Example of Exact Trade-offs in Linear Controller Design

Craig Barratt and Stephen Boyd

ABSTRACT: The design of a linear time-invariant controller for a given linear time-invariant plant, like any engineering design, involves trade-offs among many desirable qualities, such as fast response to commands without excessive overshoot, low actuator authority, robustness low controller complexity, and so on. Only for a few very special cases are analytic methods known for finding the exact form of these trade-offs. Two examples of such analytic methods are Linear Quadratic Gaussian theory, where the plant actuator and output variance can be traded off, and Nevanlinna-Pick theory, where, for instance, the achievable disturbance rejection can be traded off in two different bandwidths. In many cases, the limit of performance achievable with a linear time-invariant controller, and thus the *exact form of the trade-offs*, can be computed numerically. To demonstrate this trade-off, the paper treats the design of a regulator for a very typical plant, a double integrator with some excess phase. The trade-offs are presented for two different measures of robustness with noise sensitivity. The exact form of these trade-offs is determined numerically by using the techniques described in the appendixes.

Introduction

Given a (possibly unstable) linear time-invariant plant, a basic design problem is to find a linear time-invariant controller that will stabilize the plant and meet various design objectives. Beyond simple intuition, little is known about how the different design objectives or goals trade off.

The basic character of some of the design trade-offs is obvious from simple physical reasoning, e.g., it takes more force (actuator authority) to move a mass from one place to another more quickly (faster command response). Other trade-offs are more subtle. For example, it is well known to control engineers that robustness or stability margins trade off against closed-loop bandwidth for any plant with delay or other excess phase [1].

Recently [2], it was observed that the limit of performance achievable with a linear time-variant controller can be computed numerically, where performance reflects many practical constraints and qualities, including fast response to commands without excessive overshoot or undershoot, small and quick reactions to disturbances or noises, low actuator authority, and certain measures of robustness or insensitivity to unknown or unmodeled plant dynamics. Some design goals, such as low controller complexity and open-loop controller stability, *cannot* be included using the methods in [2].

Finding the Limit of Performance

Recent theoretical results [3] show that the set of closed-loop responses achievable with stabilizing controllers is *affine*. This means that if controllers C_1 and C_2 each stabilize a given fixed plant P, and yield closed-loop transfer functions H_1 and H_2, then, for every real number λ, there is a controller C_λ, which stabilizes P and yields the closed-loop transfer function $\lambda H_1 + (1 - \lambda) H_2$. Geometrically, this means every closed-loop transfer function on the *line* through H_1 and H_2 is achieved by some stabilizing controller. Note that, in general, the controller C_λ is *not* given by $\lambda C_1 + (1 - \lambda) C_2$. These results allow us to express simply every closed-loop transfer function achievable with stabilizing controllers.

Now consider the controller design problem, which is the problem of finding a controller that achieves particular design goals. In many cases, the design goals can be expressed as *convex constraints* on various closed-loop transfer functions. A constraint on a closed-loop transfer function is *convex* if the average value of two closed-loop transfer functions $(H_1 + H_2)/2$ satisfies the constraint whenever H_1 and H_2 do (here H_1 and H_2 are the closed-loop transfer functions yielded by controllers C_1 and C_2). Examples of convex constraints are design goals involving envelope bounds on closed-loop time responses to any input (e.g., step responses and impulse responses), sensitivity to noises, closed-loop frequency responses, and some robustness requirements. Examples of design goals that cannot be expressed as *convex* constraints on closed-loop transfer functions are restrictions on the structure or order of the controller, e.g., the number of controller poles.

In the case where *every* design goal can be expressed as a *convex* constraint on some closed-loop transfer function, the controller design problem is equivalent to finding a closed-loop transfer *matrix* (a collection of the appropriate closed-loop transfer functions) that satisfies the design goals. The simple representation of the achievable closed-loop transfer functions means that this new problem can be solved numerically to arbitrary accuracy [2].

A point on a trade-off curve can be found by constraining all but one of the design goals and minimizing the remaining one (it is assumed that all design goals are expressed so that "smaller" is "better"). The corresponding optimization problem is *convex*. This means there are no local minima. Such *convex* optimization problems can be solved numerically to arbitrary accuracy.

The Plant

We consider the design of a regulator for a single-input/single-output plant that consists of a double integrator with some excess phase,

$$P(s) = (1/s^2)\,(4 - s)/(4 + s)$$

The all-pass term $(4 - s)/(4 + s)$ approximates a 0.5-sec delay ($\exp(-s/2)$) at low frequencies. We may think of the all-pass term as accounting for any and all of a variety of sources of excess phase in a real control loop, e.g., small delays, antialias filters, equivalent excess phase contributed by a sample-and-zero-order-hold plant input, and so on. The idea of using a double integrator plant with some excess phase as a simple but realistic typical plant with which to explore control design trade-offs is taken from a study presented by Gunter Stein in [4].

The plant is discretized using a zero-order

Craig Barratt and Stephen Boyd are with the Information Systems Laboratory in the Electrical Engineering Department at Stanford University, Stanford, CA 94305.

