Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Approaches to Computationally Efficient Implementation of Gain Governors For Nonlinear Systems With Pointwise-in-Time Constraints

Ilya Kolmanovsky and Jing Sun

Abstract— **The gain governors use receding horizon optimization to adjust parameters (such as gains) in the nominal control laws. The parameters are optimized at each time instant to minimize a cost function subject to pointwise-in-time constraints and subject to the condition that the parameter values are constant over the horizon. The gain governors may be viewed as a special class of Model Predictive Control (MPC) algorithms. They provide guaranteed stability properties without terminal set conditions as well as a large degree of flexibility in accommodating the on-line computational effort. The paper reviews the properties of the gain governors and discusses different implementations allowed by the general theory with a view towards effectively accommodating the computational effort involved with the on-line optimization.**

I. Introduction

The gain governors, first proposed in [9], utilize receding horizon optimization to adjust parameters (such as gains) in the nominal control laws so that to improve closed-loop performance and satisfy pointwise-in-time constraints. The parameters are assumed to remain constant over the horizon, and they are re-optimized at each time instant to minimize a cost function defined over the horizon subject to given pointwise-in-time constraints. It was shown in [9] that if the horizon satisfies appropriate assumptions then the closed-loop stability and constraint satisfaction can be assured without explicit terminal set conditions. The underlying on-line optimization reduces to a finite number of simulations if the parameters can only assume a finite set of values.

The gain governor belongs to a larger family of suboptimal yet computationally simpler schemes for controlling systems with pointwise-in-time constraints, which also includes the reference governors (see, e.g. [1], [2], [4], [5], [6] and references therein) and feedforward governors [9]. Because these schemes do involve on-line optimization, albeit a simple one (or equivalent explicit computations of the solution [5]), the required computational effort may still be excessive for embedded applications that involve fast dynamics and low cost processors.

The present paper describes several approaches for accommodating the computational effort associated with the gain governors. For instance, it is shown that the exact minimizer is not required and the updates can be less frequent than every sampling period. Another approach involves developing an explicit implementation of the gain governor (i.e., pre-computing the solution to the optimization problem off-line and storing a functional approximation of that solution for on-line use) for a subset of the full state space (a patch) and extending it beyond this subset so that to avoid constraint violation and preserve state convergence. These and other approaches are covered in Section II. The double integrator example with input constraints is used for illustration of the main ideas in Section III. A more elaborate example is considered in Section IV. Concluding remarks are made in Section V.

II. Gain Governor and Approaches to Its Implementation

In this Section we first review the gain governor theory. We then describe various approaches to its implementation aimed at improving computational efficiency while preserving its basic properties.

A. The Gain Governor

The gain governor [9] is applied to disturbance-free, nonlinear discrete-time systems of the form,

$$x(t+1) = f(x(t), u(t)), \tag{1}$$

where $x(t)$ is the state of the system and $u(t)$ is the control input. A parameter-dependent family of nominal control laws, controlling the system to the origin is assumed to be available,

$$u(t) = U_\theta(x(t)), \tag{2}$$

where $\theta \in \Theta$ is the parameter and Θ is a given set. As one example, $\theta(t)$ may represent some or all of the gains in a given feedback law.

The pointwise-in-time constraints imposed on $x(t)$ and on $\theta(t)$ have the form:

$$(\theta(t), x(t)) \in C, \quad \forall t \in \mathbf{Z}^+, \tag{3}$$

where C is a given subset. Because of (2), the input constraints can always be recast as equivalent constraints on $x(t)$ and $\theta(t)$.

The on-line selection of $\theta(t)$ for each $t \in \mathbf{Z}^+$ is based on the minimization of a cost function,

$$J(x(t), \theta(t)) = \tfrac{1}{2}\theta(t)^{\mathrm{T}}\Psi\theta(t) + \sum_{k=0}^{k=T} Q\big(\phi_{\theta(t)}(k, x(t)), \theta(t)\big), \tag{4}$$

subject to

$$(\theta(t), \phi_{\theta(t)}(k, x(t))) \in C, k = 0, 1, \cdots, T, \tag{5}$$

where $\phi_{\bar\theta}(k, \bar x)$ denotes the solution of (1), (2) at time k given that $x(0) = \bar x$ and $\theta(t) \equiv \bar\theta$ is maintained constant.

I.V. Kolmanovsky is with Ford Motor Company, Dearborn, Michigan. Email: ikolmano@ford.com.
J. Sun is with Dept. of Naval Architecture and Marine Engineering, Ann Arbor, Michigan. Email: jingsun@umich.edu.

