Identification and Control Using Multiple Models

Matthias J. Feiler ◇ and Kumpati S. Narendra *

***Abstract*— The paper presents recent developments in an on–going effort to deal with rapid time–variations in dynamical systems. It is demonstrated that fast identification of the time–varying parameters enables stable adaptive control. The main focus is on deriving identification algorithms based on multiple models that have the required property and to provide an example from financial time–series analysis where the algorithms can be applied.**

I. INTRODUCTION

From the very beginning, the interest of adaptive control theorists has been centered around adaptation in changing environments. However, due to mathematical tractability, they confined their attention to time–invariant systems with unknown parameters. The algorithms that have proven stability properties prescribe the evolution of the parameters towards their true values in terms of a differential equation. All algorithms available in continuous and discrete time [1], [2] (including well-known regression algorithms used in statistical analysis, [3]) share this property. It is clear, though, that this may not be the fastest way to determine the parameters since adaptation thus defined is inherently incremental.

This paper deals with time–variations that are both rapid and large. It is shown how the time–varying problem can essentially be converted into a set of $N \in \mathbb{N}$ time–invariant problems which are solved in a distributed fashion by a (possibly large) set of identifiers. The framework within which this is possible is the one of multiple models, switching and tuning. The results presented are an example of how a complex problem is solved by a set of independent entities which cooperate in order to subdivide the problem into solvable pieces. The cooperation takes place in parameter space where models fan out and identify a set of distinct plants representing the time–variation. The work is related to research in the areas of node localization in sensor networks [4] and coordinated motion in networks of multi–agent systems [5]. In these problems, agents are able to measure distances to their nearest neighbors. The key point in the present work is that this critical information is unavailable. Distances in parameter space have to be inferred from observations of the system output. The system, in turn, is not constrained in any sense: its order and relative degree as well as the complexity of the time–variation are arbitrary.

◇LGT Capital Management Ltd., Asset Allocation & Research, 8808 Pfäffikon, Switzerland.

*Center for Systems Science, Yale University, P.O. Box 208267, New Haven, CT, 06520-8267 USA.

A. Problem Statement

As the title indicates, the problem addressed in this paper is two–fold (and complimentary): The first objective is to design an identification algorithm for systems with time–varying parameters. The second objective is to adaptively control such systems.

Time–variations may occur in infinitely many ways. To make the problem tractable, the time–variation is assumed to be a piecewise–constant function $\theta(k) : \mathbb{Z}^+ \to S$ where $S \subset \mathbb{R}^{2n}$ is a finite set $S = \{\theta_1, \theta_2, \ldots, \theta_N\}$ of unknown, constant parameter vectors. We denote the index set by $\Omega = \{1, 2, \ldots, N\}$ and refer to the elements of S as *plants*. The instants at which the parameters switch are unknown. If the system is constant over a long interval of time I, classical adaptive control can be used with a single model. In theory, the model takes an infinite number of steps to converge. For most practical purposes it suffices that it enters a neighborhood of the parameters. If this is achieved after $\bar{T}$ steps but the length of the interval I is smaller than $\bar{T}$, the time–variation is considered rapid. This case is the focus of the present paper. As we shall see, I must still satisfy some lower bounds for the problem to be meaningful. But it may be significantly smaller than the number $\bar{T}$ of steps required in the classic adaptive solution.

Time delays d and δ: 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 system 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 parameters affects the system 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 $\bar{N}$ adaptive models of the form

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

where $\hat{\theta}_i(k-1)$ and $\hat{y}_i(k)$ are estimates of $\theta(k-\delta) \in S$ and $y(k)$ respectively. We denote the set of models by $\bar{S}$ and the corresponding index set by $\bar{\Omega} = \{1, 2, \ldots, \bar{N}\}$. The parameters are adjusted based on the estimation error $e_i(k) = y(k) - \hat{y}_i(k)$ which satisfies

$$e_i(k) = \phi(k-d)^T [\theta(k-\delta) - \hat{\theta}_i(k-1)] \tag{4}$$

The problem is to find an algorithm that updates all models in such a way that a subset of $\bar{S}$ converges to S as $k \to \infty$.

(A) *The Identification Problem*: Determine an estimation algorithm such that for every $j \in \Omega$ there is a $i \in \bar{\Omega}$ such that $\|\hat{\theta}_i(k) - \theta_j\| \to 0$ as $k \to \infty$.

In other words, for every plant $\theta \in S$ there exists (at least) one model $\hat{\theta}_i$ that converges to it. If 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.