0272-1708/89/0100-0046 $01.00 © 1989 IEEE

hold at 10 Hz, giving the following transfer function:

$$P(z) = \frac{-0.00379(z^2 - 0.7241z - 1.1457)}{z^3 - 2.6703z^2 + 2.3406z - 0.6703} = \frac{-0.00379(z - 1.492)(z + 0.7679)}{(z - 1)^2(z - 0.6703)}$$

The objective is to investigate some trade-offs in the design of a regulator C for this plant, where the closed-loop system is shown in Fig. 1. The discrete-time inputs w and v represent actuator (input-referred process noise) and sensor noise, which are taken to be independent white noise. The discrete-time outputs u and y are the actuator and plant outputs.

LQG: An Analytically Computable Trade-off

Consider the steady-state noise at u and y in Fig. 1 due to the injection of noises at v and w. As usual, assume that v and w are zero mean, white, independent, with root-mean-square (rms) values of 1 for v and 10 for w.

The trade-off between the steady-state output variance $\lim_{k\to\infty} Ey_k^2$ and the steady-state actuator variance $\lim_{k\to\infty} Eu_k^2$ is analytically computable from Linear Quadratic Gaussian (LQG) theory [5]. Given a fixed positive ρ, the LQG optimal regulator minimizes (over all stabilizing regulators) the weighted cost function J, which is a linear combination of the actuator and output variance,

$$J = \lim_{k\to\infty} E\{y_k^2 + \rho u_k^2\}$$

Appendix B gives an example where the discrete-time LQG optimal regulator is found for the plant with ρ equal to 0.0001.

Note that for each ρ, the LQG optimal regulator gives the best noise sensitivity (i.e., the minimum J) that can be achieved by *any* linear time-invariant regulator (of any complexity or structure) that stabilizes the plant P. Thus, the rms values of u and y achieved by the LQG regulator give a point on the trade-off curve between these quantities. Let us justify this assertion. If some other linear time-invariant regulator stabilized P and achieved better rms values of *both* u and y, then it would achieve a cost J smaller than the LQG regulator (recall that ρ is greater than 0). This is not possible, since the LQG optimal regulator gives the smallest J achievable by any linear time-invariant regulator.

As the number ρ varies, the rms values of u and y achieved by the corresponding LQG optimal regulator sweep out the trade-off curve. For our plant, the trade-off between the rms values of u and y is shown in Fig. 2. The interpretation of the curve in Fig. 2 is as follows. *No linear time-invariant regulator C that stabilizes P can achieve rms values of u and y, which lie below the curve.* This is true for regulators of any order, designed by any method. We can restate this as: *every* regulator C that stabilizes P achieves rms values of u and y, which lie in the shaded region, on or above the trade-off curve. For example, the following simple lead-lag regulator stabilizes P and achieves rms values of u and y of 27.55 and 5.98, respectively.

$$C(z) = 10(z - 0.98)/(z - 0.56)$$

This property is shown in Fig. 3. This regulator achieves closed-loop performance, which lies in the shaded region of Fig. 2, as it must.

Any regulator $\tilde{C}$ that gives closed-loop performance in the shaded region in Fig. 3 will perform the same or better (in terms of rms values of u and y) than the simple lead-lag regulator C. This means that $\tilde{C}$ will achieve the same or better output regulation (rms y no worse than 5.98) *and* use the same or less actuator effort (rms u no worse than 27.55). One family of such regulators is the LQG optimal regulators with ρ between 0.00135 and 2.31. These regulators lie on the boundary of the shaded region. There are infinitely many other regulators that give closed-loop performance in the shaded region.

This idea gives us a second interpretation of the trade-off curve in Fig. 2 in terms of achievable design goals. Any *design goals* (i.e., upper bounds on the rms values of u and y) that lie in the shaded region in Fig. 2 can be met or exceeded by some linear time-invariant regulator C, which stabilizes P. Design goals outside of the shaded region cannot be achieved by any linear time-invariant regulator C, which stabilizes the plant P. For example, consider the design goal that the rms value of u should be less than 10 and the rms value of y should be less than 4. This design goal can be achieved by a linear time-invariant regulator that stabilizes the plant P, whereas the design goal that the rms value of u should be less than 3 and the rms value of y should be less than 5 cannot be achieved by any linear time-invariant regulator that stabilizes the plant P.

For ρ equal to 10^{-4}, the details of the LQG design are given in Appendix B. The LQG optimal (current estimator) regulator is shown to be

$$C(z) = \frac{45.974z^3 - 72.79z^2 + 28.138z}{z^3 - 0.8061z^2 + 0.7107z - 0.1071} = \frac{45.974z(z - 0.91305)(z - 0.6703)}{(z - 0.17897)(z - 0.3136 + 0.7072j)(z - 0.3136 - 0.7072j)}$$

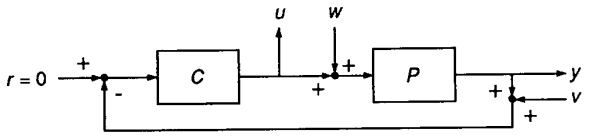

Fig. 1. Regulator system.

