2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

$\mathscr{L}_1$-adaptive Methods for Control of Patient Response to Anesthesia

M. Ralph, C. L. Beck and M. Bloom, M.D

Abstract— **In this paper, we discuss the first application of recently developed $\mathscr{L}_1$-adaptive control methods for closed-loop control of anesthesia delivery during surgery. Our initial objective, described herein, is to design controllers that are robust to inter-patient variability, such that patients follow a prespecified Bispectral Index profile. The controllers are designed from identification-based models constructed from clinical trial data.**

I. INTRODUCTION

During surgery, the anesthesiologist constantly monitors and adjusts the delivery of anesthesia to the patient in an attempt to maintain a desired level of *sedation, analgesia* and *muscle relaxation*. At the same time, the anesthesiologist maintains ventilation parameters and monitors cardiovascular and respiratory functions Additionally, intra-operative blood samples are often taken and used to observe gas concentrations, blood-sugar levels, electrolyte concentrations and coagulation parameters. In short, the anesthesiologist performs an extremely complex role, namely, that of a multivariable feedback controller. A common long term research goal in this area is thus to incorporate partially automated anesthesia delivery into the process, allowing the anesthesiologist to concentrate on urgent safety-critical events that arise during surgery.

In order to implement model-based feedback control of anesthesia delivery, two primary needs are (1) adequate and appropriate means of sensing the patient's level of sedation, analgesia and muscle relaxation, and (2) mathematical models capturing the patient response to anesthetic agents. Over the past two decades, the *bispectral index* (BIS), a statistical index based on phase and frequency relations between the component frequencies in the electroencephalograph (EEG) recordings, has found significant acceptance as a measure of sedation level (see [13], [34], [14] and the references therein). The adequacy and extent of muscle relaxation or *neuromuscular blockade* can be evaluated effectively using different modes of electrical stimulation ([20], [28]). Whereas the means and methods for establishing and monitoring adequate sedation and NMB levels are by now fairly well-accepted, there is no standardized or generally accepted method for determining the state of analgesia at this time. Under the administration of sedatives, analgesia is continually influenced by external stimuli and the administration of analgesic drugs, and interaction, or synergy, between analgesics and sedatives is, for the most part, unavoidable. In this paper, the main focus of our discussion will be on controlling the *level of sedation* via automated feedback methods, in particular implementing novel $\mathscr{L}_1$-adaptive control techniques. Our long term goals include extending these methods to the full MIMO case described above.

Our initial objective, described herein, is to design a controller such that patients under anesthesia follow a prespecified BIS profile and, simultaneously, certain vital signs are maintained in a safe range, ensuring that the proper level of sedation is maintained throughout the surgery without putting the patient at risk. The BIS value is a single dimensionless number ranging from 0 to 100, where 0 corresponds to a silent EEG, and 100 corresponds to a patient being fully awake and aware. A BIS value between 60 and 40 is considered a viable level for general anesthesia, where the patient is not aware and surgery can be performed [33]. The manner in which the BIS level of any patient responds to the infusion and/or inspiration of anesthetic agents is not linear. In fact, the standard modeling paradigm commonly used to describe the relationships between anesthetic inputs and patient outputs (or effects) is that of *pharmaco-kinetic, pharmaco-dynamic (PK-PD) compartment models*, which consist of a linear time-invariant (LTI) system cascaded with a static nonlinearity, which captures the drug concentration-to-effect relationship (see [17], [29], [25], [18] for details). The resulting (*grey box*) mathematical models are inherently single-input single-output (SISO) and contain a system of ordinary differential equations plus a nonlinear function, representing the relations between the drug input function, the concentration of drug in the various compartments, and the effect of the drug on specific patient endpoints. Unfortunately, as these models are strictly SISO, they are incapable of capturing the effects of disturbances, drug synergies, or coupling among effects in the human body.

Alternatively, in more recent previous work it has been shown that patient response to anesthesia can be adequately captured by *multivariable* piecewise-linear models, with one linear model capturing patient response around an "awake" equilibrium state, and another linear model capturing patient response around a "sedated" equilibrium state [23], [2]. These models incorporate external stimuli inputs and certain vital sign outputs. In this paper, we derive a set of models that correspond to patient response to anesthesia and external stimuli in an "awake" state, and in a "sedated" state,

M. Ralph is a M.S. student in the department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA mralph2@illinois.edu

C. L. Beck is a faculty member in the department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA beck3@illinois.edu

M. Bloom, M.D., is with the School of Medicine of the New York University, in the department of Anesthesiology, New York, NY, USA Marc.Bloom@nyumc.org

This work has been partially supported by NSF under the ECCS program grant number ECS-0725708

respectively, for seven different patients' sets of data. We note that the models we derive are *not* compartment models and thus the states in these models are not required to be positive. Using these models, we discuss control synthesis and simulation results using the recent $\mathscr{L}_1$-adaptive control methods proposed in [5], which adapt quickly, ensure stable performance in the presence of model uncertainties, and achieve the desired BIS reference tracking objectives. This represents the first application of the $\mathscr{L}_1$-adaptive control methods to the anesthesia control problem. We show that this control architecture results in very small tracking error and is robust with respect to the variability found across different patient responses to anesthesia.

A. Prior work

A number of prior control efforts for the anesthesia problem have been completed and evaluated over the past 50 years. Schwilden and colleagues used median frequencies from EEG power spectra as one measure of sedative effect to develop PK-PD model-based adaptive feedback control of the anesthetic agents propofol, methohexital, and alfentanil during both clinical studies and for surgery [32], [30], [31]. A number of model-based closed-loop anesthesia control studies also have been published by Gentilini and colleagues [9], [8], [10], [11], [12]. In [9], semi-physiological models and rule-based controllers for the regulation of respiratory functions and mean arterial pressure (MAP) under administration of isoflurane (ISO) are described. The application of model predictive control schemes to regulate MAP during delivery of isoflurane is investigated in [8]. In one of the most comprehensive control implementations completed to date, Gentilini *et al.* proposed a control scheme for the regulation of MAP and sedation level using PK-PD models of the response to isoflurane [10], [11], in which the design of a cascaded IMC (internal model control) controller to regulate the sedative affects of anesthesia via the BIS level of the patient, and a three-observer-based state feedback controller to regulate MAP are proposed; the control designs are implemented in a loop-at-a-time manner. Mortier also considered control of sedation level via BIS monitoring in [24], where PK-PD model-based adaptive control of propofol is implemented in surgeries. More recently, Haddad, Hayakawa and Bailey have completed adaptive and neural network based control designs for the regulation of unconciousness under administration of propofol [15], [16]. Although this by no means represents an exhaustive discussion of prior work on closed-loop control of anesthesia, it presents the work most closely related to that discussed in this paper. However, all of the prior and ongoing work discussed above considers the use of SISO models and control designs, whereas our long term objective is to develop MIMO models and control designs.

The rest of the paper is organized as follows. We begin by briefly stating our problem formulation in Section II. In Section III we provide an overview of the methods and theory for $\mathscr{L}_1$-adaptive control. We then discuss the construction of the models we use in the control design process, including an overview of the identification methods and data used, in Section IV. Simulation results are then presented in Section V. Section VI summarizes our results and discusses our ongoing efforts. Our notation is standard.

II. Problem Formulation

Our objective is to develop a control design platform using system identification and advanced control methods so that patients' BIS levels track a desired reference trajectory. The controller must simultaneously compensate for the variability inherent in patient response to anesthesia and disturbances. Therefore, we require a control architecture that is both adaptive and robust.

In order to achieve these objectives, we implement the recently developed $\mathscr{L}_1$-adaptive control techniques described in [6]. In this framework, we begin by considering the following ISO to BIS system model:

$$y(s) = A(s)(u(s) + d(s)) \tag{1}$$

where $y(t) \in \mathbb{R}$ is the system BIS value and $u(t) \in \mathbb{R}$ is the controllable input; in our case, percentage concentration, by volume, of the inhalational sedative isoflurane. It is assumed that $A(s)$ is an unknown transfer function that is strictly proper. The $d(s)$ term is the Laplace transform of the (time-varying) disturbances $d(t)$; it is typically assumed that $d(t) = f(t, y(t))$ where $f(t, y(t))$ is Lipschitz, with Lipschitz constant L, but is otherwise an unknown mapping. The control objective is to design an adaptive output feedback controller such that the patients' BIS level $y(t)$ tracks the BIS reference input following a desired reference model:

$$y(s) \approx M(s)r(s) \tag{2}$$

For example, in our simulations we consider a first order system of the form

$$M(s) = m/(s+m),\ m > 0. \tag{3}$$

We then rewrite (1) as

$$\begin{array}{rcl} y(s) & = & M(s)(u(s) + \sigma(s)), \text{ where} \\ \sigma(s) & = & ((A(s) - M(s))u(s) + A(s)d(s))/M(s) \end{array} \tag{4}$$

In order to control and compensate for the varying responses to isoflurane among patients, we design adaptive controllers that ensure BIS-level tracking in both the transient and steady-state, while keeping all remaining error signals bounded.

III. $\mathscr{L}_1$-Adaptive Control: Background

The $\mathscr{L}_1$-adaptive control architecture was first proposed by Cao and Hovakimyan in [4], [5], and is intended for situations in which the controller must ensure that the system output follows a given reference signal in the face of modeling uncertainties. $\mathscr{L}_1$-adaptive control techniques have seen increasing application in the last few years, particularly in aerospace applications [26], [36]. These methods are useful in the same settings in which model reference adaptive control (MRAC) techniques could be considered, but also in settings in which *fast* adaptation is required. In this setting, $\mathscr{L}_1$-adaptive controllers prevent behaviors such

as high frequency oscillations in the control channel and parameter drifts from occurring, based on the implementation of a projection operator in the adaptive law and a low pass filter in the control law. Additionally, the $\mathscr{L}_1$-adaptive control architecture guarantees *uniformly bounded* asymptotic and transient tracking for the system inputs and outputs. These bounds, which are quantifiable, improve as the adaptation rate is increased.

The basic structure of the $\mathscr{L}_1$-adaptive controller is comprised of a predictor, a parameter adaptation law, and a feedback control law. The main elements are described in [6], which we excerpt in a summarized form below.

- **Closed-loop Reference System:** Consider the following closed-loop reference system:

$$y_{ref}(s) = M(s)(u_{ref}(s) + \sigma_{ref}(s)) \tag{5}$$

$$\sigma_{ref}(s) = \frac{(A(s) - M(s))u_{ref}(s) + A(s)d_{ref}(s)}{M(s)} \tag{6}$$

$$u_{ref}(s) = C(s)(r(s) - \sigma_{ref}(s)) \tag{7}$$

  where $d_{ref}(s)$ is the Laplace transform of $f(t, y_{ref}(t))$, and $C(s)$ is a strictly proper transfer function with $C(0) = 1$. $C(s)$ is a low pass filter used to attenuate high frequency content in the control channel resulting from uncertainty. The simplest choice for $C(s)$ would be a first order system

$$C(s) = \omega/(s + \omega), \tag{8}$$

  which we use in our simulations.
- **$\mathscr{L}_1$ Stability Condition:** Choices for $M(s)$ and $C(s)$ are restricted such that

$$H(s) = \frac{A(s)M(s)}{(C(s)A(s) + (1 - C(s))M(s))} \tag{9}$$

  is BIBO stable and

$$\|G(s)\|_{\mathscr{L}_1} L < 1, \; G(s) = H(s)(1 - C(s)). \tag{10}$$

  If $M(s)$ and $C(s)$ are chosen such that (9) and (10) hold, then the closed-loop reference system in (5), (6), and (7) is BIBO stable. When $C(s)$ is chosen to be first order, then stability of $H(s)$ is equivalent to stabilization of $A(s)$ by a PI controller [6].
- **Output Predictor:** Consider the output predictor [6]:

$$\dot{\hat{y}}(t) = -m\hat{y} + m(u(t) + \hat{\sigma}(t)), \hat{y}(0) = 0 \tag{11}$$

  where the adaptive estimate $\hat{\sigma}(t)$ is used to account for unknown disturbances and uncertainty. It is governed by the adaptive law given in the following.
- **Adaptive Law:** The adaptive estimate $\hat{\sigma}(t)$ is given by

$$\dot{\hat{\sigma}}(t) = \Gamma \text{Proj}\left(\hat{\sigma}(t), -mP\tilde{y}(t)\right), \; \tilde{y} = \hat{y} - y, \; \hat{\sigma}(0) = 0, \tag{12}$$

  where $P > 0$ is arbitrary, and Γ is the adaptation rate which is subject to a fixed lower bound (given in [6]). The projection bound is

$$|\hat{\sigma}(t)| \leq \Delta \tag{13}$$

  The projection operator "Proj" is essentially of least-squares form, and ensures that the parameter estimate $\hat{\sigma}(t)$ remains inside a required compact set Σ; see [7] for details.
- **Control Law:** The control law that generates the ISO input signal is given as:

$$u(s) = C(s)(r(s) - \hat{\sigma}(s)) \tag{14}$$

  where $r(s)$ is the Laplace transform of our reference signal, $r(t)$.
- **Error Performance:** The following upper bound holds for all $t \geq 0$ [6]:

$$\|\tilde{y}(t)\|_{\mathscr{L}_\infty} \leq \frac{k}{\sqrt{\Gamma P}}, \; \forall t \geq 0 \tag{15}$$

  where $\tilde{y}(t) = \hat{y}(t) - y(t)$ and k is a constant (details for this constant can be found in [6]). As can be seen in (15) the tracking error is uniformly bounded by a constant that is inversely proportional to $\sqrt{\Gamma P}$. The higher the adaptive gain value used, the lower the tracking error achieved.

Using clinical trial data, we have constructed a set of data-based models in the form of (1). We then apply the $\mathscr{L}_1$-adaptive techniques outlined in the preceding, so that the patients' BIS levels track a prespecified reference trajectory. We first provide a brief overview of the approach used to derive the patient response models from the clinical trial data.

IV. Modeling and System Identification

As this is an initial investigation of the applicability of $\mathscr{L}_1$-adaptive techniques, we have used clinical trial data from earlier studies ([23], [2], [22]) in order to facilitate the comparison of controller performance results. The original clinical trial was designed to define the relation between clinical evaluation of the state of conciousness, explicit recall, drug concentrations and BIS effects of the anesthetic agent isoflurane when administered alone to healthy volunteers under controlled conditions. Additionally, a series of external stimuli, or disturbances, were applied to the patients (volunteers) throughout the administration of anesthesia. These stimuli included: laryngeal mask insertion and removal (LMA); evoked potential evaluations (EP); and *alertness evaluations* (EVAL) which included yelling at, shaking, and squeezing the trapezius muscle of the volunteer. Time-synchronized output measurements of the patients' BIS, MAP and HR were recorded every two seconds. For healthy individuals, normal ranges for MAP are between 70 and 110 mmHg, and the average resting HR for normal adults is around 70 beats per minute [3]. We previously developed quantitative models of the stimuli applied to the patients during the study [23] for use in system identification.

An example of a set of data taken from one subject during the clinical trial is shown in Figures 1 and 2. This data is fairly representative of the response expected from healthy volunteers to anesthesia and stimuli, however, as to be expected individual responses exhibit noticeable variation.

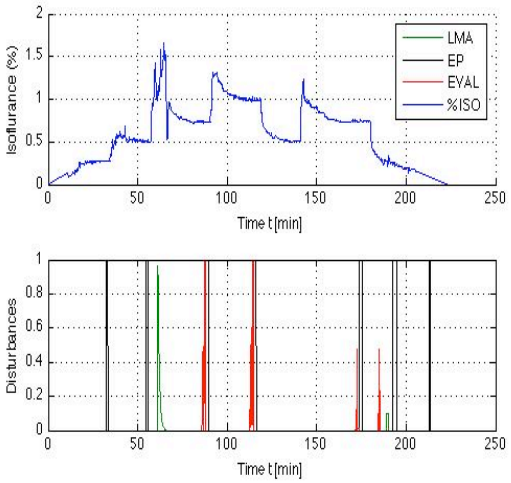

Fig. 1. Isoflurane and Stimuli Inputs versus Time

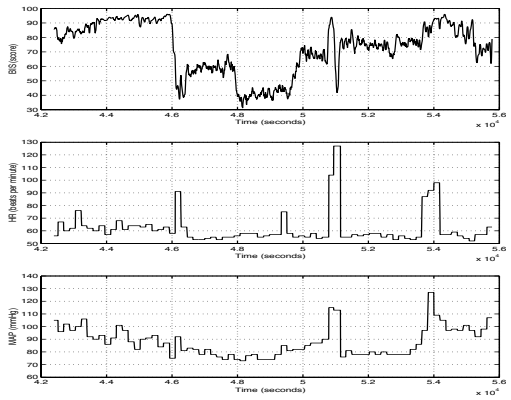

Fig. 2. BIS, HR and MAP Outputs versus Time

This inter-patient variability is one of the main motivations for considering adaptive control techniques.

For our purposes, *black-box* models were found to be sufficient for control design. Since the $\mathscr{L}_1$-adaptive control architecture relies on the transfer function, the specific structure of the A, B, and C matrices are not required to be known. Subspace identification methods were used to calculate the models for 6 patients which we have designated Patients 1, 2, 3, 5, 6, and 7. We again note that these models are not restricted to the grey-box PK-PD compartment model structures used in the earlier studies cited in the Introduction. Estimation and validation results for these models are acceptable, with an overall average normalized residual error of approximately 29.5%. One example is shown in Figure 3.

In the initial stage of the current adaptive control study, we have constructed models and adaptive controllers for SISO systems (isoflurane input to BIS output), for which we evaluate and compare the modeling and controller performance results. We specifically focus on evaluating inter-patient adaptability of the controllers, and along these lines

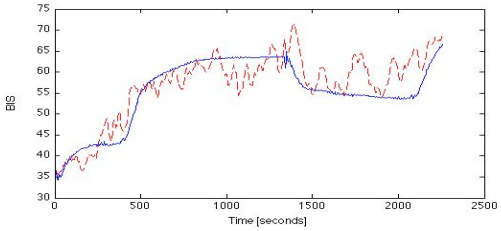

Fig. 3. Patient 5 model validation results

include performance analyses of applying controllers designed based on the model for one patient to other patient models. The second stage of this study is identification and adaptive control design evaluation for MISO systems, namely, where external stimuli are included as disturbance inputs, in addition to the controlled isoflurane input.

The final stage of this project involves the construction of MIMO models and the application of multivariable $\mathscr{L}_1$-adaptive techniques to these models. This stage includes evaluating vital sign responses (HR and MAP, for example) as well as BIS responses. We note here that our focus is on automated control of patients primarily in the sedated state. Our assumption is that the attending physician performs the initial induction from alert to the lightly-sedated state in order to closely monitor initial patient response. Upon being lightly sedated and observed for safety reasons, the patient is then switched to the proposed automated control regime.

V. Simulation Results: Analysis and Discussion

We provide details on results found for one of six control designs. Details for the remaining five control designs can be found in [27].

A. Patient 1 $\mathscr{L}_1$-Adaptive Control-Output Feedback

As noted above, models were constructed using subspace identification methods on detrended, partitioned patient data (partitioning was for estimation and validation purposes). Simulink was then used to simulate the closed-loop systems.

Tracking, no disturbances: A fourth-order system was identified for *Patient 1-sedated state.*

Note that we have assumed there is no time delay in the system in our simulations. For the $\mathscr{L}_1$-adaptive controller defined by (11), (12), and (14), we selected $P = 1$, $\Delta = 100$ and $\Gamma = 50000$ as conservative initial parameter values. The $\mathscr{L}_1$-adaptive controller is then applied to track a given reference BIS trajectory, $r(t)$. If for $M(s)$ in (3) we set $m = 1/30$ and we set $\omega = 0.001$ for $C(s)$ in (8) , we can show that $H(s)$ in (9) is strictly proper and BIBO stable as required. These values are selected using a combination of classical stability analysis methods and trial and error. Based on the selected parameter values, we expect the patient to reach the desired BIS level in approximately two minutes, with very small error between $y(t)$ and $\hat{y}(t)$, and no parameter drift. Simulation results for the $\mathscr{L}_1$ adaptive controller applied to the dynamic response model for Patient 1 verify these expectations; see Figure 4.

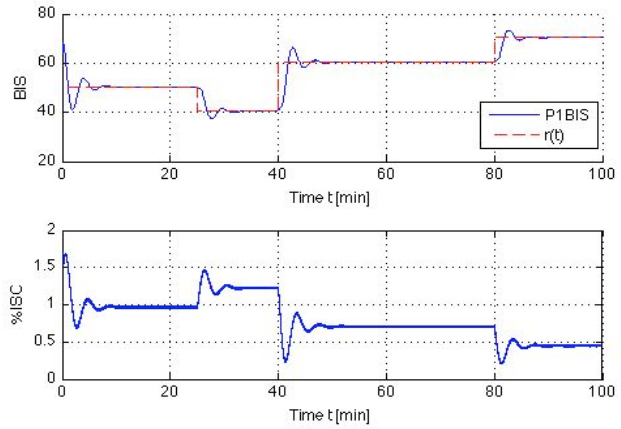

Fig. 4. Patient 1 output feedback controller with $\Gamma = 50000$

Note that the $\mathscr{L}_1$ adaptive output feedback control design achieves the desired tracking performance quickly, and the tracking error is small (Figure 4). The metric used to evaluate BIS reference tracking performance is the normalized mean square error, or residual error term,

$$r_n = \frac{\sum_{n=1}^{N}(y_n - \hat{y}_n)^2}{\sum_{n=1}^{N} y_n^2} \tag{16}$$

where the reference signal y and the patient's BIS value $\hat{y}$ contain N data points over the course of the simulation. The residual error, r_n for Patient 1 with the previously discussed parameters is 0.0019.

A second metric used to evaluate controller performance is the total amount of isoflurane required to attain the tracking performance. (Note that high total amounts of anesthetics used in surgery have been associated with negative long-term patient outcomes [1], [19]). Inhalational anesthetics are delivered as a percentage by volume to an external respiratory circuit. This gas mixture is then delivered to the lungs by a ventilation system, measured in liters/minute. So, for example if the total fresh gas flow mixture from the machine is $2-4$ l/min, and the isoflurane is 1% (at an assumed mean alveolar concentration (MAC) of 1), then $20-40$ cc/min of isoflurane gas is being delivered to the respiratory circuit. To estimate the following relative quantities of isoflurane used, we assumed an average fresh gas flow mixture delivery rate of 3 l/min at 1 MAC. The approximate isoflurane use for Patient 1 is 2.4 liters; isoflurane use for a comparable BIS reference and simulation time resulting from our previous linear parameter-varying (LPV) control designs was approximately 3.1 liters. The smaller volume of isoflurane used should lead to improved patient outcomes. We now discuss the MISO performance of our designs based on simulations run with disturbance inputs.

Tracking, with disturbances: We treat the EP, EVAL, and LMA inputs as disturbance inputs to the ISO/BIS system. Transfer functions for the disturbance models for the *Patient 1-sedated state* were also obtained through subspace identification methods.

Figure **??** shows the closed-loop performance of the $\mathscr{L}_1$ adaptive controller used on Patient 1 when disturbances are introduced. The normalized residual error computed using (16) is 0.0022 and the volume of isoflurane used is approximately 2.4 liters. For both simulations with and without disturbances, the system has a non-zero initial condition which leads to an exponentially decaying term in the control and system state signal; this does not affect the performance of the system throughout. The same design approach was followed for Patients 2,3, 5, 6 and 7, and yielded similar results.

MAP performance, with disturbances: We do not expect the mean arterial pressure to track a specific reference trajectory. Instead, during surgery, the anesthesiologist aims to maintain the patient's MAP in a given range. In our study, our goal is for the MAP to stay in the 60 to 110 range.

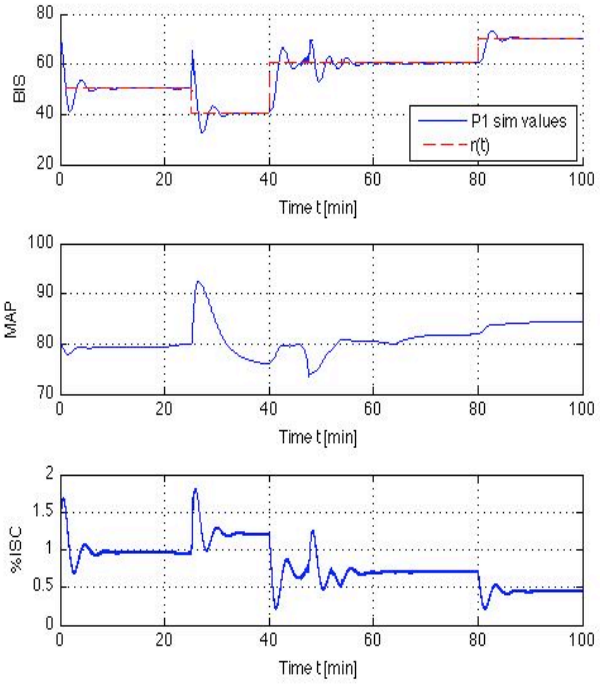

Fig. 5. BIS reference tracking: Patient 1 BIS and MAP performance

Figure 5 illustrates the BIS and MAP performance achieved when the $\mathscr{L}_1$ adaptive output feedback controller designed as described previously is applied to the model for Patient 1; the MAP values are well within the desired range throughout the entire simulation.

Robustness to Inter-patient Variability: The controller designed for Patient 1 was then simulated using the models for Patients 2, 3, 5, 6, and 7 to evaluate inter-patient design robustness. In general, the patients' BIS levels tracked the desired reference profile closely, regardless of whether the controller design was that specifically designed for that

model or not. Tables I and II detail the tracking performance and control effort when the controller designed for Patient 1 is applied to the models of Patients 2, 3, 5, 6, and 7.

TABLE I

NORMALIZED RESIDUAL ERRORS WITH P1 CONTROL, NO DISTURBANCES

Patient	Control 1
2	0.0062
3	0.0022
5	0.0016
6	0.0036
7	0.0064

TABLE II

ISOFLURANE USE IN LITERS WITH P1 CONTROL, NO DISTURBANCES

Patient	Control 1
2	2.859
3	2.232
5	2.400
6	3.292
7	2.147

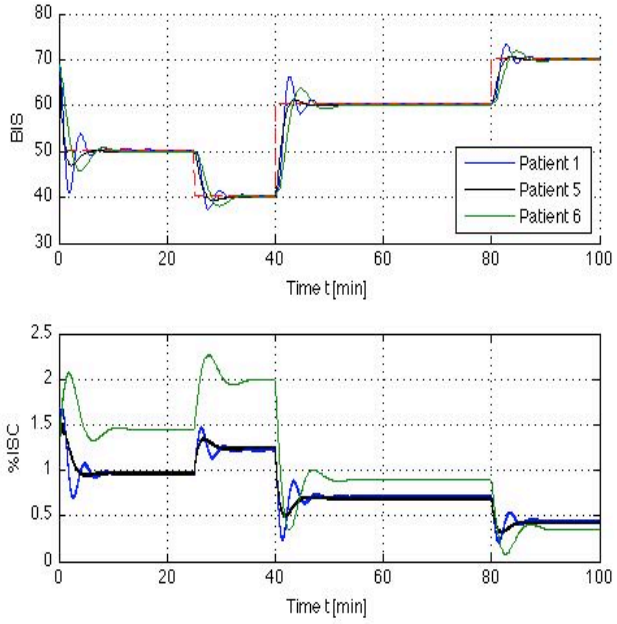

Fig. 6. BIS reference tracking: Patient 1 robustness to inter-patient variability, no disturbances

Figure 6 illustrates the robustness this controller design achieves to inter-patient variability. It is important to note that the controller used is *exactly* the same for all patients.

VI. CONCLUSIONS AND ONGOING RESEARCH

In this paper, we present the first application of recently proposed $\mathscr{L}_1$ adaptive control techniques to the anesthesia control problem. The acheived tracking performance and inter-patient robustness of these control methods, in conjunction with the relatively effective use of isoflurane, is extremely encouraging, and MIMO design efforts incorporating surgical stimuli inputs and additional performance objectives on patient vital signs are ongoing. Control designs and performance results for additional patient data sets will be incorporated into the final paper, and complete comparisons to previous design results will be fully discussed. Higher order filters in the $\mathscr{L}_1$ design process could be explored to determine if they improve performance. While our results demonstrate robustness of this control architecture, it remains to be determined what yields the best performance for the greatest number of patients. Surgical data will be used to test this control approach under harsher conditions, (i.e. greater number of anesthetic agents, disturbances, and actual surgical stimuli events).

REFERENCES

[1] J.P. Abenstein and M.A. Warner. Anesthesia providers, patient outcomes, and costs. *Anesth. Analg.*, 82:1273–1283, 1996.

[2] C. L. Beck, H-H. Lin, and M. Bloom. Modeling and control of anesthetic pharmacodynamics. In S. Tarbouriech, I. Queinnec, G. Garcia, and S. Niculescu, editors, *Biology and Control Theory: Current Challenges*, volume 357. Springer-Verlag Lecture Notes Series, 2007.

[3] R. Berne, M. N. Levy, B. M. Koeppen, and B. A. Stanton. *Physiology*. Mosby, Inc., St. Louis, MI, 4th edition, 1998.

[4] C. Cao and N. Hovakimyan. Stability margins of l_1 adaptive controller: Part ii. In *ACC*, pages 3931–3936, 2007.

[5] C. Cao and N. Hovakimyan. Design and analysis of a novel l_1 adaptive control architeture with guaranteed transient performance. *ieeeTAC*, 53(2):586–591, 2008.

[6] C. Cao and N. Hovakimyan. l_1 adaptive output feedback controller for systems of unknown dimension. *ieeeTAC*, 53(3):815–21, 2008.

[7] C. Cao, V. V. Patel, K. Reddy, N. Hovakimyan, and E. Lavretsky. Are the phase and time-delay margins always adversely affected by high-gain? In *Proc. of AIAA Guidance, Navigation and Control Conference*, AIAA-2006-6347.

[8] C. W. Frei, M. Derighetti, M. Morari, A. H. Glattfelder, and A. M. Zbinden. Improving regulation of mean arterial blood pressure during anesthesia through estimates of surgery effects. *IEEE Trans. on Biomed. Eng.*, 47(11):1456–1464, Nov. 2000.

[9] C. W. Frei, A. Gentilini, M. Derighetti, A. H. Glattfelder, M. Morari, T. W. Schnider, and A. M. Zbinden. Automation in anesthesia. In *Proc. American Control Conference*, pages 1258–1263, 1999.

[10] A. Gentilini, C. W. Frei, A. H. Glattfelder, M. Morari, T. J. Sieber, R. Wymann, T. W. Schnider, and A. M. Zbinden. Multi-tasked closed loop control in anesthesia. *IEEE Eng. in Med. and Bio. Magazine*, 20(1):39–53, Jan. 2001.

[11] A. Gentilini, M. Rossoni-Gerosa, C. Frei, R. Wymann, M. Morari, A. Zbinden, and T. Schnider. Modeling and closed loop control of hypnosis by means of bispectral index (BIS) with isoflurane. *IEEE Trans. on Biomed. Eng.*, 48(8):874–889, Aug. 2001.

[12] A. Gentilini, C. Schaniel, M. Morari, C. Bieniok, R. Wymann, and T. Schnider. A new paradigm for the closed-loop intraoperative administration of analgesics in humans. *IEEE Trans. on Biomed. Eng.*, 49(4):289–299, April 2002.

[13] P. S. Glass, M. Bloom, L. Kearse, C. Rosow, P. Sebel, and P. Manberg. Bispectral analysis measures sedation and memory effects of propofol, midazolam, isoflurane, and alfentanil in healthy volunteers. *Anesthesiology*, 86(4):836–847, April 1997.

[14] B. Guignard. Monitoring analgesia. *Best Practice & Research Clinical Anaesthesiology*, 20(1):161–180, 2006.

[15] W. M. Haddad, T. Hayakawa, , and J. M. Bailey. Adaptive control for nonnegative and compartmental dynamical systems with applications to general anesthesia. *Int. J. Adapt. Control Signal Process*, 17:209–235, 2003.

[16] T. Hayakawa, W. M. Haddad, N. Hovakimyan, and V. Chellaboina.

Neural network adaptive control for nonlinear nonnegative dynamical systems. *IEEE Trans. Neural Networks*, 16:399–413, 2005.

[17] N. H. G. Holford and L. B. Sheiner. Understanding the dose-effect relationship: Clinical application of pharmacokinetic-pharmacodynamic models. *Clinical Pharmacokinetics*, 6:429–453, 1981.

[18] J. R. Jacobs and E. A. Williams. