Proceedings of the 30th Conference
on Decision and Control
Brighton, England • December 1991

Heuristically Enhanced Feedback Control of Constrained Systems: The Minimum Time Case

M. Sznaier † *and M. J. Damborg* ‡

† Electrical Engineering, University of Central Florida, Orlando, Fl 32816

‡ Electrical Engineering, FT-10, University of Washington, Seattle, WA 98195

Abstract

Recent advances in computer technology have spurred new interest in the use of feedback controllers based upon the use of on-line optimization. Still, the use of computers in the feedback loop has been hampered by the limited amount of time available for computations. In this paper we propose a feedback controller based upon the use of on-line optimization in the feedback loop. We present theoretical results showing that this controller yields asymptotically stable systems, provided that enough computation power is available to solve on-line a constrained optimization problem considerably simpler than the original. By making use of the special structure of time-optimal systems, the constrained optimization problem is further simplified, resulting in a substantial reduction of its dimensionality.

I. Introduction.

A substantial number of control problems can be summarized as the problem of designing a controller capable of achieving acceptable performance under design constraints. This statement looks deceptively simple, but even in the case where the system under consideration is linear time-invariant, the problem is far from solved.

During the last decade, substantial progress has been achieved in the design of linear controllers. By using a parametrization of all internally stabilizing *linear* controllers in terms of a stable transfer matrix Q, the problem of finding the "best" linear controller can be formulated as an optimization problem over the set of suitable Q [1]. In this formulation, additional specifications can be imposed by further constraining the problem. However, although these methods are effective in dealing with frequency-domain constraints, they can address time-domain constraints only in an extremely conservative fashion. Hence, if the constraints are tight this approach may fail to find a solution, even if the problem is feasible

Classically, control engineers have dealt with time-domain constraints by allowing inputs to saturate, and by using "gain" scheduling" and dual mode controllers (where high gain feedback is used far from the constraints and low gain feedback is used when approaching a constraint boundary). Latter, with the appearance of the Linear Quadratic formalism, constraints have been embedded in the performance index by adjusting the penalty weights. Although these methods are relatively simple to use, they have several serious shortcomings, perhaps the most important being their inability to handle constraints in a general way. Hence, they require "ad-hoc" tuning of several parameters making extensive use of simulations.

Alternatively, the problem can be stated as an optimization problem [2]. Then, mathematical programming techniques can be used to find a solution. However in most cases the control law generated is an open-loop control that has to be recalculated entirely, with a considerable computational effort, if the system is disturbed.

Because of the difficulties with the optimal control approach, other design techniques, based upon using a Lyapunov function to design a stabilizing controller, have been suggested [3]. However, these techniques tend to be unnecessarily conservative. Moreover, several steps of the design procedure involve an extensive trial and error process, without guarantee of success (example 5.3 in [3]).

Recently, several techniques that exploit the concept of maximally invariant sets to obtain static [4–6] and dynamic [7–9] linear feedback controllers have been proposed. These controllers are particularly attractive due to their simplicity. However, it is clear that only a fraction of the feasible constrained problems admit a linear solution. Furthermore, performance considerations usually require the control vector to be on a constraint boundary and this clearly necessitates a *non-linear* controller capable of saturating.

Finally, in the last few years, there has been a renewed interest in the use of feedback controllers based upon the use of on-line minimization. Although this idea was initially proposed as far back as 1964 [10], its implementation has become possible only during the last few years, when the advances in computer technology made feasible the solution of realistically sized optimization problems in the limited time available. In [11–12] we presented a theoretical framework to analyze the effects of using on-line optimization and we proposed a controller guaranteed to yield asymptotically stable systems.

However, although these theoretical results represent a substantial advance over some previously used "ad-hoc" techniques, in some cases they are overly conservative, requiring the on-line solution of a large optimization problem. Since in most sampled control systems the amount of time available between samples is very limited, this may preclude the use of the proposed controller in many applications.

In this paper we present a suboptimal feedback controller for the minimum-time control of discrete time constrained systems. Following the approach presented in [12], this controller is based upon the solution, during the sampling interval, of a sequence of optimization problems. We will show that by making use of the special structure of time-optimal systems the proposed algorithm results in a significant reduction of the dimensionality of the optimization problem that must be solved on-line, hence allowing for the implementation of the controller for realistically sized problems.

The paper is organized as follows: In section II we present a formal definition to our problem. In section III we introduce several concepts and we review briefly the theoretical framework presented in [12]. Here we present a realistic example illustrating dimensionality problems. In section IV we present the proposed feedback controller and the supporting theoretical results. Finally, in section V, we summarize our results and we indicate directions for future research.

