2005 American Control Conference
June 8-10, 2005. Portland, OR, USA

A Sensor Management Protocol for Tracking with Diverse Sensors

Vishwesh V. Kulkarni Lucy Y. Pao

***Abstract*— We develop a sensor management protocol that a vehicle may use to track other vehicles and objects in its neighborhood using a sensor-suite, an individual element of which gives either periodic cluttered updates or aperiodic uncluttered updates. We outline how to combine the Joint Probabilistic Data Association Filter with stochastic Kalman Filters for the state estimation process. However, since such an algorithm is sensitive to modeling errors, we discuss the use of data-driven optimization algorithms to increase robustness. Finally, we discuss extensions of previously developed covariance control sensor management algorithms.**

***Index Terms*— target tracking, sensor management, data-driven controllers, covariance controller, Kalman filter, certainty equivalence principle**

I. INTRODUCTION

This paper is motivated by the problem of how a vehicle should manage a suite of diverse sensors so as to efficiently track other vehicles and obstacles in its neighborhood. The sensor management is exercised by controlling the sensor parameters such as the sampling rate or the revisit rate [1] or the emitted waveform of a radar [2]. The available sensors that are practical for this application fall into the following two classes — the class Φ_1 giving periodic but cluttered measurements, e.g., radar sensors, and the class Φ_2 giving aperiodic but uncluttered measurements that are split over multiple wireless channels, e.g., GPS receivers. We observe that the prevalent state estimation techniques using Φ_1 sensors, an example being the joint probabilistic data association filter (JPDAF) (see [3], [4], [5]), are incompatible with those using Φ_2 sensors, an example being the stochastic Kalman filter (see [6], [7]). Hence, the synthesis of a tracking protocol for a sensor-suite comprising such diverse sensors is a challenging problem. A reason behind the failure of a straight forward inclusion of a stochastic Kalman filter in the JPDAF technique is the modeling uncertainty. Much progress is possible by incorporating data-driven techniques in combining the belief-driven JPDAF and Kalman filter techniques, and we will present an overview of our approach.

Our development is around the so-called *covariance control* approach to *sensor management*, i.e., the assignment of a sensor to a target. Methods to synthesize a covariance controller to track multiple non-interacting targets using multiple Φ_1 sensors have been proposed in [8], [9] and have been extended for the case of interacting targets by [4]; a

Dr. Vishwesh V. Kulkarni and Prof. Lucy Y. Pao are with the Department of Electrical and Computer Engineering at the University of Colorado, Boulder, CO. Email: vishwesh@colorado.edu and pao@colorado.edu. This research has been supported in part by the Office of Naval Research (Grant N00014-02-1-0136), the Colorado Center for Information Storage, and the National Science Foundation (CMS-0201459).

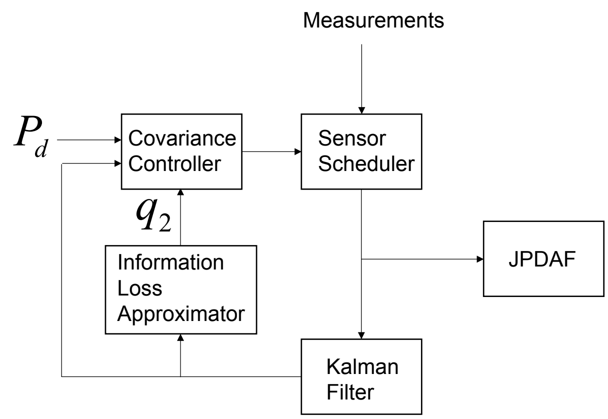

Fig. 1. A block diagram of the augmented covariance controller of [4] to track multiple targets using multiple sensors. The controller is described in further detail in Section III.

block diagram representation is given in Fig. 1. Incorporation of a Φ_2 sensor modifies the controller structure in that, since its measurement is received over multiple wireless channels which may not exhibit a strong correlation, a stochastic Kalman filter (see [6], [7]) replaces the standard Kalman filter. Recently, [6] has established upper bounds on the wireless channel losses under which the covariance update iteration of the stochastic Kalman filter converges and has synthesized an optimal linear quadratic Gaussian (LQG) controller under the assumption that the certainty equivalence principle (see [10]) holds.

The success of these techniques clearly depends on whether the certainty equivalence principle holds in practice, and, subsequently, on whether a Kalman filter-based state estimation technique is a reliable basis to synthesize the feedback controllers and to compute the allowable worst case channel loss probabilities as well. In this paper, we observe that the presence of modeling uncertainties can lead to poor performance when using such techniques. Incorporation of run-time data in the optimization process can alleviate this problem and we propose a supervisory controller with a data-driven outer loop encapsulating the belief-driven multi-sensor multi-target protocol of [9]; a conceptual block diagram is shown in Fig. 2. Further, we observe that the greedy need-based optimization algorithm implemented in the sensor scheduler of [9] can possibly be improved by casting the required optimization problem as a *market clearing* problem [11] so that some standard econometric auction algorithms may now be applied to solve the sensor management problem.