(B) *The Control Problem:* Define a non-anticipating control input $u(k)$ such that the system (1) is stable and the output $y(k)$ asymptotically follows a desired output $y^*(k)$ on a subsequence $\{k_l\}_{l \geq 0}$ where $k_l \to \infty$ as $l \to \infty$.

By non-anticipating we mean a controller that does not depend on future parameters of the system. This is also the reason why asymptotic tracking can only occur on a subsequence as will be demonstrated in section III.

B. Multiple Models, Switching and Tuning (MMST)

The methodology was introduced at the 1992 Yale Workshop [6], was proven rigorously in 1997, and later extended to the stochastic case in [7]. The basic idea comes from biology: every biological system is faced with a multiplicity of choices at any instant of time. As the environment changes, the system demonstrates an ability to rapidly modify its strategy such as to maintain optimal performance. Such an ability involves recognizing the specific situation that has arisen and taking an appropriate control action which is selected from a set of available strategies. In order to acquire a repertoire of strategies the biological system learns from past experience. Hence, it is the ability to learn and store information, and combine it with adaptation that is responsible for its satisfactory performance in rapidly time–varying situations. A crucial role in the design of MMST is played by the switching criterion which determines when to switch from one model to another and which of the models should be the new one. A number of different performance indices can be defined based on the identification error $e_i(k)$ to determine which of the models best fits the plant at any instant. These may assume the following forms:

$$\begin{array}{lllll} \text{(C i)} & J_i(k) & = & e_i^2(k) \\ \text{(C ii)} & J_i(k, T_0) & = & \sum_{\nu=0}^{T_0} e_i^2(k-\nu) \\ \text{(C iii)} & J_i(k, T_0) & = & \alpha\, e_i^2(k) + \beta \sum_{\nu=0}^{T_0} e_i^2(k-\nu) \\ \text{(C iv)} & J_i(k, T_0) & = & \alpha\, e_i^2(k) + \beta \sum_{\nu=0}^{T_0} \rho^\nu e_i^2(k-\nu) \end{array}$$

where $\alpha,\ \beta > 0$ are constants used to balance the weight of the instantaneous vs. the integral errors. In practice, a number $\bar{N}$ of models are distributed in parameter space. We assume that the space is a compact subset $K \subset \mathbb{R}^{2n}$ so that every neighborhood of a point $\theta \in K$ contains a model if $\bar{N}$ is large enough. The models simultaneously attempt to identify $\theta(k)$. At every instant of time, the performance index J_i based on the estimation error of each model is evaluated in order to determine which of the models is closest. The control input to the plant is then chosen based on that model.

II. System Identification

Identification refers to the problem of explaining how the observed response of a system (technical, biological, or social/ economic) depends on a set of input (explanatory) variables. The objective is to find parameters which minimize the error (unexplained part) between the system and the model. In regression analysis, the error term is treated as a random variable. The models typically employed belong to the system class considered in this paper and

can be written in the format (2). In the stochastic case, $\phi(k)$ is also affected by a random variable, typically a white noise sequence $\{w(k)\}$. The stochastic case does not pose a conceptually different problem. The key issue is in replacing the random entries in the regression vector by observable signals, e.g. the a posteriori identification error (see [2] for details). However, this case is not explicitly treated in this paper since our main interest is in time–varying systems. The objective of this section is to describe three identification algorithms and to assess their ability to succeed in the time–varying situation.

Assuming that the unknown parameter $\theta(k) \equiv \theta$ is constant, the identification problem may be described geometrically as follows. The parameters that achieve perfect matching of the measured output $y(k)$ lie on the hyper–plane

$$H = \{\hat{\theta}(k) \,|\, \phi(k)^T[\,\hat{\theta}(k) - \theta] = 0\} \tag{5}$$

As $\phi(k)$ changes with time, a new identification error and hence a new solution hyperplane is generated. The true parameter lies in the intersection of $2n$ such planes. If the planes are independent (not collinear) the solution is unique. $\phi(k)$ is referred to as *persistently exciting* over an interval, say T_0. This corresponds to the case where the so–called observability Gramian (see [2]) has full rank. It is then possible to uniquely determine θ from the knowledge of the identification error $e(k)$ over the interval $[k - T_0,\ k]$. Hence, in order to determine the parameter error of a model at instant k it is necessary to observe its output error over an interval of length T_0. These geometrical considerations directly motivate the following algorithm.

