. Thus the properties of the Liouville equation reflect quite closely the properties of the underlying ordinary differential equation. The example that follows uses a special case of the fact that $\det\partial f/\partial x$ is the exponential of an integral of the trace of a Jacobian.

Example 2.5.1. Using the fact that the solution of $\dot{x} = Ax + f(t)$ can be written as

$$x(t) = e^{At}x(0) + \int_0^t e^{A(t-\sigma)}f(\sigma)\, d\sigma$$

For the corresponding Liouville equation and initial density ρ_0 the solution is

$$\rho(t,x) = \frac{1}{e^{\mathrm{trAt}}}\rho_0\left(e^{-At}\left(x - \int_0^t e^{A(t-\sigma)}f(\sigma)\, d\sigma\right)\right)$$

If we have a control present, as in $\dot{x} = f(x, u)$, the Liouville equation keeps this same form but because the control is now operated on by the partial derivative operator, feedback controls and open loop controls lead to different solutions.

Example 2.5.2. Consider the scalar equation $\dot{x} = u$; $x(0) = 1$. If we let $u(t, x) = -x$ then of course $x = e^{-t}$ and $u(t) = -e^{-t}$. We get the same solution if we set $u(t, x) = -e^{-t}$. On the other hand, if we have an initial density ρ_0 the solution of the Liouville equation corresponding $u(t) = -e^{-t}$ is

$$\rho(t, x) = \rho_0(x - e^t)$$

whereas the solution corresponding to $u(x) = -x$ is

$$\rho(t, x) = e^t \rho_0(e^{-t} x)$$

2.6 Comparison with the Fokker Planck Equation

We have suggested that one interpretation of the Liouville equation is that it provides a description of the evolution of a probability density under the deterministic flow defined by $\dot{x} = f(x, u)$. Of course there is also an evolution equation for the density associated with stochastic equations containing Wiener processes, such as those of the Itô form

$$dx = f(x, u)dt + \sum g_i(x) dw_i$$

The effect of the $g_i dw_i$ terms is to introduce a diffusion, something completely absent in the model provided by the Liouville equation. For the scalar equation $dx = axdt + cdw$ the Fokker–Planck equation is

$$\frac{\partial \rho(t, x)}{\partial t} = -\frac{\partial ax\rho}{\partial x} + \frac{1}{2} c^2 \frac{\partial^2 \rho(t, x)}{\partial x^2}$$

If the initial density is Gaussian, $\rho_0(x) = \left(1/\sqrt{2\pi s(0)}\right) e^{(x-\bar{x}(0))^2/2s(0)}$ then the solution of this equation remains Gaussian for all time

$$\rho(t, x) = \frac{1}{\sqrt{2\pi s(t)}} e^{(x-\bar{x})^2/2s(t)}$$

where $\bar{x}$ is $e^{at}\bar{x}(0)$ and $s(t)$ is the solution of the variance equation given $s(0)$

$$\dot{s} = 2as + c^2$$

This can be compared with the solution of the Liouville equation corresponding to $c = 0$ which, for the same initial condition, is

$$\rho(t, x) = \frac{1}{\sqrt{2\pi t e^{2at}}} e^{(x-\bar{x})^2/2e^{at}}$$

2.7 Sample Problems Involving the Liouville Equation

In the case where

$$\dot{x}(t) = f(x) + \sum u_i(t) g_i(x(t))$$

the Liouville equation is

$$\frac{\partial \rho(t,x)}{\partial t} = -\left\langle \frac{\partial}{\partial x}, f(x)\rho(t,x) \right\rangle - \sum \left\langle \frac{\partial}{\partial x}, u_i(t,x) g_i(x) \rho(t,x) \right\rangle$$

When solving a standard control problem modeled as $\dot{x} = f(x,u)$ one seeks a control u defined on the interval of interest. Often u will be found through the use of variational principles and may be found as a function of t or as a function of the optimal trajectory x. Whether u is expressed as a function of t alone or as some combination of t and x is regarded as being of secondary importance. However, the situation is quite different for the Liouville equation because now $f(x,u)$ is acted on by the partial derivatives with respect to x. The value of $\rho(t,x)$ depends on whether u is expressed as an open loop function ($u = u(t)$ or as a closed loop function ($u = u(t,x)$.

We now briefly describe a number of problems which can be phrased in terms of the Liouville equation even though they fall outside the usual theory of optimal control.

Problem 2.7.1. *The regulator in a box:* Just as one of the basic examples in quantum mechanics is the charged particle in a square well potential, we can consider control problems where the domain of interest is limited for technological reasons to a sharply defined interval. Suppose that there exist limitations such that values of x and u that lie outside a certain range are of no interest. We seek a control that has good performance and is easy to implement. Building on our earlier example, we consider the scalar control problem $\dot{x} = u$ with the distribution of initial conditions given by a density $\rho_0(x)$ which is uniform on $[-1, 1]$ and zero outside that interval. Find u as a function of x such as to minimize

$$\eta = \int_0^\infty \int_{-1}^1 \rho(t,x) a x^2 dx dt + \int_{\mathbb{R}} \left(\frac{\partial u}{\partial x} \right)^2 dx$$

Problem 2.7.2. *Maximizing the Domain of Attraction:* Consider the system $\dot{x} = f(x,u)$ with $f(0,0) = 0$ having the property that the solution $(x,u) = (0,0)$ is unstable. It is often of interest to determine u so as to make the null solution of $\dot{x} = f(x,u(x))$ asymptotically stable and to make its domain of attraction as large as possible. We can formulate this in terms of a controlled Liouville equation in the following way. For the equation

$$\frac{\partial\rho(t,x)}{\partial t} = -\left\langle \frac{\partial}{\partial x}, f(x,u)\rho(t,x)\right\rangle$$

Find u as a function of x so as to minimize

$$\eta = \int_0^\infty \int_{\mathbb{R}^n} \tanh(k\|x\|)\rho(t,x)\ dxdt$$

Notice that for large positive values of k this assigns zero cost to trajectories that go to zero as t goes to infinity.

