Attention as a Performance Measure for Control System Design

2002 ACC

Anchorage, Alaska

May 8, 2002

Roger Brockett

Engineering and Applied Sciences

Harvard University

A More Inclusive Definition of Optimal Control

- Set up an optimization theory framework that will include both the quality of the performance and the cost of the controller.

2. Include the costs for communication and computation together with more traditional trajectory based performance terms.

3. We will be led to a quantitative measure of attention, with the word being used in somewhat the same way as it is psychology.

4. As presented here, both open loop control and closed loop control require attention. The definitions will be such that low attention solutions will turn out to be less “expensive” to implement.

- The best known work in attention is associated with cognitive and sensory attention, priming, etc., but there is other work as well..

Starting from the Familiar

Consider the standard problem of finding u(t,x) to optimize

$$\dot{x} = Ax + bu$$

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

Compare this with finding k(t) to optimize the same performance measure

$$\dot{x} = Ax + bk(t)cx$$

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

The solution of the first problem can be expressed in feedback form in such a way as to be independent of the initial condition. This is not possible in the second case.

Variations on the Problem

By specifying the way in which the control is allowed to depend on x we may make the control law easier to implement but at the same time introduce a dependence on x(0). Compare this with other ways to make the control law easier to implement such as limiting dk/dt or including the square of du/dt in the performance measure.

$$\dot{x} = Ax + bk(t)cx$$

$$\eta = \int_0^T x^2 + u^2 + (\dot{k})^2 dt$$

The solution of this problem also depends on x(0). Assuming a density for x(0) makes optimization possible.

Families of Trajectories and Densities: Liouville's Equation

Closely related to the differential equation dx/dt = f(x,u) is a partial differential equation, known as the Liouville equation.

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

Instead of describing how a single initial condition evolves, it describes, in one breath, the evolution of a density of initial states.

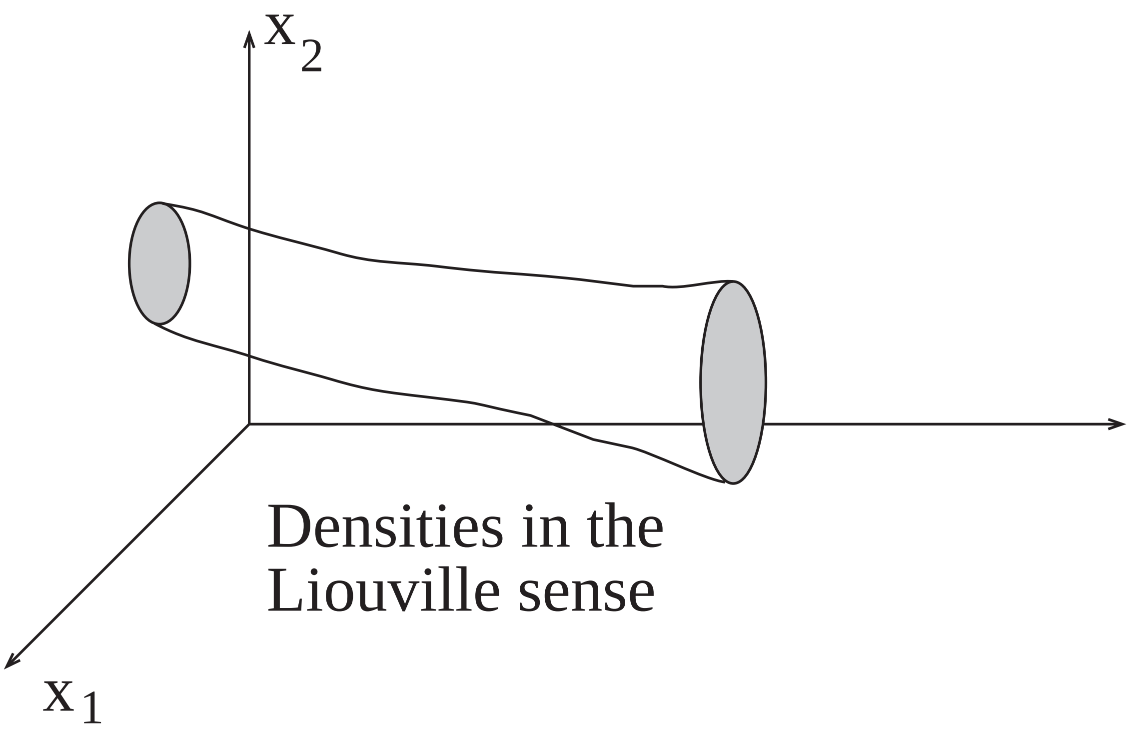

Example of a Solution of Liouville's Equation

$$\dot{x} = -k(t)x \;\; ; \;\; \rho_0(x) = given$$

$$\frac{\partial \rho(t,x)}{\partial t} = \frac{\partial}{\partial x} k(t) x \rho(t,x)$$

$$\rho(t,x) = e^{a(t)} \rho_0(e^{a(t)} x) \;\; ; \; a(t) = -\int_0^t k(\tau) d\tau$$

Much more general cases can be solved explicitly, including the general linear time invariant dx/dt =Ax+Bu.