Algorithm 1: Sequential Projections

A single model $\hat{\theta}$ is used to identify $\theta(k)$. The model is updated by subsequently projecting it on the hyperplanes that appear when new data points are added. While the location of the planes is computed the estimation error $e(k)$, the projection is along the direction of the regression vector $\phi(k-1)$. The update equation reads:

$$\hat{\theta}(k+1) = \hat{\theta}(k) + \eta\,\frac{\phi(k-d)}{1 + \phi(k-d)^T\phi(k-d)}\; e(k) \tag{6}$$

where $0 < \eta < 2$ is the adaptation step–size. If the vector $\phi(k-d)$ is orthogonal to all its previous values $\phi(1), \ldots, \phi(k-d-1)$ then the parameters converge in a finite number ($2n$) of steps. In general, this is not the case and convergence takes an infinite number of steps. In practice, on the other hand, it often suffices that models enter a small neighborhood of the plant. Due to the sequential nature of the algorithm, time-variations may be tracked up to a neighborhood where the size of the neighborhood depends on the speed of the time-variation.

Algorithm 1 can be used to adapt several models. If the plants are separated by a known distance D an ad–hoc way of identifying them is to start with an equal grid of models and adapt only the nearest model at every instant of time. The success critically depends on the fact that there is a model already close to every plant and that this model is always closest to its respective plant, so also the models are separated by some minimum distance. To ensure this, the grid size must be adapted to D. But the location of the plants, and hence D, are unknown so the problem has only been shifted. Another simple approach would be to increase the number of models whenever a chosen grid–size proves to be insufficient (based e.g. on the estimation error). This means that the number of models becomes infinite when $D \to 0$. In practice one may stop adding models if the errors become small. But then there is no convergence in the strong mathematical sense. From a theoretical viewpoint there is a huge difference between zero and "small". In particular, the stability proof collapses when certain signals do not go to zero asymptotically. The following two algorithms ensure (strong) convergence without assuming any prior knowledge as to distances among the plants.

Algorithm 2: Simultaneous Identification

Algorithm 2 represents the most elegant and compelling solution to the identification problem (A). The number of models required is only $\bar{N} = N$. In its formulation, the authors were influenced by earlier work by Kawato et al. [8] who presented a way of measuring the extent to which different models contribute to observed system behavior in a problem of motor learning and control. Unlike our predecessors, our interest is in proving convergence of an identification algorithm based multiple models. The performance of a given model relative to all the other models is used in the adaptive law (6) to determine the step-size by which this model is updated. In particular,

$$\hat{\theta}_i(k+1) = \hat{\theta}_i(k) + \eta_i(k)\,\frac{\phi(k-d)}{1 + \phi(k-d)^T\phi(k-d)}\; e_i(k) \tag{7}$$

where

$$\eta_i(k) = \frac{\frac{1}{J_i^2(k,T_0)}}{\frac{1}{J_1^2(k,T_0)} + \frac{1}{J_2^2(k,T_0)} + \cdots + \frac{1}{J_N^2(k,T_0)}} \tag{8}$$

While the projection vector $\phi(k-d)$ in equation (7) is the same for all models, its length varies according to the error $e_i(k)$ and a step–size $\eta_i(k)$ associated with every model $\hat{\theta}_i(k)$, $i \in \bar{\Omega}$. $J_i(k,T_0)$ measures the performance of a model in terms of its estimation error which, in turn, relates to its distance from the plant as discussed above. The closest model is updated by a large amount and, hence, its advantage over the other models increased. Notice that the nonlinear function (8) does not define a "winner takes it all"– policy since all models move at every instant of time. This is why the algorithm is referred to as *simultaneous identification*. The *sum* of all step–sizes is equal to one, and the step–size of any other model $\eta_l(k)$, $l \neq i$, is zero only in the limit when $e_i(k) \rightarrow 0$ as $k \rightarrow \infty$ (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 without the help by a supervisor/ coordinator. This is possible since every model has full information about the performance of all its competitors.

A salient feature of the algorithm is that convergence takes place in a non–monotonic fashion except for some initial conditions where the models are already close to the plants. In phase space, the models seem to approach a saddle point which attracts all solutions first and then repels them towards the region of attraction of the other invariant set, that we are interested in, when all models converge. Detailed simulations and a partial analysis of the algorithm have been presented in [9], [10]. In all the simulation studies produced over the last five years, the algorithm invariably converges, independently of the number of unknown plants, the choice of initial conditions or the nature of the time–variation (periodic or random). Due to the non–monotonicity and the ensuing difficulties in finding a Lyapunov function, however, the proof of convergence has not been completed yet.