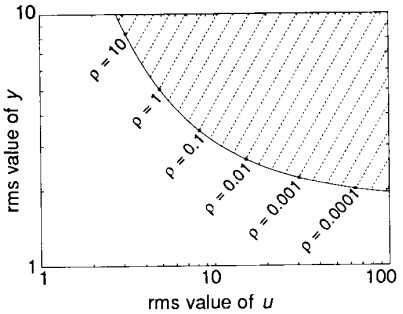

Fig. 2. Trade-off between achievable input and output rms noise sensitivities, on a log-log scale.

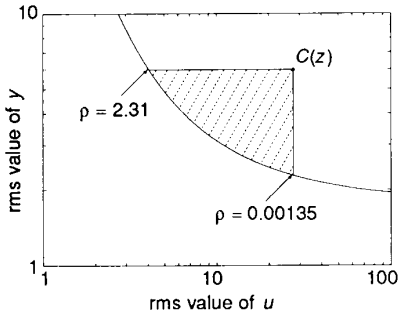

Fig. 3. Simple regulator C, which achieves performance above the trade-off curve.

For this regulator, the rms values of u and y are 63.73 and 2.0184, respectively. The value of the cost function J is 4.48.

Two Measures of Robustness

Tolerance of Additive Loop Perturbations

One measure of robustness of a control system, which combines the gain and phase margins, is the M-circle radius, defined as the minimum distance from the Nyquist plot of the loop gain $PC(\exp j\Omega)$ to the critical point -1. Mathematically, this can be written as

$$M = \text{dist}\,(PC(\exp j\Omega), -1) = \min_{0 \le \Omega \le \pi} |1 + PC(\exp j\Omega)|$$

The M-circle radius gives a measure of the sensitivity to additive loop perturbations. If the distance M is small, then slight variations in the loop gain PC could change the number of net clockwise encirclements of -1 by the loop gain, resulting in an unstable closed-loop system.

The M-circle radius is related to the *maximum sensitivity* of the control system. The M-circle radius is the inverse of the maximum sensitivity

$$M = 1/\|S\|_\infty$$

where $S = 1/(1 + PC)$ is the *sensitivity transfer function*, and the maximum of the sensitivity transfer function over all frequencies is

$$\|S\|_\infty = \max_{0 \le \Omega \le \pi} |S(\exp j\Omega)|$$

(In this notation, $\|H\|_\infty$ denotes the maximum magnitude over real frequencies of a transfer function H.) Thus, a small M-circle radius corresponds to *peaking* of the sensitivity function at some frequency. It also follows that constraining the maximum magnitude of the closed-loop transfer function S to be at most 1 over $M_{\min}$ is equivalent to constraining the M-circle radius to be at least $M_{\min}$.

For example, the $\rho = 10^{-4}$ LQG regulator gives $\|S\|_\infty = 3.34$, so that $M = 0.30$ is the closest the Nyquist plot comes to the critical point -1. This can be seen from the Nyquist plot in Fig. 4. The M-circle of radius 0.30 is also shown in Fig. 4. The magnitude of the sensitivity transfer function S for the $\rho = 10^{-4}$ LQG regulator is shown in Fig. 5. Note that its maximum is 10.5 dB (which is equal to $-20 \log M$).

Tolerance of Additive Plant Perturbations

The second measure of robustness we will consider concerns the ability of the regulator

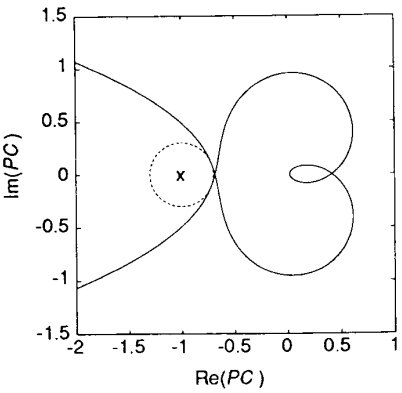

Fig. 4. Nyquist plot of the LQG regulator. Note: $\rho = 10^{-4}$ and the $M = 0.30$ circle is shown.

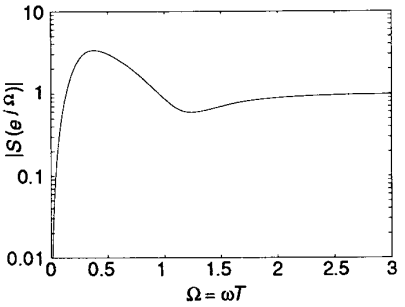

Fig. 5. Magnitude of sensitivity transfer function of LQG regulator. Note: $\rho = 10^{-4}$.

to maintain closed-loop stability in the face of stable additive plant perturbations ΔP, as shown in Fig. 6. The plant perturbation could represent errors in modeling the plant, component variations, or deliberately ignored plant dynamics.

Intuitively, if ΔP is small at all frequencies, then we would expect that if the regulator C stabilizes the plant P, then C should stabilize $P + \Delta P$ as well. This intuition is indeed correct, a consequence of the small gain theorem [6], [7].