Problem 2.7.3. *Trajectory Confinement:* In most models concerned with discretized control signals, asymptotic stability is not possible. In [5] we discussed the possibility of confinement to a region about the target value. This can be restated as requiring that the support of the density should be limited to some neighborhood of the target. If the target is $x = 0$ we might also reformulate the problem in terms of minimizing a measure such as

$$\eta = \int_X x^2\rho(T,x)\ dx$$

Problem 2.7.4. *Enhancing Controllability:* As we have seen, identical linear systems are not ensemble controllable in any reasonable sense. Yet with nonlinear feedback they can become so. We can ask about the nonlinearities that make the linear system ensemble controllable. Of course we need the Lie algebra generated by $Ax + bg(x), b$ to have enough independent components so as to achieve controllability. Moreover, we would like $[Ax, b(g \cdot b)]$ to be "strongly independent" in some sense, probably involving an average over the domain of interest. As noted in the example above, replicated systems are not controllable along the walls of the cone defined by the planes characterizing equality of components, but in the interior they can be.

Problem 2.7.5. *Restricted Range Feedback:* In our paper [5] we discussed the possibility of controlling a linear system with outputs that are generated by a finite state machine. The idea was to model the feedback controller as a Markov process and to adjust the transition rates of the Markov process in such a way as to achieve control. This can be contrasted with the older idea of pulse-width modulated control, commonly used in less sophisticated control systems, which operates in an on-off mode, with the switching times synchronized with a clock.

2.8 Controllability

Suppose that $\dot{x} = f(x,u)$ and that there is a density of initial conditions for x with support of ρ_0 being the set X_0. Suppose, further, that we would like to find $u(t,x)$ so as to steer ρ from its initial value to $\rho_1(x)$ whose support is X_1. For example, if, in fact, we have a regulator problem then X_1 could be a small set containing 0. If we have a cost function involving u we could arrive at a problem of the form

$$\frac{\partial\rho(t,x)}{\partial t} = -\left\langle \frac{\partial}{\partial x},\ f(x,u)\rho(t,x)\right\rangle;\quad \rho(0,x) = \rho_0;\quad \text{Supp}\ \rho(T,x) \in X_1$$

$$\eta = \int_0^T \int_X L(u(x))\rho(t,x)dxdt$$

with the goal of minimizing η.

In other situations the final density might be completely specified or it might be that certain linear functionals of it are to satisfy some inequalities. It might happen that L depends on x as well as u, etc. Some concrete examples appear elsewhere in these notes.

Let X be an oriented differentiable manifold with a fixed, nondegenerate, volume form dv. Let $\phi : X \to X$ be a diffeomorphism. If ρdv is a nonnegative measure on X and if ϕ is orientation preserving, then ϕ acts on densities according to

$$\rho(\cdot) \mapsto \rho(\phi^{-1}(\cdot))/\det \mathrm{J}_\phi$$

where J_ϕ is the Jacobian of ϕ. In this sense $\mathrm{Diff}_O(X)$, the set of orientation preserving diffeomorphisms, generates an orbit through a given ρ.

If the manifold is compact and we restricted discussion to strictly positive densities then this action is transitive, see Moser [6] and Dacorogna and Moser [7]. If the densities are only assumed to be nonnegative the situation is much more complicated.

A natural question to ask is then, given two nonnegative densities, ρ_0 and ρ_1, each of which integrates to one, does there exist a control vector $u(t,x)$ defined on $[0,T]$ that steers ρ_0 to ρ_1? From the point of view that the Liouville equation defines an evolution equation on $L_1(\mathbb{R}^n)$, It might be expected that in considering this question the Lie algebra generated by the first order linear operators

$$\mathscr{L} = \left\{-\left\langle\frac{\partial}{\partial x}, f(x)\rho(t,x)\right\rangle, \left\langle\frac{\partial}{\partial x}, u_i(t,x)g_i(x)\rho(t,x)\right\rangle\right\}_{LA}$$

should play a role. However, because the bracket

$$\left[\left\langle\frac{\partial}{\partial x}, u_i(t,x)g_i(x)\rho(t,x)\right\rangle, \left\langle\frac{\partial}{\partial x}, u_j(t,x)g_j(x)\rho(t,x)\right\rangle\right]$$

involves the partial derivatives of $u(x)$, and deeper brackets involve successively higher partial derivatives, this line of attack leads to complications.

Of course the set of operators of the form

$$L = \left\langle \frac{\partial}{\partial x}, u_i(t,x) g_i(x) \right\rangle$$

as u varies over the set of $\mathscr{C}^\infty$ functions of x is an infinite dimensional set. We could reformulate the problem in this way. Let $\psi_i(x) \in \mathscr{C}^\infty(\mathbb{R}^n)$ be a basis for some subset of $\mathscr{C}^\infty$ and consider vector fields of the form

$$L_j = \left\langle \frac{\partial}{\partial x}, \sum \psi_j(x) g_i(x) \right\rangle$$

This is to be compared with the controllability of the system

$$\dot{x} = f(x) + \sum u_i g_i(x)$$

for which the relevant Lie algebra is

$$\mathscr{L} = \{f, g_1, g_2, \ldots, g_m\}_{LA}$$

In our paper [8] we studied the problem of controllability of the density equation associated with linear systems. More recently Agrachev and Caponigro [9] published a study phrased in terms of controlling diffeomorphisms, not restricted to linear systems.

2.9 Optimization with Implementation Costs

Not surprisingly, the addition of an implementation term usually complicates the mathematics required to solve a trajectory optimization problem.

