Proceedings of the 45th IEEE Conference on Decision & Control
Manchester Grand Hyatt Hotel
San Diego, CA, USA, December 13-15, 2006

Simultaneous Identification and Control of Time–Varying Systems

Matthias J. Feiler[◇] and Kumpati S. Narendra[*]

Abstract— **The paper addresses two complementary problems. The first one is to derive a rigorous formulation of adaptive control when the parameters of the plant are time-varying. To make the problem tractable, the time–variation is approximated by a piecewise–constant function assuming a finite number N of unknown values. The instants at which the parameters switch are also unknown. A proof of stability is presented assuming that an algorithm for simultaneous identification of N values exists. In the second part of the paper, an algorithm is designed which causes N adaptive models to converge to N different locations in parameter space in a self–organizing fashion. Work has been in progress but the proof of convergence has defied solution for over three years. The paper discusses the theoretical difficulties encountered, solutions obtained in special cases, new mathematical tools that are being investigated, and potential applications of the approach.**

I. Introduction

It is well known that while the main reason for using adaptive control in practical applications is to compensate for time–variations of the parameters, a coherent theory exists only when the plant is time–invariant. The generally accepted philosophy is that if an adaptive system is fast and accurate when the plant parameters are constant but unknown, it would also prove satisfactory when the parameters vary with time, provided the latter occurred on a relatively slower time–scale.

This paper deals with the case of unknown, rapidly time–varying parameters. It is shown how the time–varying problem can be converted into a set of $N \in \mathbb{N}$ time–invariant problems which are solved simultaneously. The framework within which this is possible is the one of multiple models, switching and tuning [1]. The solution presented is an example of how a nonautonomous problem can be transformed in an autonomous one by enlarging the state space of the system. This is in line with well–known techniques for eliminating harmonic time–varying disturbances by incorporating an internal model into the controller. However, since the time–variation is assumed to be arbitrary, the procedure is substantially more complex.

[◇]Chair of Systems Design, ETH Zurich, CH-8092 Zurich, Switzerland. mailto: mfeiler@ethz.ch; web: http://www.sg.ethz.ch

[*]Center for Systems Science, Yale University, P.O. Box 208267, New Haven, CT 06520-8267 USA. kumpati.narendra@yale.edu

A. Motivation for the problem

We assume that the plant parameters $\breve{\theta}(k) : \mathbb{Z}^+ \to \mathbb{R}^{2n}$ evolve in an unknown, bounded region of the parameter space $\mathbb{R}^{2n}$ (n is the order of the system). $\breve{\theta}(k)$ is approximated by a piecewise constant function $\theta(k) : \mathbb{Z}^+ \to \mathbb{S}$ where $S \subset \mathbb{R}^{2n}$ is a finite set $S = \{\theta_1, \theta_2, \ldots, \theta_N\}$ of unknown, constant parameter vectors. It is clear that as $N \to \infty$ the approximation error tends to zero. Let us fix some finite value $N \in \mathbb{N}$. If the norm of the error is larger than a prescribed value, N is increased at the discretion of the designer. For our purposes, we assume that an appropriate N is known. Furthermore, only those values are approximated that the function $\breve{\theta}(k)$ assumes repeatedly, i.e. all transient errors will be tolerated.

B. The Problem

From these practical considerations the following assumptions about the plant P are abstracted:

(A1) P switches randomly between N unknown values according to an ergodic Markov chain.
(A2) P is constant over an interval T of minimum length $d+1$ where d is the relative degree of the system.
(A3) The plant has asymptotically stable zero dynamics. This is a standard prerequisite for a solution to the inverse control problem to exist.

