Robust Control with an H_2 Performance Objective

Andrew Packard [♯ ♭] John C. Doyle [♮ ♭]

Univ. of California, Berkeley
Mechanical Engineering
Berkeley, CA 94720

[♮] California Institute of Technology
Electrical Engineering 116-81
Pasadena, CA 91125

[♭] Honeywell SRC MN65-2500
3660 Technology Drive
Minneapolis, MN 55418

Abstract

This paper considers the problem of designing robust controllers with an H_2 performance objective. A modified version of μ-synthesis is proposed and compared with two alternative schemes.

I. Introduction

H_∞ methods have gained great popularity recently, in part because, unlike H_2 methods (e.g. LQG), they offer a single framework in which to study both performance and robust stability [F1]. Furthermore, extensions using the structured singular value, μ, provide methods for treating robust performance with structured uncertainty [D1]. Important questions remain, however, about the suitability of weighted H_∞ as a performance measure. We will not try to resolve this issue here, but simply make the claim that in some situations an H_2 performance objective is natural.

The relative merits of the H_∞ versus H_2 frameworks for robust performance problems is also less straightforward than is often assumed. It is well-known that H_2 optimal controllers have no a priori guaranteed robustness characteristics [D2], except in very special cases [S1]. This is not in the least remarkable since no optimal method should be expected to provide any characteristic that is not explicitly included in its cost function. Likewise, H_∞ optimal controllers have no particular robustness characteristics unless they are penalized. To obtain robust performance one must resort to something like μ-synthesis, which is still experimental. For a practically motivated example which illustrates these issues, see [D3].

Robust performance with an H_2 performance objective, but with H_∞ norm bounded uncertainty, is even less developed than for H_∞ performance. Some early attempts [D4], though primitive, have been refined into rather elaborate design methodologies (e.g. [ACC '86 invited session on LQG/LTR]). This paper will compare some alternatives, which though ad hoc, aim directly at robust H_2 performance. The first method is by Bernstein and coworkers [B1], who replace uncertain parameters with multiplicative white noise and solve a generalized LQG problem. The second method uses additive instead of multiplicative noise and is similar in spirit to [D4] and [S2]. The third method is a refinement motivated by μ, and seems quite promising. An attempt is made, with much hand-waving, to justify it. The methods are compared on the example from [D2].

II. Method 1: Multiplicative Noise

Method 1 uses a generalization of LQG, called the *Optimal Dynamic Compensation Problem* (ODCP), [B1], to achieve robustness and H_2 performance. In addition to the two stochastic processes which enter additively as measurement noise and state driving noise, the ODCP also allows for additional stochastic processes in the A, B, and C matrices which define the system. It is these processes that can, in some sense, model uncertain parameters. The following outline is taken from [B1], where details and other references can also be found.

Consider the system

$$\dot{x} = \left(A + \sum_{i=1}^{p} v_i A_i\right) x + \left(B + \sum_{i=1}^{p} v_i B_i\right) u + J\eta$$

$$y = \left(C + \sum_{i=1}^{p} v_i C_i\right) x + M\xi \quad , \quad e = \begin{bmatrix} Lx \\ Ru \end{bmatrix}$$

where η, ξ, v_1, v_2 , . . . , v_p are uncorrelated, unit intensity, zero mean, white noise processes. The ODCP method attempts to minimize $\lim_{t\to\infty} \mathrm{E}[e^T(t)e(t)]$. The optimality conditions yield 4 equations: two Riccati like equations, and two Lyapunov like equations which are all coupled together via the matrices A_i , B_i , and C_i. For nonzero A_i , B_i , and C_i, the equations are solved iteratively to generate the ODCP controller. If these are all zero, then the solution reduces to standard LQG.

To apply this technique for robust H_2 performance, begin with a standard LQG problem, call it $P1$, with uncertain parameters in the A, B, and C matrices. Ignoring this uncertainty, and solving the straight LQG with nominal values, may give a closed loop system whose stability is arbitrarily sensitive to parameter variations. To reduce this sensitivity, define a second problem, $P2$, which fits in the ODCP framework by treating the unknown parameters of

$P1$ as *stochastic processes* with mean equal to the nominal value and variance proportional to the size of the interval uncertainty.