Algorithm 3: Swarms of Models

The number of models is substantially larger than in the previous algorithm. The idea is to generate swarms of models in the neighborhoods of the plants.

1) Distribute a large number $\bar{N}$ of models uniformly over the compact space K.
2) The estimation errors are evaluated according to one of the performance criteria e.g. (C ii) $J_i(k,T_0) = \sum_{\nu=0}^{T_0} e_i^2(k-\nu)$. The models are ordered based on this criterion. Let $[\hat{\theta}](k) = [\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_M, \hat{\theta}_{M+1}, \ldots, \hat{\theta}_N]$ be an ordering of the models.
3) Let $\hat{\theta}^* = \hat{\theta}_1$. Create a swarm of models around $\hat{\theta}^*$ consisting of the M closest models $\hat{\theta}_2, \ldots, \hat{\theta}_M$. The models are distributed uniformly within a sphere centered at $\hat{\theta}^*$ with radius $r = \gamma\, J_1(k,T_0)$ where $\gamma > 0$ is a constant. This distribution is periodically updated after T_0 instants of time.
4) $\hat{\theta}^*$ is fixed for all $\bar{k} > k$. This means that if a model has once been assigned the '*' it is no longer adapted. If there are μ such models in the list $[\hat{\theta}_1 ... \hat{\theta}_M]$ they are skipped and the list is prolonged up to $\theta_{M+\mu}$.

A central assumption is that the criterion applied in 2 reflects the distance of model i to the plant:

Theorem 1: Assume that the models can be ordered according to their parametric distances by a criterion based on the estimation error. The algorithm 1-4 described above solves the identification problem (A) with probability one, i.e.

For every $j \in \Omega$ there is a $i \in \bar{\Omega}$ such that

$$\lim_{k\rightarrow\infty} \|\hat{\theta}_i(k) - \theta_j\| = 0 \quad \text{a.s.}$$

We outline a sketch of the proof.

Case 0: First, let us assume that S consists of a single element. This is the time-invariant identification problem $\theta(k) \equiv \theta$. Let $\hat{\theta}_i$ be the closest model at instant k. A swarm will be generated around $\hat{\theta}_i$ with radius $\gamma\, J_i$. Let p be the probability, that one model $\hat{\theta}_j$ contained in the swarm is closer to the plant that $\hat{\theta}_i$. By assumption this can be detected using criterion (ii). After T_0 instants of time all entries appearing in the sum of the criterion have been replaced and the '*' is passed from model i to model j and a new entourage is created around $\hat{\theta}_j$ with smaller radius γJ_j. We have that

$$\|\hat{\theta}_j - \theta\| < \|\hat{\theta}_i - \theta\| \quad \text{with probability } p \qquad (9)$$

Suppose that no such model $\hat{\theta}_j$ was contained in the swarm around $\hat{\theta}_i$. The distribution is renewed after T_0 instants. Hence $p \rightarrow 1$ as time $k \rightarrow \infty$.

Now assume that S consists of $N > 1$ elements. There are two cases.

Case 1: Swarms do not overlap. Any time the system switches to a new value, a new model ordering $[\theta]_{new}(k)$ is in place. The set Q_{new} consisting of the first M

components of the ordering $[\theta]_{new}$ does not intersect any of the previous sets Q. Since there is a finite number of plants, there exist a finite number of such non-intersecting sets. From the perspective of the ith plant the problem is one of time–invariant identification on the subsequence $\{k_i\}$ over which the ith plant exists. Notice that by choosing the number of models $\bar{N}$ large enough the conditions of Case 1 can always be met. However, there is no guarantee and we have to consider the following case.

Case 2: Swarms overlap. The overlap may either occur at the top level (i.e. $\hat{\theta}_1$ of different orderings coincide, one model is closest to more than one plant) or at the swarm level (i.e. one or more $\hat{\theta}_i$, $i > 1$ coincide) or both. If $\hat{\theta}_1$ is close to more than one plant the performance index J_1 cannot be zero for all plants. When $J_1 \neq 0$, the model will form an entourage of non–zero radius. As in the time–invariant case, there will be a closer model with probability p. Since $p \to 1$, the top models of the swarm will eventually be two different ones. The fact that the different top models share the same swarm poses no problem since, by step 4 of the algorithm, the number of members in the swarm may be increased. This means that the problem has been seperated into N time–invariant identification problems and convergence takes place as in case 0. □