New Problem Formulation

To formulate a optimization problem that involves both trajectory terms and implementation terms, select an L so as to reflect implementation costs. An example:

$$\dot{x} = -u(t,x) \;\; ; \;\; \rho_0(x) = given$$

$$\frac{\partial \rho(t,x)}{\partial t} = \frac{\partial}{\partial x} u(t,x)\rho(t,x)$$

$$\eta = \int_0^T \int_{-\infty}^{\infty} (x^2 + u^2)\rho(t,x) + L(u)\,dx\,dt$$

This is a trajectory optimization problem with a fixed initial condition where we think of the dynamics being defined by the given Liouville equation.

Implementation Cost Factors: Choice of L

1. Number of quantization levels (12 bit vs. 16 bit, single precision, double precision, etc.)
2. Sampling rate, 30 Hz, 100 Hz, …
3. Tolerance to delay 20 millisecond latency, 60 millisecond latency…
4. Computational complexity of the control law
5. Quality of sensors (speed and accuracy).

Nothing takes less attention than letting u be constant. Adjusting u, as x and t change, requires attention.

An implementation must be good for the full range of initial conditions and disturbances that will be confronted. The implementation is not adjusted in accordance with the current trajectory. (That is, any "adaptation" is considered to be part of the controller.)

Implementation Cost Factors: Choice of L

Two examples of meaningful attention functionals or given below. Each reflects the desirability of minimizing the change in the contro.l

$$\eta = \int_0^\infty \int_{\mathbb{R}^n} \alpha \left(\frac{\partial u}{\partial t}\right)^2 + \beta \left(\frac{\partial u}{\partial x}\right)^2 dxdt$$

$$\eta = \int_0^\infty \int_{\mathbb{R}^n} \alpha \left(\frac{\partial u}{\partial t}\right)^2 + \beta \left(\frac{\partial u}{\partial x}, f(x,u)\rangle\right)^2 dxdt$$

Each of these separates nicely into an open loop part weighted by α and a closed loop part weighted by β.

Trajectory and Implementation Costs Formulated Jointly

We are interested in combining two types of terms:

1. A performance term that will insure stability, hitting the target, conserving resources, minimizing time, etc. as dictated by the problem.

2. A implementation term that assures that the control is not excessively sensitive to small changes in the measurements, small errors in the clock, does not require a high sampling rate or ultra fine quantization. (Some type of regularization term.)

The inclusion of the second term will complicate the mathematics but can give control laws that saturate for large values and are more easily approximated, thus giving more flexibility in their implementation.

What kind of solution can we expect?

We can get some feeling from exploring an example.