The function Q is the incremental cost and $\Psi = \Psi^{\mathrm{T}}$ is a parameter penalty matrix. If C admits an inequality characterization

$$C = \{(\theta, x) : \; g_j(\theta, x) \leq 0, \; j = 1, \cdots, q\},$$

then the constraint (5) reduces to

$$g_j(\theta, \phi_\theta(k, x(t))) \leq 0, \; j = 1 \cdots, q; k = 0, \cdots, T. \quad (6)$$

The constraint (5) can be also restated equivalently as

$$(\theta(t), x(t)) \in O_T, \; O_T \stackrel{\Delta}{=} \{(\theta, x) : \; (\theta, \phi_\theta(k, x)) \in C, \\ k = 0, 1, \cdots, T\}. \quad (7)$$

The rationale for the gain governor is easy to understand in the case of systems with input constraints and it is similar to that for the multi-mode controller [7], [8]. Specifically, the gain governor can lower the gains when it becomes necessary to avoid violating the input constraints; the gain governor can increase the gains when there is no danger of constraint violation and doing so improves the performance. In comparison to the multi-mode controller, the gain governor has an advantage in that it explicitly incorporates a cost function in deciding which gains to use.

Under reasonable assumptions (see [9]), the desirable response properties of the gain governor, including the asymptotic convergence of $x(t)$ and $\theta(t)$ to the origin if starting with a feasible initial state, can be rigorously guaranteed. The needed assumptions for convergence of $x(t)$ include compactness of C and Θ, asymptotic stability of (1), (2) with $\theta(t)$ fixed, continuity and positive-definiteness properties of the function Q ($Q(a,b) \geq 0$, for all a, b, $Q(a,b) > 0$ if $a \neq 0$), $\Psi \geq 0$, and an adequate horizon, T. Specifically, T must satisfy the condition that if $t > T$, then $Q(\phi_\theta(t,x), \theta) \leq q \cdot Q(x, \theta)$ for some $0 \leq q < 1$ and all $(\theta, x) \in C$, $\theta \in \Theta$. This condition is reasonable in view of asymptotic stability of (1), (2) with $\theta(t)$ fixed and compactness of C and Θ. In addition, T must satisfy the property that if $(\theta, x) \in C$, $\theta \in \Theta$, $\phi_\theta(t, x) \in C$ for $t = 0, 1, \cdots, T$ then $\phi_\theta(t, x) \in C$ for $t > T$. Under additional technical assumptions which include convexity of Θ, $0 \in int\Theta$, and $\Psi > 0$, the convergence of $\theta(t)$ to the origin can be guaranteed.

Numerical procedures are available for estimating an acceptable value of T [9]. One procedure is based on computing two quantities, $L_1(k)$ and $L_2(k)$, $k \in \mathbf{Z}^+$:

$$L_1(k) = \max_{j=1,\cdots,q, \theta \in \Theta, (\theta, x) \in C} g_j(\theta, \phi_\theta(k, x)),$$

$$L_2(k) = \max_{\theta \in \Theta, (\theta, x) \in C} \frac{Q(\phi_\theta(k, x), \theta)}{Q(x, \theta)},$$

where g_j is defined in (6). The $L_1(k)$, $L_2(k)$ are, respectively, the maximum constraint violation and the minimum decay in the incremental cost due to $\phi_\theta(k, x)$ as x and θ vary ($\theta \in \Theta$, $(\theta, x) \in C$). An acceptable T must satisfy the conditions $L_1(k) \leq 0$ and $L_2(k) \leq q$ for all $k \geq T$ and some $0 \leq q < 1$. Either off-line numerical optimization or multiple off-line simulations of the model for different x and θ can be used to estimate $L_1(k)$ and $L_2(k)$. An acceptable T can be easily picked from the graphical plots of $L_1(k)$ and $L_2(k)$ versus k. We note that the resulting T is a numerical approximation to the required horizon and not a guaranteed upper bound.

Note that as compared to more general MPC schemes in which a control sequence needs to be computed over the specified horizon, the condition $\theta(t+k) = \theta(t)$, $k = 0, \cdots, T$, makes the optimization problem (4), (5) lower dimensional; in fact, its dimension does not depend on the horizon, T. In the case when Θ has a finite number of elements (which is allowed by the theory in [9]), the optimization reduces to a finite number of simulations. In that case, the gain governor provides a stability preserving and performance improving mechanism for switching between a finite family of controllers.

The gain governor permits a large degree of flexibility in incorporating the on-line optimization to improve the performance and enforce the constraints.