Example 2.9.1. Consider the problem of minimizing the quantity

$$\eta = \int_0^\infty x^2 + u^2 dt$$

while steering the solution of $\dot{x} = -x + u$ from $x(0) = 10$ to $x(1) = 0$. Of course a variational argument implies immediately that

$$\ddot{x} - 2x = 0$$

and together with the boundary conditions on x this determines the optimal trajectory. But another way to solve this problem is to find a solution of the Riccati equation

$$\dot{k} = 2k - 1 + k^2$$

on the interval [0, 1] and to make the substitution $u = v - kx$. It then follows that the original trajectory optimization problem is equivalent to a modified one for which the evolution equation is

$$\dot{x} = (-1 - k)x + v \;\; ; \;\; \eta = \int_0^1 v^2 \, dt$$

and the performance measure is

$$\eta = \int_0^1 v^2 \, dt$$

The optimal v is then expressible in terms of the controllability Gramian W associated with the new system. Matters being so, optimizing v leads to an expression for u. In more detail,

$$v(t) = e^{\int_0^t (-1-k(\tau))d\tau} p \implies u = -k(t)x(t) + v(t)$$

This solution has both open loop, and closed loop terms. Their relative size depends on which solution of the Riccati equation is chosen. The above construction works for any solution of the Riccati equation and includes the possibility that we choose an equilibrium solution. This choice could be made with the goal of minimizing some functional of the form

$$\eta = \int_0^1 L(\dot{k}, \dot{v}) \, dt$$

such as

$$\eta = \int_0^1 (\partial u/\partial t)^2 \, dt$$

Example 2.9.2. As an example of a problem in this setting that is solvable in special cases, consider

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

For this system

$$\frac{\partial \rho(t,x)}{\partial t} = -\left\langle \frac{\partial}{\partial x}, (f(x) + g(x)u)\rho(t,x) \right\rangle$$

with a given initial probability density $\rho_0(x)$. Suppose we consider a trajectory term

$$\eta_p = \int_0^{\infty} \int_{\mathbb{R}^n} (x^T L x + u^T u)\rho(t,x)dxdt$$

and an implementation penalty that favors a linear control.

$$\eta_i = \int_0^{\infty} \int_{\mathbb{R}^n} \left\| \frac{\partial u}{\partial x} x - u(x) \right\|^2 \rho(t,x)\ dxdt$$

It is obvious that in the special case where $f(x) = Ax$ and $g(x) = b$ the optimal solution is

$$u = -B^T K x$$

where K satisfies $A^T K + KA - L + KBB^T K = 0$. More interesting is the suggestion that if $\| f(x) - Ax \|$ and $\| g(x) - b \|$ are not too large in the region of interest then we can use the known solution as the initial guess in a successive approximation scheme.

In the context of this example there are three distinct aspects of the linear case are worth noting. i) The pure trajectory optimization is solvable in feedback form, ii) the implementation term adds no additional cost at precisely at the optimal feedback control, and iii) the form of the initial distribution is irrelevant. Generalizing the problem in such a way as to take away any of these will yield more interesting solutions.

2.10 Controlling the Variance

We now turn our attention to questions involving the simultaneous use of open loop and closed loop terms to shape the first and second moments of the density. This can be thought of as part of the larger problem of controlling the Liouville equation. For linear stochastic systems this amounts to controlling the mean and the variance and represents a compromise between controlling one individual trajectory associated with $\dot{x} = f(x,u)$ and controlling the entire density. It is, perhaps, the simplest set of problems illustrating how the parametrization of the control as a sum of an open loop part plus a closed loop part can provide additional controllability beyond what is available using open loop control alone. For simplicity, we suppose that the uncontrolled system is linear and time invariant; the extension to the time varying case presents little additional difficulty.

Consider the stochastic system

$$dx = Axdt + Budt + Gdw$$

Let $\bar{x}$ and Σ denote the corresponding mean and variance so that with the control law $u(t) = K(t)x(t) + u_0(t)$ we have

$$\frac{d}{dt}\bar{x} = (A + BK)\bar{x} + Bu_0$$
$$\dot{\Sigma} = (A + BK(t)\Sigma + \Sigma(A + BK(t))^T + Q$$

with $Q = GG^T$. We now investigate the set of reachable values for $\bar{x}$ and Σ, considering K and u_0 to be controls.

In thinking about controlling the variance, it is helpful to keep in mind that the set of positive semidefinite matrices is both a cone and an additive semigroup and that any vector field of the form $F(\Sigma) = A\Sigma + \Sigma A^T$ maps this cone into itself. Moreover, the general linear group acts transitively on the set of positive definite matrices in accordance with the group action

$$(T, Q) \mapsto TQT^T$$

Of course there is a large literature devoted to the steady state solution of the variance equation, going back to Wiener's work on filtering and continuing with the celebrated linear-quadratic-Gaussian theory developed in the context of modern control theory . Much of this work is devoted to questions about how to minimize the variance through the choice of constant K. Here we are interested in treating $K(t)$ as a control and focusing on the transient behavior.

Remark 2.10.1. As motivation consider the following type of problem. Suppose that an athlete has an objective such as placing the ball with a tennis serve or gaining a certain height as a pole vaulter. The penalty for missing the objective may be highly nonlinear and the number of tries limited. Thus the best policy typically involves a tradeoff between controlling the mean and controlling the variance. If the only uncertainty enters through the initial state, the problem can be phrased in the terms described above.

The feedback gains K enter the variance equation multiplicatively and hence this is an example of what has come to be called bilinear control. The presence of the bias term Q and the constraint imposed by the fact that the variance is automatically positive semidefinite sets this problem apart from much of the literature. In the appendix we give some results on the general bilinear problem but here we focus on the variance equation itself. We will make use of the idea that when studying controllability for systems with a drift term, if the drift vector field generates a periodic motion then the effect of moving backwards along the drift vector field can be achieved by letting the system flow along the drift vector field for something less than a full period. This idea was used by Jurdjevic and Sussmann [10] in the context of control on Lie groups and later, without the Lie group hypothesis, in [11].

Lemma 2.10.2. *Let A be a real n-by-n matrix and let B be n-by-m. If A, B is a controllable pair in the sense that the rank of $[B, AB, \ldots, A^{n-1}B]$ is n then, considering K as a time varying control, the system*

$$\dot{\Sigma} = (A + BK(t)\Sigma + \Sigma(A + BK(t))^T + Q; \quad \Sigma(0) \geq 0$$

has the property that the reachable set from any $\Sigma(0) > 0$ has nonempty interior in the space of symmetric positive definite matrices.