Given the linear time–varying system

$$y(k) = \sum_{\nu=d}^{d+n-1} \Big[\, a_\nu(k-\delta)\, y(k-\nu) + b_\nu(k-\delta)\, u(k-\nu) \,\Big] \tag{1}$$

we define the plant parameter vector by

$$\theta(k) = [a_0(k), \ldots, a_{n-1}(k), b_0(k), \ldots, b_{n-1}(k)]^T$$

and the regression vector by

$$\phi(k) = [y(k), \ldots, y(k-n+1), u(k), \ldots, u(k-n+1)]^T$$

and write (1) in the following compact form:

$$y(k) = \phi(k-d)^T \theta(k-\delta) \tag{2}$$

A time–variation of the plant parameters affects the plant output after δ instants of time whereas the input affects the output after d instants of time, where d is the relative degree of the system. We define N adaptive models of the form

$$\hat{y}_i(k) = \phi(k-d)^T \hat{\theta}_i(k-1) \qquad i = 1 \dots N \tag{3}$$

where $\hat{\theta}_i(k-1)$ and $\hat{y}_i(k)$ are estimates of $\theta \in S$ and $y(k)$ respectively. The parameters are updated using an algorithm of the form

$$\hat{\theta}_i(k) = \hat{\theta}_i(k-1) + g(k)\phi(k-d)e_i(k) \tag{4}$$

where $e_i(k) = y(k) - \hat{y}_i(k)$ is the estimation error which satisfies

$$e_i(k) = \phi(k-d)^T [\theta(k-\delta) - \hat{\theta}_i(k-1)] \tag{5}$$

The problem is to find an adaptive gain function $g(k)$ such that every one of the models $\hat{\theta}_i(k)$, $i = 1, \dots, N$ converges to a different element of the set S as $k \to \infty$.

(B) *Problem Statement*: Define $g(k)$ such that for every $j \in \Omega$ there is a $i \in \Omega$ such that $\|\hat{\theta}_i(k) - \theta_j\| \to 0$ as $k \to \infty$.

In other words, each model identifies exactly one element of the set S as the plant switches among all the elements of S. Each parameter vector $\hat{\theta}_i(k)$, $i = 1 \dots N$ has to be adapted in such a way that it converges to one of the $\theta \in S$ or (in the absence of a persistently exciting input) to the solution manifold corresponding to θ_i. If for every $\theta \in S$ there exists a model $\hat{\theta}_i$ that converges to it and the control is based at every instant on the correct model, then the adaptive control problem is no different from the time–invariant case. This is due to the fact that if the switching from one model to another occurs synchronously with the time–variation of the environment, then —from the viewpoint of one model— the environment is constant. This is referred to as recasting the time-varying problem as N time–invariant ones.

The term "correct" used above refers to the fact that if an element θ_j of S is estimated by a particular model $\hat{\theta}_i$ (i.e. we ultimately have that $\hat{\theta}_i \to \theta_j$) then the certainty equivalence control input must be computed from $\hat{\theta}_i$, whenever θ_j is in existence. At every instant the model that best approximates the plant according to one of the following criteria is chosen for control.

$$\begin{array}{lll} \text{(i)} & J_i(k) & = e_i^2(k) \\ \text{(ii)} & J_i(k) & = \sum_{\nu=0}^{k} e_i^2(\nu) \\ \text{(iii)} & J_i(k) & = \alpha\, e_i^2(k) + \beta \sum_{\nu=0}^{k} e_i^2(\nu) \\ \text{(iv)} & J_i(t) & = \alpha\, e_i^2(k) + \beta \sum_{\nu=0}^{k} \rho^{k-\nu} e_i^2(\nu) \end{array}$$

II. Analysis

We now take a closer look at the instants at which the controller switches from one model to another. It is useful to define the a posteriori identification error as

$$e_p(k) \;=\; \phi(k-d)^T [\theta(k-\delta) - \hat{\theta}_{i^*}(k)] \tag{6}$$

The function $i(k)^* : \mathbb{Z} \to \Omega$, $\Omega = \{1, \dots N\}$ yields the index of the best model according to the criterion. The qualifier "a posteriori" refers to the fact that the switching to the best model occurs *after* the performance of all the models has been assessed.