B. The use of non-exact minimizers

In minimizing (4) it is not necessary to obtain the exact optimizer to preserve the state convergence. Indeed, suppose a feasible $\theta(t-1)$ has been computed at a time instant $t-1$ and suppose $\theta(t)$ has been determined at the time instant t using numerical optimization (e.g., with $\theta(t-1)$ as the initial guess) so that the condition,

$$J(x(t), \theta(t)) \leq J(x(t), \theta(t-1)), \quad (8)$$

is satisfied. Then,

$$\begin{aligned} & J(x(t), \theta(t)) \leq J(x(t), \theta(t-1)) \\ & = J(f(x(t-1), U_{\theta(t-1)}(x(t-1))), \theta(t-1)) \\ & \leq J(x(t-1), \theta(t-1)) - (1-q) Q(x(t-1), \theta(t-1)) \\ & \leq J(x(t-1), \theta(t-1)). \end{aligned}$$

Consequently, if (8) is satisfied for all t and given that J and Q are non-negative, it follows that $Q(x(t), \theta(t)) \to 0$ as $t \to \infty$, and by continuity and positive-definiteness of Q with respect to the first argument, it then follows that $x(t) \to 0$. For unconstrained problems or for problems where the constraints are incorporated into the cost through the penalty function, the condition (8) can be satisfied by the usual line search methods, for which the gradient of $J(x(t), \theta(t))$ with respect to $\theta(t)$ can be easily computed. If the line search iterations at a particular time instant do not produce a value for $\theta(t)$ satisfying (8) (e.g., because of running out of available time), then $\theta(t) = \theta(t-1)$ can be used instead to preserve the state convergence.

C. Dealing with large number of constraints

The number of constraints in (5) or (6) grows with T and can complicate on-line optimization if T is large. As in the reference governor case [6], the use of a simple off-line functional characterization of a subset, $M \subset O_T$, in place of O_T provides an alternative. If

$$M = \{(\theta, x) : \; V(x, \theta) \leq 0\} \subset O_T,$$

then multiple inequalities in (5) or (6) can be replaced by a single inequality, $V(x(t), \theta(t)) \leq 0$. As in the reference governor case, the methods to construct V can be based on parameter-dependent closed-loop Lyapunov functions or classification/pattern recognition techniques. When M is used in place O_T, a situation may occasionally arise that no feasible $\theta \in \Theta$ exists for a time instant t, i.e., $V(x(t), \theta) > 0$ for all $\theta \in \Theta$. In this case, as in [6], $\theta(t)$ is set equal to $\theta(t-1)$. This procedure, under the same assumptions as in Section II-A, guarantees that the constraints are satisfied and $x(t) \to 0$.

D. The use of less frequent updates

Further, the updates of θ can be less frequent than at every time instant, for example, at every $t \in I_u = \{0, n, 2n, 3n, \cdots; n \in \mathbf{Z}^+, n > 1\}$. Whenever $t \notin I_u$, $\theta(t)$ can be kept constant, $\theta(t) = \theta(t-1)$, while the on-line optimization is performed only for $t \in I_u$. The time interval, $k \cdot n \leq t < (k+1) \cdot n$, provides n sampling periods to calculate the optimal value of $\theta((k+1)n)$ which is useful in case these computations cannot be completed within a single sampling period. Note that the value of $x((k+1)n) = \phi_{\theta(kn)}(n, x(kn))$ can be predicted via on-line simulations assuming that $\theta(t) = \theta(kn)$ for $kn \leq t \leq (k+1)n$. The drawback of less frequent parameter updates is in lost opportunities for transient performance improvements and, in practice, potentially degraded performance due to unmodelled disturbances.

In a common situation when (1), (2) represent a discrete-time approximation of a continuous-time system and Δ is the physical time period between two subsequent parameter updates, it is usually the underlying continuous-time dynamics that dictate an acceptable value for $T \cdot \Delta$. In particular, selecting larger Δ (i.e., using less frequent parameter updates) can lead to smaller T and reduce complexity of the on-line optimization problem. The drawback of using large Δ is cruder enforcement of constraints for the original continuous-time system. This drawback can be addressed by enforcing constraints on a finer time grid wherein (5) is replaced by

$$\phi_{\theta(t)}(n\delta, x(t)) \in C, \; n = 0 \cdots, N,$$

where $\phi_{\theta(t)}(n\delta, x(t))$ is the predicted state of the continuous-time system at time $n\delta$, where $\delta < \Delta$ and $N\delta > T\Delta$. If the approach of Section II-C, with $M \subset O_T$, and V is used, the number of constraints in the resulting optimization problem may not necessarily be large. Note also that the computational effort to simulate the continuous-time model to a desired level of accuracy does not decrease with larger Δ.

E. The use of terminal set and terminal penalty conditions

As an alternative to constraining the horizon, T, based on our assumptions terminal set and terminal penalty conditions may also be used. In this case, a terminal penalty function term, $F(\phi_{\theta(t)}(T, x(t)))$, where $F(0) = 0$, is added to (4) and a terminal set condition $\phi_{\theta(t)}(T, x(t)) \in \Gamma$ is imposed. The terminal set Γ should be positively invariant *for all* $\theta \in \Theta$ (i.e., $\phi_\theta(t, x) \in \Gamma$ if $x \in \Gamma$) and constraint-admissible (i.e., $(\theta, \phi_\theta(t, x)) \in C$ for all $x \in \Gamma$, $\theta \in \Theta$ and $t \in \mathbf{Z}^+$); the terminal penalty function F must satisfy $F(\phi_\theta(1, x)) - F(x) \leq -Q(\phi_\theta(1, x), \theta)$ for all $x \in \Gamma$, $\theta \in \Theta$.

F. Explicit implementation of the gain governor

An explicit implementation of the gain governor provides another mechanism for reducing the computational burden. In the explicit implementation, the optimal values of $\theta = \theta^*(x)$ are first pre-computed off-line for different x and then they are used to develop a functional approximation, $\bar{\theta}^*(x)$ of $\theta^*(x)$. The $\bar{\theta}^*(x)$ can be applied during the on-line operation of the system thereby eliminating the need for on-line optimization.

Suppose that such an explicit solution is available for $x \in \Pi$ where Π is a set (referred to as a patch) such that $0 \in int\Pi$. As long as $x(t) \in \Pi$, $\bar{\theta}^*(x(t))$ is defined. If the trajectory of x starts in Π but exits Π at some time t, then $\bar{\theta}^*(x(t))$ is not defined but $\theta(t)$ can be set to the value of $\theta(\tilde{t}) = \theta^*(x(\tilde{t}))$, where $\tilde{t} < t$ is the last time instant for which $x(\tilde{t}) \in \Pi$. This procedure guarantees that the constraints are satisfied, because if T is selected consistently with the assumptions in [9], $x(\tilde{t}) \in \Pi$ implies $(\theta^*(t), \phi_{\theta^*(t)}(k, x(t))) \in C$ for all $k \in \mathbf{Z}^+$. The procedure also preserves, under the assumptions of [9], the cost non-increase condition $J(x(t+1), \theta^*(t+1)) \leq J(x(t), \theta^*(t))$ and thus the convergence of $x(t)$ to 0. Even if $x(t)$ exits Π at a time instant t, the condition $0 \in int\Pi$ and asymptotic stability of (1), (2) with $\theta(t)$ maintained constant guarantee that $x(t)$ must re-enter Π in finite-time where $\bar{\theta}^*(x(t))$ can again be applied.

The appropriate selection of the set Π over which the functional approximation to the explicit solution is developed and deployed provides a mechanism for decreasing the complexity of this functional approximation and for improving its accuracy. We note that this simple mechanism may not be available with more general MPC schemes since for them the explicit implementation typically retains the information only about the first element of the optimal control sequence and discards the rest of this sequence.

III. Double Integrator Example

We consider an application of a gain governor to the double integrator system, $\dot{x}_1 = x_2, \; \dot{x}_2 = u$, under an input saturation constraint, $|u| \leq 1$. The nominal control law has the form, $u = -\omega_n^2(x_1 - r) - 2\zeta\omega_n x_2$, where $\zeta = 0.5$ and $\omega_n = \omega_{n,0} + \theta, \quad \omega_{n,0} = 10$. The continuous-time system is discretized assuming the sampling period of $\Delta T = 0.01$ sec.

We consider two cases. In the first case, $\theta(t)$ is selected from a continuous interval, $\Theta = [-9, 0.5]$. In the second case, $\theta(t)$ is selected from a discrete set,

$$\Theta = \{-9, -8, \cdots, -1, 0, 0.1, 0.2, 0.3, 0.4, 0.5\}.$$

The cost is (4) with $\Psi_\theta = 10^{-4}$ and $Q = 10 \cdot (x_1)^2 + 0.1 \cdot (x_2)^2$. The negative values in Θ provide a mechanism to slow down the response while the positive values can speed up the response. The set C in (3) reflects the input constraint, and additional constraints, $\; -1 \leq x_1 \leq 1, \; -2 \leq x_2 \leq 2$, were added to make it compact. The numerical procedure described in Section II-A was used to estimate an adequate horizon as $T = \frac{3.5}{\Delta T}$.