Proof. Step 1: Clearly the Lie algebra generated by the matrices A and BK for all possible constant K, contains A and every matrix whose range space is contained in the range space of B. It also contains all matrices of the form $ABK - BKA$. However, the range space of BKA is contained in the range space of B and so we see that the Lie algebra in question contains all matrices of the form ABK. It also contains all matrices whose range space is AB as well as those whose range space is contained in the range space of B. Continuing with $[A, ABK] = A^2BK - ABKA$, etc. we see the Lie algebra contains all matrices whose range space is contained in the sum of the ranges of $B, AB, \ldots, A^nB$ which is the entire Lie algebra of n-by-n matrices.

Step 2: In the case where $Q = 0$ and K is piecewise constant on $[0, t]$ we have

$$\Sigma(t) = M\Sigma(0)M^T$$

where

$$M = e^{(A+BKr)t_r} e^{(A+K_{r-1})t_{r-1}} \cdots e^{A+BK_1)t_1}$$

Thus with $Q = 0$ the given equation is controllable on the space of symmetric matrices with rank and a signature matching that of $\Sigma(0)$, provided that the matrix equation $\dot{X} = (A+BK)X(t)$ is controllable on the space of nonsingular matrices. In particular, it is controllable on the space of symmetric, positive definite matrices.

Step 3: The effect of Q is simply to offset the solution in accordance with the variation of constants formula

$$\Sigma(t) = \Phi(t,0)\Sigma(0)\Phi^T(t,0) + \int_0^t \Phi(t,\tau)Q\Phi^T(t,\tau)d\tau$$

and thus even with $Q \neq 0$ the reachable set retains the property of containing an open set. □

Remark 2.10.3. Theorem 1 of [12] provides a complete characterization of the Lie algebra generated by A and bc^T, under the assumption that (A, b, c) is controllable and observable. In particular, it is established there that if the trace of $A + \alpha bc^T$ is nonzero for some α and if $c^T(I(s+\alpha)-A)^{-1}B$ is not equal to $c^T(-I(s+\alpha)-A)^{-1}b$ for any α then the Lie algebra generated by A and bc^T is the set of all of n-by-n matrices. Observe that in the present situation we can choose K such that the trace of BK is nonzero and by virtue of the controllability assumption we can select a rank one matrix K such that $BK = bc^T$ meets these requirements. We give a general result later (Theorem 2.10.6) but perhaps a concrete example will be helpful at this point.

Example 2.10.4. Consider the two-by-two variance equation associated with

$$\begin{bmatrix} dx_1 \\ dx_2 \end{bmatrix} = u \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} dt + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u + \begin{bmatrix} 0 \\ dw \end{bmatrix}$$

If we let $u = k_1x_1 + k_2x_2$ the corresponding variance equation is

$$\frac{d}{dt}\begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ k_1 & k_2 \end{bmatrix}\begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{bmatrix} + \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{bmatrix}\begin{bmatrix} 0 & k_1 \\ 1 & k_2 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

We want to show that this equation is controllable on $\Sigma > 0$.

Write the equations in component form

$$\begin{aligned} \dot\sigma_{11} &= 2\sigma_{12} \\ \dot\sigma_{12} &= k_1\sigma_{11} + k_2\sigma_{12} + \sigma_{22} \\ \dot\sigma_{22} &= 2k_1\sigma_{12} + 2k_2\sigma_{22} + 1 \end{aligned}$$

Positive definiteness can be characterized by $\sigma_{11} > 0$ and $\sigma_{11}\sigma_{22} > \sigma_{12}^2$. Observe that given (u_1, u_2), the simultaneous equations

$$\begin{bmatrix} u \\ v \end{bmatrix} = u \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ 2\sigma_{12} & 2\sigma_{22} \end{bmatrix}\begin{bmatrix} k_1 \\ k_2 \end{bmatrix} dt + \begin{bmatrix} \sigma_{22} \\ 0 \end{bmatrix}$$

can be solved for (k_1, k_2) in the set $\Sigma > 0$ and if we make the corresponding replacements we have

$$\begin{aligned} \dot\sigma_{11} &= 2\sigma_{12} \\ \dot\sigma_{12} &= u \\ \dot\sigma_{22} &= v \end{aligned}$$

Now the first two of these equations depend on u alone; $\ddot\sigma_{11} = 2u$ and $\sigma_{12} = \dot\sigma_{11}/2$. It is clear, for example from the classical treatment of the time-optimal control of $\ddot x = u$, the point $(\sigma_{11}(0), \sigma_{12}(0))$ can be steered to any point in the half-plane $\sigma_{11} > 0$

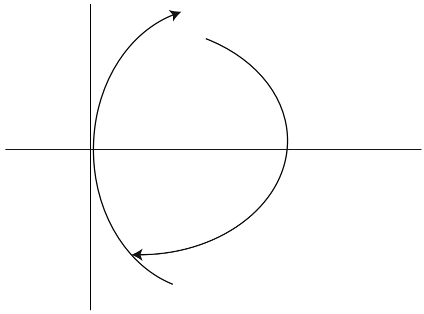

Fig. 2.2 Showing possible trajectories respecting $\sigma_{11} > 0$ in the $(\sigma_{11}, \sigma_{12})$-plane

without leaving that half-plane. Suppose we select some u that accomplishes this transfer in T units of time (Fig. 2.2). Define h by

$$\frac{\sigma_{12}^2(T)}{\sigma_{11}(T)} - \frac{\sigma_{12}^2(0)}{\sigma_{11}(0)} = h/T$$

Finally choose v to be the time derivative of

$$\sigma_{22}(t) = \sigma_{22}(0) + \frac{\sigma_{12}^2(0)}{\sigma_{11}(0))} + \frac{t}{T}\left(\frac{\sigma_{12}^2(t)}{\sigma_{11}(t)} - \frac{\sigma_{12}^2(0)}{\sigma_{11}(0))}\right)$$