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

$$u(t, x) = \frac{t}{1 + t^2} \tanh x$$

$$u_t = \left(\frac{1}{1 + t^2} - \frac{2t^2}{(1 + t^2)^2} \right) \tanh x$$

$$u_x = \frac{1}{1 + t^2} \frac{1}{\cosh^2 x}$$

The Change in Performance

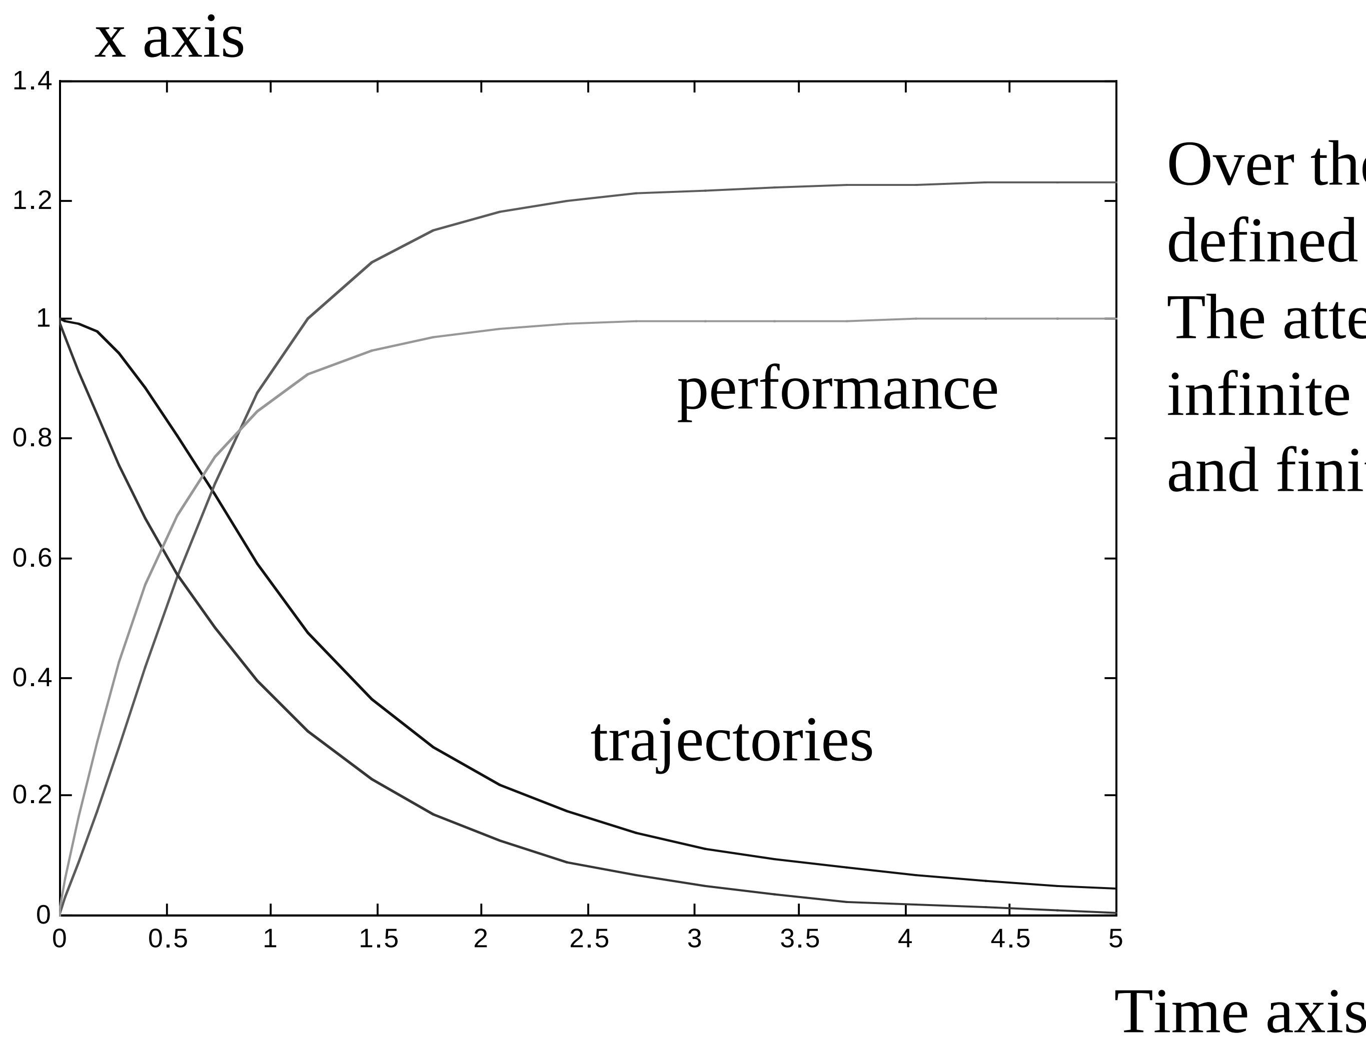

Over the rectangle defined by |x|<1, t>0 The attention is infinite in one case and finite in the other.

$u(t,x) = -2t/(1+t^2)\ \tanh(x)$ restricted to a rectangle

Solving a Special Case

Consider keeping on the closed loop loop penalty and asking that the initial density be a delta function supported at x(0)=1

$$\dot{x} = -a(t)x$$

$$\eta = \int_0^\infty x^2(t) + \beta \left(\frac{\partial u}{\partial x}\right)^2 dt$$

In this case we can solve for x in terms of the integral of a and pass directly to an Euler-Lagrange equation for A

$$A(t) = \int_0^t a(\tau)d\tau$$

$$\beta\ddot{A} - e^{-2A} = 0$$

The Optimality of a/(t+b)

The rest of the calculation goes through smoothly

$$A(t) = \int_0^t a(\tau)d\tau$$

$$\beta \ddot{A} - e^{-2A} = 0$$

$$a(t) = \frac{1}{t + \sqrt{\beta}}$$

This establishes a sense in which the 1/t type gains showing up in stochastic approximation are actually optimal.

What can we learn from other fields about choosing L

Books on neuroscience, such as those of M. Ito on the cerebellum, contain many references to feedback control and have many block diagrams. There are also a number of disclaimers because the feedback loops depend on higher level cognitive processes. The reader from control will conclude that our theories, while useful, do not have enough generality to provide the right tools.

Robotics has been a successful application of control theory at the hardware level, but the more important problem is software and we have not yet provided much in the way of theory for language driven systems.

The inspiration will not come from "cost is no issue" solutions but rather the "just enough control" solutions as found in in washing machines and humans.