It seems reasonable to expect that solving $P2$ via ODCP might produce a controller more robust for $P1$ than the LQG controller, but at the expense of some degradation in nominal performance. Indeed, sometimes this is the case. Unfortunately, it is possible for this scheme to actually degrade robustness as well, since the stochastic performance objective in $P2$ is only indirectly related to the robust H_2 performance of $P1$. In fact, it is possible for $\lim_{t\to\infty} \mathbf{E}[e^T(t)e(t)]$ to be finite (e.g. stochastic stability) but have the system be unstable for all fixed values of the parameters. While this phenomena is interesting in its own right ([K1]), it naturally raises questions about the appropriateness of the ODCP approach to problems with uncertain, rather than white noise, parameters. Nevertheless, as an ad hoc scheme for improving H_2 robust performance, it has possibilities.

III. Modeling Assumptions for Methods 2 and 3

The general framework to be used in Method 2 and Method 3 is illustrated in the diagram in Figure 1a. Any linear interconnection of inputs, outputs, commands, perturbations, and a controller can be rearranged to match this diagram. G will be taken to be a linear, time-invariant, lumped system and be represented by a rational transfer function. The interconnection structure G can be partitioned so that the transfer function from d to e can be expressed as the linear fractional transformation

$$e = F_u(G,\Delta)\, d = \left[G_{22} + G_{21}\Delta(I - G_{11}\Delta)^{-1}\, G_{12}\right]\, d.$$

The external input d is an additive signal entering the system and is typically used to model disturbances, commands, and noise. In addition, the system model itself typically has uncertainty which can have a significant impact on system performance. This uncertainty is a consequence of unmodeled dynamics and parameter variations and is modeled as the perturbations Δ to the nominal interconnection structure G.

IV. Method 2: Additive Noise

Method 2 is quite simple. We assume that the interconnection structure, N, with its perturbation inputs and outputs, disturbances and errors, and measurements and control inputs is given. The perturbation inputs and the actual disturbances are lumped together into one large disturbance (and likewise for the errors) to give a new interconnection structure $\tilde{N}$ with no uncertainty, just additive noise (see figure 2). The H^2 solution for *this* problem is the Method 2 controller.

We call this the "additive noise" method, in contrast with the "multiplicative noise" method, because where perturbations enter the loop, new noise sources and errors are added to produce the modified problem. The rationale is simple - if a particular closed loop transfer function g_{cl} is not explicitly in the cost functional, it is possible that the minimization process will radically affect it, making $\|g_{cl}\|_\infty$ large, consequently giving poor robustness at that point in the loop. Conversely, the optimization should not do great violence to transfer functions in the cost (it certainly keeps a handle on their $\|\cdot\|_2$ norm).

The most serious drawback of this method is that the *structure* of the perturbation is totally ignored. Any two problems with the same interconnection structure $N(s)$, regardless of their respective uncertainty sets $\mathbf{\Delta}$, will be solved identically. This means that this method can be very conservative.

V. Preliminaries to Method 3

The third method ("mμs" - modified μ synthesis) is a beefed up version of method 2 and eliminates some of its deficiencies. The method is heavily structured singular value based, [D5] and [D1]; a brief review of definitions and notation for μ is in [D6] (this session).

The *Robust Performance* theorem of [D1] implies that

$$\inf_K \sup_{\omega\in\mathbf{R}} \mu_{\tilde{\mathbf{\Delta}}}(F_l(N,K)(j\omega)) \qquad (1)$$

where K is stabilizing, is an important control problem, since a smaller μ implies both increased stability robustness and better guaranteed H_∞ performance (the structure $\tilde{\mathbf{\Delta}}$ is an augmented block structure, diag $\left(\mathbf{\Delta}, \hat{\mathbf{\Delta}}_{r+1}\right)$ where $\hat{\mathbf{\Delta}}_{r+1}$ is a full block and is dimensioned compatibly with $N_{11_{22}}$). The complete solution to this is not known, but combining the DMD^{-1} upper bound and H_∞ techniques leads to

