OPTIMAL LOOP-SHAPING IN THE CEPSTRAL DOMAIN FOR ACTIVE NOISE CANCELING HEADSETS

Jean Laroche
jeanl@atc.creative.com

Creative Advanced Technology Center
1500 Green Hills Rd., Scotts Valley, CA 95067

ABSTRACT

Simple active noise-canceling systems (such as noise canceling headphones) use a feedback mechanism whereby the signal recorded by microphone placed near the loudspeaker is phase-inverted and sent back through the loudspeaker via a feedback filter. Such systems can be implemented using analog circuitry, are fairly inexpensive, and achieve reasonably good performance at low frequencies (below 500Hz). The feedback filter is designed to achieve a high gain at low-frequencies (for best noise reduction) while maintaining closed-loop stability under various conditions. To design the feedback filter, fairly ad-hoc techniques were traditionally used under the broad denomination of "loop-shaping" until the advent in the 80's of optimal H_∞ algorithms. The design technique outlined in this paper shares the intuitive approach of traditional loop-shaping methods, but removes most of the guess-work while producing optimal filters that guarantee closed-loop stability. The main idea is to represent the filter in the cepstral domain, where gain and phase constraints take a convenient linear form, and use linear programming to design the optimal feedback filter.

1. INTRODUCTION

The past few years have seen an increasing interest in active noise attenuation, particularly as it applies to noise-canceling headphones. Many manufacturers now offer low-to-medium price headsets equipped with active noise attenuation and a few papers have appeared that give fairly detailed descriptions of such setups [1, 2, 3, 4]. Because of cost and battery-life issues, most if not all noise canceling headsets use analog circuitry to implement a negative feedback algorithm. The headset loudspeaker is used to send the normal audio signal as well as the "secondary path" feedback signal responsible for the noise attenuation. A microphone is placed near the loudspeaker and its output is filtered by a feedback filter $C(z)$, phase inverted and sent back to the loudspeaker, forming a feedback path as depicted in Fig. (1). The amount of noise attenuation achieved by the system is given by the magnitude of the so-called "sensitivity function" [5]:

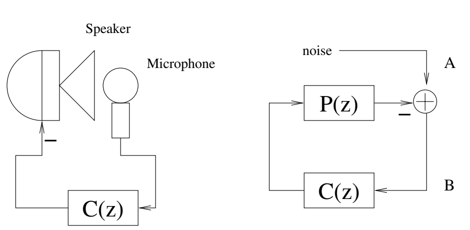

Fig. 1. Typical negative feedback ANR setup and equivalent signal flow. $P(z)$ represents the electro-acoustic path and $C(z)$ is the feedback filter.

$$S(z) = \frac{1}{1 + C(z)P(z)} \tag{1}$$

in which $P(z)$ is the "plant" (in control theory parlance) i.e., the electro-acoustic path between the input of the loudspeaker and the output of the microphone, and $C(z)$ is the feedback filter (the "controller" in the control theory literature). $S(z)$, also called the closed-loop transfer function is the transfer function between points A and B. Eq. (1) shows that good noise attenuation is achieved when $C(z)P(z)$ has a large gain, but because $S(z)$ is an IIR filter, its stability isn't guaranteed. As we will see, too large a gain for $C(z)P(z)$ usually yields an unstable $S(z)$ and the design of the feedback filter $C(z)$ for a given plant $P(z)$ involves a tradeoff between performance on the one hand, and stability and robustness to variations of $P(z)$ on the other.

Historically, feedback filters were designed using an iterative ad-hoc procedure called "synthetic loop-shaping" [6, 7] consisting of progressively adding lead and lag compensators to the feedback filter, making sure closed-loop stability was preserved. The Nyquist criterion [7] was typically used to check closed-loop stability, and the resulting feedback filter was, in most case, suboptimal. Some amount of robustness to "plant variation" (variations in $P(z)$ resulting from changes to the physical setup - for example taking the headset off the head) was obtained by using phase and gain margins in the design [5]. During the 80's, H_∞ robust control techniques were introduced [8, 9] which provided a solid basis for the design of optimal or near-optimal feedback filters, and robustness was characterized in a more accurate manner that was previously possible [5]. The two main techniques (the model matching [8] and state-space [9] formulations) however, have several drawbacks. They are complex, require a model (either a rational transfer function or a state-space model) for the plant $P(z)$, depend on the careful selection of two rational weighing transfer functions to achieve good performance and robustness and can produce very high-order filters [1]. Nonetheless, H_∞ control theory has been put to good use in the context of noise-canceling headsets (see for example [4]). The alternative filter design technique presented in [1] has several advantages: the filter order is selected in advance and there is no need for a model of the plant (the technique works from direct experimental measurements of the headset). Furthermore, robustness is easily incorporated in the algorithm, and the technique allows a direct control of the performance tradeoff between good noise attenuation in low frequencies and noise amplification at higher frequencies (the usual "waterbed" effect [5] inherent in non minimum-phase plants). Unfortunately, because the solution is obtained via a non-linear constrained optimization technique, the algorithm only finds local minima and there is a possibility that a filter exists that achieves better performance. The design technique presented in this paper fits the broad category of "loop-shaping" techniques in the sense that the algorithm focuses on the characteristics of the open-loop filter $C(z)P(z)$ (as opposed to the sensitivity function $S(z)$). The algorithm yields a feedback filter $C(z)$ that achieves the best noise rejection (in the H_∞ sense) in a specified frequency area, while maintaining closed-loop stability. As in [1], the algorithm does not require a plant model, and uses direct frequency-domain measurements of the plant. Robustness to plant variations is achieved either by use of traditional gain and phase margins, or by using several sets of measurements that correspond to various positions of the headset, or by a combination of the two. The key to the design is to express the feedback filter $C(z)$ in the cepstral domain, where magnitude and phase constraints can be expressed linearly in terms of the unknown cepstral coefficients. The design then translates into a convex optimization problem which can be solved by fast and well known algorithms, producing a unique optimal solution. For practical implementations, the resulting filter must be transformed in a subsequent stage into an analog filter ($B(s)/A(s)$), which can be done using straightforward filter transformation techniques [10]. The second part of this paper briefly summarizes the theory of negative feedback systems. The third part presents the optimal design technique, and the fourth part concludes with a design example.

2. NEGATIVE FEEDBACK ANR

Fig. (1) above presents a typical setup for negative feedback Active Noise Reduction. It is easy to see that the transfer function between the incoming noise (point A) and the attenuated noise (point B) is given by Eq. (1) above. Because $S(z)$ is an IIR filter, stability is an issue. In general, it is not easy nor simple to derive a condition on the open-loop filter $P(z)C(z)$ to guarantee the stability of $S(z)$. However, the Nyquist stability theorem provides an answer. It can be shown [11] that the closed-loop filter $S(z)$ is stable if the Bode plot of the open-loop filter $P(z)C(z)$ does not encircle the $z = -1$ point in the Z-plane (the Bode plot of a filter $L(z)$ is obtained by plotting the complex $L(z)$ for $z = \exp(j\omega)$ for $0 \leq \omega < 2\pi$). A simple and traditional [7] design strategy derived from the Nyquist criterion consists of making sure that the magnitude of the open-loop filter $L(z)$ is lower than 1 when its phase crosses $-\pi$. It is easy to see that if this is true, the $z = -1$ point is never encircled by the Bode plot. Note however that this design strategy is not a necessary condition for closed-loop stability.

3. OPTIMAL DESIGN TECHNIQUE

3.1. Design template

The design strategy discussed above can be expressed in a way that is reminiscent of the Parks McClellan FIR filter design algorithm. Let us assume for simplicity that $P(z) = 1$, we will see later that accounting for non-unity $P(z)$ is easy. One can create a template characteristics that represents the expression to be maximized and the constraints to be satisfied as a function of frequency, as shown in Fig. (2). In the H_∞ framework, we seek to maximize the minimum amplitude response $A_{\min}$ (to achieve maximum noise reduction) over a designer-specified low frequency region (up to frequency F_0) while the phase is constrained to remain above $-\pi$. We constrain the amplitude to reach a negative dB level at frequency F_1, at which point the phase is allowed to dip below $-\pi$, which ensures the stability of the closed-loop filter. As with the Remez exchange algorithm, there can be "don't care" regions where we allow the filter response to be essentially free. Given our design strategy, the question is: is there a global optimum, and what algorithm can be used to determine it?

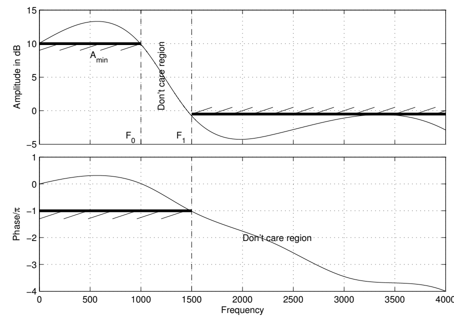

Fig. 2. Design template for the open-loop filter $C(z)$. The goal is to maximize the minimum open-loop gain $A_{\min}$ while satisfying the magnitude and phase constraints.

3.2. Using the cepstral domain

Because the design objective and constraints affect both the magnitude and phase of $C(z)$, it is natural to express $C(z)$ in terms of its (complex) cepstral coefficients. Given an impulse response $c(n)$, the cepstral coefficients of the filter are obtained by computing the Fourier Transform of $c(n)$, taking the complex logarithm of the result, and then the inverse Fourier Transform [10]. In other words, the cepstral coefficients $\hat{c}(n)$ are the inverse Fourier transform of the logarithm of the filter's frequency response:

$$\sum_{n=-\infty}^{\infty} \hat{c}(n) e^{-j\omega n} = \log(C(\omega)) = \log|C(\omega)| + j\angle C(\omega) \quad (2)$$

where $C(\omega)$ represents the complex-valued frequency response of the filter. Eq. (2) shows that both the filter's log-magnitude frequency response and the filter's phase response are linear functions of the filter's cepstral coefficients. A practical constraint is that the filter $c(n)$ should be causal. If we further constrain $c(n)$ to be minimal-phase, then this is equivalent to constraining the cepstral coefficients to be right-sided [10], and the sum above to be restricted to positive values of n. We will say more about that later.

3.3. Linear programming

Our design task now consists of maximizing a linear function of the unknown parameters $\hat{c}(n)$ (the filter's low-frequency gain), subject to constraints that are also linear in the unknowns (constraints on the phases and the magnitudes above F_1). This is a standard convex optimization problem for which there is either no solution, a unique global maximum or an infinity of maxima, and the well-known linear programming algorithm [12] can be used to compute the solution(s). One form of the linear programming (or "simplex") algorithm maximizes the expression

$$\mathbf{a^t}.\hat{\mathbf{c}} \quad \text{subject to the constraints} \quad \mathbf{B}.\hat{\mathbf{c}} \leq \mathbf{c} \quad (3)$$

where $\hat{\mathbf{c}}$ is a vector containing the N unknown parameters $\hat{c}(n)$, $\mathbf{a}$ is a $N \times 1$ vector, $\mathbf{B}$ is a $L \times N$ matrix expressing L linear constraints and c is a vector used in the constraints. Here, we use the short-hand notation that $\mathbf{x} < \mathbf{y}$ if and only if $x_i < y_i$ for all the components of vectors $\mathbf{x}$ and $\mathbf{y}$. The linear programming optimization technique fits our design problem very well: We wish to maximize level $A_{\min}$ on Fig. (2) under the constraint that the filter's log-magnitude response should be above $A_{\min}$ in the frequency range $[0 \;\; F_0]$, should be below 1 in the frequency range $[F_1 \;\; S_r/2]$, (S_r is the sampling frequency) and the phase response should be above $-\pi$ in the frequency range $[0 \;\; F_0]$. To express this in matrix form, we need to discretize the frequency axis: we define M frequency points $\omega_k = k\pi/M$ between 0 and π (although the sampling needs not be uniform), and the inverse Fourier transforms used to calculate the magnitudes and phases from the unknown cepstral coefficients are expressed in a matrix form as

$$\begin{aligned} |C(\omega_k)| &= \sum_{n=0}^{N-1} \hat{c}(n) \cos \omega_k n = \sum_{n=0}^{N-1} \hat{c}(n) C_n^k \\ \angle C(\omega_k) &= \sum_{n=0}^{N-1} -\hat{c}(n) sin \omega_k n = \sum_{n=0}^{N-1} \hat{c}(n) S_n^k \end{aligned} \tag{4}$$

The linear programming setup of Eq. (3) is used, with the following parameters:

$$\hat{\mathbf{c}} = (A_{\min}, \hat{c}(0), \hat{c}(1), ..., \hat{c}(N-1))^t \tag{5}$$

$$\mathbf{a} = (1, 0, 0, ..., 0)^t \tag{6}$$

$$\mathbf{B} = \left[\mathbf{B}^{\mathbf{pass}} | \mathbf{B}^{\mathbf{stop}} | \mathbf{B}^{\mathbf{phase}}\right]^{\mathbf{t}} \tag{7}$$

$$\mathbf{c} = \left[\mathbf{c}^{\mathbf{pass}} | \mathbf{c}^{\mathbf{stop}} | \mathbf{c}^{\mathbf{phase}}\right]^{\mathbf{t}} \tag{8}$$

where B^{pass}, B^{stop} and B^{phase} are the constraints related to the magnitudes in the passband, in the stopband, and the phases in the passband. Specifically:

$$\begin{aligned} \mathbf{B_i^{pass}} &= \left(1, -C_0^i, -C_1^i, ..., -C_{N-1}^i\right) \quad \text{and} \\ \mathbf{c^{pass}} &= (0, 0, ..., 0)^t \quad \text{for} \quad 0 \le i < L_0 \end{aligned} \tag{9}$$

$$\begin{aligned} \mathbf{B_i^{stop}} &= \left(0, C_{L_1}^i, C_{L_1+1}^i, ..., C_{N-1}^i\right)^t \quad \text{and} \\ \mathbf{c^{stop}} &= (g, g, ..., g)^t \quad \text{for} \quad 0 \le i < M - L_1 \end{aligned} \tag{10}$$

$$\begin{aligned} \mathbf{B_i^{phase}} &= \left(0, -S_0^i, -S_1^i, ..., -S_{N-1}^i\right)^t \quad \text{and} \\ \mathbf{c^{phase}} &= (\alpha, \alpha, ..., \alpha)^t \quad \text{for} \quad 0 \le i < L_1 \end{aligned} \tag{11}$$

in which L_0 and L_1 are the indexes corresponding to frequencies F_0 and F_1 in Fig. (2). Note that the added dummy unknown $A_{\min}$ in $\hat{\mathbf{c}}$, allows us to maximize the minimum gain below F_0. This a standard technique for min-max optimization. It is easy to verify for example that $\mathbf{B_i^{pass}}.\hat{\mathbf{c}} < \mathbf{c^{pass}}$ ensures that the filter's magnitude frequency response will be above $A_{\min}$ in the passband below frequency F_0. In the equations above, g is the gain margin and should be such that $g < 1$, and α is the phase margin and should be such that $\alpha < \pi$. Adding gain and phase margins is a traditional way to ensure some amount of robustness. We will see below that there is another way.
Note that the design template isn't restricted to piecewise constant constraints. Any arbitrary frequency-dependent function can be used both in the area where the magnitude is to be maximized and in the constraints. The linear programming algorithm (for example, the linprog function in matlab) is then run using the above setup and yields the usually unique set of cepstral coefficients $\hat{c}(n)$ that maximizes the passband minimum gain $A_{\min}$ while satisfying the required constraints on the magnitudes and phases. Once the cepstral coefficients $\hat{c}(n)$ are obtained, the impulse response $c(n)$ is calculated by an inverse cepstral transform, for example using the recursion [10]:

$$c(n) = \begin{cases} 0 & \text{for} \quad n < 0 \\ \exp(\hat{c}(0)) & \text{for} \quad n = 0 \\ \frac{1}{n} \sum_{k=0}^{N-1} k\hat{c}(k)c(n-k) & \text{for} \quad n > 0 \end{cases}$$

Note that even though $\hat{c}(n)$ has a finite length, the impulse response $c(n)$ is IIR.

3.4. Adding the contribution of the plant $P(z)$

In the discussion above, we assumed that $P(z) = 1$ which of course is never true in practice (if only because of the acoustic delay between the loudspeaker and the microphone). Because we design $C(z)$ in the cepstral domain, the cepstrum of $C(z)P(z)$ is simply the sum of the two cepstrums. We can simply measure $P(z)$, then calculate its magnitude and phase spectra and modify our constraints in Eqs. (11) to account for the contribution of $P(\omega)$. For example, $\mathbf{c^{phase}}$ would now read:

$$\mathbf{c^{phase}} = (\alpha + \angle P(\omega_0), \alpha + \angle P(\omega_1), ..., \alpha + \angle P(\omega_{L_1-1}))^t$$

In addition, robustness to plant variations can be achieved by including additional measurements of the plant $P(z)$, for example when the headset is off the user's head. At each frequency ω_i, the smallest measured $\angle P(\omega_i)$ would be used in the equation above, to guarantee stability in the worst scenario. The same holds for magnitudes: the vector $\mathbf{c^{stop}}$ would need to be modified to include the largest (worst case) gain contributed by $P(\omega)$. Including multiple measurements into the constraints provides a simple way to account for expected real-world plant variations. This stands in contrast with H_∞ robust control where robustness is guaranteed via the use of a rational weighing transfer function W_2 that must be carefully adjusted [5].
As mentioned above, we added the constraint that filter $c(n)$ be minimum phase as a way to ensure its causality, and one can question whether this artificial constraint hurts the optimality of the solution. However, it is easy to see that if a causal filters $c(n)$ satisfies the design constraints and achieves a minimum passband gain of $A_{\min}$, its equal-amplitude minimum-phase equivalent will also satisfy the constraints (because its phase response will be larger or equal to that of $c(n)$) and achieve the same gain. If the solution of the problem is unique (which is normally the case), we can conclude that the optimal causal solution has to be minimum-phase.

4. PRACTICAL CONSIDERATIONS, RESULTS AND DISCUSSION

Contrary to what happens in standard filter design tasks, increasing the number N of cepstral coefficients $\hat{c}(n)$ does not necessarily yield better results. This is because the phase constraint in Fig. (2) prevents the filter from exhibiting too sharp a transition between the passband and the stopband, regardless of its length. When the number N of cepstral coefficients is increased beyond a certain limit, the results remain the same.
We now present a practical design example for a case where $P(z)$ is a $125\mu s$ delay corresponding to a distance of about 4.1cm between the microphone and the loudspeaker. The sample rate is $16kHz$. The constraints are shown as dashed lines on Fig. (3), and the resulting filter is shown as a solid line. The constraints were slightly offset vertically in the figures, for the sake of clarity. The design concentrates

the attenuation between 0 and 400Hz with a maximum at 140Hz. F_1 is set to 1500Hz, a phase margin α of about $\pi - 0.12\pi$ was used, as well as a gain margin g of 15dB. We used $N = 256$ cepstral coefficients and $M = 300$ frequency sampling points. As is usu-

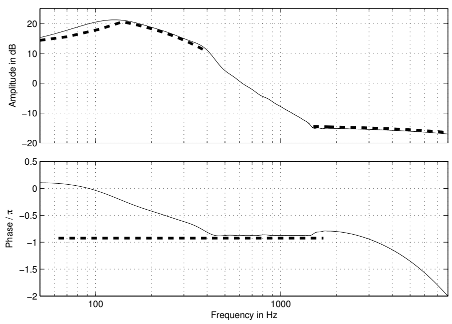

Fig. 3. Design template (dotted line) and resulting optimal magnitude and phase responses (solid lines).

ally the case in linear convex optimization, the optimum meets some of the constraints: the filter's responses touch the design boundaries at some frequencies. Finally, Fig. (4) presents a 6th-order (3 pairs of complex conjugate poles and 3 pairs of complex conjugate zeros) rational approximation of the filter of Fig. (3) obtained by the Steiglitz-McBride algorithm (the stmcb() function in Matlab. As ex-

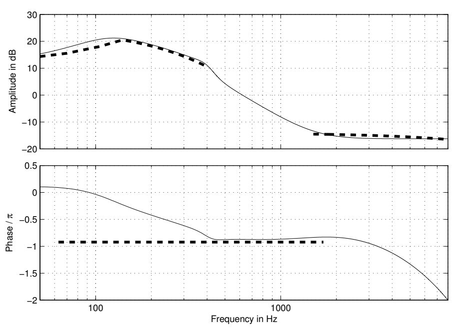

Fig. 4. 6th-order rational approximation of the design of Fig. (3).

pected the rational transfer function is no longer guaranteed to satisfy the constraints, and indeed, its magnitude response slightly cuts the constraint around 1.5kHz, but because of the gain margin, the closed-loop filter is still stable. In all other respects however, the filter performance is practically identical to its FIR counterpart. This confirms that the design technique presented here only provides an approximate sub-optimal solution to the problem of fixed-order controller design. Note however that provided the order is large enough, the solution is very close to the optimal filter.