More generally, consider the variance equation associated with an n^{th} order system with scalar control. Let e_i denote the standard basis vectors in $\mathbb{R}^n$ $dx = (A + e_n k^T)x dt + e_n dw$. If we partition the variance in terms of blocks compatible with e_n and its complement we have

$$\frac{d}{dt}\begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12}^T & \sigma_{22} \end{bmatrix} = \begin{bmatrix} S & 0 \\ k^T & k_n \end{bmatrix}\begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12}^T & \sigma_{22} \end{bmatrix} + \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12}^T & \sigma_{22} \end{bmatrix}\begin{bmatrix} S^T & k \\ 0 & k_n \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

Using a linear transformation and a suitable offset for k we can arrange matters so that A and $e_n k^T$ take the form

$$A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & 1 \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix}; \quad e_n k^T = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & 0 & \cdots & 0 \\ k_1 & k_2 & k_3 & \cdots & k_n \end{bmatrix}$$

Observe that if Σ is positive definite then the equation

$$u = \Sigma k + b$$

can be solved for k and that in terms of u the variance equation can be written as

$$\begin{aligned}\dot{\Sigma}_{11} &= S\Sigma_{12} + \Sigma_{11}S^T \\ \dot{\Sigma}_{12} &= S\Sigma_{12} + u \\ \dot{\sigma}_{22} &= u\end{aligned}$$

This is a more general formulation of the example. In this notation the problem is that of showing that for

$$\dot{\Sigma} = A\Sigma + \Sigma A^T + e_n u^T + u e_n^T$$

it is possible to steer Σ from $\Sigma_0 > 0$ to $\Sigma_1 > 0$.

The proof of the following theorem shows that this is possible for a general controllable linear system.

Theorem 2.10.5. *Let (A, B) be a controllable pair and let Σ satisfy*

$$\dot{\Sigma} = (A + BK(t))\,\Sigma + \Sigma\,(A + BK(t))^T + Q; \quad \Sigma(0) > 0; \quad Q \geq 0$$

Considering K to be a control, any $\Sigma_1 > 0$ can be reached from any $\Sigma(0) > 0$.

Proof. Clearly the variance equation is linear and the operator mapping real symmetric matrices into real symmetric matrices defined by

$$L(\Sigma) = (A + BK)\Sigma + \Sigma(A + BK)^T$$

has eigenvalues which are all possible pairs of the form $\lambda_i + \lambda_j$ where λ_i and λ_j are eigenvalues of $A + BK$. Thus if there exists a K such that the eigenvalues of $(A + BK)$ are integer multiples of $\mu\sqrt{-1}$ then $\exp L$ is periodic and Theorem A1 of the appendix applies, provided that $e^{(A+BK_0C)t}$ is periodic for some choice of K_0 $\square$

Theorem 2.10.6. *Assume that (A, B) is a controllable pair. The system of equations*

$$\begin{aligned}\dot{x}(t) &= (A + BK(t))x + Bu(t) \\ \dot{\Sigma}(t) &= (A + BK(t))\Sigma(t) + \Sigma(t)(A + BK(t))^T\end{aligned}$$

is controllable in the sense that given any two pairs (x_0, Σ_0) and (x_1, Σ_1) with $\Sigma_0 = \Sigma_0^T > 0$ and $\Sigma_1 = \Sigma_1^T > 0$ and given any time $T > 0$ there exists a control (u, K) defined on $[0, T]$ steering the system from (x_0, Σ_0) to (x_1, Σ_1).

Proof. Select K in accordance with Theorem 1 so as to steer Σ to the desired state. Having selected K, select u by standard controllability arguments to steer x.

Going beyond controllability, there are a variety of optimization questions that arise this context. The most basic are the extensions of the problem considered in the previous section involving the minimization of

$$\eta = \int_0^T L(u_0, k, \dot{k})\, dt$$

while using the control law $u = u_0 + kx$ to force the solution of $\dot{x} = f(x, u)$ to move from $x(0) = x_0$ to $x(T) = x_1$. Extending this idea to a stochastic setting, we can, for example, consider controlling the mean and variance equation as in Theorem 2.10.6, while minimizing

$$\eta = \int_0^T L(u_0, K, \dot{K})\, dt \qquad \square$$

2.11 Ensembles, Symmetric Functions and Thermodynamics

This section is adapted from our paper [2]. It can be seen as taking the idea of simultaneous control of the mean and variance in a new direction.

Let u be a m-dimensional vector, let x_i for $i = 1, 2, \ldots, k$ be a n-dimensional vector, and let y be a p-dimensional vector. Consider a system consisting of k copies of a first order model, each with the same input

$$\dot{x}_i = f(x_i, u); \quad i = 1, 2, \ldots, k$$

We limit our attention to outputs of the form

$$y = c(x_1, x_2, \ldots, x_k)$$

with c being a symmetric function in the sense that for any permutation of the index set $\{1, 2, \ldots, k\} \to \{\pi(1), \pi(2), \ldots, \pi(k)\}$ we have $c(x_1, x_2, \ldots, x_k) = c(x_{\pi(1)}, x_{\pi(2)}, \ldots, x_{\pi(k)})$. If the system is stochastic we replace this model with a family of Itô equations of the form

$$dx_i = f(x_i, u)dt + g(x_i, u)dw_i; \quad i = 1, 2, \ldots k$$

with the Wiener processes $w_1, w_2, \ldots, w_k$ being independent.

Of course there are significant limitations that arise in the control of such systems because u acts on each system in the same way and y is constrained to be a symmetric function. In particular, linear systems of this type are never controllable or observable if $k > 1$.