We may ask, what is the smallest (in the sense of maximum magnitude of frequency response) stable additive plant perturbation ΔP that will destabilize the system in Fig. 6? Let us define the quantity D as the smallest plant perturbation ΔP that will destabilize the closed-loop system,

$$D = \min_{\Delta P \text{ destabilizes system}} \|\Delta P\|_\infty$$

It is not difficult to derive that the inverse of D is equal to the maximum of the closed-loop transfer function from the reference in-

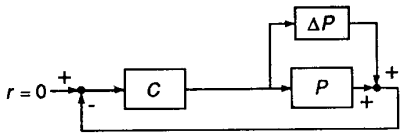

Fig. 6. Additive plant perturbation.

put r to the actuator output u,

$$1/D = \|C/(1 + PC)\|_\infty$$

This observation is due to Doyle and Stein [6].

The positive number D can be interpreted as the largest size of stable additive plant perturbation the control system can be guaranteed to withstand. A regulator that yields small values of D corresponds to a system that can be destabilized by a small stable additive plant perturbation. It also follows that constraining the maximum magnitude of the closed-loop transfer function $C/(1 + PC)$ to be at most $1/D_{\min}$ is equivalent to constraining the additive plant perturbation tolerance D to be at least $D_{\min}$.

For the $\rho = 10^{-4}$ LQG optimal regulator, the peak of the closed-loop transfer function $C/(1 + PC)$ is 83.02; so $D = 0.0121$. Hence, there are stable plant perturbations ΔP that destabilize the system in Fig. 6, with $\|\Delta P\|_\infty$ as small as 0.0121. One destabilizing perturbation is shown for which $\|\Delta P\|_\infty = 1/80$, which is greater than D.

$$\Delta P(z) = \frac{0.2152 \times 10^{-3}(1 - z)}{z^2 - 0.985z + 0.9801}$$

This ΔP destabilizes the system in Fig. 6. The frequency-response magnitude of ΔP is plotted in Fig. 7. The transfer function ΔP might represent a mechanical resonance, which was ignored when the plant P was modeled for the LQG design.

In general, the robustness requirements that M and D be large are independent. A system may have good margins (i.e., large M) but be quite sensitive to additive plant perturbations (i.e., small D) and vice versa.

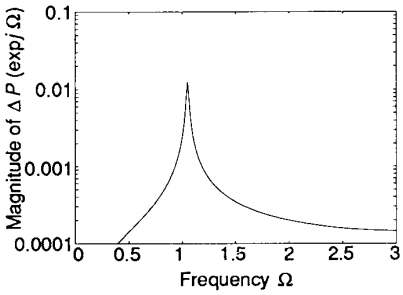

Fig. 7. Frequency response of a destabilizing ΔP.

Trade-off Curves Involving Noise Sensitivity and Robustness

We wish to minimize the noise sensitivity of the system and determine how this trades off with the two different robustness requirements described in the previous section. Specifically, how does the noise sensitivity J trade off against M when we require $D \geq D_{\min}$, and how does it trade off against D when we require $M \geq M_{\min}$?

The design problem can be expressed mathematically as

$$J_{\min} = \min_{\substack{C \text{ stabilizes } P \\ 1/M = \|1/(1+PC)\|_\infty \leq 1/M_{\min} \\ 1/D = \|C/(1+PC)\|_\infty \leq 1/D_{\min}}} J$$

Unlike LQG, no exact analytical solution of this minimization problem is known, although some work has been done [8]. Nevertheless, this minimization problem can be solved numerically since it can be cast as an (infinite-dimensional) convex optimization problem [2]. In general, the optimizing controllers are of high order. In general, no method is known for finding low-order controllers that achieve close to the optimal performance.

For three fixed values of $D_{\min}$, the trade-off between $J_{\min}$ and $1/M_{\min}$ is shown in Fig. 8. The $\rho = 10^{-4}$ LQG regulator is also shown. Since this regulator achieves $1/D = 83.02$, it lies below the $1/D_{\min} = 10$ curve and on the $1/D_{\min} = 83.02$ curve. All trade-off curves with $1/D_{\min} \geq 83.02$ will pass through the LQG performance point and be horizontal to the right of it.

Note the interesting fact that by allowing the M-circle radius to be less than 0.5, only modest improvement in the noise response is gained, with the same tolerance D to additive plant perturbations.

For four fixed values of $M_{\min}$, the trade-off between $J_{\min}$ and $1/D_{\min}$ is shown in Fig. 9. The $\rho = 10^{-4}$ LQG regulator is also shown. Since this regulator achieves $1/M = 3.34$, it lies below the $1/M_{\min} = 2.0$ curve and on the $1/M_{\min} = 3.34$ curve. All trade-off curves with $1/M_{\min} \geq 3.34$ will pass through the LQG performance point and be horizontal to the right of it.

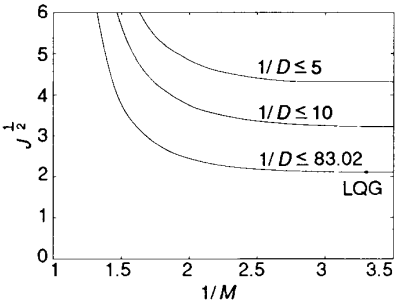

Fig. 8. Trade-off between rms noise and M-circle radius.