Comment: The swarm algorithm has the effect of increasing the density of models in the neighborhoods of the plants. The assumption that the models can be rank–ordered according to their parametric distance to the plant is essential. It is clear, however, that this cannot be guaranteed over the T_0 instants following a variation (switch) of the plants. The effect is that some models will be placed at the location of a (fictious) plant having an output sequence that corresponds to the mixed output of the two plants involved in a transition. Since the number of models is assumed to be large, this has no adverse effect on the performance of the algorithm – not so in the case of algorithm 2, where the models may be diverted from the true plant locations during a transition period. The swarm algorithm is both faster and more robust as can be observed in the following example.

Application: Replication of Financial Time–Series

The performance of the different algorithms is studied using a real–world example. We begin by motivating and describing the problem.

In the analysis of financial time–series, identification problems are ubiquitous. Recent interest has been in the problem of index tracking, i.e. of creating a synthetic asset which duplicates the composite returns of a specific financial market. This falls under the category of *passive* management: the investor follows the index instead of outperforming it thereby avoiding the costs of research (stock picking, market timing) entailed in active management. There are very few human inputs involved in the construction of the portfolios. The holdings are identified by statistically sampling the market and regressing the observed data on the set of representative securities. One way of looking at this approach is that it mechanically extracts the investment decisions of the average investor in a given market and uses this insight to construct a "market" portfolio. In other words, the method identifies what the market is doing.

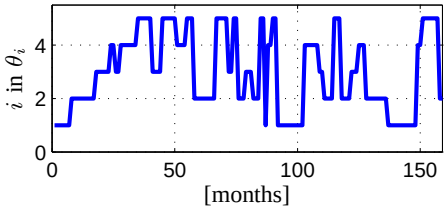

Fig. 1. Dynamic asset allocation

In the following example, we create a (fictitious) index reflecting the returns of a long/ short equity portfolio and we aim to replicate the index by investing in (small and large market capitalization) stocks. Not all of our money needs to be invested, part of it can be retained as (risk–free) cash. The index represents a (hedge fund–like) trading style in which positions may be shifted very rapidly (once per month) from one asset to another. It is assumed that all the assets in the replicating portfolio can also be traded at least monthly. The objective is to construct the replicating portfolio based on monthly updates on the index performance. The model used for identifying the portfolio weights is given by

$$\hat{r}_I(t) = \theta_{SC} r_{SC}(t) + \theta_{LC} r_{LC}(t) + \theta_C r_C(t) \qquad (10)$$

where r_I is the index return at instant t, r_{SC}, r_{LC} and r_C are the returns of small cap stocks, large cap stocks and cash (the money market) respectively. The portfolio weights satisfy $-1 \leq \theta_{xC} \leq 1$ and $\theta_{SC} + \theta_{LC} + \theta_C = 1$. The time t is measured in months and the return at time t is simply equal to the price change of a given security $r(t) = p(t)/p(t-1)$. The index portfolio consists of the