II. Statement of the problem.

In this paper we consider linear, time invariant, controllable discrete time systems modeled by the difference equation:

$$\underline{x}_{k+1} = A\underline{x}_k + B\underline{u}_k, \qquad k = 0, 1, \dots \tag{S}$$

with initial condition $\underline{x}_o$, and the constraints

$$\begin{gathered} \underline{u}_k \in \Omega \subset R^m, \underline{x}_k \in \mathcal{G} \subseteq R^n \\ \mathcal{G} = \{\underline{x} : |G\underline{x}| \leq \underline{\gamma}\}, \ \Omega = \{\underline{u} : |W\underline{u}| \leq \underline{\omega}\} \end{gathered} \tag{C}$$

where $\underline{\gamma} \in R^p, \underline{\omega} \in R^q$, γ_i, $\omega_i > 0$, $G \in R^{pn}, W \in R^{qm}$ with full column rank, $\underline{x}$ indicates x is a vector quantity and where the inequalities (C) should be interpreted in a component by component sense. Furthermore, we will assume as usual that A^{-1} exists. The objective is to find a sequence of admissible feedback controls, $\underline{u}_k[\underline{x}_k]$, that minimizes the *transit time* to the origin. Thus the problem has the form of the following constrained optimization problem:

$$\min_{\underline{u}\in\Omega} N$$

subject to:

$$\begin{aligned} \underline{x}_{k+1} &= A\underline{x}_k + B\underline{u}_k \\ \underline{x}_k &\in \mathcal{G}\ \forall\, k,\ \underline{x}_N = 0 \end{aligned} \qquad (P)$$

Throughout this paper we will assume that problem (P) is well posed in the sense of having a solution for *any* initial condition $\underline{x}_o \in \mathcal{G}$.

III. Definitions and Background Theoretical Results.

In order to analyze the proposed controller we need to introduce some definitions and background theoretical results. We begin by introducing some definitions that deal with the controllability aspects of the problem and in particular, with the effect of quantizing the control space.

3.1 Definitions

- **Def. 1:** The *Null Controllable domain* of (S) is the set of all points $\underline{x} \in \mathcal{G} \subseteq R^n$ that can be steered to the origin by applying a sequence of admissible controls $\underline{u}_k \in \Omega \subset R^m$, such that $\underline{x}_k \in \mathcal{G}, k = 0,1\ldots$. The Null Controllable domain of (S) will be denoted as C_∞. The Null Controllable domain in j or fewer steps will be denoted as $C_j \subseteq C_\infty$.

- **Def. 2:** Consider a closed set $\Omega \subseteq R^m$. A (uniform) *quantization* Ω_s of Ω is defined as the set $\Omega_s \stackrel{\Delta}{=} \{\underline{u} \in \Omega: u_i = n_i/s\}$ where u_i is the i^{th} coordinate of $\underline{u}$, n_i is integer and s is a scaling factor. The quantity $1/s$ will be called the norm of the quantization.

- **Def. 3:** Let O be a convex open set containing the origin and such that for all the optimal trajectories starting out in O, the constraints (C) are not effective, and let $J_o(\underline{x})$ be the optimal cost–to–go from the state $\underline{x}$. A function $g: R^n \rightarrow R$ such that:

$$\begin{aligned} 0 \leq g(\underline{x}) &\leq J_o(\underline{x})\ \forall\ \underline{x} \in \mathcal{G} \\ g(\underline{x}) &= J_o(\underline{x})\ \forall\ \underline{x} \in O \end{aligned}$$

will be called an *underestimate* relative to the set O.

3.2 Background Theoretical Results

In this section we establish the background theoretical framework required to support the proposed controller. We begin by showing that problem (P) can be *exactly* solved by solving a sequence of suitable approximations.

- **Theorem 1:** Let O be the set introduced in Def. 3 and let $\underline{x}^u_k(\underline{\xi})$ be the (unconstrained) optimal trajectory corresponding to the initial condition $\underline{\xi} \in O$. Finally let $g(\underline{x}): R^n \rightarrow R$ be an underestimate relative to O. Consider the following optimization problems:

$$\min_{\mathcal{U}} \left\{ J(\underline{x}) = \sum_{k=1}^{N} 1 \right\} \qquad (1)$$

$$\min_{\mathcal{U}} \left\{ J_m(\underline{x}) = \sum_{k=1}^{m-1} 1 + g(\underline{x}_m) \right\} \qquad m < N \qquad (2)$$

subject to (P), with $\underline{u}_k \in \Omega \subset R^m$ and where $\mathcal{U} = \{\underline{u}_0, \underline{u}_1, \ldots\}$. Then an optimal trajectory, $\underline{x}^o_k, k = 1,2\ldots m$ which solves (2), extended by defining $\underline{x}^o_k = \underline{x}^u_k(\underline{x}^o_m), k = m+1\ldots N$, is also a solution of (1) provided that $\underline{x}^o_m \in O$.

Proof: The theorem corresponds to a special case of Theorem 1 in [12], with $L_k(\underline{x}_k, \underline{u}_k) \equiv 1$.

It follows that problem (P) can be solved by using the sampling interval to solve a sequence of optimization problems of the form (2) until a number m and a trajectory $\underline{x}_k$ such that $\underline{x}_m \in O$ are obtained. However this approach presents the difficulty that the asymptotic stability of the resulting closed loop system can not be guaranteed when there is not enough time to reach the region O. This difficulty can be solved by imposing an additional constraint, which does not affect feasibility, upon (P). To show this, we first introduce a constraint–induced norm in the set $\mathcal{G}$. We then show that there exists at least one admissible control sequence such that this norm defines a Lyapunov function for the system, hence guaranteeing asymptotic stability. This result will be applied to generate a stabilizing controller.