Suppose that all elements of the set S have been identified correctly using multiple models, i.e. $\hat{\theta}_i \equiv \theta_j$ for N *different* pairs $(i, j) \in \Omega \times \Omega$. Criterion (i) can be used in this case in order to decide which model best approximates the plant. We obtain that

$$e_p(k) \equiv 0 \qquad \text{for all } k > 0 \tag{7}$$

This can be seen as follows. Denote the instant of time at which the plant parameters switch from some arbitrary θ_{i_1} to θ_{i_2}, by s. In view of equation (2), this change of the environment affects the output only δ instants of time later. At instant of time $k = s + \delta$ the (a priori) identification error of the ith model is equal to

$$\begin{array}{rcl} e_i(s+\delta) & = & \phi(s+\delta-d)^T [\theta(s) - \hat{\theta}_i(s+\delta-1)] \\ & = & \phi(k+\delta-d)^T [\theta_{i_2} - \theta_{i_1}] \\ & \neq & 0 \end{array} \tag{8}$$

since $\theta(s) = \theta_{i_2}$ whereas $\hat{\theta}_i(s+\delta-1) = \theta_{i_1}$. This error is used to detect the time–variation of the plant. At the same instant $k = s + \delta$ a new model can be selected, to obtain $\hat{\theta}(s+\delta) = \theta_{i_2}$. It follows that the a posteriori error satisfies $e_p(s+\delta) = 0$. But this means that $e_p(k) \equiv 0$, $\forall\ k > 0$ in the exact matching case.

It is worth emphasizing that the instant s at which a time–variation occurs and the instant at which it can be detected at the output differ by δ steps. Let us now examine e_c. As in standard adaptive control we have the relationship between the identification error $e(k)$ and the control error $e_c(k)$

$$e_c(k) = e_p(k) + \phi(k-d)^T [\hat{\theta}(k) - \hat{\theta}(k-d)] \tag{9}$$

Due the detection problem, $\hat{\theta}(k)$ switches to the new model only at instant $k = s+\delta$. In view of equation (9), this means that $e_c(k) \neq 0$ for $k = s+\delta, \ldots, s+\delta+d-1$. More precisely, it satisfies

$$e_c(k) = \phi(k-d)^T[\theta_{i_2} - \theta_{i_1}] \tag{10}$$

for $k \in I_c = [s+\delta, \ldots, s+\delta+d-1]$. The interval I_c is referred to as the inherent control error interval, which is due to relative degree of the plant.

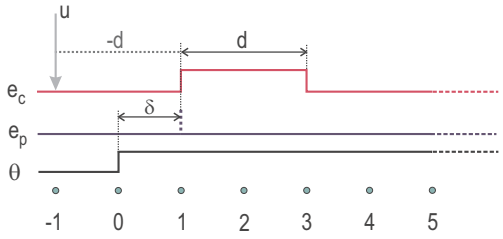

Fig. 1. Binary representation of the inherent control error

Figure (1) illustrates the sequence of events causing the control error e_c to be nonzero of the interval I_c. To avoid the error, $\theta(k)$ would have to be known, i.e. the switching instant s and the "new" environment $\theta(s) = \theta_{i_2}$ must be known for every s. We then have to apply the new control input (i.e. corresponding to θ_{i_2}) d instants before the change affects the output of the plant, i.e. at instant of time $k = s+\delta-d$ to cause $e_c(k) \equiv 0$, $k > 0$.

III. Proof of Stability

Solving the simultaneous identification problem (B) is equivalent to solving the adaptive control problem in the presence of large and rapid time–variations. This is demonstrated in the following paragraph.

Theorem 1: Assume that for any $j \in \Omega$ there is a $i \in \Omega$ such that $\|\hat{\theta}_i(k) - \theta_j\| \to 0$ as $k \to \infty$. Further assume that $T \geq d+1$. Then the certainty equivalence controller obtained by solving

$$y^*(k) = \phi(k-d)^T\hat{\theta}_i(k-d) \tag{11}$$

asymptotically stabilizes the time–varying system (2) in the sense that $\|\phi(k)\| < \infty$ for all $k > 0$.

Proof:

From the properties of the parameter adaptation algorithm (4) we obtain the following well–known inequality (see [2]).

$$\frac{|e_c(k)|}{[1+\phi(k-d)^T\phi(k-d)]^{\frac{1}{2}}} \leq \frac{|e_i(k)|}{[1+\phi(k-d)^T\phi(k-d)]^{\frac{1}{2}}} \\ \ldots + \frac{\|\phi(k-d)\|\|\hat{\theta}(k-1) - \hat{\theta}(k-d)\|}{[1+\phi(k-d)^T\phi(k-d)]^{\frac{1}{2}}} \tag{12}$$

The term on the left hand side corresponds to the "normalized" control error which we denote by $\bar{e}_c(k)$ for future reference. Similarly, $\bar{e}_i(k)$ denotes the normalized identification error due to the ith model, which we find as the first term on the right hand side of the above inequality. Also we introduce the normed distance between two parameter estimates, $\bar{\Delta}(k) = \|\hat{\theta}(k-1) - \hat{\theta}(k-d)\|$.