The constraints described in this paper correspond to H_∞ optimization (in the sense that we maximize the smallest loop gain in the desired frequency range), but could easily be modified to maximize the sum of the loop log-magnitudes in the frequency range instead, because the domain of acceptable filters (i.e., filters that yield stable sensitivity functions $S(z)$) is convex. Convexity guarantees that the maximum is global (if it exists) and can be obtained by well-known linear programming algorithms. Finally it is worth mentioning that the choice of F_1 (the frequency where the phase is allowed to cross $-\pi$) is somewhat arbitrary, and different values will yield different filters with better or worse performance in terms of the achieved value of $A_{\min}$. With respect to the ultimate goal of maximizing noise attenuation over a frequency range while maintaining closed-loop stability, having to choose F_1 is an unnecessary constraint. Because the design procedure is fairly rapid, it is not unreasonable to imagine testing a range of values for F_1, then selecting the value that yields the best filter (the largest value for $A_{\min}$).

5. CONCLUSION

The main advantages of the technique presented in this paper are its simplicity and its ability to yield truly optimal filters in the min-max (H_∞) sense. As in [1], direct measurements of the headset can be used without additional modeling stages, and robustness to plant variation is accounted for in a very intuitive and realistic manner. The design algorithm can incorporate the traditional loop shaping concepts of gain and phase margins and provides the designer with intuitive parameters such as passband gain and bandwidth to achieve the desired noise reduction performance, getting rid of most of the guess-work. One drawback of the algorithm is that it does not provide an optimal solution for fixed-order feedback filters. Finding fixed-order H_∞ optimal controllers is still an unsolved problem [13].

6. REFERENCES

[1] S.Y. Yu and J.S. Hu, "Controller design for active noise cancellation headphones using experimental raw data," *IEEE Trans. Mechatronics*, vol. 6, no. 4, pp. 483–490, Dec 2001.

[2] Y. Song, Y. Gong, and S.M. Kuo, "A robust hybrid feedback active noise cancellation headset," *IEEE Trans. Speech and Audio Processing*, vol. 13, no. 4, pp. 607–617, Jul 2005.

[3] W.S. Gan and S.M. Kuo, "An integrated audio and communication active noise control headset," *IEEE Trans. Consumer Elec.*, vol. 48, pp. 242–247, May 2002.

[4] M. Bai and Dunjay Lee, "Inplementation of an active headset by using the H_∞ robust control theory," *J. Acoust. Soc. Am.*, vol. 102, no. 4, pp. 2184–2190, Oct 1997.

[5] J. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory*, Macmillan Publishing Co., Boston, 1990, also at http://www.control.utoronto.ca/people/profs/francis/dft.html.

[6] H.W. Bode, *Network Analysis and Feedback Amplifier Design*, D. Van Nostrand, Princeton, NJ, 1945.

[7] S.J. Elliott and P.A. Nelson, "Active noise canceling," *IEEE Signal Processing Mag.*, vol. 9, no. 10, pp. 12–35, Oct 1993.

[8] B.A. Francis and G. Zames, "On H_∞-optimal sensitivity theory for SISO feedback systems," *IEEE Trans. Automat. Contr.*, vol. AC-29, no. 1, pp. 9–16, Jan 1984.

[9] J.C. Dolye, K. Glover, P.P. Khargonekar, and B.A. Francis, "State-space solutions to standard H_2 and H_∞ control problems," *IEEE Trans. Automat. Contr.*, vol. 34, no. 8, pp. 831–847, Aug 1989.

[10] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, Prentice Hall, Englewood Cliffs, New Jersey, 1989.

[11] P.A. Nelson and S.J. Elliott, *Active Control of Sound*, Academic Press, New York, 1992.

[12] W. H. Press, *Numerical Recipes, The Art of Scientific Computing*, Cambridge University Press, Cambrige, 1986.

[13] K.H. Keel and S.P. Bhattacharyya, "A linear programming approach to controller design," in *Proc. 36th conf. Decision & Contr.*, San Diego, CA, 1997, pp. 2139–2148.