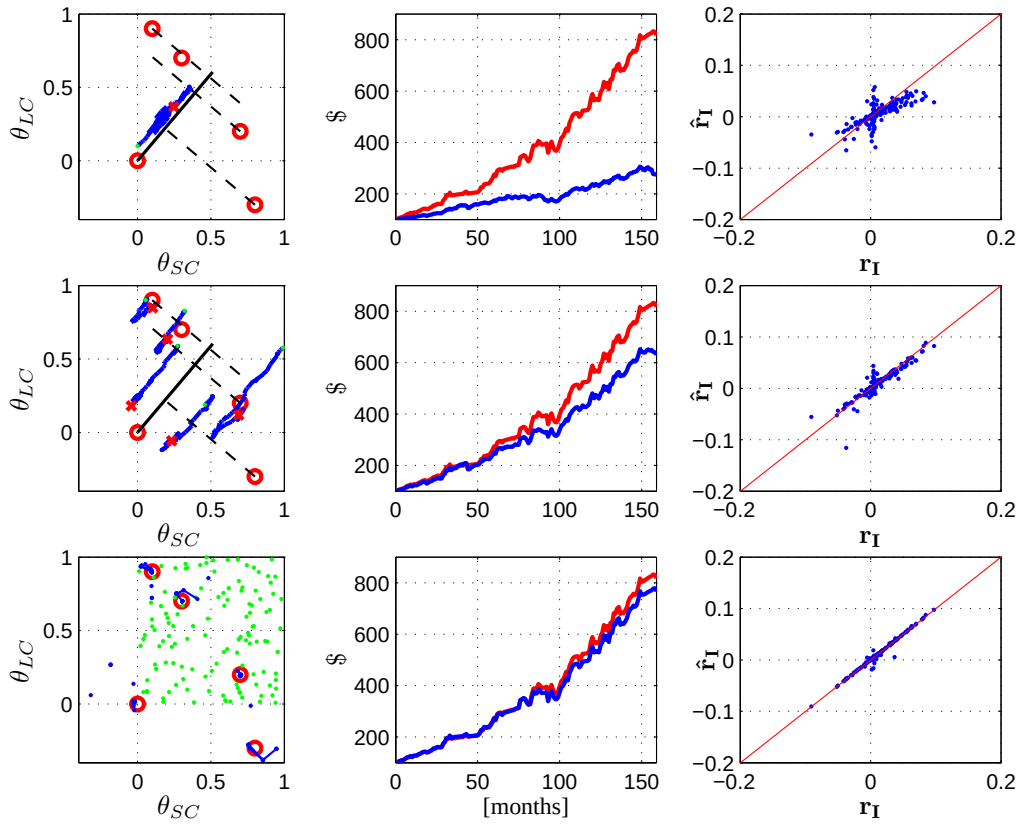

Fig. 2. Index tracking by identification of the time-varying positions in an equity portfolio. Top row: projection algorithm using a single model. Middle row: simultaneous identification algorithm using 5 models. Bottom row: clustering algortihm using a swarm of models (> 100).

five positions:

$$\theta_1 = \begin{bmatrix} 70\% \\ 20\% \end{bmatrix} \quad \theta_2 = \begin{bmatrix} 10\% \\ 90\% \end{bmatrix} \quad \theta_3 = \begin{bmatrix} 80\% \\ -30\% \end{bmatrix}$$
$$\theta_4 = \begin{bmatrix} 30\% \\ 70\% \end{bmatrix} \quad \theta_5 = \begin{bmatrix} 0\% \\ 0\% \end{bmatrix} \tag{11}$$

The remainder (100% minus the quotas) is invested in cash. The representative investor is rapidly switching among these positions in order to obtain the optimal (according to some return/risk performance criterion) allocation of assets in the portfolio. The allocations displayed in figure 1 correspond to the investment decisions. Figure 2 displays the simulated tracking of the index using the three different identification algorithms described above.

In the first row, algorithm 1 is applied. It performs poorly both because the regression vector does not form an orthogonal sequence (to the contrary: small and large cap equity indices are highly correlated) and because the time–variation is rapid. The parameters converge to an average portfolio (average over time, not over the market). In the second row, algorithm 2 is employed. It performs better, since the identification was distributed among $\bar{N} = 5$ models. However, since convergence is non–monotonic, models which were once close to the index allocation move away from that allocation and return to it only eventually. This digression costs real money (about \$200 in column 2) of the figure. Finally, in row 3, we observe the swarm algorithm which starts with $\bar{N} = 150$ possible allocations (models) but rapidly collapses them into a few very similar ones. The figure clear displays the successful determination of the index allocations as well as a close–to–optimal performance of the index tracker (the loss is $<$ \$50) over a period of 150 months.

III. Control

Solving the identification problem (A) is equivalent to solving the adaptive control problem in the presence of

large and rapid time–variations. This is demonstrated in the following section.

A. Analysis

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 $(i,j) \in \bar{\Omega} \times \Omega$. A performance criterion is used to decide which model best approximates the plant. Denote the instant at which the system switches from a plant 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. The instant s at which a time–variation occurs and the instant at which it can be detected at the output are offset by δ steps. At $k = s+\delta$ the identification error is equal to

$$\begin{aligned} e(s+\delta) &= \phi(s+\delta-d)^T[\theta(s) - \hat{\theta}(s+\delta-1)] \\ &= \phi(k+\delta-d)^T[\theta_{i_2} - \theta_{i_1}] \\ &\neq 0 \end{aligned} \tag{12}$$

since $\theta(s) = \theta_{i_2}$ whereas $\hat{\theta}(s+\delta-1) = \theta_{i_1}$. This error is used to detect the time–variation of the plant. After some time $T \leq T_0$ the detection is complete and a new model is selected, $\hat{\theta}(s+\delta+T) = \theta_{i_2}$. This means that even in the case of perfect matching of models and plants, the identification error is non–zero over the interval $I_{id} = [s+\delta,\ s+\delta+T)$. We define the (infinite) set of detection instants $Q_1 = \{k \mid s+\delta \leq k < s+\delta+T\}$.