Step 1: Define a subsequence $\{k_n\}_{n\geq 0}$ with $k_n \to \infty$ as $n \to \infty$ where $k_n \in \mathbb{Z}^+ \setminus Q_1$ and $Q_1 = \{k \,|\, k = s+d\}$ for all switching instants $s > 0$. In view of the discussion in section II, Q_1 corresponds to the set of detection instants. We obtain that the ith identification tends to zero on a subsequence, i.e.

$$\bar{e}_i(k_n) \to 0 \quad \text{for all } i \in \Omega \text{ as } n \to \infty \tag{13}$$

Step 2: As in step 1, we define a subsequence $\{k_m\}_{m\geq 0}$ with $k_m \to \infty$ as $m \to \infty$ where $k_m \in \mathbb{Z}^+ \setminus Q_2$ and $Q_2 = \{k \,|\, s+\delta+1 \leq k \leq s+\delta+d-1\}$ for all $s > 0$. We obtain,

$$\bar{\Delta}(k_m) \to 0 \quad \text{as } m \to \infty \tag{14}$$

Notice that $\bar{\Delta}(k) = \|\hat{\theta}(k-1) - \hat{\theta}(k-d)\|$ cannot be zero when $k \in Q_2$ since $\hat{\theta}(k-1)$ and $\hat{\theta}(k-d)$ belong to different models during the instants following a switching (the one that currently best approximates the plant and the one that was used for control at instant $k-d$).

From step 1, 2 and inequality (12) it follows that

$$\bar{e}_c(k_l) = \frac{|e_c(k_l)|}{[1+\phi(k_l-d)^T\phi(k_l-d)]^{\frac{1}{2}}} \to 0 \quad \text{as } k_l \to \infty \tag{15}$$

where $\{k_l\}_{l\geq 0}$ is the subsequence tending to infinity as $l \to \infty$ and $k_l \in L = \mathbb{Z}^+ \setminus Q_1 \cup Q_2 = \mathbb{Z}^+ \setminus I_c$.

Every interval of length $T \geq d+1$ contains at least one instant $k = k_l \in L$, i.e. there exists a subsequence $\{k_l\}$ as defined above on which $\bar{e}_c(k_l) \to 0$. This implies that $\phi(k_l-d)$ grows faster than $e_c(k_l)$ if indeed both signals grow in an unbounded fashion. But, by assumption (A3), the input of the plant cannot grow faster than its output and hence $\phi(k)$ cannot grow faster than $e_c(k)$, not even

on a subsequence. Hence, the assumption that $\phi(k)$ grows without bound results in a contradiction. ■

The last part of the proof is analogous to the proof of stability in the time–invariant case. We obtain

Corollary 2:

$$e_c(k_l) \to 0 \text{ as } k_l \to \infty \tag{16}$$

i.e. the actual control error (not only the normalized one, cf. eq. (15)) tends to zero on a subsequence.

IV. Identification Algorithm

We now turn to the second problem addressed in this paper. An identification algorithm is defined by setting the adaptive gain $g(k)$ in equation (4) equal to

$$g(k) = \frac{\eta_i(k)}{\phi(k-d)^T \phi(k-d)} \tag{17}$$

While the projection vector $\phi(k-d)$ is the same for all models, its length varies according to the error $e_i(k)$ and step–size $\eta_i(k)$ associated with every model $\hat{\theta}_i(k)$, $i \in \Omega$. The step–size is a function of the performance of $\hat{\theta}_i(k)$ relative to all the other models:

$$\eta_i(k) = \eta_0 \frac{\frac{1}{e_i^2(k)}}{\frac{1}{e_1^2(k)} + \frac{1}{e_2^2(k)} + \cdots + \frac{1}{e_N^2(k)}} \qquad 0 < \eta_0 \leq 1 \tag{18}$$