- **Lemma 1:** Let:

$$\|\underline{x}\|_{\mathcal{G}} \stackrel{\Delta}{=} \max_{1\leq i\leq p} \left\{ \frac{|G\underline{x}|_i}{\gamma_i} \right\} = \|\Gamma^{-1} G\underline{x}\|_\infty \qquad (3)$$

where $\Gamma = diag(\gamma_1 \ldots \gamma_p)$. Then $\|.\|_{\mathcal{G}}$ is a norm in R^n.

Proof: The proof follows by noting that since G has full column rank, $\|\Gamma^{-1}G\underline{x}\|_\infty = 0$ iff $\underline{x} = 0$. The additional properties of norms follow immediately from the definition.

- **Theorem 2:** Consider problem (P) and assume that the following condition holds:

$$\min_{\underline{u}\in\Omega} \{\|A\underline{x} + B\underline{u}\|_{\mathcal{G}}\} < 1 \qquad \forall\ \|\underline{x}\|_{\mathcal{G}} \leq 1 \qquad (4)$$

Then, there exists an admissible control sequence $\mathcal{U} = \{\underline{u}_0, \ldots\}$, $\underline{u}_k \in \Omega$ such that:

$$\|\underline{x}_{k+1}\|_{\mathcal{G}} < \|\underline{x}_k\|_{\mathcal{G}},\ k = 0,1\ldots \forall \underline{x}_k \in \mathcal{G} - O \qquad (5)$$

Proof: $\mathcal{G}$ satisfies the constraint qualification conditions of [12]. Therefore the theorem reduces to a special case of Theorem 2 therein.

Corollary: Consider the problem (P') defined as:

$$\min_{\underline{u}\in\Omega} N$$

subject to:

$$\begin{aligned} \underline{x}_{k+1} &= A\underline{x}_k + B\underline{u}_k \\ \|\underline{x}_{k+1}\|_{\mathcal{G}} &< \|\underline{x}_k\|_{\mathcal{G}}, \quad k = 0,1\ldots\ \forall\ \underline{x}_k \in \mathcal{G} - O \\ \underline{x}_k &\in \mathcal{G}\ \forall\, k,\ \underline{x}_N = 0 \end{aligned} \qquad (P')$$

If (4) holds, then (P') is feasible.

3.3 Model Algorithm (Algorithm M)

Consider now the following control algorithm applied to (P'):

Begin:

1) Let $\underline{x}_k$ be the current state of the system, k the current time instant and T the sampling interval. Then:

   i) If $\underline{x}_k \in O$ the solution coincides with that of the unconstrained problem.

   ii) If $\underline{x}_k \notin O$, solve a sequence of optimization problems of the form (2), with the additional constraint (5), until a number m such that $\underline{x}_m \in O$ is found. Use as next control law, the first element of the control sequence corresponding to this solution.

iii) If there is no more computation time available for searching and the region O has not been reached, use the minimum partial cost trajectory that has been found.

2) Repeat step 1 until the origin is reached.

End.

From Theorem 2 and its Corollary it follows that the application of algorithm M to problem (P') yields a control law such that $\|\underline{x}_k\|_{\mathcal{G}}$ is monotonically decreasing in $\mathcal{G}-O$. Hence the system is guaranteed to reach the region O. But, since the solution to (P) is exactly known in this region, it follows that the cost–to–go, $g(\underline{x})$, is a Lyapunov function for the system in O. Thus the closed–loop system resulting from the application of the feedback control law $\underline{u}_k = M(\underline{x}_k)$ is asymptotically stable. This result is summarized in the next theorem:

- **Theorem 3:** [12] The closed loop system resulting from the application of algorithm M to problem (P') is asymptotically stable, *provided* that there is enough computational power available to solve a problem of the form (2) with $m = 1$ during the sampling interval.

Note that algorithm M is a "conceptual" algorithm, in the sense that it can not be implemented until a finite procedure to perform the optimization required by step 1-ii is specified. In our previous work [11–12] we solved this optimization problem by partitioning the control space Ω into a finite set Ω_s. The attainable domain from the initial condition, using controls in Ω_s, can be represented now as a tree with each node corresponding to one of the attainable states. Hence the original optimal control problem is recast as a tree search, with the approximation resulting from the control quantization. The resulting tree can be scanned efficiently for minimum cost paths using heuristic search techniques, based upon an underestimate of the cost-to-go [13].

In the remainder of this section, we summarize, for completeness, the features of the algorithm (see [12] for a complete description) and we illustrate, through the use of a realistic example, potential difficulties with this approach. We begin by showing that if (P) is feasible, then (P') is also feasible, even when the controls are restricted to an appropriate quantization Ω_s of Ω.

- **Theorem 4:** [11] Assume that the system (S) satisfies the condition (4). Let:

$$\Lambda = \min_{\underline{x}\in\mathcal{G}-O}\left\{\lambda : \frac{\underline{x}}{\lambda} \in \partial\mathcal{G}\right\}$$
$$\mathcal{S} = \left\{ s \in R : \max_{\underline{u}\in B_s} \|B\underline{u}\|_\infty \le \Lambda \min_{\underline{y}\in\partial\mathcal{G}} \left\{1 - \min_{\underline{u}\in\Omega} \|A\underline{y} + B\underline{u}\|_{\mathcal{G}}\right\}\right\}$$
$$B_s = \left\{\underline{u} \in R^m : \|\underline{u}\|_\infty \le \frac{1}{s}\right\} \tag{6}$$

where $\partial\mathcal{G}$ denotes the boundary of $\mathcal{G}$. Finally, let $s_o = \min\{s \in \mathcal{S}\}$. Then the constraint (5) can be satisfied for *all* $\underline{x} \in \mathcal{G} - O$ with $\underline{u}$ restricted to *any* quantization Ω_s of Ω with $s \ge s_o$

Remark: Note that condition (4) guarantees that $\mathcal{S} \ne \emptyset$. Moreover, $\mathcal{S}$ is closed, bounded below. Hence s_o is well defined.