Let us now examine the control error $e_c(k) = y(k) - y^*(k)$. As in standard adaptive control, we have the following relationship between the identification error $e(k)$ and the control error $e_c(k)$

$$e_c(k) = e(k) + \phi(k-d)^T[\hat{\theta}(k-1) - \hat{\theta}(k-d)] \tag{13}$$

Due to the detection problem, $\hat{\theta}(k)$ switches to the new model only at instant $k = s+\delta+T$. In view of equation (13), this means that $e_c(k) \neq 0$ for $k \in I_c = [s+\delta, \ldots, s+\delta+T+d)$. The interval I_c is referred to as the inherent control error interval. The set of error instants which is due to relative degree is defined as $Q_2 = \{k \mid s+\delta+T \leq k < s+\delta+T+d\}$. We have $I_c = Q_1 \cup Q_2$.

Figure 3 illustrates the sequence of events causing the control error e_c to be nonzero over the interval I_c. To avoid the error, the time–variation would have to be known, i.e. the switching instant s and the "new" environment $\theta(s) = \theta_{i_2}$ must be anticipated. If the new control input (corresponding to θ_{i_2}) was applied d instants before the time-variation affects the output of the plant, i.e. at instant of time $k = s+\delta-d$, the control error would satisfy $e_c(k) \equiv 0$ for all $k > 0$.

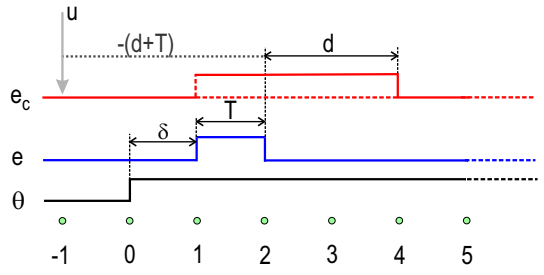

Fig. 3. Binary representation of the inherent identification and control error when $\delta = 1$, $d = 2$, and $T = 1$.

B. Stable Adaptive Control

Theorem 2: Suppose the following is satisfied:

- (1) has asymptotically stable zero–dynamics.
- $\theta(k) \equiv \theta_j$ for some $j \in \Omega$ over an interval $I > d+T$.
- For any $j \in \Omega$ there is a $i \in \bar{\Omega}$ such that $\|\hat{\theta}_i(k) - \theta_j\| \to 0$ as $k \to \infty$.

The certainty equivalence controller obtained by solving $y^*(k) = \phi(k-d)^T\hat{\theta}_i(k-d)$ asymptotically stabilizes the time–varying system (2) in the sense that $\|\phi(k)\| < \infty$ for all $k > 0$.

The identification algorithms 2 and 3 in section II satisfy assumption three of the theorem. For a detailed proof of stability when algorithm 2 is used, see [11]. In this paper, the focus is on algorithm 3.

Proof: Observe that

$$|e(k)| \leq \|\phi(k-d)\|\|\tilde{\theta}(k-\delta)\| \tag{14}$$

where $\|\tilde{\theta}(k-\delta)\|$ is the distance between the swarm leader ($\hat{\theta}_1$ of the ordering $[\theta](k-\delta)$) and the plant $\theta(k-\delta)$. By assumption three of the theorem, $\|\tilde{\theta}\| \to 0$ for a set of model–plant pairings $\hat{\theta}_i$ and θ_j where $\hat{\theta}_i$ refers to the respective swarm leader (i.e. $[i] = 1$ when the models are ordered).

Define a subsequence $\{k_n\}_{n\geq 0}$ with $k_n \to \infty$ as $n \to \infty$ where $k_n \in \mathbb{Z}^+ \setminus Q_1$. Using equation (14) and the properties of the swarm algorithm (Theorem 1) we have that, with probability one, the normalized identification error satisfies,

$$\frac{|e(k_n)|}{\|\phi(k_n-d)\|} \xrightarrow{\text{a.s.}} 0 \quad \text{as } n \to \infty \tag{15}$$