The paper is organized as follows. Section II presents the system description and the problem formulation. Section

III reviews relevant background results on the covariance controllers and the Kalman filter modifications. Section IV presents a counter-example to demonstrate the shortcomings of the traditional techniques to solve the estimation/control problems of interest and demonstrates the utility of a data-driven supervisory controller in overcoming these shortcomings. The paper is concluded in Section VI after a brief discussion in Section V on the use of auction algorithms to improve the performance of the sensor scheduler proposed by [9]. Some control theoretic notions on data-driven control are noted in the Appendix.

II. System Description and Problem Formulation

Every vehicle is equipped with Φ_1 and Φ_2 sensors. The target vehicle dynamics are assumed to be

$$x_t(k) = F_t x(k-1) + w_t(k-1),$$

and the measurement process is assumed to be

$$\begin{array}{ll} z_{jt}(k) = H_{jt}x(k) + v_{jt}(k) & \text{if the sensor } \in \Phi_1 \\ z_{jt}(k) = \gamma_j(k)\left(H_{jt}x(k) + v_{jt}(k)\right) & \text{else,} \end{array}$$

where $x_t(k)$ represents the state of the target t at time k, F_t represents the state transition matrix, z_{jt} represents the measurement of target t from the j-th sensor, H_{jt} represents the measurement matrix that creates z_{jt}, w_t represents unmodeled dynamics, v_{jt} represents the measurement noise, and $\gamma_{(\cdot)}$ is i.i.d. Bernoulli with mean λ. It is assumed that both w_t and v_{jt} are zero-mean white Gaussian with covariance matrices $Q_t \geq 0$ and $R_j > 0$, respectively. It is assumed that the measurement z_{jt} associated with a Φ_2 sensor has a representation $z_{jt} = [z_{j1t}^T \;\; z_{j2t}^T]^T$ where $z_{jt} = [H_{j1t}^T \;\; H_{j2t}^T]^T x_t + [v_{j1t}^T \;\; v_{j2t}^T]$ with

$$R \doteq \begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix},\; p(v_{jt}|\gamma_{jt}) = \left\{ \begin{array}{ll} N(0, R_{jj}), & \text{if } \gamma_{jt} = 1 \\ N(0, \sigma_{jt}^2 I), & \text{else,} \end{array} \right.$$

where $\sigma_{jt} \to \infty$ if the measurement is lost.

Given a vehicle equipped with Φ_1 and Φ_2 sensors, a tracking protocol can be implemented by incorporating the modified Kalman filter of [6] in the covariance controller of [9]. However, [6] assumes an i.i.d. Bernoulli distribution for the packet losses on a wireless link, and a perfect knowledge of the process matrices. In practice, the process matrices are liable to change in response to vehicle maneuvers and operating conditions, and, furthermore, the packet losses can rarely be described by a Bernoulli process and are correlated across the wireless links. Therefore, a problem of interest is to determine the robustness properties of a controller synthesized on the lines of [6] based on the state estimation using the covariance conroller of [9].

III. Background Results and Extensions

Since measurements from a Φ_1 include true target measurements as well as clutter measurements, these are often processed by a data association algorithm such as a joint probabilistic data association filter (JPDAF) [5]. The JPDAF filter tries to resolves the ambiguity using weighted averages of the presumed target-originated measurements and the clutter measurements, the weights being a function of the normalized distance of each measurement from the predicted target location [5, Ch. 6]. Since w_t and v_{jt} are zero mean white Gaussian, the estimated target state at the next time step is given by $\widehat{x}_t(k|k-1) = F_t\, \widehat{x}_t(k-1|k-1)$ and the estimated measurement state at the next time step is given by $\widehat{z}_{jt}(k) = H_{jt}\, \widehat{x}_t(k|k-1)$. The *innovation* $\nu_{tj\ell}$ is given by $\nu_{tj\ell}(k) \doteq z_{j\ell}(k) - \widehat{z}_{jt}(k)$ and the Mahalanobis distance [5] of each innovation is used to gate measurements that are likely to have originated from the target. The *combined innovation* is given by $\nu_{tj}(k) \doteq \sum_{\ell=0}^{m_{kj}} \beta_{tj\ell}(k)\; \nu_{tj\ell}(k)$ where $\beta_{tj\ell}$ is the probability that measurement $z_{j\ell}(k)$ is the true measurement of target t from sensor j, m_{kj} is the number of gated measurements from sensor j at time k, and $\ell = 0$ signifies the event that none of the gated measurements is the true measurement. The state prediction covariance and the innovation prediction covariances are

$$P_t(k|k-1) = F_t P_t(k-1|k-1) F_t' + Q_t(k-1) \quad (1)$$
$$S_{tj}(k) = H_j P_t(k|k, j-1) H_j' + R_j(k) \quad (2)$$