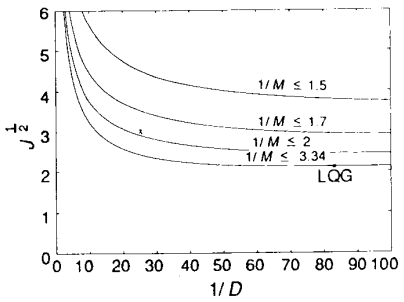

Fig. 9. Trade-off between rms noise and additive plant sensitivities.

From Fig. 9 we can draw some interesting conclusions. Consider the $1/M \leq 3.34$ curve, which corresponds to regulators that yield the same or larger M-circle radius as the $\rho = 10^{-4}$ LQG regulator. The curve is relatively flat for $1/D \geq 20$, meaning that D can be increased to about 0.05 with a relatively small increase in rms noise response and the same or larger M-circle radius (0.3). For the $\rho = 10^{-4}$ LQG regulator, this represents an increase in additive plant perturbation tolerance D of a factor of 4.

Of course, all regulators that stabilize P yield a noise sensitivity $J \geq J_{\text{LQG}}$, so that all curves lie on or above the horizontal asymptote $J_{\text{LQG}}^{1/2} = 2.117$. Imposing further constraints on the two measures of robustness naturally will increase the minimum noise sensitivity J achievable with linear time-invariant stabilizing regulators. What is neither intuitively obvious nor analytically computable is *how much* the minimum noise sensitivity J must increase when we impose various constraints on the two measures of robustness. Figures 8 and 9 show this trade-off precisely.

Let us give two examples of specific conclusions we may draw from Figs. 8 and 9. First, the following design goals can be achieved with a linear time-invariant regulator that stabilizes P (this point is marked "x" in Fig. 9).

$$J^{1/2} \leq 3, \quad M \geq 0.5, \quad D \geq 0.04$$

These goals represent an increase in noise response over the LQG regulator (as any achievable goal must), a moderate improvement in M-circle radius over the LQG controller, and a substantial improvement in D, tolerance to additive plant perturbation, over the LQG controller:

$$J_{\text{LQG}}^{1/2} = 2.117, \quad M_{\text{LQG}} = 0.3,$$
$$D_{\text{LQG}} = 0.0121$$

Another important conclusion we may draw from Figs. 8 and 9 is that the following design goals are not achievable by any linear time-invariant regulator that stabilizes P.

$$J^{1/2} \leq 3, \quad M \geq 0.5, \quad D \geq 0.08$$

Consider the (easily) achievable design goals

$$J^{1/2} \leq 3, \quad M \geq 0.35, \quad D \geq 0.04$$

It is interesting to consider whether we can achieve these goals using an LQG design with some appropriate weight $\rho > 0$. The process of "tuning" the LQG weights (ρ in our case) in an attempt to achieve particular design goals is typical of how LQG is used in practice.

For each fixed ρ, the LQG optimal regulator stabilizes the plant P and achieves particular values of J, D, and M (note that J denotes the fixed weighted noise sensitivity, $J = \lim_{k\to\infty} E\{y_k^2 + 10^{-4}u_k^2\}$, and is not a direct function of ρ). These values are plotted in Fig. 10. As expected, the minimum value of J is achieved for $\rho = 10^{-4}$, since the other LQG regulators each minimize a slightly different cost function.

The preceding design goals are also shown in Fig. 10. As can be seen from Fig. 10, no

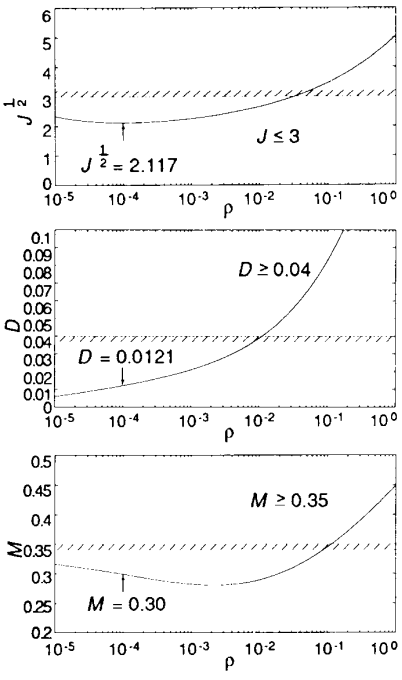

Fig. 10. Performance of a family of LQG regulators.

value of ρ (i.e., no amount of "tweaking" of ρ) will yield a regulator that achieves the preceding goals. However, it is clear from Fig. 9 that these design goals can be achieved by some linear time-invariant stabilizing regulator.

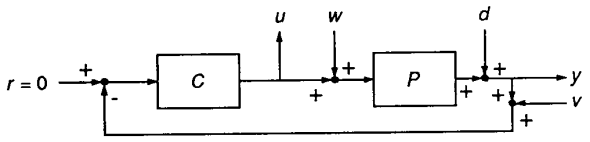

Fig. A1. Regulator system setup for `qdes`.

Conclusion

We have given an example demonstrating how the methods of [2] can be used to determine numerically the various trade-off curves for competing design goals or objectives. It should be emphasized that the trade-off curves shown in Figs. 8 and 9 represent *fundamental trade-offs:* they do not merely show the best we can do; they show the best anybody can do with a linear time-invariant regulator of any form or complexity, designed by any scheme or method. We believe that this information can be very useful to the designer.

No control engineer would be surprised by the general shape of the trade-off curves shown in Figs. 8 and 9: It is intuitively obvious that some improvement in robustness can be realized at the cost of some degradation in root-mean-square noise sensitivity. However, even for this very simple, typical plant, the numerical values of the trade-offs—how much improvement in robustness can be "bought" for a given increase in noise response—are not at all obvious. Naturally, for a complicated multi-input/multi-output plant and a much larger set of design goals, specifications, and objectives, the trade-offs would be considerably less obvious, and such computations correspondingly more valuable.

Appendix A: Using qdes to Find the Limits of Performance

Points near the exact trade-off curve were found using the computer program `qdes`, described in [2], which should be consulted for full details. The program `qdes` translates a description of the control design problem into an approximation of the corresponding convex infinite-dimensional optimization problem (over a parameterization of all controllers that stabilize the plant P). The accuracy of the approximation can be improved arbitrarily at the cost of computation time. The source file for `qdes` contains the specification of the design objective (to be minimized) and the design constraints.

The discrete-time plant was set up with two exogenous inputs v and w. To enable the specification of a bound on $\| 1/(1 + PC) \|_\infty$, a third fictitious exogenous input d was used. The three exogenous inputs were then v, w, and d. The two regulated outputs were the controller output u and the plant output y. The corresponding block diagram is shown in Fig. A1.

These exogenous inputs and regulated outputs were referred to symbolically by `V, W, D, U,` and, `Y` respectively, in the `qdes` source file by use of the `qdes #define` facility.

A nominal controller (which stabilized the plant) was found, and the Q parameterization was applied to yield transfer-function matrices T_1, T_2, and T_3. Every closed-loop transfer matrix from $[v\ w\ d]^T$ to $[u\ y]^T$ achievable by a controller that stabilizes P is of the form shown for some stable Q.

$$T_1 + T_2 Q T_3$$

Here, T_1 is a 2×3 transfer matrix, T_2 is 2×1, T_3 is 1×3, and Q is a single-input/single-output transfer function.

To set up the problem for `qdes`, the impulse responses of T_1, T_2, and T_3 (which are always stable) were truncated at 200 samples.

Since the noises v and w are independent and white, the weighted noise sensitivity J can be expressed as a linear combination of the sums of squares of impulse responses. For example, if w acted alone, then it is easy to show that

$$\lim_{n\to\infty} Ey_n^2|_w = Ew^2 \sum_{i=0}^{\infty} h_{yw}(i)^2$$

$$= Ew^2 \| h_{yw} \|_2^2$$

where $h_{yw}(i)$ is the impulse response from w to y at time i. We see that the steady-state variance of y is the variance of w times $\| h_{yw} \|_2^2$, the sum of the squares of the impulse response from w to y.

Similarly, the steady-state variance of y due to sensor noise is given by $Ev^2 \| h_{yv} \|_2^2$. Since v and w are independent, the total steady-state variance or power in y is simply the sum of these powers,

$$\lim_{n\to\infty} Ey_n^2 = Ew^2 \| h_{yw} \|_2^2 + Ev^2 \| h_{yv} \|_2^2$$

$$= 100 \| h_{yw} \|_2^2 + \| h_{yv} \|_2^2$$

In `qdes`, the quantity $\| h_{yw} \|_2^2$ is referred to by `norm_h_sqr[Y] [W].`

After the preamble, the design problem was expressed to `qdes` as (the objective is the sum of each of the terms)

```
minimize {
  norm_h_sqr[Y] [V];          /* output pwr due to sensor noise */
  100 * norm_h_sqr[Y] [W];    /* output pwr due to process noise */
  0.0001 * norm_h_sqr[U] [V]; /* actuator pwr due to sensor noise */
  0.01 * norm_h_sqr[U] [W];   /* actuator pwr due to process noise */
}

subject_to {
  max_Mag_H[U] [V] <= 1/D_MIN; /* robust to additive plant perturbation */
  max_Mag_H[U] [D] <= 1/M_MIN; /* M-circle radius constraint */
}
```

D_MIN and M_MIN were specified using the #define facility in qdes. A single point on the trade-off curve was found by running qdes with the constraints D_MIN and M_MIN fixed. The qdes objective gave the value of the constrained minimum noise sensitivity J. A trade-off curve was swept out by varying one of the constraints D_MIN or M_MIN.

We expect the value of J computed by qdes to be close to the minimum J achievable by any stabilizing controller that meets the specified robustness constraints. To provide confidence in the values computed, the two approximations in qdes were refined steadily. These approximations are the number of taps in Q and the number of frequencies at which frequency-response constraints are sampled.

For each point on each trade-off curve, the number of taps in Q was increased (by roughly a factor of 2) and the objective function J was recomputed. This was repeated until J changed negligibly. We found a negligible change in the qdes results for the number of taps between 65 and 200, and, in many cases, for the number of taps between 20 and 200. (It is interesting to note that we found a very *nonnegligible* change in total qdes computation time, however! The computation times were on the order of 60 min on a Sun 3/260 for a 200-tap run.) Similarly, the number of points at which the frequency-response constraints were sampled was also increased, also with negligible change in qdes output. We comment that even very crude approximations (e.g., 10 taps) often yielded values of J only a few percent above those yielded for very fine approximations.

As another test of the validity of the qdes results, the nominal controller was changed and the procedure described earlier was repeated. This also produced negligible change in the computed results.

We believe that the trade-off curves shown in Figs. 8 and 9 are within, at most, a few percent of the true trade-off curves, and probably much closer.

Appendix B: Details of LQG Design

A state-space realization of the plant with input u, input-referred process noise w, and sensor noise v is

$$x_{k+1} = Ax_k + bu_k + bw_k$$

$$y_k = cx_k + v_k$$

where

$$A = \begin{bmatrix} 0.6703 & 0 & 0 \\ 0.08242 & 1 & 0 \\ 0.004395 & 0.1 & 1 \end{bmatrix}$$

$$b = \begin{bmatrix} 0.08242 \\ 0.004395 \\ 0.0001512 \end{bmatrix}$$

$$c = [0 \quad -1 \quad 4]$$

The process and sensor noise variances are

$$W = Ew_k^2 = 100$$

$$V = Ev_k^2 = 1$$

We wish to minimize the weighted noise at the plant's input and output

$$J_{LQG} = \lim_{k\to\infty} E\{y_k^2 + \rho u_k^2\}$$
$$= \lim_{k\to\infty} E\{x_k^T Q x_k + u_k^T R u_k\}$$

where

$$Q = c^T c = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & -4 \\ 0 & -4 & 16 \end{bmatrix}$$

$$R = \rho = 10^{-4}$$

Solving the discrete-time LQG problem gives an optimal estimator gain l and an optimal state feedback gain k, where

$$l = \begin{bmatrix} 0.0 \\ 0.2 \\ 0.14 \end{bmatrix}$$

$$k = [10 \quad 89.96 \quad 199.9]$$

A state-space realization of the optimal current estimator LQG controller is

$$A_{LQG} = A - bk - Alc + bklc$$

$$= \begin{bmatrix} -0.154 & -11.2 & -1.317 \\ 0.03846 & 0.6025 & -0.8702 \\ 0.002882 & 0.2394 & 0.3576 \end{bmatrix}$$

$$b_{LQG} = (A - bk)\, l = \begin{bmatrix} -3.789 \\ -0.002066 \\ 0.153 \end{bmatrix}$$

$$c_{LQG} = k - klc = [10 \quad 135.9 \quad 15.97]$$

$$d_{LQG} = kl = 45.974$$

The transfer function of the controller is then

$$C_{LQG}(z)$$
$$= c_{LQG}(zI - A_{LQG})^{-1} b_{LQG} + d_{LQG}$$
$$= \frac{45.974z^3 - 72.79z^2 + 28.138z}{z^3 - 0.8061z^2 + 0.7107z - 0.1071}$$

Acknowledgments

Research sponsored in part by the Office of Naval Research under N00014-86-K-0112, by the National Science Foundation under ECS-85-52465, by an IBM faculty development award, and by Bell Communications Research. The authors thank Art Bryson and Gene Franklin for their numerous comments and helpful suggestions. Thanks also to P. Gill, W. Murray, M. Saunders, and M. Wright of the Stanford Optimization Laboratory for providing the convex optimization code LSSOL [9], [10].

References

[1] B. A. Francis and G. Zames, "On H_∞-Optimal Sensitivity Theory for SISO Feedback Systems," *IEEE Trans. Auto. Contr.*, vol. AC-29, no. 1, pp. 9–16, Jan. 1984.

[2] S. Boyd, V. Balakrishnan, C. Barratt, N. Khraishi, X. Li, D. Meyer, and S. Norman, "A New CAD Method and Associated Architectures for Linear Controllers," *IEEE Trans. Auto. Contr.*, vol. AC-33, no. 3, pp. 268–283, Mar. 1988.

[3] M. Vidyasagar, *Control System Synthesis: A Factorization Approach*, MIT Press, 1985.

[4] D. Flamm, S. Boyd, G. Stein, and S. Mitter, Tutorial Workshop on H_∞-Optimal Control, Preconference Workshop, Conference on Decision and Control, Dec. 1987.

[5] A. E. Bryson and Y. C. Ho, *Applied Optimal Control*, New York: Hemisphere Publishing, 1975.

[6] J. C. Doyle and G. Stein, "Multivariable Feedback Design: Concepts for a Classical/Modern Synthesis," *IEEE Trans. Auto. Contr.*, vol. AC-26, no. 1, pp. 4–16, Feb. 1981.

[7] G. Zames, "On the Input-Output Stability on Nonlinear Time-Varying Feedback Systems," *IEEE Trans. Auto. Contr.*, Parts I and II, vol. AC-11, pp. 228–238, 1966.

[8] D. S. Bernstein and W. M. Haddad, "LQG Control with an H_∞ Performance Constraint: A Riccati Equation Approach," in *Proc. Amer. Contr. Conf.*, pp. 796–802, 1988.

[9] P. E. Gill, S. J. Hammarling, W. Murray, M. A. Saunders, and M. H. Wright, *User's Guide for LSSOL (Version 1.0): A FORTRAN Package for Constrained Least-Squares and Convex Quadratic Programming*, Tech. Rept. SOL 86-1, Operations Research Dept., Stanford Univ., Stanford, CA, Jan. 1986.

[10] P. E. Gill, W. Murray, and M. Wright, *Practical Optimization*, London: Academic Press, 1981.

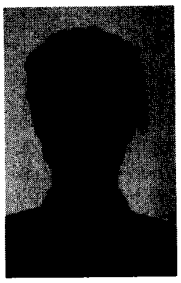

Craig Barratt received the B.Sc. degree in mathematics and physics in 1983 and the B.E. degree (Hons.) in electrical engineering in 1985 from the University of Sydney, Sydney, Australia, and the M.S. degree in electrical engineering from Stanford University. He is currently pursuing a Ph.D. degree in electrical engineering at Stanford.

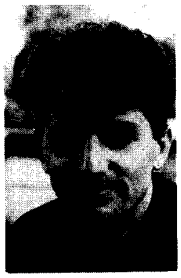

Stephen Boyd received the A.B. degree in mathematics from Harvard University in 1980 and the Ph.D. in EECS from the University of California, Berkeley, in 1985. Since 1985 he has been an Assistant Professor in the Information Systems Laboratory, Electrical Engineering Department, Stanford University.

1990 IFAC World Congress

The 11th World Congress of the International Federation of Automatic Control (IFAC) will be held August 13–17, 1990, in Tallinn, USSR. The theme of the Congress is "Automatic Control in Service of Mankind." The program will include plenary, paper, and discussion sessions, as well as case studies, all devoted to theory, applications, and industrial experience.

The Congress will take place at the Tallinn Technical University. Accommodations will be available in several easy-to-reach hotels. Social programs, sight-seeing tours, and technical visits will be organized. There will be many possibilities to meet and discuss with control experts from all over the world.

Tallinn is the capital of the Estonian SSR. It has half a million inhabitants. Tallinn is an important industrial, scientific, and cultural center, and one of the oldest and most interesting cities in the Soviet Union, famous for its medieval architecture. Tallinn can be reached by air via Moscow and Leningrad International Airports or by train, and also by ship from Helsinki, Finland. Tallinn lies in the temperate zone and has a humid climate. The average daytime temperature is 18–23°C in August.

IFAC's scope of interest deals with the science and technology of automatic control, both in theory and broad applications for all systems. This wide variation is indicated by its technical committees: Applications; Biomedical Engineering; Components and Instruments; Computers; Developing Countries; Economic and Management Systems; Education; Manufacturing Technology; Mathematics of Control; Social Effects of Automation; Space; Systems Engineering; Terminology and Standards; and Theory.

The provisional deadline for full-paper submission is July 1989. However, as soon as possible, prospective authors should obtain "Author Kits" from the IFAC Secretariat, Schlossplatz 12, A-2361, Laxenburg, Austria. For further information, contact: IFAC Tallinn 1990, Institute of Cybernetics, Akadeemia tee 21, 200108 Tallinn, USSR.

Workshop on Spectral Analysis

A Workshop on Higher-Order Spectral Analysis, sponsored by the Office of Naval Research and the National Science Foundation and in cooperation with the IEEE Control Systems Society, will be held at The Lodge at Vail, Vail, Colorado, on June 28–30, 1989. The objective of the workshop is to provide a forum for discussion of new theories and techniques in the area of higher-order spectral analysis and their application to practical problems. The intent is to attract researchers from this field to exchange ideas and debate current and future trends. The emphasis will be on technology transfer issues, i.e., how the theory and properties of polyspectra (or cumulants) can be used to solve important signal and systems analysis problems. Both digital and optical signal processing problems will be addressed.

The workshop will feature four types of sessions: (1) plenary sessions with keynote speakers; (2) tutorial sessions; (3) sessions with contributed presentations that will emphasize interdisciplinary contributions; and (4) panel discussions. The workshop also will include working social activities and will offer opportunities for those from similar disciplines to meet and discuss in smaller groups. Attendance at the workshop will be by both invitation and application. Copies of the Proceedings of the workshop will be made available to all attendees of the workshop.

Prospective authors should submit four (4) copies of a two-page summary including authors' names, addresses, affiliations, and telephone numbers prior to January 15, 1989, to: Prof. C. L. Nikias, CDSP Research Center, 416 Dana Research Bldg., Northeastern University, 260 Huntington Ave., Boston, MA 02115; phone: (617) 437-3352.