$$\inf_{K,D(\omega)\in\mathbf{D}} \|DF_l(N,K)D^{-1}\|_\infty \qquad (2)$$

which is currently the best approach to μ *synthesis* [D1]. This is why the upper bound is so important–synthesizing μ directly is beyond existing theory. The minimization in (2) is carried out by alternately minimizing over $K(s)$ and over D. While each separate minimization achieves the global minimum (with the other variable fixed) the problem is not convex jointly in both variables, and may have local minima which are not global minima, even in the constant case. This *is* a problem, and research to guarantee convergence to the global minimum continues (e.g. [S3]). A detailed design example using μ synthesis is in [D3]. Other synthesis methods involving μ and nonlinear, time varying controllers and plants are also being developed. [D6]

VII. Method 3: mμs

The basic idea of mμs is to do μ synthesis with the $\|\cdot\|_2$ norm replacing $\|\cdot\|_\infty$. Two main points will emerge: in a limited sense, the method achieves robust H_2 performance; computationally, it is relatively simple, involving a sequence of H_2 solutions. We begin with a outline of the method.

Assume that the interconnection structure N is given, and for simplicity, all the uncertainty blocks are full blocks ($m = 0$). Suppose that there are r perturbations and one additional block for performance, so

$$\tilde{\mathbf{\Delta}} = \left\{ \text{block} \quad \text{diag}[\Delta_1, \Delta_2, \ldots, \Delta_r, \hat{\Delta}_{r+1}] : \Delta_i \in \mathbf{C}_{j_i \times j_i} \right\}.$$

If $\hat{d}_1(s), \ldots \hat{d}_r(s)$ are stable, minimum phase, rational functions, let $D(s) = diag(\hat{d}_1(s) I_{j_1}, \ldots, \hat{d}_r(s) I_{j_r}, I_{j_{r+1}}$ and define $\tilde{N}_D$ as in figure 3.

The Method 3 iteration is as follows:

1. Design H_2 optimal controller, K, for $\tilde{N}_D$.
2. For each $\omega \in \mathbf{R}$, find $D_\omega \in \mathbf{D}$ to minimize

   $$\|D_\omega F_l(N,K)(j\omega) D_\omega^{-1}\|_F.$$

   Scale D_ω so $d_{r+1}(\omega) = 1$.
3. Fit each element $d_i(\omega)$ with a stable, minimum phase $\hat{d}_i(s)$, i.e. $\mid \hat{d}_i(j\omega) \mid \approx d_i(\omega)$ for all $\omega \in \mathbf{R}$.
4. This defines a new $D(s)$, hence a new $\tilde{N}_D$. Go to 1 and repeat.

Note that existing μ synthesis is the same iteration, with two changes: in step (1) the controller would be H_∞ optimal for $\tilde{N}_D$, and in step (2) the maximum singular value is minimized rather than $\|\cdot\|_F$.

VI. Standing Assumptions for Method 3

There are two assumptions we will use below in "justifying" this scheme. Both "assumptions" are, strictly speaking, false. Thus they must be assumptions rather than theorems. We are willing to accept them because counterexamples are difficult to construct and would be extremely unlikely to occur in any typical examples, especially if practically motivated. The first assumption is that, for our problems, H_2 optimal controllers perform well in an H_∞ sense. That is, define

$$K_2 = \underset{K}{\text{argmin}} \|F_l(P,K)\|_2$$

and

$$\beta = \inf_K \|F_l(P,K)\|_\infty.$$

Then we assume the ratio $\frac{\|F_l(P,K_2)\|_\infty}{\beta}$ is not "too large". In general, this is false. In fact, the ratio can be made arbitrarily large as is shown in the appendix. However, trying reasonable examples will usually lead one to believe the assumption. Our experience has been that realistic, physically motivated problems lead to ratios less than 2.