for all $j = 1, \ldots, N_{s_i}$, respectively, where $Q_t(k)$ is the process noise covariance for the t-th target and $R_j(k)$ is the measurement noise covariance from the j-th sensor, and N_{s_i} is the number of sensors in the i-th combination of sensors available to the multi-sensor manager. A sequential algorithm runs a separate filter for each sensor in the combination, propagating its state estimate to the next filter:

$$\begin{array}{rcll} \widehat{x}_t(k|k,1) & = & \widehat{x}_t(k|k-1) + K_{t1}(k)\nu_{t1}(k) & (3) \\ \widehat{x}_t(k|k,j) & = & \widehat{x}_t(k,j-1) + K_{tj}(k)\nu_{tj}(k) & (4) \\ \widehat{x}_t(k|k) & = & \widehat{x}_t(k|k,N_s) & (5) \\ \text{with } K_{tj}(k) & \doteq & P_t(k|k,j-1) H_j' S_{tj}^{-1}(k). & (6) \end{array}$$

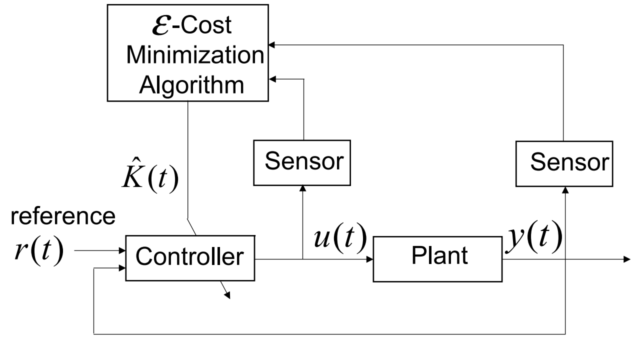

Fig. 2. The data-driven supervisory controller in the outer loop modifies the parameters of belief-driven controllers implemented in the inner loop of the system Σ. An example of the algorithms that may be implemented in the supervisory controller is the ε-cost optimization algorithm, described in this paper.

The state covariance update for each filter is given by [9]:

$$
\begin{aligned}
P_t(k|k,j) &= P(k|k,j-1) + \widetilde{P}_{tj}(k) \\
&\quad - (1-\beta_{tj0}(k))\, K_{tj}(k) H_j P_t(k|k,j-1) \quad (7)\\
P_t^i(k|k) &= P_t(k|k,N_{s_i}) \quad (8)\\
\widetilde{P}_{tj}(k) &\doteq K_{tj}(k) \sum_{\ell=1}^{m_{kj}} \beta_{tj\ell}(k)\nu_{tj\ell}(k)\nu'_{tj\ell}(k) \\
&\quad - K_{tj}(k)\nu_{tj}(k)\nu'_{tj}(k)K'_{tj}(k). \quad (9)
\end{aligned}
$$

Once the state and covariance estimates have been updated, they are fed back into the algorithm and the entire process is repeated for the measurement update at the next time step. To efficiently account for clutter, the sensor management algorithm computes the state prediction covariance and the innovation prediction covariance as in a sequential multisensor Kalman Filter algorithm. It then estimates the loss of information parameter q_2 by calculating the expected number of clutter measurements as a linear function of the gate volume and by using that information to compute the loss of information parameters, which are then passed on to the covariance controller. The covariance controller then ranks each target based on the need n_t, given by $n_t = -\psi_t \min\{\text{eig}\,(P_d(t) - P_t(k|k-1))\}$ where $P_d(t)$ is the desired covariance for the target t and ψ_t represents the target priority, with larger ψ_t indicating higher priority. The target with the largest need is then selected and the covariance that would result from each feasible sensor is calculated [9]. The sensor that maximizes

$$\min\left\{\text{eig}\left(P_d(t) - E\left[P_{tj}(k|k)|Z^{k-1}, P(k|k-1)\right]\right)\right\}, \quad (10)$$

where Z^{k-1} is the set of all measurements before time k, is selected, the covariance is updated $P_t(k|k,j) \rightarrow P_t(k|k-1)$, and the need is updated as $n_t = -\psi_t \min\{\text{eig}\,(P_d(t) - P_t(k|k,j))\}$. The remaining neediest target is selected and the process is repeated until either the need of each target is non-positive or the total tracking capacity of the sensors has been exhausted. An inclusion of Φ_2 sensors brings along random measurement losses so that the error covariance matrix and Kalman filter updates (1)—(9) are now stochastic. The detailed stochastic Kalman filter equations, along with bounds on the channel loss probabilities that ensure a bounded error for the state estimates, are derived in [6] and [7]. These assume an i.i.d. Bernoulli distribution for the packet losses on a wireless link and a perfect knowledge of the process matrices. A straight forward extension of the JPDAF algorithm (see [5], [3], [4]) and the covariance controller (see [9], [4]) to include the sensors from the class Φ_2 can be obtained by appropriately substituting the stochastic Kalman filter updates for the standard Kalman filter updates. The question of interest is whether such an extension will work in practice. In practice, the process matrices have an associated uncertainty; the factors contributing to the uncertainty in a single-lane lead-follower collision avoidance application (see [13]) are shown in Table I. Now, a direct extension