If the elemental systems are linear then the overall system obtained by applying feedback $u = \sum Cx_i$ is described by

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \cdots \\ \dot{x}_k \end{bmatrix} = \begin{bmatrix} A+BC & BC & \cdots & BC \\ BC & A+BC & \cdots & BC \\ \cdots & \cdots & \cdots & \cdots \\ BC & BC & \cdots & A+BC \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \cdots \\ x_k \end{bmatrix}$$

The matrix on the left is similar to a block triangular matrix with $n-1$ diagonal blocks of the form $A+BC$ and hence can not be controllable. A similar limitations apply to stochastic models of the form

$$dx_i = Ax_i dt + Budt + Gdw_i; \quad dy = C(x_1 + x_2 + \cdots + x_k)dt + d\nu$$

We discuss these in Theorem 2.11.1 after giving a few additional definitions. In a probabilistic setting it is meaningful to discuss statistical properties such as the mean and variance. In the case of many copies of a given system we can consider various averages taken over the ensemble. Of course the sample statistics, as opposed to the statistics themselves, are random variables. In the present situation, with $x_1, x_2, \ldots, x_k$ being described by identical probability laws, we have an interest in a particular type of sampling leading to what can be termed *ensemble sample-statistics*. This refers to averages over the variables $x_1, x_2, \ldots x_k$. For example, we refer to the random variable

$$a(t) = \frac{x_1(t) + x_2(t) + \cdots + x_k(t)}{k}$$

as the ensemble sample-mean.

We say that a homogeneous function $\phi(x_1, x_2, \ldots x_k)$ is *centered* if the sum of its partial derivatives vanishes, i.e.,

$$\sum_{i=1}^{k} \frac{\partial \phi}{\partial x_i} = \sum_{i=1}^{k} \begin{bmatrix} \partial\phi/\partial x_{i1} \\ \partial\phi/\partial x_{i2} \\ \cdots \\ \partial\phi/\partial x_{in} \end{bmatrix} = 0$$

Theorem 2.11.1. *Consider the linear stochastic ensemble*

$$dx_i = Ax_i dt + Budt + Gdw_i; \quad i = 1, 2, \ldots, k$$

The application of feedback in the form $u = \alpha(x)$ does not change the evolution equation of any centered homogeneous function of $x_1, x_2, \ldots x_k$.

Proof. Let ϕ be homogeneous and centered. Applying the Itô rule to ϕ we see that

$$d\phi(x) = \sum \langle \nabla\phi, Ax_j dt + Budt + Gdw_j \rangle + \frac{1}{2} \sum \left\langle \frac{\partial^2 \phi}{\partial x_{ij}}, GG^T \right\rangle$$

Clearly under the given hypothesis the effect of u disappears. $\square$

Corollary. The ensemble sample variance

$$\Sigma_{esv} = \frac{1}{n}\sum\left(x_i - \frac{1}{n}(x_1 + x_2 + \cdots + x_k)\right)\left(x_i - \frac{1}{n}(x_1 + x_2 + \cdots + x_k)\right)^T$$

associated with the system

$$dx_i = Ax_i dt + Budt + Gdw_i; \quad i = 1, 2, \ldots k$$

is not altered by feedback of the form $u = \phi(x)$.

Proof. It is easy to see that each term in the sum defining Σ_{esv} is homogenous and centered and therefore the sum is as well.

These results show that it is necessary to go beyond linear theory if we are to find any benefit from the use of control in the ensemble setting. The following theorem applies to multiplicative control. □

Theorem 2.11.2. *Consider the ensemble*

$$dx_i = Ax_i + uBx_i dt + Gdw_i; \quad i = 1, 2, \ldots, k$$
$$y = \frac{1}{n}(x_1^T Lx_1 + x_2^T Lx_2 + \cdots + x_k^T Lx_k)$$

with $L = L^T > 0$. If there exists a symmetric matrix Q such that $QB + B^TQ$ is negative definite and the eigenvalues of Q all have the same sign, then there exists $\beta > 0$ such that for any real c between 0 and β there is a feedback control law $u = \phi(y)$ which stabilizes the trace of the variance of the sample variance at c.

Proof. The variance of the sample variance, i.e.

$$\Sigma_{esv} = \sum\left(x_i - \frac{1}{n}(x_1 + x_2 + \cdots + x_k)\right)\left(x_i - \frac{1}{n}(x_1 + x_2 + \cdots + x_k)\right)^T$$

satisfies the equation

$$\dot{\Sigma}_{esv} = (A + uB)\Sigma_{esv} + \Sigma_{esv}(A + uB)^T + GG^T$$

Let $QB + B^TQ = R < 0$. Then $Q(A+uB)+(A+uB)^TQ$ is negative definite for suitable choice of u. Thus we see that for a semi-infinite range of u the eigenvalues of $A + uB$ have negative real parts. In fact, $Q(A + uB) + (A + uB)^TQ$ can be made more negative definite than $-\alpha I$ for any α and so the eigenvalues of $A + uB$ can be placed to the left of any vertical line in the complex plane. This means that a steady state variance exists and satisfies

$$0 = (A + uB)\Sigma_{esv} + \Sigma_{esv}(A + uB)^T + GG^T$$

Over the range of u for which the system is stable let $ly^{-1}_{A+uB}(GG)$ denote the solution of this equation. Clearly $\mathrm{tr}(ly^{-1}_{A+uB}(GG))$ goes to zero as the eigenvalues of $A + uB$ go to minus infinity and so, as u varies $\mathrm{tr}\Sigma_{esv}$ sweeps out a range of values of the form $(0, c)$ as required. We can let the feedback control law be a constant, independent of y. $\square$

The control of heat engines and provides a good example of ensemble control. The mathematical description consists of a family of identical scalar linear systems with multiplicative control driven by independent Brownian motion terms.

$$dx_i = (u_1 - u_2)x_i(t)dt + u_3 dw_i; \quad i = 1, 2, \ldots, k$$

$$y = \sum_{i=1}^{k} x_i^2$$

In this case the ensemble equations are supplemented by two auxiliary equations which complete the description and serve to distinguish u_1 from u_2. These are

$$\dot{x}_{k+1} = u_1$$

$$\dot{x}_{k+2} = u_1 y$$

The physical interpretation is as follows. The x's represent (one dimensional) velocities of individual particles in the ensemble. The controls represent the time rate of change of the volume occupied by the gas (u_1), the type of contact the gas has with the available heat sources (u_2), and the selection of a heat source with a particular temperature (u_3). Further details will emerge from the discussion. If we had the services of a Maxwell demon we could observe each of the x_i individually but in reality only certain ensemble averages are observable. Likewise, if we had access to a demon we could generate individual controls for each of the state variables but in reality we can only apply controls which influence all elements of the ensemble in the same way. In the context of the elementary thermodynamics of gases, we are able to change the volume of the gas by moving a piston and to alter the internal energy of the gas by adding or removing heat. Such actions translate into choices of u_1, u_2, u_3 in the above model. The objective of the control action might be, for example, to cause the development of a given quantity of work over a period $[0, P]$. Here work corresponds to the integral

$$w = \int_0^P u_1(t)y(t)dt$$

The relevant summary of the behavior of the population is, in this case, provided by the ensemble sample variance. From the equations for $x_1, x_2, \ldots, x_k$ we see that the ensemble sample variance satisfies the stochastic differential equation

$$d\sigma_{esv} = 2(u_1 - u_2)\sigma_{esv}dt + u_3\frac{1}{n}\sum dw_i + u_3^2 dt$$