Figure 1 shows the time response of x_1 with and without the gain governor. The response (a) of the nominal controller with $\omega_n(t) \equiv \omega_{n,0} = 10$ is very fast if there is no saturation, but it behaves poorly with the saturation, see the response (b). The use of a controller with a fixed lower gain $\omega_n(t) \equiv 1.0$ avoids control input saturation (see the response (c)) but it significantly slows down the system thereby sacrificing the performance. Finally, with the gain governor the system avoids control constraint violation, and the response is much faster than both (c) and (b), see responses (d) and (e) in Figure 1. The response (d) with θ selected from the interval is faster than the response (e) with θ selected from the discrete set. However, the selection from a discrete set can be implemented just using on-line model simulations.

Figure 2 illustrates the behavior of $\theta^*(t)$. Note that the response of $\theta^*(t)$ is non-monotonic so that the system is first slowed down to prevent violation of the input constraint, and then made faster once close to the desired equilibrium; ultimately, $\theta^*(t)$ settles to zero in finite time. Figure 3 shows that the control constraints are satisfied by the gain governor.

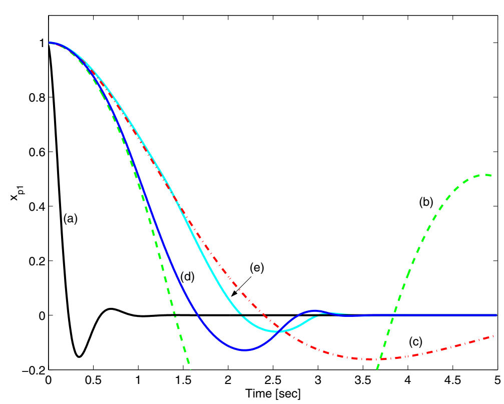

Fig. 1. Time response of x_1 without the gain governor and with the gain governor: (a) With high gain controller and no input saturation; (b) With high gain controller and input saturation; (c) With low gain controller and input saturation; (d) With the gain governor and θ taking values from an interval; (e) With the gain governor and θ taking values from a discrete set.

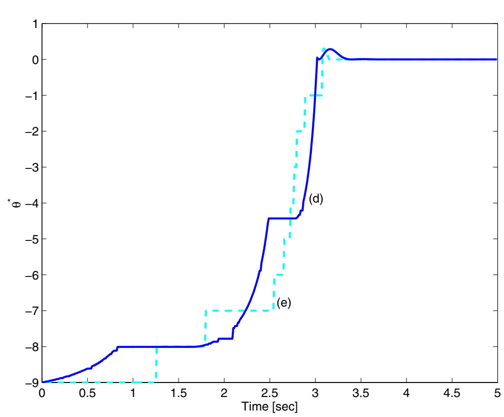

Fig. 2. Time response of $\theta(t) = \theta^*(t)$ in the control law: (d) With θ taking values from an interval; (e) With θ taking values from a discrete set.

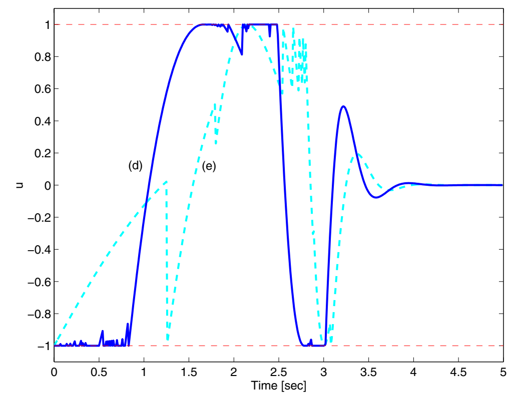

Fig. 3. Time history of the control input u with the gain governor: (d) With θ taking values from an interval; (e) With θ taking values from a discrete set.

The explicit implementation of the gain governor was carried out next. The $\theta^*(x)$ was pre-computed off-line over a mesh of the set (a patch) $\Pi = [-0.5, 0.5] \times [-0.5, 0.5]$, see Figure 4. Note that the area of this patch is only one eighth of that allowed by the constraints. Two functional approximations were developed. A refined approximation was implemented using a 50×50 look-up table. Then a coarse approximation was implemented using a 10×10 look-up table. The look-up tables utilized linear interpolation in-between the mesh points. The responses of the explicitly implemented gain governor are shown and compared to the response of the on-line optimization-based gain governor in Figures 5-7. Both refined approximation and coarse approximation come very close to fully enforcing the input constraints and the response of x_1 is only mildly affected by the approximation errors. The constraint violation may be avoided even though the approximation errors are present if the gain governor is redesigned by assuming that the constraints are slightly tighter than they really are.

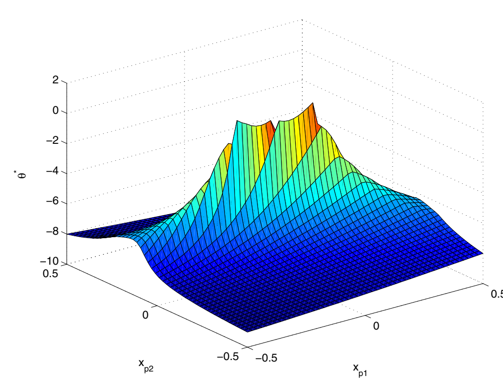

Fig. 4. The function $\theta^*(x)$ pre-computed off-line and used in explicit implementation of the gain governor.

IV. A HIGHER ORDER EXAMPLE

In this section we consider a more elaborate example of the gain governor application application to a model arising in engine control. In the engine which we consider, an

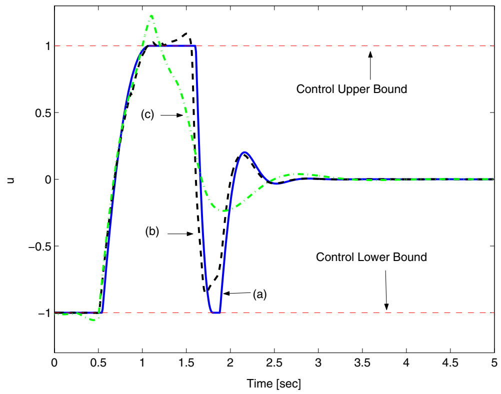

Fig. 5. The time history of $u(t)$ with (a) the on-line optimization-based θ; (b) with explicit implementation based on the refined approximation; (c) with explicit implementation based on the coarse approximation.

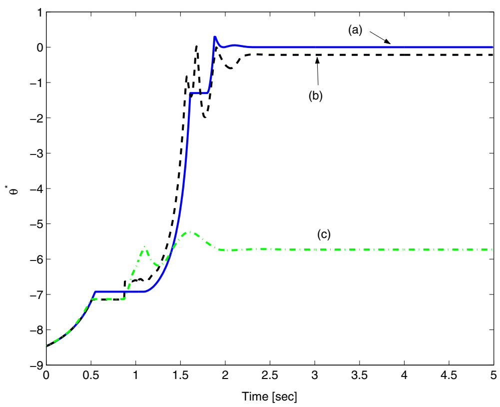

Fig. 6. The time history of $\theta(t)$ with (a) the on-line optimization-based θ determination; (b) with explicit implementation based on the refined approximation; (c) with explicit implementation based on the coarse approximation.

electronic throttle is used to control the flow of air into the engine while the adjustment of cam phasing is used to reduce residuals at low loads (to improve combustion stability) and to increase residuals at medium loads in order to reduce oxides of nitrogen emissions. Since the residuals displace fresh air charge in the cylinders, the cam phasing transitions can influence the transient cylinder air flow behavior and the engine torque response. In particular, if the cam phasing transitions occur too fast, the cylinder flow and the engine torque may undershoot; if this transition is too slow, the cylinder flow and the engine torque may overshoot. See Figure 8. While spark timing can, in principle, help mitigate the overshoot or undershoot in the engine torque, the spark timing authority is limited and using spark timing may degrade fuel economy. To ensure the monotonic cylinder flow response, an on-line selection procedure for the cam phasing transition rate has been developed in [10]. Here we approach this problem using the gain governor and synergistically treat both the electronic throttle and cam phasing actuation.

The engine breathing dynamics, in simplified form as compared to [10], have the following form in continuous

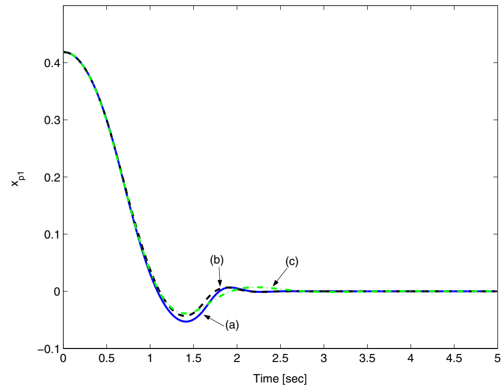

Fig. 7. The time history of $x_1(t)$ with (a) the on-line optimization-based θ determination; (b) with explicit implementation based on the refined approximation; (c) with explicit implementation based on the coarse approximation.

time:

$$\begin{aligned} \dot{p} &= c_m\left(k_1 \cdot u_{th} \cdot \sqrt{p - p^2} - W\right), \\ W &= k_2 \cdot p \cdot \left(1 - \frac{\alpha}{90}\right), \\ \dot{\alpha} &= -\tau(\alpha - \alpha_e), \\ \ddot{u}_{th} &= -2\zeta\omega_n \dot{u}_{th} - \omega_n^2(u_{th} - u_{th,e}), \end{aligned}$$

where p is the intake manifold pressure, u_{th} is the throttle angle, W is the cylinder flow, α is the cam phasing angle, and the subscript e signifies the equilibrium value of a variable. Note that

$$p_e = \frac{1}{1 + \left(k_2 \cdot \frac{1 - \frac{\alpha_e}{90}}{k_1 u_{th,e}}\right)^2}, \quad W_e = k_2 \cdot p_e \cdot \left(1 - \tfrac{\alpha_e}{90}\right).$$

The constants are $c_m = 0.0414$, $k_1 = 4.0$, $k_2 = 30.0$, $\omega_n = 24.5$. The governed parameters are τ and ζ so that $\tau = 8 + \theta_1$, $\zeta = 1.2 + \theta_2$, where $\theta \in \Theta = [-7.54, 25.33] \times [-1, 0.8]$. They determine, respectively, the speed of cam phasing transitions and the damping ratio in the throttle position response. These parameters are updated every $\Delta T = 0.05$ sec by the gain governor. The incremental cost function Q in the form of (4) is

$$Q = q_1 \cdot (W - W_e)^2 + q_2 \cdot (\alpha - \alpha_e)^2 + q_3 \cdot (u_{th} - u_{th,e})^2 + q_4 \cdot \dot{u}_{th}^2,$$

where $q_1 = 100$, $q_2 = 0.01$, $q_3 = 10^{-4}$, $q_4 = 10^{-4}$ while $\Psi_\theta = diag(0.001, 0.001)$. The cost emphasizes the cylinder flow response to provide better engine responsiveness and drivability.

Assuming the command is to increase the cylinder flow, the constraint which ensures the monotonic cylinder flow response is $\dot{W}(t) \geq 0$. Strictly speaking, the theory in [9] does not permit the treatment of constraints in this form because the equilibrium is on the boundary of the feasible set. We therefore relax the constraint to $\dot{W}(t) \geq -0.1$.

Figure 8 demonstrates that the gain governor is able to coordinate throttle and cam phasing to produce a monotonic cylinder flow response. Figures 9-10 indicate that the gain governor creates an initial overshoot in throttle response (by decreasing the damping ratio) to increase air flow and mitigate the increase in the residuals. It initially adjusts the cam phasing position slowly and then speeds it up.

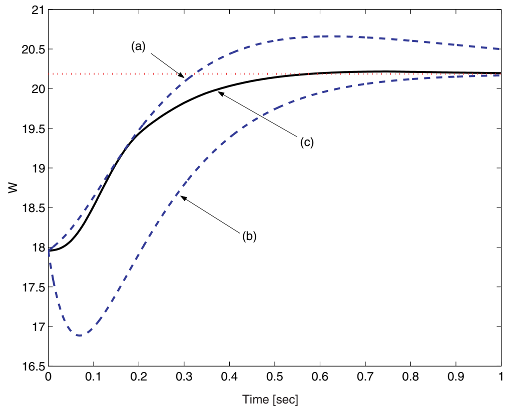

Fig. 8. Time history of cylinder flow with (a) slow cam phasing transition and nominal throttle transition; (b) fast cam phasing transition and nominal throttle transition; (c) cam phasing transition and throttle transition controlled by the gain governor.

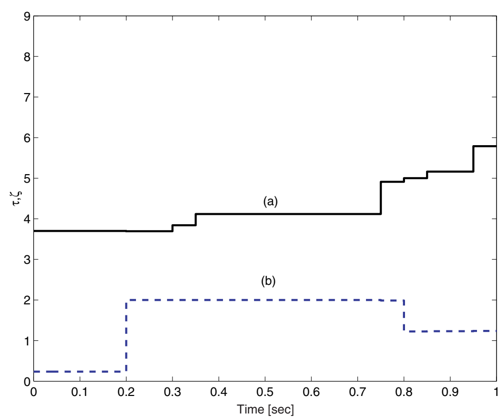

Fig. 9. Time histories of $\theta_1(t) = \tau(t)$ (trajectory (a)) and $\theta_2(t) = \zeta(t)$ (trajectory (b)) prescribed by the gain governor.

V. Concluding Remarks

The paper discussed several approaches to the implementation of the gain governors. The gain governors use receding horizon optimization for on-line adjustment of parameters in nominal control laws so that to avoid violation of pointwise-in-time state and input constraints, and to improve transient performance. The adjustable parameters remain constant over the prediction horizon. Thus the dimensionality of the optimization problem being solved does not depend on the horizon and is equal to the number of parameters. The paper has demonstrated that a large degree of flexibility exists in accommodating the on-line optimization required to implement the gain governor. For example, the exact minimizer is not required or need not be computed at every sample time instant. Key results hold even if the parameter values are restricted to a finite set; in this case the optimization reduces to a finite number of on-line simulations. In addition, an explicit implementation (wherein the solution to

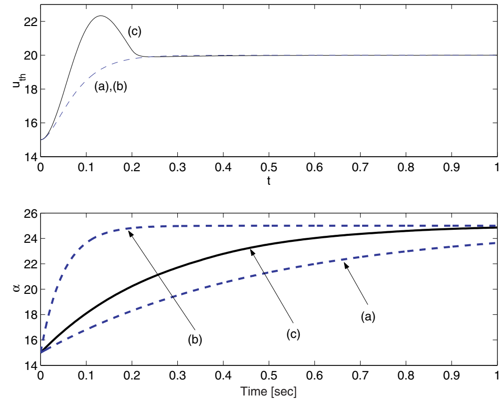

Fig. 10. Time histories of u_{th} (upper subplot) and α (lower subplot) with (a) slow cam phasing transition and nominal throttle transition; (b) fast cam phasing transition and nominal throttle transition; (c) cam phasing transition and throttle transition controlled by the gain governor.

the receding horizon optimization problem is pre-computed off-line and its functional approximation is applied on-line) can be generated for a subset of the state space and then extended in a simple way beyond this subset while preserving the state convergence. The modest computational effort may make gain governors and more general parameter governors [9] practical *embedded optimization* controllers for systems with fast dynamics and limited computational resources.

References

[1] Bemporad, A., (1998). Reference governor for constrained nonlinear systems, *IEEE Transactions on Automatic Control, vol. 43, no. 3,* pp. 415-419.

[2] Bemporad, A., Casavola, A., and Mosca, E., (1996). A nonlinear command governor for constrained control systems, *Proceedings of 13th IFAC World Congress,* San Francisco, pp. 473-478.

[3] Bemporad, A., Casavola, A., and Mosca, E., (1997). Nonlinear control of constrained linear systems via predictive reference management, *IEEE Transactions on Automatic Control*, vol. 42, no. 3, pp. 340-349.

[4] Gilbert, E., Kolmanovsky, I., and Tan, K.T., (1995). Discrete time reference governors and the nonlinear control of systems with state and control constraints, *International Journal of Robust and Nonlinear Control,* vol. 5, pp. 487-504.

[5] Gilbert, E.G., and Kolmanovsky, I., (1999). Fast reference governors for systems with state and control constraints and disturbance inputs, *International Journal of Robust and Nonlinear Control,* vol. 9, no. 15, pp. 1117-1141.

[6] Gilbert, E., and Kolmanovsky, I., (2002). Nonlinear tracking control in the presence of state and control constraints: A generalized reference governor, *Automatica,* vol. 38, no. 12, pp. 2063-2073.

[7] Kolmanovsky, I., and Gilbert, E., (1997). Multimode regulator for systems with state and control constraints and disturbance inputs, in *Control Using Logic-Based Switching,* edited by Morse, A.S., Springer-Verlag, pp. 104-117.

[8] Kolmanovsky, I., and Gilbert, E., (1998). Theory and computation of disturbance invariant sets for discrete-time linear systems, *Mathematical Problems in Engineering,* vol. 4, pp. 317-367.

[9] Kolmanovsky, I., and Sun, J., (2004). Parameter governors for discrete-time nonlinear systems with pointwise-in-time state and control constraints, *Proceedings of 2004 American Control Conference,* Boston, MA, pp. 3075-3081, June.

[10] Stefanopoulou, A., and Kolmanovsky, I., (1999). Analysis and control of transient torque response in engines with internal exhaust gas recirculation, *IEEE Transactions on Control Systems Technology,* vol. 7, no. 5, pp. 555-566.