In a similar fashion, define a subsequence $\{k_m\}_{m\geq 0}$ with $k_m \to \infty$ as $m \to \infty$ where $k_m \in \mathbb{Z}^+ \setminus \bar{Q}_2$ and $\bar{Q}_2 = \{k \mid s+\delta+T < k < s+\delta+T+d\}$. Since the swarm radius γJ_i tends to zero almost surely we obtain,

$$\|\hat{\theta}(k_m - 1) - \hat{\theta}(k_m - d)\| \xrightarrow{\text{a.s.}} 0 \quad \text{as } m \to \infty \tag{16}$$

Notice that $\|\hat{\theta}(k-1) - \hat{\theta}(k-d)\|$ cannot be zero when $k \in \bar{Q}_2$ since $\hat{\theta}(k-1)$ and $\hat{\theta}(k-d)$ belong to different models during the $d-1$ instants following the detection of a time-variation. Using the two results in order to bound the absolute value of the r.h.s. in equation (13) it follows that

$$\frac{|e_c(k_l)|}{\|\phi(k_l - d)\|} \xrightarrow{\text{a.s.}} 0 \text{ as } k_l \to \infty \tag{17}$$

where $\{k_l\}_{l\geq 0}$ is a subsequence tending to infinity as $l \to \infty$ and $k_l \in L = \mathbb{Z}^+ \setminus Q_1 \cup \bar{Q}_2$.

Every interval of length $I > d + T$ contains at least one instant $k = k_l \in L$, i.e. there exists a subsequence $\{k_l\}$ as defined above on which the normalized control error in equation (17) tends to zero with probability one. This implies that $\phi(k_l - d)$ must grow faster than $e_c(k_l)$ if both signals grow in an unbounded fashion. But, by assumption one of the theorem, 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.

□

Corollary 3:

$$e_c(k_l) \xrightarrow{\text{a.s.}} 0 \text{ as } \; k_l \to \infty \tag{18}$$

i.e. the actual control error (not only the normalized one, cf. eq. (17)) tends to zero on a subsequence with probability one.

The paper demonstrates how a complex time–varying problem can be subdivided into a set of time–invariant identification and control problems. The last part of the above proof is in fact analogous to the proof of stability in the time–invariant case.

REFERENCES

[1] NARENDRA, K.S.; ANNASWAMY, A.M.: *Stable Adaptive Systems.* Prentice–Hall, Englewood Cliffs, NJ, 1989.

[2] G. C. GOODWIN, K.S. SIN: *Adaptive Filtering, Prediction and Control.* Prentice–Hall, 1984.

[3] LJUNG, L.: *System Identification – Theory for the User*. Prentice–Hall, Englewood Cliffs, NJ, 1999.

[4] VIRRANKOSKI, R.; SAVVIDES, A.: *Topology Adaptive Spatial Clustering for Sensor Networks.* In: Proceedings of the 2nd IEEE Conference on Mobile Ad Hoc and Sensor Systems, Washington DC, Nov 2005.

[5] JADBABAIE, A.; LIN, J.; MORSE, A.S.: *Coordination of groups of mobile autonomous agents using nearest neighbor rules*, IEEE Transactions on Automatic Control, vol. 48, no. 6, 2003.

[6] NARENDRA, K.S.; BALAKRISHNAN, J.: *Performance Improvement in Adaptive Control Systems Using Mulitple Models and Switching.* In: Proceedings of the Seventh Yale Workshop on Adaptive Learning Systems, New Haven CT, May 1992.

[7] NARENDRA, K.S.; XIANG, C.: *Adaptive Control of Discrete–Time Systems Using Multiple Models.* IEEE Transactions on Automatic Control, vol. 45, no. 9, Sep. 2000.

[8] HARUNO, M.; WOLPERT, D.M. AND KAWATO M.: *MOSAIC Model for Sensorimotor Learning and Control.* Neural Computation, 13: 2201–2220, 2001.

[9] FEILER, M.J.; NARENDRA, K.S.: *Adaptive Control of Rapidly Time-Varying Systems*. In: Proceedings of the 12th Yale Workshop on Adaptive and Learning Systems, p. 33-42, Yale University, New Haven, CT, USA, 2003.

[10] FEILER, M.J.; NARENDRA, K.S.: *Simultaneous Identification and Control of Linear Time–Varying Plants using Multiple Models*. Technical Report No. 0401, Center for Systems Science, Yale University, New Haven, CT, USA, March 2004.

[11] FEILER, M.J; NARENDRA, K.S.:*Simultaneous Identification and Control of Time–Varying Systems* In: Proceedings of the 45th IEEE Conference on Decision and Control, 2006.