From Theorem 4 it follows that (P') can be solved by using quantized controls in the region $\mathcal{G}-O$ and switching to non-quantized controls inside O. However, selecting the norm of the partition to satisfy equation (6) may prove overly conservative leading to a very large tree. This phenomenon is illustrated in the following example:

3.4 A Realistic Problem.

Consider the minimum time control of an F-100 jet engine. The system at sea level static and PLA = 83° can be represented by:

$$A = \begin{pmatrix} 0.8907 & 0.0474 & -0.0980 & 0.2616 & 0.0689 \\ 0.0237 & 0.9022 & -0.0202 & 0.1057 & 0.0311 \\ 0.0233 & -0.0149 & 0.8167 & 0.2255 & 0.0295 \\ 0.0 & 0.0 & 0.0 & 0.7788 & 0.0 \\ -0.0979 & 0.3532 & 0.3662 & 0.6489 & 0.0295 \end{pmatrix} \quad \underline{\gamma} = \begin{pmatrix} 50.0 \\ 64.0 \\ 20.0 \\ 5.0 \\ 18.1 \end{pmatrix}$$

$$B = \begin{pmatrix} 0.0213 & -0.3704 \\ 0.0731 & -0.1973 \\ -0.0357 & -0.5538 \\ 0.2212 & 0.0 \\ 0.0527 & -3.9068 \end{pmatrix} \tag{7}$$

$$G = I; \Omega = \{\underline{u} \in R^2 : |u_1| \le 31.0; |u_2| \le 200.0\}$$

The sampling time for this system is 25 msec. In this case, condition (6) yields $\sim 10^3$ points for the quantization Ω_s, which is clearly impractical.

IV. Proposed Control Algorithm

As we illustrated with the example of the last section, even though the assumptions of Theorem 4 are not very restrictive, selecting the size of the quantization from (6) may result in an extremely large number of possible control actions to be considered. In this section we indicate how the special structure of time–optimal systems can be used to eliminate most of these candidate control actions. Specifically, we use a *modification of the Discrete Time Minimum Principle to show* that the points that satisfy a *necessary* condition for optimality are the *corners* of a subset of Ω. Hence, only these points need to be considered by the optimization algorithm.

4.1 The Modified Discrete Time Minimum Principle.

- **Theorem 5:** Consider the problem (P'') defined as:

$$\min_{\substack{\underline{x}_k\in\mathcal{G} \\ \underline{u}_k\in\Omega}} S(\underline{x}_N) \tag{8}$$

subject to:

$$x_{k+1} = f(x_k, u_k), \qquad x_o, N \text{ given} \tag{9}$$
$$\|\underline{x}_{k+1}\|_{\mathcal{G}} < \|\underline{x}_k\|_{\mathcal{G}} \tag{10}$$

where $\mathcal{G}$ and Ω are the compact, convex polyhedrons defined in (C), and where S is Frechet differentiable. Define

$$F_k = \frac{\partial f(\underline{x}_k, \underline{u}_k)}{\partial \underline{x}'} \tag{11}$$

Let the co–states ψ_k be defined by the difference equation:

$$\psi'_k = \psi'_{k+1} F_{k+1}$$
$$\psi'_{N-1} = \frac{\partial S(\underline{x}_N)}{\partial \underline{x}'} \tag{12}$$

Finally, define the Hamiltonian as:

$$H(\underline{x}_k, \underline{u}_k, \psi_k) = \psi'_k f(\underline{x}_k, \underline{u}_k) \tag{13}$$

Then, a necessary condition for optimality is:

$$H(\underline{x}^*_k, \underline{u}^*_k, \psi^*_k) = \min_{\underline{u}\in O\subseteq\Omega_1} H(\underline{x}^*_k, \underline{u}, \psi^*_k),\ k = 1, \ldots, N \tag{14}$$

where

$$\Omega_1(\underline{x}_k, k) = \{\underline{u} \in \Omega : \|\underline{x}_{k+1}\|_{\mathcal{G}} \leq (1-\epsilon)\|\underline{x}_k\|_{\mathcal{G}}\} \tag{15}$$

where O is some neighborhood of $\underline{u}$ and where $\epsilon > 0$ is chosen such that Ω_1 is not empty.

Proof: From Theorem 2 it follows that there exist $\epsilon > 0$ such that Ω_1 is not empty. From the definition of Ω_1 it follows that for any $\underline{u}_k \in \Omega_1(\underline{x}_k, k)$ there exist a neighborhood $O \subseteq \Omega_1$, not necessarily open, where (10) holds. Hence, if $\underline{x}_k \in \mathcal{G}$, $\underline{x}_{k+1} = f(\underline{x}_k, \underline{u}_k) \in \mathcal{G} \; \forall \, \underline{u}_k \in O$. We complete the proof using induction:

a:) Let $\tilde{\underline{x}}_k$ denote a non-optimal feasible trajectory obtained by employing the non-optimal control law $\tilde{\underline{u}}_{N-1}$ at stage $k = N-1$. Consider a neighborhood $O \subseteq \Omega$ of $\underline{u}^*_{N-1}$ such that the state constraints are satisfied for *all* the trajectories generated employing controls in O. For any such trajectory $\tilde{\underline{x}}, \tilde{\underline{u}}$ we have:

$$S(\underline{x}^*_N) \leq S(\tilde{\underline{x}}_N) \tag{16}$$

or

$$S(\tilde{\underline{x}}_N) - S(\underline{x}^*_N) \geq 0 \tag{17}$$

Hence:

$$\frac{\partial S}{\partial \underline{x}'}\Big|_{\underline{x}^*_N} \Delta \underline{x}_N = \frac{\partial S}{\partial \underline{x}'}\Big|_{\underline{x}^*_N} (f(\underline{x}^*_{N-1}, \tilde{\underline{u}}_{N-1}) - f(\underline{x}^*_{N-1}, \underline{u}^*_{N-1})) \geq 0 \tag{18}$$

$$\begin{aligned} H(\underline{x}^*_{N-1}, \tilde{\underline{u}}_{N-1}, \psi^*_{N-1}) &= \frac{\partial S}{\partial \underline{x}'}\Big|_{\underline{x}^*_N} f(\underline{x}^*_{N-1}, \tilde{\underline{u}}_{N-1}) \\ &\geq \frac{\partial S}{\partial \underline{x}'}\Big|_{\underline{x}^*_N} f(\underline{x}^*_{N-1}, \underline{u}^*_{N-1}) \\ &= H(\underline{x}^*_{N-1}, \underline{u}^*_{N-1}, \psi^*_{N-1}) \end{aligned} \tag{19}$$

b:) Consider a neighborhood $O \subseteq \Omega$ of $\underline{u}^*_k$ such that *all* the trajectories obtained by replacing $\underline{u}^*_k$ by any other control in O satisfy the constraints and assume that (14) does not hold for some $k < N-1$. Then, there exists at least one trajectory $\tilde{\underline{x}}, \tilde{\underline{u}}$ such that:

$$H(\underline{x}^*_k, \tilde{\underline{u}}_k, \psi^*_k) < H(\underline{x}^*_k, \underline{u}^*_k, \psi^*_k) \tag{20}$$

Therefore:

$$0 > \psi^{*'}_k (f(\underline{x}^*_k, \tilde{\underline{u}}_k) - f(\underline{x}^*_k, \underline{u}^*_k)) = \psi^{*'}_k \Delta \underline{x}_{k+1} \tag{21}$$

Hence:

$$\begin{aligned} H(\tilde{\underline{x}}_{k+1}, \underline{u}^*_{k+1}, \psi^*_{k+1}) - H(\underline{x}^*_{k+1}, \underline{u}^*_{k+1}, \psi^*_{k+1}) &= \\ \psi^{*'}_{k+1} f(\tilde{\underline{x}}_{k+1}, \underline{u}^*_{k+1}) &- \psi^{*'}_{k+1} f(\underline{x}^*_{k+1}, \underline{u}^*_{k+1}) \\ &= \psi^{*'}_{k+1} \left(\frac{\partial f}{\partial \underline{x}'}\Big|_{\underline{x}^*_{k+1}}\right) \Delta \underline{x}_{k+1} \\ &= \psi^{*'}_k \Delta \underline{x}_{k+1} < 0 \end{aligned} \tag{22}$$

or:

$$H(\tilde{\underline{x}}_{k+1}, \underline{u}^*_{k+1}, \psi^*_{k+1}) < H(\underline{x}^*_{k+1}, \underline{u}^*_{k+1}, \psi^*_{k+1}) \tag{23}$$

Using the same reasoning we have:

$$\begin{aligned} H(\tilde{\underline{x}}_{k+2}, \underline{u}^*_{k+2}, \psi^*_{k+2}) &< H(\underline{x}^*_{k+2}, \underline{u}^*_{k+2}, \psi^*_{k+2}) \\ &\vdots \\ H(\tilde{\underline{x}}_{N-1}, \underline{u}^*_{N+1}, \psi^*_{N+1}) &< H(\underline{x}^*_{N+1}, \underline{u}^*_{N+1}, \psi^*_{N+1}) \end{aligned} \tag{24}$$

From (24) it follows that

$$0 > \psi^{*'}_{N-1} \Delta \underline{x}_N = \frac{\partial S}{\partial \underline{x}'}\Big|_{\underline{x}^*_N} \Delta \underline{x}_N = \Delta S \tag{25}$$

against the hypothesis that $S(\underline{x}^*_N)$ was a *local* minimum $\diamond$.

4.2 Using the Modified Discrete Minimum Principle.

In this section we indicate how to use the results of Theorem 5 to generate a set of points that satisfy the necessary conditions for optimality. In principle, we could apply the discrete minimum principle to problem (P') by taking $S(\underline{x}_N) = \|\underline{x}_N\|_2^2$ and solving a sequence of problems, with increasing N, until a trajectory $\underline{x}^*$ and a number N_o such that $\underline{x}^*_{N_o} = 0$ are found. However note that Theorem 5 does not add any information to the problem since:

$$\psi^{*'}_{N-1} = \frac{\partial S(\underline{x}^*_N)}{\partial \underline{x}'}\Big|_0 = \underline{x}\Big|_0 = 0 \tag{26}$$

It follows that $\psi_k = 0 \; \forall k$, and hence the optimal trajectory corresponds to a "singular arc". Therefore, nothing can be inferred a priori about the controls. In order to be able to use the special structure of the problem, we would like the co-states, ψ, to be non-zero.

Consider now the special case of problem (P'') where $S(\underline{x}_N) = \frac{1}{2}\|\underline{x}_N\|_2^2$ (with fixed terminal time N). Let n be the dimension of the system (S) and assume that the initial condition $\underline{x}_o$ is such that the origin *can not* be reached in N stages. Then, it follows from (12) that:

$$\begin{aligned} \psi^{*'}_k &= \psi^*_{k+1}{}' A \\ \psi^*_{N-1}{}' &= \underline{x}_N \neq \underline{0} \end{aligned} \tag{27}$$

It follows (since A was assumed regular) that $\psi^*_k \neq 0 \; \forall k$. Furthermore, since (S) is controllable, C_n has dimension n [11]. It follows that, by taking N large enough will have $\underline{x}_N \in C_n$. Hence and approximate solution to (P') can be found by solving (P'') for N such that $\underline{x}_N \in C_n$ and by using Linear Programming to find the optimal trajectory from $\underline{x}_N$ to the origin. This idea is the basis of the proposed algorithm.

- **Theorem 6:** The optimal control sequence $\mathcal{U} = \{\underline{u}^*_0 \cdots \underline{u}^*_{N-1}\}$ that solves problem (P'') for the case of linear time invariant dynamics and linear inequality constraints, is always in the boundary of the set Ω_1. Further, the control sequence can always be selected to be a corner point of such set.

  Proof: Since the constraints are linear and $\psi^*_k \neq \underline{0}$, it follows that the control $\underline{u}^*_k$ that solves (14) belongs to the boundary of the set $\Omega_1(\underline{x}_k, k)$. Further, except in the case of degeneracies, the *control* $\underline{u}^*_k$ *must* be a corner point of the set. In the case of degeneracies, *all* the points of the boundary parallel to the co-state yield the same value of the Hamiltonian and therefore the optimal control $\underline{u}_k$ can be selected to be a corner of Ω_1.

4.3 Algorithm H_{MP}.

In this section we apply the results of Theorem 6 to obtain a suboptimal stabilizing feedback control law. From Theorem 6 it follows that problem (P'') can be solved by using the following algorithm, a modification of the Heuristically Enhanced Control idea proposed in [11-12]:

Algorithm H_{MP} (Heuristically Enhanced Control using the minimum principle)

Begin.

1) Determine ϵ for equation (15). Note that since the constraints are assumed to be linear, the maximum value of ϵ can be determined off-line using Linear Programming. Let $O = C_1$, null controllability region in 1 step, and determine an underestimate $g(.)$ relative to O.

2) Let $\underline{x}_k$ be the current state of the system:

   2.1) If $\underline{x}_k \in C_n$, null controllability region in n steps, solve problem (P) exactly using Linear Programming.

2.2) If $\underline{x}_k \notin C_n$ generate a tree by considering all possible trajectories starting at $\underline{x}_k$ with controls that lie in the corners of the polytope $\Omega_1(\underline{x}_k, k)$. Search the tree for a minimum cost trajectory to the origin, using heuristic search algorithms and $g(.)$ as heuristics.

2.3) If there is no more computation time available for searching and the region O has not been reached, use the minimum partial cost trajectory that has been found.

3) Repeat step 2 until the region the origin is reached.

End.

Remark: Note that by solving problem (P'') instead of (P) we are relinquishing optimality, since the trajectory that brings the system closer to C_n is *not* necessarily the trajectory that will yield minimum transit time to the origin. However, for any "reasonable" problem, we would expect both trajectories to be close in the sense of yielding approximately equal transit times (in the next section we will provide an example where this expectation is met).

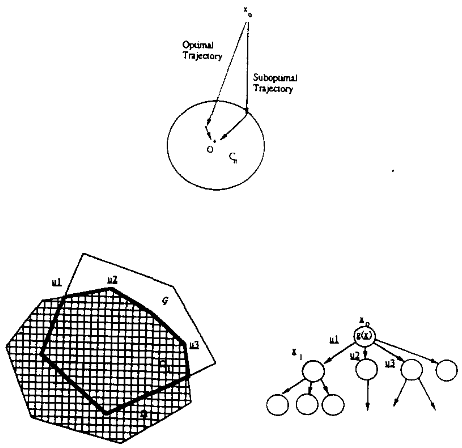

Fig 1. Using the Discrete Minimum Principle to Limit the Search

4.4 The Heuristic for Algorithm H_{MP}.

In order to complete the description of algorithm H_{MP} we need to provide a suitable underestimate $g(\underline{x})$. In principle, an estimate of the number of stages necessary to reach the origin can be found based upon the singular value decomposition of the matrices A and B, using the same technique that we used in [12]. However, in many cases of practical interest such as the F-100 jet engine of section 3.4, the limitation in the problem is essentially given by the state constraints (i.e. the control authority is large). In this situation, this estimate yields an unrealistically low value for the transit time, resulting in poor performance.

The performance of the algorithm can be improved substantially by considering an heuristic based upon experimental results. Recall that optimality depends on having, at each time interval, an underestimate $g(\underline{x})$ of the cost-to-go. Consider now the Null Controllability regions (C_k). It is clear that if they can be found and stored, the true transit time to the origin is known. If the regions are not known but a supraestimate C_k^s such that $C_k \subseteq C_k^s$ is available, a suitable underestimate $g(\underline{x})$, can be obtained by finding the largest k such that $\underline{x} \in C_k^s$ and $\underline{x} \notin C_{k-1}^s$. However, in general these supraestimates are difficult to find and characterize [11]. Hence, it is desirable to use a different heuristic, which does not require the use of these regions. From the convexity of Ω and $\mathcal{G}$ it follows that the regions C_k are convex. Therefore, a *subestimate* C_{sk} such that $C_{sk} \subseteq C_k$ can be found by finding points in the region C_k and taking C_{sk} as their convex hull. Once a subestimate of C_k is available, an estimate $\tilde{g}(\underline{x})$ of the cost-to-go can be found by finding the largest k such that $\underline{x} \in C_{sk}$ and $\underline{x} \notin C_{sk-1}$. Note that this estimate *is not* an underestimate in the sense of Def. 4. Since $C_{sk} \subseteq C_k$ then $\underline{x}_k \in C_k \not\Rightarrow \underline{x}_k \in C_{sk}$ and therefore $\tilde{g}(\underline{x}_k)$ is *not* necessarily $\leq k$. Thus, Theorem 1 that guarantees that once the set O has been reached the true optimal trajectory has been found is no longer valid. Nevertheless, we expect that if enough points of each region are considered so that the subestimates are close to the true Null Controllability regions, then the control law generated by algorithm H_{MP} should be close to the true optimal control. This idea is illustrated in the following section.

4.5 Application to the Realistic Example.

In this section we will argue the soundness of the approximation introduced in the last section by comparing the results of applying algorithm H_{MP}, with the heuristic resulting from the use of subestimates of the controllability regions, against the true optimal controller for the problem of section III.

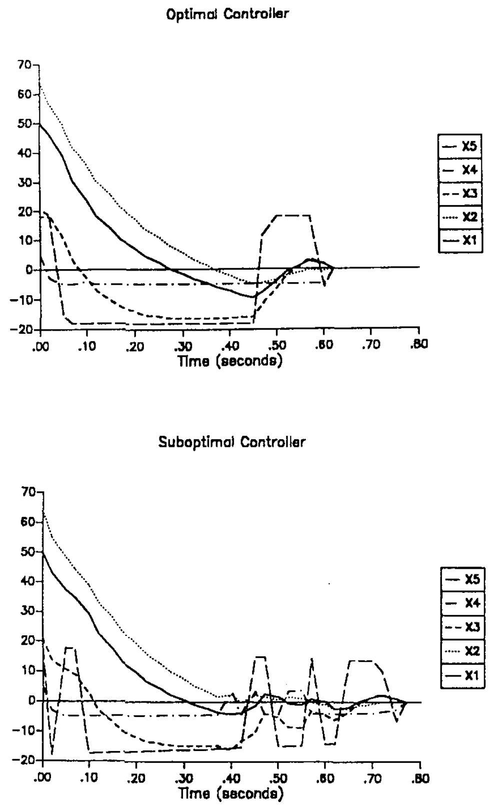

Fig. 2. Optimal Control vs. Algorithm H_{MP} for the Example of Section 3.4

Figure 2 shows a comparison between the trajectories for the optimal control law and algorithm H_{MP}. In this particular case, the optimal control law was computed off-line by solving a sequence of linear programming problems, while algorithm H_{MP} was limited to computation time compatible with an on-line implementation. The value of ϵ was set to 0.01 (using linear programming it was determined that the maximum value of ϵ compatible with the constraints is 0.025), and each of the regions C_{sk} was found as the convex hull of 32 points, using optimal trajectories generated off-line. Note that in spite of being limited to running time roughly two orders of magnitude smaller than the computation time used off-line to find the true optimal control solution, algorithm H_{MP} generates a solution that takes only 25% more time to get to the origin (25 vs. 31 stages).

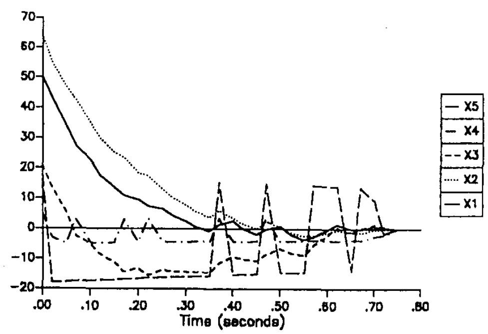

Fig. 3. Algorithm H_{MP} with perfect information

Figure 3 shows the results of applying algorithm H_{MP} when the heuristic is perfect (i.e the exact transit time to the origin is known). By comparing figures 3 and 4 we see that most of the additional cost comes from the approximation made in Theorem 6, while the use of an estimate of the cost-to-go based upon the subestimates C_{sk} (rather than a "true" underestimate as required by Theorem 1) adds only 1 stage to the total transit time.

V. Conclusions

Most realistic control problems involve some types of constraints. However, although there presently exist efficient techniques for addressing frequency domain constraints, satisfactory techniques for systematically handling time–domain constraints have started to appear recently. Following the idea presented in [11-12], in this paper we propose to address time–domain constraints by using a feedback controller based upon the on–line use of a dynamic–programming approach to solve a constrained optimization problem. Theoretical results are presented showing that, with the addition of a rather modest condition, this controller yields asymptotically stable systems, *provided* that the solution to a optimization problem, considerably simpler than the original, can be computed in real-time.

In the first part of the paper we presented a realistic example, illustrating potential difficulties with dynamic- programming approaches. These difficulties can be circumvented by applying a suitably modified discrete time minimum principle, which allows for checking only the vertices of a polytope in control space. The proposed approach results in a substantial reduction of the dimensionality of the problem (two orders of magnitude for the case of the example presented in section 3.4). Hence, the proposed algorithm presents a significant advantage over previous approaches that use the same idea, specially for cases, such as Example 3.4, where the time available for computations is very limited.

In the second part of the paper, we addressed the problem of finding an underestimate of the cost, required by the algorithm, for the case where the limitations in the problem come essentially from the state constraints. We showed that by relinquishing theoretical optimality, an heuristic estimate can be found based upon experimental data obtained off-line. Through simulation results we showed that even though theoretically we are giving up optimality, the behavior of the system is practically unaltered by this choice. Furthermore, the closed loop system obtained by combining this estimate with the approximation of the first part, presents very good performance when compared to the true optimal solution found using an off-line procedure.

We believe that the algorithm presented in this paper shows great promise, especially for cases where the dimension of the system is not small. Note however, that the algorithm requires the real time solution of two non-trivial computational geometry problems in R^n: determining the inclusion of a point in a convex hull and finding all the vertices of a polytope. Recent work on trainable non-linear classifiers such as artificial neural nets and decision trees may prove valuable in solving the first problem.

Perhaps the most serious limitation to the theory in its present form arises from the implicit assumption that the model of the system is perfectly known. Since most realistic problems involve some degree of uncertainty, clearly this assumption limits the domain of application of the proposed controller. We are currently working on a technique, patterned along the lines of the Norm Based Robust Control framework introduced by Sznaier [6], to guarantee robustness margins for the resulting closed–loop system. A future paper is planned to report these results.

References.

[1]. Boyd S. et. al., "A New CAD Method and Associated Architectures for Linear Controllers," *IEEE Trans. Autom. Contr.*, Vol 33, 3, pp. 268–283, 1988.

[2]. Frankena J. F. and R. Sivan, "A Nonlinear Optimal Control Law for Linear Systems," *Int. J. Contr.*, Vol. 30, 1, pp 159-178, 1979.

[3]. Gutman P. O. and P. Hagander, " A New Design of Constrained Controllers for Linear Systems," *IEEE Trans. Automat. Contr.*, Vol. 30, pp. 22-33, 1985.

[4]. Vassilaki M. et. al., "Feedback Control of Discrete-Time Systems Under State and Control Constraints," *Int. J. Control*, Vol. 47, 6, pp 1727-1735, 1988.

[5]. Benzaouia A. and C. Burgat, "Regulator Problems for Linear–Discrete Time Systems with Non–Symmetrical Constrained Control," *Int. J. Control*, Vol. 48, 6, pp. 2441–2451, 1988.

[6]. Sznaier M., "Norm Based Robust Control of Constrained Discrete–Time Linear Systems," *Proc. of the 29th IEEE CDC*, Hawaii, Dec 5–7, pp. 1925–1930, 1990

[7]. Sznaier M. and A. Sideris, "Norm Based Optimally Robust Control of Constrained Discrete–Time Linear Systems", *Proc. 1991 ACC.*, Boston, Massachusetts, June 26–28, pp. 2710–2715, 1991.

[8]. Sznaier M. and A. Sideris, "Norm Based Robust Dynamic Feedback Control of Constrained Systems", *Proc of the First IFAC Symposium on Design Methods for Linear Systems*, Zurich, September 4–6, 1991, to appear.

[9]. Sznaier M., "Set Induced Norm Based Robust Control Techniques," to appear in *Advances in Control and Dynamic Systems*, C. T. Leondes Editor, Academic Press, 1992.

[10]. Dreyfus, S. E., "Some Types of Optimal Control of Stochastic Systems," *J. SIAM Control A*, Vol 2, 1, pp. 120–134, 1964.

[11]. Sznaier M., "Suboptimal Feedback Control of Constrained Linear Systems," *Ph.D. Dissertation*, University of Washington, 1989.

[12]. Sznaier M. and M. J. Damborg, "Heuristically Enhanced Feedback Control of Constrained Discrete–Time Linear Systems," *Automatica*, Vol 26, 3, pp. 521–532, 1990.

[13]. Winston P. H., *Artificial Intelligence*, 2nd edition, Addison-Wesley, Massachusetts, Chapter 4, 1984.