As stated earlier, the performance index $J_i(k)$ is used to determine which model is closest to a given plant θ_j. In general, this cannot be concluded from the instantaneous error $e_i(k)$ as in equation (18), but requires observing a sequence of estimation errors $\{e_i(k), \ldots, e_i(k-T_0)\}$ over an interval of sufficient length T_0. This is a technical detail which is not developed here (see [3]). We simply assume that e_i^2 provides an unambiguous measure of the parameter error, i.e. the model with the smallest instantaneous error corresponds to the model that is closest to θ_j in parameter space. The closest model is updated by a large amount and, hence, its advantage over the other models consolidated. Notice that the nonlinear function (18) does not define a "winner takes it all"–policy since the *sum* of all step–sizes equals η_0, i.e. $\eta_l(k)$, $l \neq i$, is zero only in the limit when $e_i(k) \to 0$ as $k \to \infty$ (if indeed $\hat{\theta}_i$ converges to θ_j). The rule induces competition among the models for the largest adaptive gain. It causes the models to "scatter" in a self–organizing fashion, such that ultimately, there is only one model in the vicinity of every plant.

Extensive simulation studies have demonstrated that the algorithm successfully causes the N models to converge to the N different elements of the set S. In figure (2) we observe that $N = 5$ models converge to 5 different plant locations even as they are initialized far from the plants (and close to each other). As expected from our analysis in section II, this causes the control error $e_c = y - y^*$ to approach zero on a subsequence $k_l \in \mathbb{Z}^+ \setminus I_c$ where $I_c = \{s+2, s+3\}$ and s: switching instant, see figure (3). The lower part of the figure shows that the switching of the models also synchronizes with that of the environment (plants).

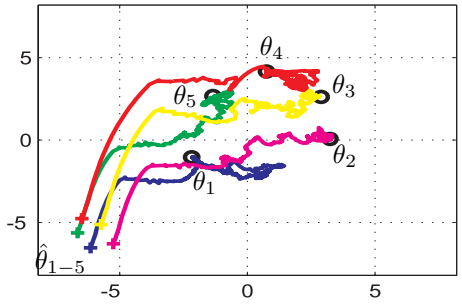

Fig. 2. Self–organization and convergence of $N = 5$ models

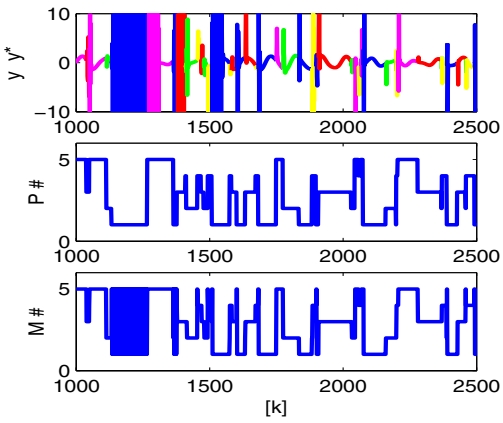

Fig. 3. Convergence of the control error of an unknown time–varying 2nd–order–system

Despite the efforts of the last three years, the algorithm is surprisingly resistent to a succinct mathematical analysis. While the low–order case ($N = 2$) displays relatively simple convergence properties, for $N \geq 3$, a very large number of special cases had to be analyzed separately.

A. Convergence Analysis

A first reduction of the problem is to assume that an algorithm can be determined by which the norms of all

the error vectors are reduced at every instant according to the equation

$$\tilde{\theta}_i(k+1) = \tilde{\theta}_i(k) - \eta_i(k)\tilde{\theta}_i(k) \qquad (19)$$

where the step–size $\eta_i(k)$ is defined in (18) with $e_i = \|\tilde{\theta}_i\|$. All the critical features of the algorithm are retained. In particular, $\tilde{\theta}_i$ is not used to update the parameters directly, i.e. by setting $\eta_i(k)$ to one. Instead, the self–organization of the models due to the time–varying $\eta_i(k)$ is studied. It is easy to see that the equilibrium set of the system is given by

$$E = \{e_{ij} \,|\, e_{f(i)\,j} = 0 \text{ whenever } f(i) = j, \text{ for all } i, j \in \Omega\}$$

where $f(\cdot) : \Omega \to \Omega$ is a rearrangement of the model indices such that every model has the same index as the plant to which it converges. Note that since *any* model may converge to *any* plant (depending on the initial conditions) there are $N!$ possible rearrangements of the model indices.