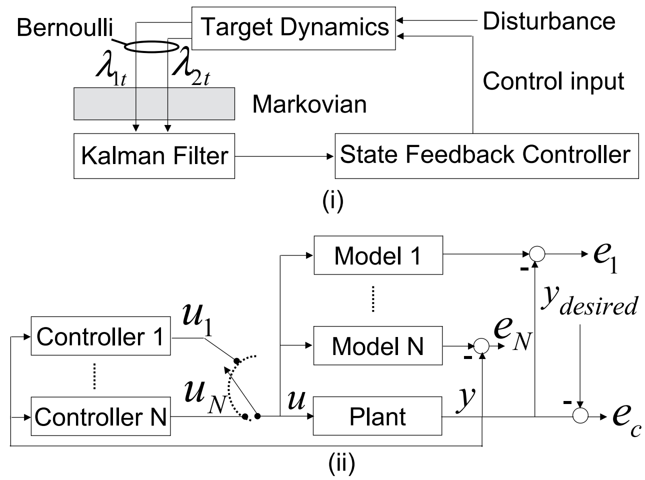

Fig. 3. (i) The full linear quadratic Gaussian (LQG) regulator of [6]. A target measurement is split into components and is sent to a remote controller over separate wireless channels, each of which incurs packet losses as an i.i.d. Bernoulli process. (ii) If a more realistic Markovian structure were imposed on the packet losses and/or if modeling uncertainties were considered, a multi-model adaptive controller (MMAC) of [12] is one of the standard solutions.

of the counterexample in [14] shows that, for even two channels having no packet losses, the state feedback controller in [6] has vanishingly small robustness margins. A first-cut extension to account for packet losses that cannot be described by a Bernoulli process is to assume that a wireless link can be described by a Markovian process with finitely many operating states, each described by a Bernoulli process (see Fig. 3). Then, the traditional approach to synthesizing the state-feedback controller is a multi-model adaptive controller (MMAC) [12]. We next present an example to demonstrate some practical shortcomings.

IV. ILLUSTRATIVE EXAMPLE

The following prototype counterexample demonstrates how the standard use of multiple models fails whereas the ε-cost minimization algorithm, given in the Appendix, succeeds in finding a stabilizing controller. For simplicity, let the Φ_2 measurement be available on a single wireless channel. Furthermore, suppose the wireless channel admits no packet losses, i.e., it can be described by a Bernoulli process with $\lambda = 1$. Let the structure of plant models and controllers be the same as in [12] with parameters $\left(\beta_0, \beta_1^T, \alpha_0, \alpha_1^T\right)^T$ for plant models and $\left(k, \theta_1^T, \theta_0, \theta_2^T\right)^T$ for controllers. Two candidate plant models and their corresponding controllers are designed so that their parameters are far from those of the true plant P^* and its corresponding controller C^*, with the parameter vectors as follows:

$$
\begin{aligned}
&P^* : (1,\ 0,\ -2,\ 0)^T; \quad C^* : (1,\ 0,\ 2,\ 0)^T;\\
&P_1 : (2,\ 0,\ 4,\ 0)^T; \quad C_1 : (0.5,\ 0,\ -2,\ 0)^T;\\
&P_2 : (1,\ 0,\ -6,\ 0)^T; \quad C_2 : (1,\ 0,\ 6,\ 0)^T.
\end{aligned}
$$

Such representations can be justified by introducing the uncertainties in Table I in the lead-follower models of [15], and will be discussed in detail in our future work. The control specification is assigned via the reference model

Variable	Distribution	Mean	Stdev	Minimum	Maximum
Reaction Time	Gaussian	1.1 sec	0.305 sec	0 sec	2 sec
System Delay	Impulse	0.2 sec	0 sec	0.2 sec	0.2 sec
Braking Deceleration	Gaussian	0.6 g	0.1 g	0.3 g	0.8 g

TABLE I

Factors contributing to modeling uncertainty in single-lane lead-follower systems. These models were parameterized in [13] based on the distributions used by NHTSA to test advanced collision warning systems [13].

$W_m(s) = 1/(s+3)$, while the unknown plant is $W_p(s) = 1/(s+5)$. The input is a step signal. The simulations are carried out with a dwell time of 0.001 sec. All initial conditions are zero. As in [12], the cost function to be minimized is $J(t) = e^2_{I_j}(t) + \int_0^t e^{-\nu(t-\tau)} e^2_{I_j}(\tau)d\tau, \quad j = 1,2$, where $e_I(t)$ is the identification error and $\nu = 0.05$. Fig. 4(i) shows the on-line values of the cost function for both identifiers, when either controller C_1 or C_2 is initially in the loop. C_2 is switched into the loop since it has smaller cost value than C_1 from the very beginning. However, C_2 is destabilizing, as can be confirmed by the analysis of the stability margins whereas C_1 is stabilizing [16]. The adaptive control method in [12] based on minimizing $J(t)$ fails to pick the stabilizing controller in this case and the cost for both controllers quickly explodes regardless of which controller is in the loop initially. Next, we used the ε-cost minimization algorithm with the cost function

$$J(t) = \max_{\ell \in (0,t)} \left\{ \frac{\tilde{e}_i^2(\ell) + \int_0^\ell e^{-\lambda(\ell-\tau)} \cdot \tilde{e}_i^2(\tau)d\tau}{\int_0^\ell e^{-\lambda(\ell-\tau)} \cdot \tilde{r}_i^2(\tau)d\tau} \right\}, \ i = 1,2, \tag{11}$$

where $\int_0^\ell e^{-\lambda(\ell-\tau)}\tilde{r}_i^2(\tau)d\tau \neq 0$, and $\tilde{r}_i, \tilde{e}_i$ are the fictitious reference signal and the fictitious error, respectively (see Appendix). At time $t = 0$, a controller is chosen arbitrarily and is put in the loop. The stabilizing controller C_2 is quickly switched into the loop. The parameter ε is set to 0.001. Fig. 4(ii) shows the simulation result of the unfalsified cost for both controllers: the cost of C_1 is much smaller than that of C_2, regardless of which controller is initially in the loop, and, hence, gets switched into the loop. The stabilizing controller C_1 is successfully chosen.

Remark 1: This counter-example demonstrates the utility of the unfalsified control methodology in using run-time data to synthesize a stabilizing feedback controller for a prototype system, which could not be stabilized using a traditional MMAC controller. The counter-example rests on the fact that the true process is not captured by any of the candidate models so that, as a result, the certainty equivalence principle is violated. The methodology can be expected to provide a robust state-feedback controller, which is an objective of [6], and to fine-tune the combined innovation parameter. □

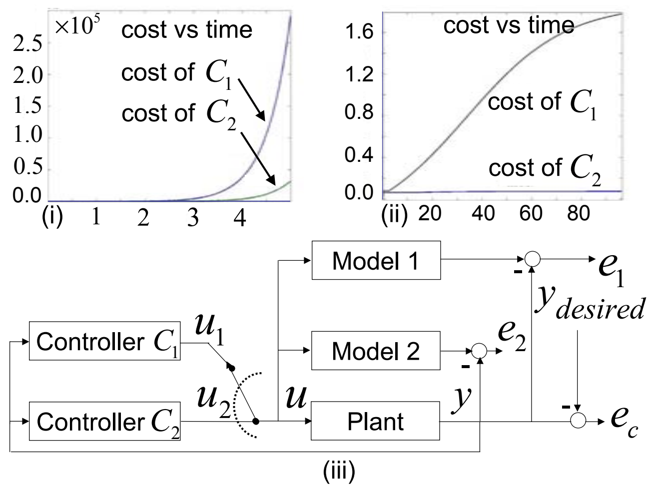

Fig. 4. (i) Cost of C_1 and C_2 under the MMAC method of [12]. (ii) Cost of C_1 and C_2 under the unfalsified method [16]. (iii) Plant-controller decomposition for the prototype system.

V. Sensor Management Algorithm Extensions

Auction algorithms on the lines of [11] and the Hungarian method [17] can improve on the need-based sensor-target assignment algorithm of [9], as follows. Consider the vehicle network collection as a market with N number of buyers, i.e., targets, and N number of goods, i.e., sensors. Suppose a target i is initially assigned a_{ij} amount of the sensor j with the total amount of the sensor j in the market as $a_j \doteq \sum_i a_{ij}$, and $a \doteq \sum_j a_j$. Let v_{ij} be the utility of the sensor j on the target i; an example of v_{ij} is the inverse of the predicted covariance of the estimate of target i when the sensor j is assigned to it. Suppose p_i denotes the price of the sensor i. Then, the scheduler would like to assign only the sensors that maximize v_{ij}/p_j. Let x_{ij} denote the amount of sensor j available with the target i. Let $\underline{\mathrm{P}}$ denote the price vector and let $\underline{\mathrm{X}}$ denote the matrix of x_{ij} assignments. The pair $(\underline{\mathrm{X}}, \underline{\mathrm{P}})$ forms a market equilibrium iff (a) there is no surplus or deficiency of any good; (b) all traders receive the goods that maximize their utility per money spent. The prices $\underline{\mathrm{P}}$ are known as the market clearing prices and $\underline{\mathrm{X}}$ is known as the equilibrium assignment. Then, the conditions for the market equilibrium are as follows:

$$\begin{cases} \sum_i x_{ij} = a_j,\ \forall j;\quad \sum_j x_{ij}p_j = \sum_j a_{ij}p_j,\ \forall i, \\ x_{ij} > 0 \text{ implies } v_{ij}/p_j \geq v_{ik}/p_k,\ \forall k;\quad x_{ij} \geq 0, p_j \geq 0. \end{cases}$$

Since the sensor prices p_j can be presumed to be fixed for most purposes, the market equilibrium conditions can be written as the solution to the following linear program:

$$\begin{aligned} \text{minimize } & \sum_i \sum_j a_{ij} p_j \alpha_i + \sum_j \alpha_j \beta_j \\ \text{subject to } & \alpha_i p_j + \beta_j \geq v_{ij},\ \forall i, j. \end{aligned}$$

An auction algorithm to realize the approximate market clearing is then as follows [11]. At any stage in the auction, each sensor j is available at two prices: $p_j/(1+\epsilon)$ and p_j, where $\epsilon > 0$ is a small fixed quantity chosen at the initialization stage. Let y_{ij} be the amount of sensor j assigned to the target i at the price $p_j/(1+\epsilon)$ and let h_{ij} be the amount assigned to the target i at the price p_j. Define the demand set D_i of the target i as $D_i = \arg\ \max_j v_{ij}/p_j$ and define the surplus r_i left with the target i as:

$$r_i = \sum_j a_{ij} p_j - \sum_j y_{ij} \frac{p_j}{1+\epsilon} - \sum_j h_{ij} p_j$$

with the total surplus $r = \sum_i r_i$. Let $a_* = \min_j a_j$ and $a = \sum_j a_j$. The sensor j is said to be unassigned if $\sum_i x_{ij} < a_j$ and assigned otherwise. It is said to be available at price p if its current price $p_j = p$ and $\sum_i h_{ij} < a_j$.

To start with, $p_j = 1,\ \forall j$ and $y_{ij} = h_{ij} = x_{ij} = 0,\ \forall i, j$. A target with positive surplus acquires the sensors in its demand set. If a sensor in the demand set is still unassigned, it is acquired at unit price. If a sensor j is available at its current price p_j, it is acquired by outbidding another target which has been assigned the sensor at a lower price $p_j/(1+\epsilon)$. If the sensor j is not available at its current price p_j, its price is increased by a factor of $1+\epsilon$ to make it available. The process continues until either the surplus of all targets becomes sufficiently small or all the sensors are assigned. The algorithm has been proven to terminate in $O(\frac{1}{\epsilon^2} N^2 \log(\frac{p^* a}{\epsilon a_*}) \log p^*)$ rounds [11], where $\max_j p_j \leq p^*$. The algorithm can be readily extended for collaborative sensor-target assignments in vehicular networks.

VI. Conclusion

We have described a sensor management protocol which a vehicle may use to track other vehicles and objects in its neighborhood using a sensor-suite, an individual element of which gives either periodic cluttered updates or aperiodic uncluttered updates. The use of Kalman filters has been a key component in the prevalent solutions to this tracking problem. The reliance of such techniques on the certainty equivalence principle was exploited to produce a counterexample demonstrating inadequacies of such techniques to provide the desired performance in the face of modeling uncertainties. The use of data-driven optimization algorithms in overcoming this difficulty was discussed. Such algorithms may be used in fine-tuning the augmented covariance controller of [9] and improving on the state feedback controller of [6]. The use of auction algorithms in the sensor-target assignments has also been overviewed.

VII. Acknowledgment

We thank Prof. Michael Safonov and Margareta Stefanovic (USC) for unfalsified control correspondence.

Appendix. Data-Driven Control Concepts

This section reviews some relevant notions for the problem of data-driven discovery of controllers that fit control goals, as outlined in [18], [19], [20], and [21]. A given phenomenon (plant, process) produces elements (outcomes) that reside in some set $\mathcal{Z}$ (universum). A subset $\mathcal{B} \subseteq \mathcal{Z}$ (behavior of the phenomenon) contains all *possible* outcomes. The mathematical model of the phenomenon is the pair $(\mathcal{Z}, \mathcal{B})$. The set $\mathcal{T}$ denotes an underlying set that describes the evolution of the outcomes in $\mathcal{B}$ (usually, the time axis). We distinguish between manifest variables $z_{manifest} \in \mathcal{Z}$ that completely describe the phenomenon, and latent (auxiliary) variables $z_{latent} \in \mathcal{Z}$; e.g., the plant input and output may serve as the manifest data $\{(u, y) \in L_{2e} \times L_{2e} \subset \mathcal{Z}\}$. We define the linear truncation operator $P_\tau : \mathcal{Z} \mapsto \mathcal{Z}_\tau$ as:

$$(P_\tau z)(t) = \begin{cases} z_{manifest}(t), & \forall t \leq \tau \\ o, & \text{else.} \end{cases}$$

This definition differs slightly from the usual definition of the truncation operator in that the truncation is performed with respect to both time and signal vector z. The *measured data set* [19] contains the observed (measured) samples of the manifest plant data, and is defined as $\{z_{data}\} = \{(y_{data}, u_{data})\} \subset B_{p_{true}}$ where $B_{p_{true}}$ is the behavior of the true plant. The available plant data at time τ is denoted $P_\tau(z_{data}) \subset P_\tau(B_{p_{true}})$. The set $\mathcal{K}$ denotes a finite set of candidate controllers. The *fictitious reference signal* $\tilde{r}(K, P_\tau z_{data}, \tau)$ is the reference signal that would have exactly reproduced the measured signals $P_\tau(z_{data})$ had the controller K been in the loop when the data was collected. Any adaptive control algorithm can be viewed as a mechanism that chooses a particular controller that minimizes a cost function. In a multiple-model/multiple-controller switching scheme, this function has a role of ordering candidate controllers according to the chosen criterion. A data-driven cost-minimization paradigm used here implies that the ordering of the controllers is based on the available plant data.

Definition 1: The *cost functional* is a mapping $V : P_\tau \mathcal{Z} \times \mathcal{K} \times \mathcal{T} \to \mathbb{R}_+$ for the given controller $K \in \mathcal{K}$, measured data $P_\tau z_{data} \in P_\tau \mathcal{Z}$ and $\tau \in \mathcal{T}$. □

Definition 2: The *true cost* $V_{true} : \mathcal{K} \mapsto \mathbb{R}_+ \cup \{\infty\}$ is defined as $V_{true}(K) \doteq \sup_{z \in B_{p_{true}}, \tau \in \mathcal{T}} V(K, P_\tau z, \tau)$. □

The true cost represents, for each K, the maximum cost that would be incurred if it were possible to perform a worst case experiment, for all possible experimental data. Let $[y_{data}, u_{data}]$ represent the output signals of the supervisory feedback adaptive system $\Sigma : L_{2e} \rightarrow L_{2e}$ in Fig. 2. It is assumed that all components of the system give zero output when subject to zero input.

Definition 3: A system with input w and output z is said to be *stable* if $\limsup_{\tau \to \infty} \|z\|_\tau / \|w\|_\tau < \infty$ holds for all $w \in L_{2e}, w \neq 0$. If, in addition, $\sup_{w \in L_{2e}, w \neq 0} (\|z\|_\tau / \|w\|_\tau) < \infty$, the system is said to be *finite-gain stable*; otherwise, it is said to be *unstable*. □

Definition 4: A robustly stabilizing and performing controller K_{RSP} is a controller that *stabilizes* the given plant and *minimizes* the true cost V_{true}. □

A *data-driven adaptive control law* is an algorithm that selects at each time τ a controller $\hat{K}_\tau$ dependent on experimental data. There are many ways to choose a controller (see [20] and [22]), and an example is as follows.

The ε-Cost Minimization Algorithm

$$\begin{cases} \text{1. Initialize: Let } t = 0, \tau = 0; \text{ choose } \varepsilon > 0. \\ \text{Let } \hat{K}_t \in \mathcal{K} \text{ be the first controller in the loop.} \\ \text{2. } \tau \leftarrow \tau{+}1. \\ \text{If } V\left(\hat{K}_t, P_\tau z, \tau\right) > \min_{K \in \mathcal{K}} V\left(K, P_\tau z, \tau\right) + \varepsilon \\ \text{then} \\ t \leftarrow \tau \text{ and } \hat{K}_t \leftarrow \arg\min_{K \in \mathcal{K}} V\left(K, P_\tau z, \tau\right) \\ \text{3. } \hat{K}_\tau \leftarrow \hat{K}_t; \text{ return } \hat{K}_\tau; \\ \text{4. go to step 2.} \end{cases}$$

The time instant t is the time of the last controller switch. The switch occurs only when the current unfalsified cost related to the currently active controller exceeds the minimum of the current unfalsified cost by at least ε. Here, ε serves to limit the number of switches to a finite number, and so prevents the possibility of limit cycle types of instability that may occur when there is a continuous switching between two or more stabilizing controllers. It also ensures a non-zero dwell time between switches. We assume that the candidate controller set $\mathcal{K}$ contains at least one robustly stabilizing and performing controller and that the performance cost functional V has a *monotone non-decreasing cost* in the sense that for all τ_1, τ_2 such that $\tau_2 \geq \tau_1$, and $\forall K \in \mathcal{K}$, $\forall z_{data}$ with which K is consistent, $V(K, P_{\tau_2} z, \tau_2) \geq V(K, P_{\tau_1} z, \tau_1)$. It may be observed that when V is monotonically non-decreasing in time, its optimal value $\min_{K \in K} V(K, P_\tau z, \tau)$ is monotonically non-decreasing in time and uniformly bounded from above by $V_{true}(K_{RSP}) \doteq \min_{K \in K} V(K, P_{\tau_1} z, \tau_1) \leq \min_{K \in K} V(K, P_{\tau_2} z, \tau_2)$, $\forall \tau_2 > \tau_1$ for all $z \in \mathcal{Z}$.

Definition 5: Given $K \in \mathcal{K}$ and measured data $[y_{data}, u_{data}]$ we say that the stability of the system given in Fig. 2 is *falsified* if

$$\exists \tilde{r}\,(K, z_{data}) \text{ s.t. } \limsup_\tau \left(\frac{\|[y_{data}, u_{data}]\|_\tau}{\|\tilde{r}\|_\tau} \right) = \infty.$$

Otherwise, it is said to be unfalsified. □

References

[1] G. Watson, W. Blair, and T. Rice, "Enhanced electronically scanned array resource management through multisensor integration," in *Proc. SPIE Conf. Signal Processing of Small Targets*, 1997, pp. 329–340.

[2] T. Kirubarajan, Y. Bar-Shalom, W. Blair, and G. Watson, "IMMPDAF for radar management and tracking benchmark with ECM," *IEEE Trans. Aerospace and Electronic Systems*, vol. 34, no. 4, pp. 1115–1132, Oct 1998.

[3] M. K. Kalandros, L. Trailović, L. Y. Pao, and Y. Bar-Shalom, "Tutorial on multisensor management and fusion algorithms for target tracking," in *Proc. of American Control Conference*, Boston, MA, June/July 2004.

[4] M. K. Kalandros and L. Y. Pao, "Covariance control strategies for reducing bias effects in interacting target scenarios," *IEEE Trans. Aerospace and Electronic Systems*, vol. 41, no. 1, Jan 2005.

[5] Y. Bar-Shalom and T. Fortman, *Tracking and Data Association*. San Diego, CA: Academic Press, 1988.

[6] X. Liu and A. Goldsmith, "Kalman filtering with partial observation losses," in *Proc. of American Control Conference*, Boston, MA, June/July 2004.

[7] B. Sinopoli, L. Schenato, M. Francescheti, K. Poolla, M. Jordan, and S. Sastry, "Kalman filtering with intermittent observations," in *Proc. Conference on Decision and Control*, Maui, HI, Dec 2003.

[8] M. K. Kalandros and L. Y. Pao, "Covariance control for multisensor systems," *IEEE Trans. Aerospace and Electronic Systems*, vol. 38, no. 4, pp. 1138–1156, Oct 2002.

[9] ——, "Covariance control for sensor management in cluttered tracking environments," *Journal of Guidance, Control, and Dynamics*, vol. 27, no. 3, pp. 493–496, March 2004.

[10] H. D. Water and J. Willems, "The certainty equivalence principle in stochastic control theory," *IEEE Trans. Automatic Control*, vol. 26, no. 5, pp. 1080–1087, 1981.

[11] R. Garg and S. Kapoor, "Auction mechanism for market equilibrium," in *Symposium on Theory of Computing*, Chicago, IL, June 2004.

[12] K. Narendra and J. Balakrishnan, "Adaptive control using multiple models," *IEEE Trans. Automatic Control*, vol. 42, no. 2, pp. 171–181, Feb 1997.

[13] L. Yang, J. Yang, E. Feron, and V. Kulkarni, "Development of a performance-based approach for a rear-end collision warning and avoidance system for automobiles," in *IEEE Proc. Intelligent Vehicles*, June 2003.

[14] J. Doyle, "Guaranteed stability margins of LQG regulators," *IEEE Trans. Automatic Control*, vol. 23, no. 4, pp. 756–757, 1978.

[15] X. Liu and A. Goldsmith, "Effect of communication delay on string stability in vehicle platoons," submitted to IEEE Trans. Automatic Control.

[16] M. Stefanovic, R. Wang, and M. Safonov, "Stability and convergence in adaptive systems," in *Proc. American Control Conference*, Boston, MA, June/July 2004, pp. 1923–1928.

[17] H. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, pp. 83–97, 1955.

[18] P. Brugarolas and M. Safonov, "A data-driven approach to learning dynamical systems," in *Proc. Conference on Decision and Control*, Las Vegas, NV, Dec 2002, pp. 4162–4165.

[19] J. Willems, "Paradigms and puzzles in the theory of dynamical systems," *IEEE Trans. Automatic Control*, vol. 36, no. 3, pp. 259–294, March 1991.

[20] T. Tsao, "Set theoretic adaptor systems," Ph.D. dissertation, University of Southern California, Los Angeles, CA, 1994.

[21] M. Safonov and F. Cabral, "Fitting controllers to data," *System and Control Letters*, vol. 43, no. 4, p. 2001, July 1998.

[22] M. Safonov and T. Tsao, "The unfalsified control concept and learning," *IEEE Trans. Automatic Control*, vol. 42, no. 6, pp. 2819–2824, 1997.