Because we have assumed that the w_i are independent the sum appearing here may be replaced by $dw/\sqrt{k}$ without changing the statistical properties of the solution. That is, we may as well adopt the model

$$d\sigma_{esv} = 2(u_1 - u_2)\sigma_{esv}dt + u_3\frac{1}{\sqrt{k}}dw + u_3^2 dt$$

This stochastic differential equation represents a sample statistic obtained from k samples. If we are dealing with a mole of gas then $k \approx 6 \times 10^{23}$! If we assume that the Brownian motion term is insignificant we are led to the set of deterministic equations

$$\begin{aligned}\frac{d}{dt}\sigma_{esv}(t) &= 2(u_1 - u_2)\sigma_{esv}(t) + u_3^2\\ \dot{x}_{k+1} &= u_1\\ \dot{x}_{k+2} &= u_1 y\end{aligned}$$

as a reduced model for the stochastic ensemble.

For the sake of simplicity we rename the variables $(\sigma_{esv}, x_{k+1}, x_{k+2})$ as (x_1, x_2, x_3). The control terms enter these equations in such a way as to define three vector fields as brought out by the notation

$$\begin{bmatrix}\dot{x}_1\\ \dot{x}_2\\ \dot{x}_3\end{bmatrix} = u_1\begin{bmatrix}2x_1\\ 1\\ 2x_1\end{bmatrix} - u_2\begin{bmatrix}2x_1\\ 0\\ 0\end{bmatrix} + u_3^2\begin{bmatrix}1\\ 0\\ 0\end{bmatrix}$$

subject to the constraint that u_2 should be nonnegative. The three vector fields appearing here are

$$A = 2x_1\frac{\partial}{\partial x_1} + \frac{\partial}{\partial x_2} + 2x_1\frac{\partial}{\partial x_3};\quad B = -2x_1\frac{\partial}{\partial x_1};\quad C = \frac{\partial}{\partial x_1}$$

Together with the pair

$$D = x_1\frac{\partial}{\partial x_3};\quad E = \frac{\partial}{\partial x_3}$$

they obey the commutation relations

$$\begin{aligned}&[A, B] = 4D;\quad [A, C] = -2C - 2E;\quad [B, C] = -2C\\ &[A, D] = -[B, D] = 2D;\quad [A, E] = 0\end{aligned}$$

and these five vector fields define a basis for a Lie algebra. This algebra is a solvable subalgebra of the algebra corresponding to the three dimensional affine group. One might say that it is *the* Lie algebra of the Carnot cycle. Constraining $u_2(t)$ to be nonnegative, this system generates admissible flows.

Appendix

We collect here a few results on the bilinear controllability putting in a larger context the result of Sect. 2.10 on the control of the mean and variance. There is a large literature on this subject and we only touch a few points. References [10–15] are relevant.

As is well known, if the off-diagonal elements of $A(t)$ are nonnegative for all t then the solutions of the system $\dot{x}(t) = A(t)x(t)$ leave the positive orthant invariant. Thus if b is a vector with nonnegative entries and $\dot{x}(t) = (A(t) + U(t))x(t) + b(t)$ with U diagonal but otherwise unconstrained, and $A(t)$ is nonnegative off the diagonal then the positive orthant is an invariant set. This can be seen as a being a consequence of the direction of the vector field along the boundary of the positive orthant. In a similar way, if a symmetric matrix X satisfies $\dot{X} = A(t)X(t) + X(t)A^T(t) + B(t)$ with $B(t) = B^T(t)$ nonnegative definite then the cone of nonnegative definite matrices is an invariant set.

Thus, in the case of the scalar system $\dot{x} = (a + k)x + b$ with $b > 0$ the set $\{x|x > 0\}$ is positively invariant and x cannot leave the positive half-line. It is controllable there in the sense that any point in $\{x|x > 0\}$ can be steered in positive time to any other point in the set. In higher dimensions the situation is more complicated. For example, if b_1 and b_2 are positive then solutions of the system

$$\frac{d}{dt}\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} a + k_1 & 1 \\ 1 & b + k_2 \end{bmatrix}\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}$$

can never leave the first orthant regardless of the choice of (k_1, k_2) but if $b \neq 0$ the system

$$\frac{d}{dt}\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 0 & a + k \\ -a - k & 0 \end{bmatrix}\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}$$

can be steered from any initial state in $\mathbb{R}^2$ to any final state.

In studying the controllability of an n-dimensional system

$$\dot{x} = (A + BK(t)C)x(t)$$

it is natural to appeal to Lie algebraic methods. In [12] it is shown that if (A, b, c) is a minimal triple in the sense that $[b, Ab, \ldots, A^{n-1}b)]$ and $[c; cA; \cdots cA^{n-1}]$ are both of rank n then the Lie algebra generated by A and bc is either $gl(n), sl(n), sp(n/2)$ or $sp(n/2) \oplus I$, and is $gl(n)$ unless $g(s) = c(Is - A)^{-1}b$ has a reflection symmetry in the form $g(s + \sigma) = g(-s - \sigma)$ for some real number σ and $\mathrm{tr}A = cb = 0$. Adapting that result to the present situation, we see that the Lie algebra generated by A and matrices of the form $BE_{ij}C$ is $gl(n)$ unless $CB = 0$, tr$A = 0$, and $C(Is - A - \sigma I)^{-1}B = C(-Is - A + \sigma I)^{-1}B$ for some real number σ. However, because the system has an irreversible drift term $\dot{x} = Ax$ the Lie algebra does not tell the whole story.

Theorem A1. *Let A, B, C be constant matrices with $A, B)$ being controllable and (A, C) being observable. Consider the system evolving in the space of n-by-n nonsingular matrices with positive determinant.*

$$\dot{X} = (A + BKC)X$$

with B, K, C being n-bym, m-by-p and p-by-n, respectively. Assume that the solution of $\dot{X} = (A + BKC)X;\ X(0) = I$ is periodic for some $K = K_0(\cdot)$. Then given any pair of nonsingular matrices with positive determinant there exists a K that steers one to the other provided that CB and tr$A$ are not both zero and $C(Is - A)^{-1}B$ does not have the reflection symmetry described above.

Proof. First of all the system is controllable in the sense that it is possible to reach an open set of nonsingular matrices because A and the possible values of BKC generate the entire Lie algebra $gl(n)$. Second, as is well known, form early work on controllability on Lie groups, if $\dot{X} = (A + BKC)X$ has a periodic solution with X nonsingular periodic and we have local controllability then we have global controllability. □

Remark A1. Let A be n-by-n and b be 1-by-n. Observe that

$$M(t) = \exp\begin{bmatrix} At & bt \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} e^{At} & e^{\int_0^t e^{A(t-\sigma)}d\sigma} b \\ 0 & 1 \end{bmatrix}$$

If e^{At} is periodic with period T then its eigenvalues lie on the imaginary axis. If none are zero then A is invertible and

$$e^{\int_0^T e^{A(T-\sigma)}d\sigma} b = A^{-1}(I - e^{-AT})b = 0$$

Thus M is periodic. If 0 is in the spectrum of A then the explicit form of the integration is not available. However, if b lies in the range space of A then we can write b as Av so that

$$e^{\int_0^t e^{A(t-\sigma)}d\sigma} b = e^{\int_0^t e^{A(t-\sigma)}d\sigma} Av = (I - e^{-AT})b = 0$$

and M is periodic. When restricted to evolution equations on $\mathbb{R}^n$ the conditions for controllability simplify because $sp(n/2)$ acts transitively on $\mathbb{R}^n$.

Theorem A2. *Let A, B, C be constant matrices with $A, B)$ being controllable and (A, C) being observable. Consider the system evolving in $\mathbb{R}^n$*

$$\dot{x} = (A + BKC)x + b$$

with B, K, C being n-by$m$, m-by-p and p-by-n, respectively. Assume that the solution of $\dot{X} = (A + BKC)X$; $X(0) = I$ is periodic for some $K = K_0(\cdot)$. Then the system is controllable in the sense that any pair $x_0 \neq 0$ and $x_1 \neq 0$ can be joined by a solution of the given equation.

References

1. R.W. Brockett, Minimum Attention Control, in *Proceedings of the 36th IEEE Conference on Decision and Control* (1997), pp. 2628–2632
2. R. Brockett, *Control of Stochastic Ensembles*, ed. by B. Wittenmark, A. Rantzer. Åström Symposium on Control (Studentlitteratur, Lund, 1999), pp. 199–216
3. R. Brockett, N. Khaneja, in *System Theory Modeling, Analysis, and Control* ed. by T. Djaferis, I. Schick. On the Stochastic Control of Quantum Ensembles (Kluwer Academic, Norwell, 1999), pp. 75–96
4. R. Brockett, On the control of a Flock by a leader. Proc. Skeklov Inst. Math. **268**(1), 49–57 (2010)
5. R.W. Brockett, Reduced Complexity Control Systems, in *Proceedings of the 17th World Congress*, The International Federation of Automatic Control, Seoul, 2008
6. J. Moser, On the volume elements of a manifold. Trans. Am. Math. Soc. **120**, 286–294 (1965)
7. B. Dacorogna, J. Moser, On a partial differential equation involving the Jacobian determinant. Ann. Inst. Henri Poincare C **7**(1), 1–26 (1990)
8. R. Brockett, in *Proceedings of the International Conference on Complex Geometry and Related Fields*, ed. by Z. Chen et al. Optimal Control of the Liouville Equation (American Mathematical Society, Providence, 2007), pp. 23–35
9. A.A. Agrachev, M. Caponigro, Controllability on the group of diffeomorphisms. Ann. Inst. Henri Poincare C Non Lin. Anal. **26**(6), 2503–2509 (2009)
10. V. Jurdjevic, H.J. Sussmann, Control systems on Lie groups, J. Differ. Equat. **12**(2), 313–329 (1972)
11. R. Brockett, Nonlinear Systems and Differential Geometry, in *Proceedings of the IEEE*, vol 64 (1976), pp. 61–72
12. R.W. Brockett, Linear feedback systems and the groups of Lie and Galois. Lin. Algebra Appl. **50**, 45–60 (1983)
13. Yu.L. Sachkov, *in Positive orthant scalar controllability of bilinear systems*. Springer Mathematical Notes, Mat. Zametki, **58**(3), pp. 666–669 (1995)
14. Y.L. Sachkov, On positive orthant controllability of bilinear systems in small codimensions. SIAM J. Contr. Optim. **35**(1), 29–35, (1997)
15. W.M. Boothby, Some comments on positive orthant controllability of bilinear systems. SIAM J. Contr. Optim. **20**(5), 634–644 (1982)

Springer

http://www.springer.com/978-3-642-27892-1

Control of Partial Differential Equations
Cetraro, Italy 2010, Editors: Piermarco Cannarsa, Jean-Michel Coron
Alabau-Boussouira, F.; Brockett, R.; Glass, O.; Le Rousseau, J.; Zuazua, E. - Cannarsa, P.; Coron, J.-M. (Eds.)
2012, XIII, 344 p. 66 illus., 49 illus. in color., Softcover
ISBN: 978-3-642-27892-1