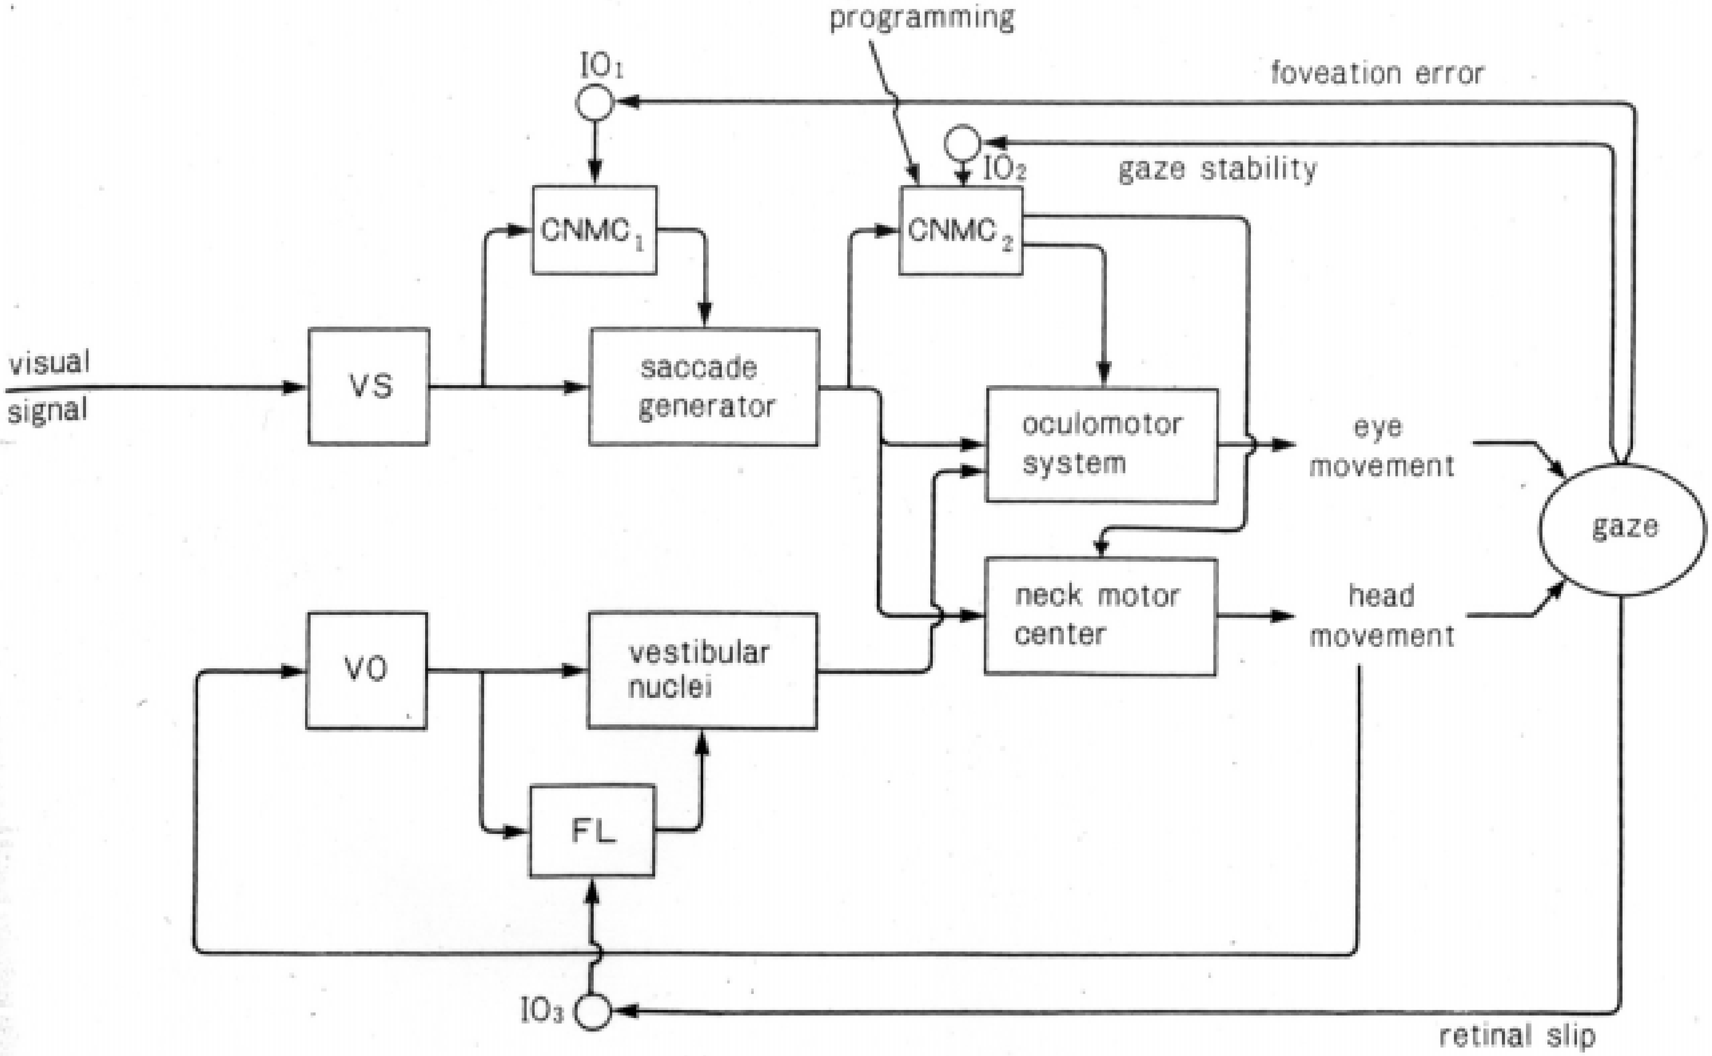

FIG. 195. Hypothetical block diagram for cerebellar contribution to eye–head coordination. VS, visual system; VO, vestibular organ; FL, flocculus; IO_1, IO_2, IO_3, areas of the inferior olive.

Hermann von Helmholtz, 1821-1894

"Theoretical natural science must, therefore, if it is not to rest content with a partial view of the nature of things, take a position in harmony with the present conception of simple forces and the consequences of this conception. Its task will be completed when the reduction of phenomena to simple forces is completed…"

Helmholtz, 1894

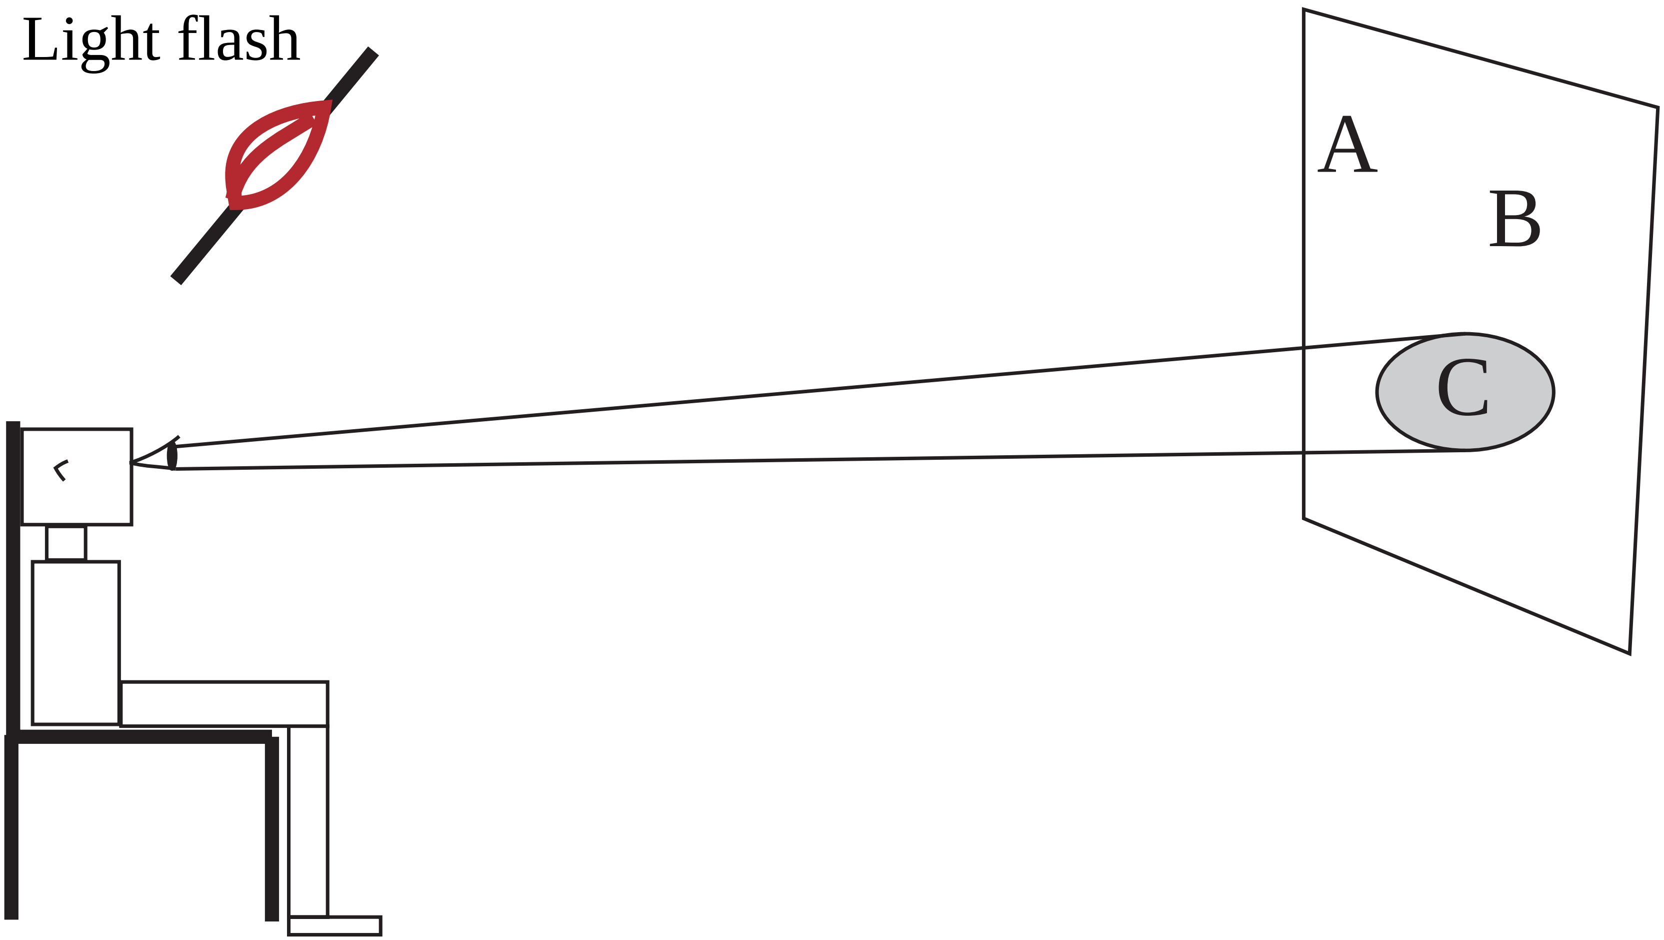

The famous Holmholtz experiment showing that humans can direct visual attention without physical motion.

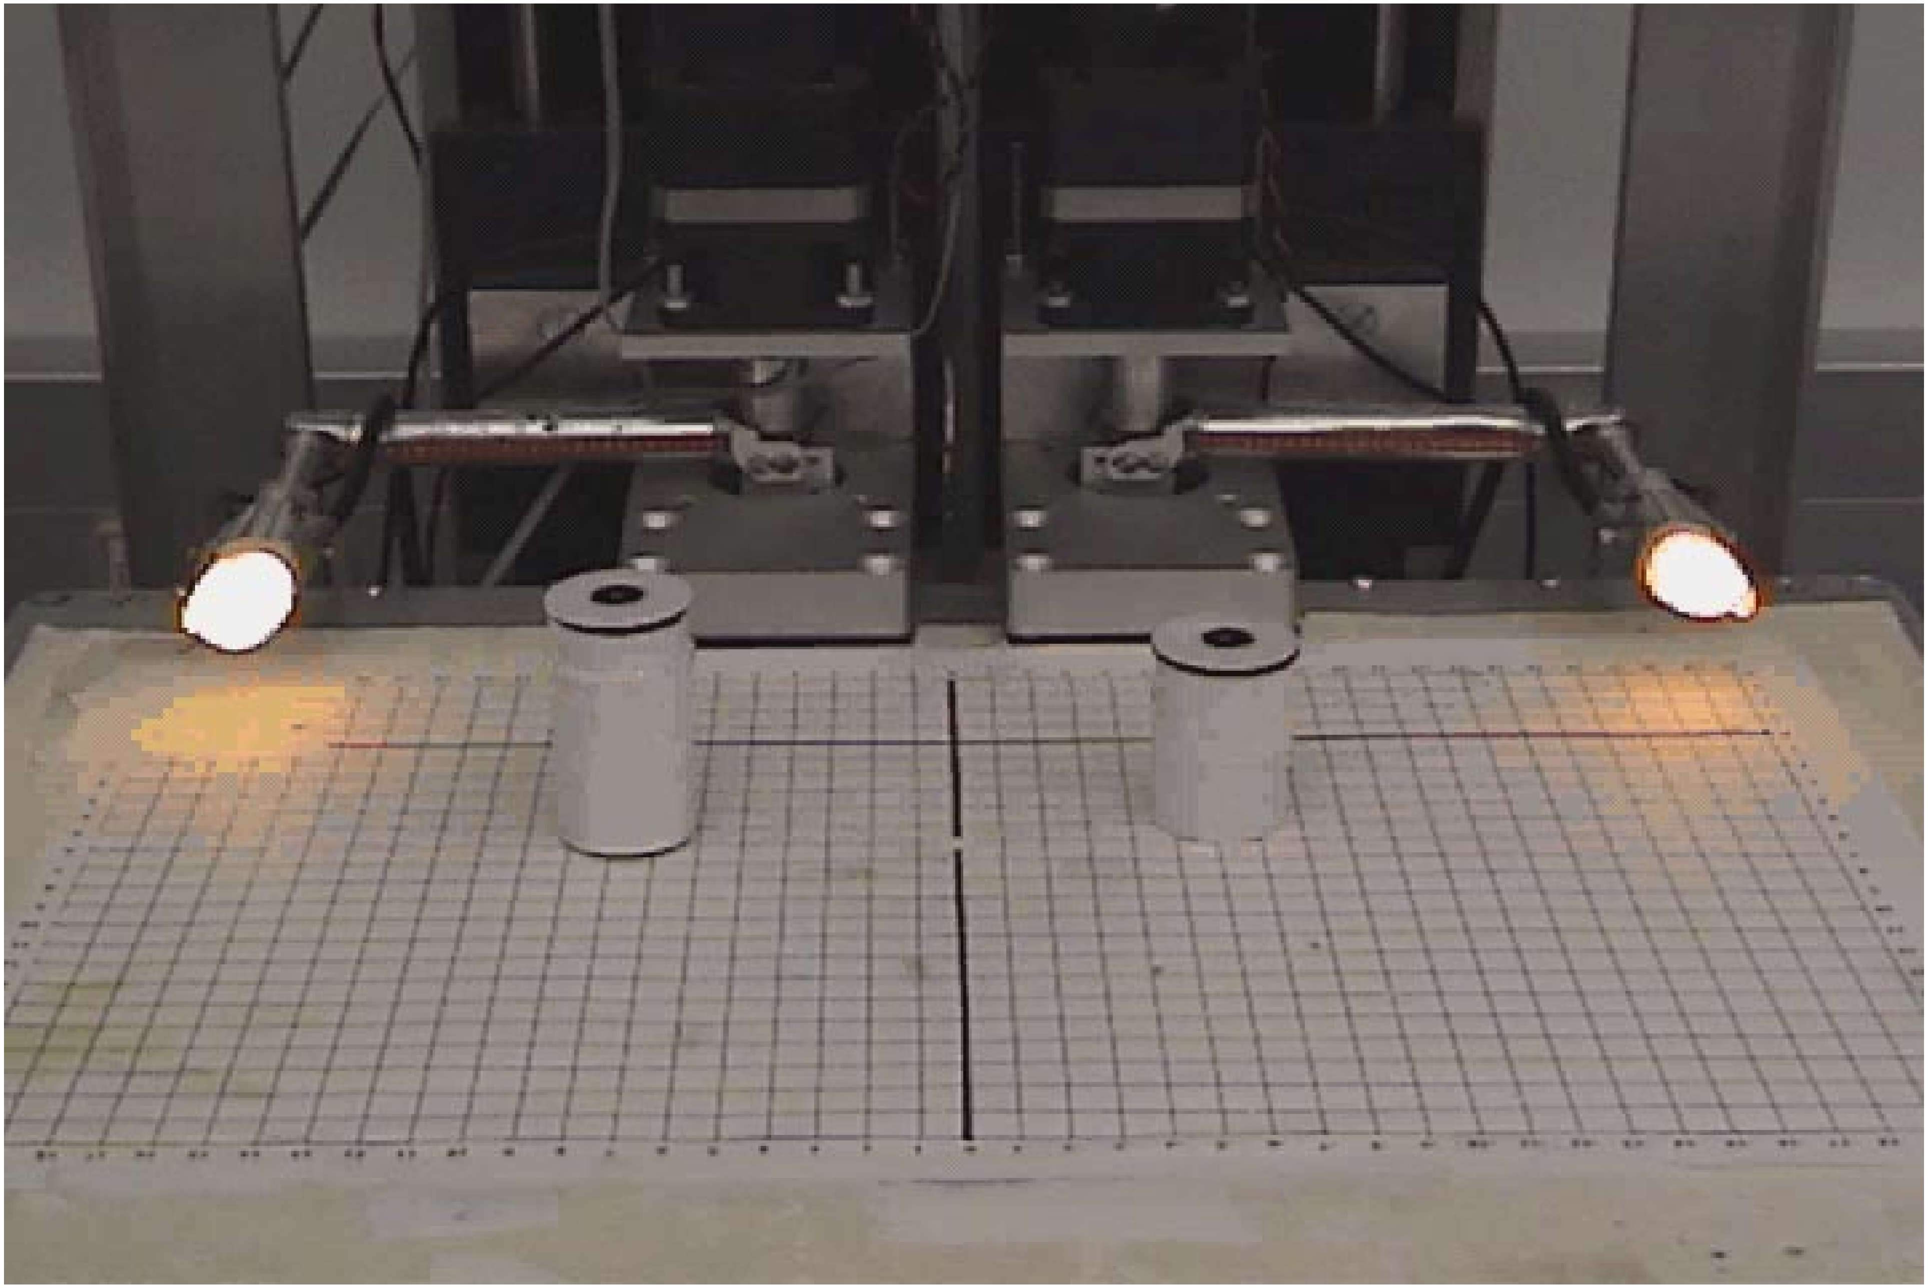

Representation of the values of u(t,x) and the tiling of space-time that is implicit in any implementation of computer control.

Bridging the Gap Between the Continuous and the Discrete

If $u(y,t)$ is the desired control then the magnitudes of the partial derivatives $u_x(x,t)$ and $u_t(x,t)$ give an indication of how hard it will be to approximate u with a piecewise constant function.

We can save resources with little loss in performance by non-uniform quantization when u(t,x) changes slowly outside the normal range of the variables.

Relationship with Stochastic Control

Suppose that the model is a stochastic differential equation. In this case the density does not evolve according to a Liouville equation but rather according to a refinement, the Fokker-Planck equation.

$$dx = f(x,u)dt + gdw$$

$$\frac{\partial \rho(t,x)}{\partial t} = -\langle \frac{\partial}{\partial x}, f(x,u)\rho(t,x)\rangle + \frac{1}{2}\sum \frac{\partial}{\partial x}^T gg^T \frac{\partial}{\partial x}$$

In the case where x can be observed perfectly, the problem might be to pick u(t,x) so as to minimize

$$\eta = \int_0^T \int_{-\infty}^{\infty} (x^2 + u^2)\rho(t,x)dxdt$$

Multiple Trajectories and the Fokker-Planck Equation

If we add noise to the equation dx/dt = f(x,u) then the partial differential equation for the evolution of the density of states is

$$\frac{\partial \rho(t,x)}{\partial t} = -\nabla^T f(x,u)\rho(t,x) + \frac{1}{2}\nabla^T g g^T \nabla$$

The presence of the second order derivative leads to additional smoothness as characterized by diffusion processes.

Adriaan Fokker

Max Planck

Controlling the Whole Set of Behaviors

$$\frac{\partial \rho(t,x)}{\partial t} = -\nabla^T f(x,u)\rho(t,x) + \frac{1}{2}\nabla^T g g^T \nabla$$

The performance measure is

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

The boundary condition are

$$\rho(0,x) = \rho_0(x) \;\; ; \;\; u(t,0) = 0 \;\; ; \;\; u(0,x) = 0$$

The Singular Limits

The variational equation associated with the square of the gradient is just Laplace's Equation. This corresponds to the limit when the trajectories are ignored completely.

The variational equation corresponding to trajectory term with no importance being associated to the attention term yields the usual optimal trajectory problem.

These two limits do not seem to offer much insight. In the first case the performance is not just bad, it is infinitely bad. In the second case the attention required is infinite.

Limitations of Linearity

Because of saturation we can assert that there are no linear systems. Even so, linear models are very useful and their properties offer considerable insight. We can design a control law as if linear implementation were possible and let the hardware do the truncation. Or, we can acknowledge the limitations it in the design process.

It makes little sense to model many on-off systems such as those that find wide use in low tech control such as the electric valves in dish washers and gasoline pumps as linear systems.

The outcome of a measurement is often a go/no-go decision. In such cases it makes little sense to regard the measurement as being linear.

The default assumption of linearity can be misleading either because of saturation or because of discontinuity, or both.

A Way to Think about Learning and Practice

The optimization problem posed here involves a trade-off between The quality of the trajectory and the implementation costs. The latter involves a trade-off between open loop cost measured by u_t and closed loop costs measured by u_x. A model for what happens when one practices a task is to imagine that the weighting shifts from the open loop term to the closed loop term and from the implementation cost to the trajectory cost.

Conclusions

1. We have framed the problem of optimizing the implementation cost in terms of an optimization problem involving variational problems on (t,x)-space.

2. The solution of such problems will generally lead to less control "activity", saturating control laws, and require nonlinearity. .

3. The smallness of the partial derivatives implies that the control law will change slowly, be relatively insensitive to error, and lend itself to roughly quantized, slowly sampled implementations.

4. These ideas have been useful to us in thinking about the how achieve good control with limited resources in terms of communication and computation.

Acknowledgements

These ideas have developed over a period of years with inputs from many people. These include

John Baillieul, Boston University
Dimitri Hristu, University of Maryland
Navin Khaneja, Harvard University
P. S. Krishnaprasad, University of Maryland
Magnus Egerstedt, Georgia Tech
Wing Wong, Chinese University of Hong Kong

With Krishnaprasad providing particular advice on the presentation given today