The key question in understanding the convergence properties of the SIC algorithm is how an elementary sorting operation contributes to global convergence. By an elementary sorting operation we mean an update action that increases the relative distances of all models towards one of the plants and, hence, the order between the models as seen from the perspective of that plant. Such a step follows a very simple rule, namely

$$\left[\frac{e_i}{e_j}\right]_{k+1} = \underbrace{\frac{1-\eta_i}{1-\eta_j}}_{<1} \left[\frac{e_i}{e_j}\right]_k \qquad e_i < e_j \qquad (20)$$

From the perspective of the other plants, however, the order may or may not increase depending on the location of these plants. This dependency on the geometry of the model–plant configuration is the essential difficulty in the analysis of the algorithm.

B. Numerical Experiments

Let us illustrate this point using the above example with $N = 3, n = 1$. The central question is whether, at each step, the model configuration approaches the plant configuration in some sense. A natural measure of the progress of self–organization is given by the area of the triangles formed by the model and the plants. The difference between the two areas may be taken as a Lyapunov candidate. However, numerical experiments demonstrate that the candidate function does not decrease along the trajectories of the models (see figure 4). This is due to the fact that all models are drawn towards a unique point M^* (located near the center of gravity of the triangle) before they expand and converge. This suggests the existence of a saddle point in the (3×2 dimensional) state–space of the system. The models first follow a stable direction leading to M^* and then are repelled from M^* along the unstable direction. The models $[\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3]$ also collectively approach M^* and expand only after attaining a neighborhood of M^*. The existence of a "contraction" and "separation" phase during the convergence of the models has been observed in all the numerical experiments that have been performed over the last three years.

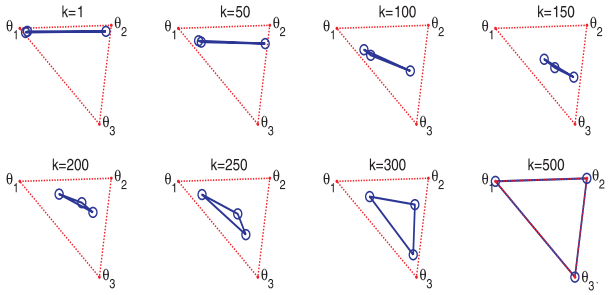

Fig. 4. Self Organization in the 3/3 case: The area of the inner triangle (as well as its perimeter) decreases initially and then increases monotonically up to the point where the inner and outer triangle match.

The same non–monotonic behavior is observed if we collapse the triangle to a straight line, on which all models and plants are initialized. However, if there are only $N = 2$ plants (at a distance L) and we define $x(k) = e_{21}(k) - e_{11}(k) = e_{12}(k) - e_{22}(k)$ we find that $L - x(k)$ **is** a Lyapunov function. The last equal sign is due to the fact that the models and plants are on the same line. It is easy to see that x increases no matter whether θ_1 or θ_2 is in existence and that $\Delta x = 0$ iff $x = \|\theta_1 - \theta_2\| = L$. Hence, $\{x(k)\}$ is a monotonically increasing sequence bounded above by L and $\lim_{k\to\infty} x(k) = L$. □

C. A new mathematical framework

The numerical experiments display what is sometimes called an *emergent* property [5]. A simple form of interaction (among models) produces a complex large–scale effect, which cannot be explained by simply aggregating the properties of each individual model. The dynamics of the system is defined by a set of nonlinear, non–autonomous, coupled difference equations, (19) and (18), where the determining factor for the qualitative behavior of solutions is the coupling, not the eigen-dynamics of every model. The fact that the coupling is nonlinear

together with the typically large number of models makes analysis of the system an extremely complex one. Notice also that the update equation (19) is different every time the plant switches. Even in the simple case of a periodic switching (with period NT) we have to compose the corresponding update steps to obtain the final equation for the evolution of models. The problem here is that there is no reasonably simple algebraic expression for the r.h.s. of such an equation.