The second assumption has to do with an approximation to the upper bound for μ. Let $M \in \mathbf{C}^{n \times n}$ and let $\|M\|_F^2 := \sum_{i,j=1}^{n} \mid m_{ij} \mid^2$. Recall the definition of $\mathbf{D}$ in the review of μ and define

$$D_2 := \underset{D \in \mathbf{D}}{\text{argmin}} \|DMD^{-1}\|_F \tag{1}$$

$$\alpha := \inf_{D \in \mathbf{D}} \bar{\sigma}\left(DMD^{-1}\right). \tag{2}$$

Then it is easily shown that

$$\frac{\bar{\sigma}\left(D_2 M D_2^{-1}\right)}{\alpha} \leq \sqrt{\text{rank}(M)}. \tag{3}$$

Since $\bar{\sigma}\left(D_2 M D_2^{-1}\right)$ is very easy to compute, it provides a cheap alternative to α. Unfortunately, (3) gives us little guarantee of the quality of this upper bound, but as in the above H_2/H_∞ situation, the bound that occurs in practice seems much better than the theoretical worst case. For the case of n 1×1 blocks, a tighter bound seems possible and extensive computational experience suggests that a practically reliable bound is approximately 1.2, independent of the number of blocks. The argument to derive (3) does not take advantage of the structure that is present in the problem and can probably be refined. Therefore, for the justification of method 3, we will assume that the ratio in (3) is very close to 1. Given these assumptions, we be jammin'.

Justification of Method 3

Consider both the above iteration and the μ synthesis iteration. Let K_2, D_2 and K_∞, D_∞ be the controller and D scales that the each of the methods stop at. If the two assumptions in the last section are true, then the ratio

$$\frac{\|D_2 F_l(N,K_2) D_2^{-1}\|_\infty}{\|D_\infty F_l(N,K_\infty) D_\infty^{-1}\|_\infty}$$

is not a lot larger than 1. Consequently, if the μ synthesis controller has reasonable stability robustness properties, then so will the method 3 controller.

The important question is how Method 3 performs in the $\|\cdot\|_2$ norm. Conceptually, the iteration is trying to solve

$$\inf_K \inf_{\substack{D(\cdot) \\ D_\omega \in \mathbf{D}}} \|DF_l(N,K)D^{-1}\|_2 \tag{4}$$

By the second assumption above, this is "nearly" equivalent to

$$\inf_K \int_{-\infty}^{\infty} \left(\inf_{D_\omega \in \mathbf{D}} \bar{\sigma}\left(D_\omega F_l(N,K)(j,\omega) D_\omega^{-1}\right)\right)^2 d\omega \tag{5}$$

which is approximately

$$\inf_K \|\mu(F_l(N,K))\|_2 \tag{6}.$$

The situation is illustrated in figure 4. There are three curves – a flat μ plot for the μ synthesis controller, and a μ plot and nominal disturbance to error plot from the mμs controller. Consider a perturbation $\Delta \in \mathbf{\Delta}$ with $\bar{\sigma}(\Delta) < \frac{1}{\beta}$. The *robust performance* theorem implies that both closed loops are robustly stable to this perturbation, and furthermore, each maintain an H_∞ disturbance to error performance level less than β. The $\|\cdot\|_2$ norm performance of the mμs controller can also be bounded: a simple application of μ implies that for each ω, $\bar{\sigma}\left(F_u\left(F_l(N,K_2)(j\omega),\Delta\right)\right) \leq \mu(F_l(N,K_2)(j\omega)$. That is, in the presence of Δ, the disturbance to error transfer function (nominally the lower curve) can at most degrade up to the $\mu(F_l, K_2)$ curve. Hence, for all $\Delta \in \mathbf{\Delta}$ with $\bar{\sigma}(\Delta) < \frac{1}{\beta}$

$$\int_{-\infty}^{\infty} [\bar{\sigma}\left(F_u(F_l(N,K_2),\Delta)\right)]_2 \, d\omega \leq \|\mu(F_l(N,K_2))\|_2 \quad (7)$$

Modulo the difference between $\|\cdot\|_F$ and $\bar{\sigma}(\cdot)$, this gives a bound on worst case 2 norm performance, and the bound is the the quantity being minimized.

There are two potential problems. First, if in minimizing (4), the peak value of $\mu(F_l(N,K_2)(j\omega))$ does get very large compared to the μ synthesis peak, then all we are guaranteeing is better robust H_2 performance for a smaller class of perturbations. Clearly, this is not desirable, and highlights the importance of the two assumptions. Second, and apart from the above assumptions, the bound in (7) is conservative. It would only be achieved in a worst case by an acausal perturbation. These are not considered, hence controllers that do not minimize (6) may perform as well or better than controllers that do.

VIII. Example

This section briefly describes the sample application of the three methods. The problem is taken from [D2], and shows that *LQG* solutions can have arbitrarily small gain margins. The matrices are

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \qquad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \qquad J = \sqrt{\sigma}\begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

$$C = [1 \quad 0] \qquad L = \sqrt{\rho}[1 \quad 1] \qquad R = M = 1$$

The *LQG* solution has arbitrarily small gain margins in both directions as these weighting parameters σ and ρ get large. At $\rho = \sigma = 60.0$, the gain margin is about 1 percent and the transfer function from $\binom{\eta}{\xi}$ to e has a $\|\cdot\|_2$ norm of 90. The goal is to increase this margin while maintaining low H^2 cost.

Method 1 proceeds by treating the (2,1) entry in the B matrix as a stochastic process, with mean equal to 1 and variance b [B1]. b is a free parameter that will be adjusted for different designs of varying robustness. Hence in the ODCP framework $p = 1$, $B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$, and $B_1 = \begin{bmatrix} 0 \\ b \end{bmatrix}$. For "additive noise" and mμs, the uncertainty is modeled as a multiplicative perturbation as shown in the interconnection structure, Fig. 5. The scalar constant α is used to vary the robustness among designs.

Results:

Method 1 was done at four conditions, $b = 0.05, 0.10, 0.15$, and 0.20. Each controller is 2nd order, as is inherent in the procedure. Method 2 also produces controllers with 2 states. 10 controllers with α taking on values from 1 to 20 were calculated. Method 3 was carried out for values of α between 1 and 30. In each design, the iteration was done twice – that is, first an H^2 minimization with $D(s) = I$, solve Frobenious minimization to get new D, absorb this into N, and design the controller. A second order $\hat{d}(s)$ was used to fit $d(\omega)$. Consequently, the interconnection $\tilde{N}_D$ was 6'th order, as were the controllers. Truncated balanced realizations reduced each controller to 4 states.

The first plot shows nominal performance .vs. gain margin. Gain margin is defined as the largest γ such that for all $\delta \in (-\gamma, \gamma)$ the closed loop system with B matrix $\begin{bmatrix} 0 \\ 1+\delta \end{bmatrix}$ is stable. Nominal performance is the H^2 norm of the disturbance to error transfer function with B matrix equal to $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$. This plot does not show the possible degradation of performance as δ changes, and hence is only a part of the story, however it does qualitatively show all methods doing what was expected – improving robustness at the expense of nominal cost.

A word about the choice of *real* gain margin versus a *complex* gain margin: Method 1 is set up to give robustness to real perturbations, but it is probably just as likely to give good complex margins; The mμs method is definitely geared toward complex margins, as that is what the perturbation set $\mathbf{\Delta}$ represents; in this example the choice didn't matter as the above plot for complex gain margins looked nearly identical.

Next are the robust performance plots. Each curve corresponds to a different controller and it plots the $\|\cdot\|_2$ norm of the disturbance to error transfer function .vs. the perturbation δ, as δ varies in the gain margin interval. The solid line are mμs controllers and the dashed lines are from Method 1. The "additive noise" plots are not included as their performance is quite poor compared to these two. The plots are fairly self explanatory. Method 1 controllers do better at the left edge of the gain margin interval, while the mμs controllers are superior elsewhere.

Conclusions

This paper explored the issue of robust H_2 performance in the presence of H_∞ norm bounded uncertainties. Two very different underlying theories, a stochastic approach and a deterministic approach based on a generalization of singular values, give rise to three ad hoc schemes that address the problem. We outlined the methods and applied them to a well known LQG gain margin counterexample. All the methods performed as we hoped, at least qualitatively, and given the theoretical differences between them, remarkably similar. It is difficult to extrapolate these results to more complicated problems (several uncertainties), so more examples need to be tried. In addition to examples, though, it is obvious that basic research in these "mixed norm" problems is necessary.

Acknowledgements

This work has been supported by Honeywell IRAD Funding, ONR Research Grant N00014-82-C-0157, AFOSR Research Grant F49620-86-C-0001, and the National Science Foundation, a NSF graduate fellowship and Grant ECS-8451519.

References

[B1] D. Bernstein, S. Greeley, *IEEE Trans. Auto. Cont.*, July 1985.

[D1] J. C. Doyle, *IEEE CDC Proceedings*, Ft. Lauderdale FL., 1985.

[D2] J.C. Doyle, *IEEE Trans. Auto. Cont.*, Aug. 1978.

[D3] J.C. Doyle, K. Lenz, A. Packard, *IEEE CDC Proceedings*, Athens, Greece, 1986.

[D4] J.C. Doyle, G. Stein, *IEEE Trans. Auto. Cont.*, Feb., 1981.

[D5] J.C. Doyle, *IEE Proceedings*, part D, Nov., 1982.

[D6] J.C. Doyle, *ACC Proceedings*, Minneapolis, MN., 1987.

[F1] B.A. Francis, J.C. Doyle, *SIAM J. Cont.*, to appear July, 1987.

[J1] E. Jonckheere, M. Safonov, L. Silverman, *IEEE CDC Proceedings*, 1981.

[K1] D. Kleinman, *IEEE Trans. Auto. Cont.*, pp. 673-677, 1969.

[S1] M. Safonov, M. Athans, *IEEE Trans. Auto. Cont.*, April, 1977.

[S2] M. Safonov, A. Laub, G. Hartmann, *IEEE Trans. Auto. Cont.*, Feb., 1981.

[S3] M. Safonov, V. Le, submitted to *Systems and Controls Letters*, Feb., 1987.

Appendix

For a bunch of counterexamples, consider stabilizing a unstable plant robustly to additive plant uncertainty. In this case, $\|\frac{c}{1+pc}\|_\infty$ should be minimized. We can construct plants for which the optimal $\|\cdot\|_2$ controller does arbitrarily bad in a H_∞ sense. The procedure is as follows:

1. Let n be any integer that is a bad enough ratio to be convincing. Choose p so that the hankel norm of $g(s) = \sum_{k=1}^{n} \frac{p^{2k}}{s-p^{2k}}$ is less than 1. This can always be done [J1]. Note that $\|g\|_\infty = n$.
2. Define polynomials a and b by $\frac{a(s)}{b(-s)} := g(s)$.
3. The pair $\left(d_g := \frac{b(-s)}{b(s)}, n_g := \frac{a(s)}{b(s)}\right)$ is a coprime factorization (over $\mathbf{R}H_\infty$) of $g(s)$. Let (v, u) be a coprime factorization of a stabilizing compensator for g.
4. Define a plant $p(s) := \frac{u(s)}{d_g(s)}$. The set of all FDLTI compensators which stabilize $p(s)$ is given by

$$\left\{ c(s) = \frac{n_c(s)}{d_c(s)} : q \in \mathbf{R}H_\infty, n_c = n_g - q d_g, d_c = v + qu \right\}.$$

Then for $\alpha = 2$ or ∞

$$\inf_{c,stabilizing} \|\frac{c}{1+pc}\|_\alpha = \inf_{q,stable} \|g(s) - q(s)\|_\alpha.$$

The $\|\cdot\|_2$ is minimized by the unique choice of $q(s) = 0$. This gives the H_∞ cost of the H_2 controller to be $\|g\|_\infty = n$. The optimal $\|\cdot\|_\infty$ performance level is the hankel norm of g. By construction, this is less than 1, hence the ratio is greater than n as promised.

The plants resulting from this construction are rather pathological looking. For instance, setting $n = 4$ and $p = 4$ results in

$$P(s) = 371601.5\frac{(s-20.343)(s-393.83)(s-8809.5)}{(s-16)(s-256)(s-4096)(s-65536)}$$

but a ratio of only 4.244. It seems likely that extremely large ratios are only possible with even wilder plants. There are probably simple conditions that can be placed on general interconnection structures so that H_2 controllers have some guaranteed level of H_∞ performance. More research is needed here, as a result like this could reduce the computations required in some H_∞ problems.

Figure 1a General Framework

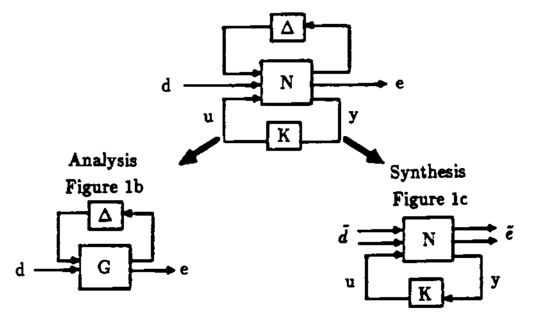

$F_u(G,\Delta) =$

$G_{22} + G_{21}\Delta(I - G_{11}\Delta)^{-1}G_{12}$

$e = F_u(G,\Delta)d$

$F_l(N,K) =$

$N_{11} + N_{12}K(I - N_{22}K)^{-1}N_{21}$

$\tilde{e} = F_l(N,K)\tilde{d}$

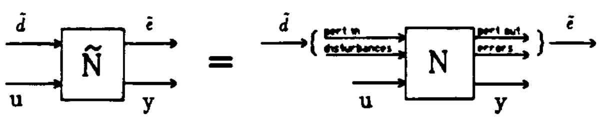

Figure 2

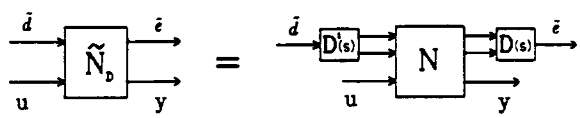

Figure 3

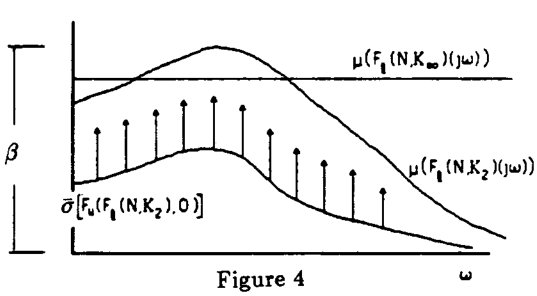

Figure 4

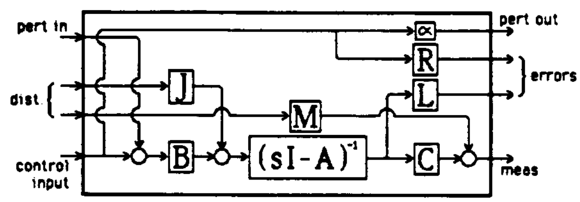

Figure 5

Gain Margin .vs. Nominal Cost

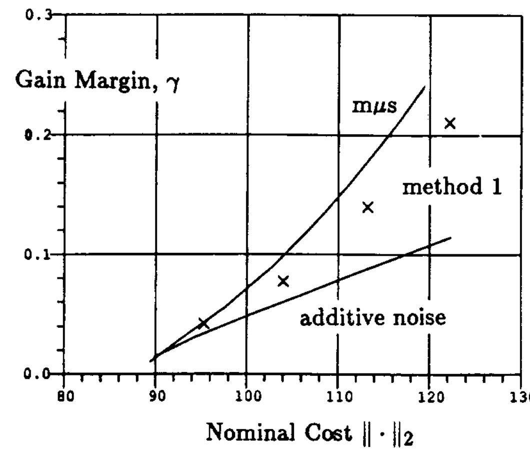

Robust Performance