From analysis we know that the key to all convergence results is monotonicity. Similarly, in the context of dynamical systems, we rely on the monotonic decrease of a Lyapunov function along the orbits of the system in order to make any statements as to the long–term behavior of the system. The search for a Lyapunov function $V(x)$ is often guided by considering sample orbits of the system and relating them to the vectorfield $f(x)$ as defined by the right hand side of the differential equation. In order to solve the inequality $\nabla V^T f(x) < 0$, the algebraic expression for $f(x)$ must offer some "structure" in order to obtain a closed–form expression for V. In the present case, the vectorfield fails to display such a simple structure. On the other hand, if we look at sample paths of the system, we observe a rather simple qualitative dynamic pattern including a "saddle-type" equilibrium.

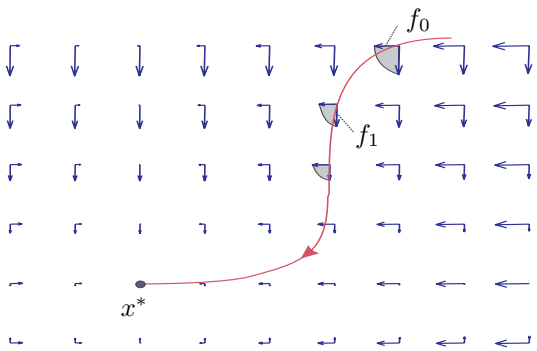

Fig. 5. Evolution of the solution of Σ inside an envelope formed by the vectorfields of two simpler systems Σ_0, Σ_1. The solution is almost tangent to f_0 for one point of the state space and tangent to f_1 for another. x^* denotes the equilibrium.

The idea is to find a relation between the given system Σ and a simpler dynamical system defined in the same space. Let Σ_0 and Σ_1 be two such systems. Suppose we can enclose the solutions of Σ in a region formed by convex combinations of the vectorfields f_0 and f_1, it is clear that the solution can be characterized by the two simpler systems since it evolves inside the envelope defined by the two "border" vectorfields. This is shown in figure (5). Notice that the vectorfield of Σ does not itself correspond to a convex combination of f_0 and f_1.

The fact that the r.h.s. is very complex does not in itself mean that the dynamic behavior of the system is complex. In a seminal paper [4], Smale emphasized that dynamical systems should be classified according to the complexity of their qualitative dynamical behavior, rather than, the more traditional way of classification by the complexity of the algebraic form of the differential equation. It is in this spirit that we have to find a correspondance between the given system and a (set of) simpler systems with "similar" trajectories. This relation should be weaker than a topological conjugacy, since the latter will be just as complex as the system itself.

D. Application

The construction of an entirely new mathematical framework is justified by the broad spectrum of conceivable applications of the algorithm. In fact, the methodology presented here is not limited to the adaptive control of time–varying systems. It is applicable whenever multiple agents are involved in cooperative problem solving such as distributed search, swarms, information retrieval, traffic control, and the study of social and economic systems. The algorithm has in fact a game theoretic background in that the performance of the system is optimal if all the models $\hat{\theta}_i$ learn to be team players. The reward J_i allotted to $\hat{\theta}_i$ decreases if the model tries to be close to more than one point. In such a case, another model will have the chance to step in. The interesting point to note here is that this is directly a consequence of geometry and not due to the (probabilistic) choice of an update action. In game theory, the actions usually belong to a finite set which is why they can be analyzed using the framework of stochastic automata. In the present paper, we have formulated a game having a continuous space of update actions (moves) $\eta_i\tilde{\theta}_i$. Moreover, the game is deterministic in the sense that if the time–variation is periodic, the outcome is completely determined by initial conditions.

REFERENCES

[1] NARENDRA, K.S.; BALAKRISHNAN, J.: *Adaptive Control Using Multiple Models*. IEEE AC–42, No. 2, Feb. 1997.
[2] G. C. GOODWIN, K.S. SIN: *Adaptive Filtering, Prediction and Control.* Prentice–Hall, 1984.
[3] FEILER, M.J.: *Adaptive Control in the Presence of Disturbances*. http://tumb1.biblio.tu-muenchen.de/publ/diss/ei/2004/feiler.html, Diss. TU Munich, 2004.
[4] S. SMALE: *Differentiable Dynamical Systems*. Bull. American Math. Soc. No. 73, 1967.
[5] R. AXELROD: *The Complexity of Cooperation*. Princeton Studies in Complexity 1997.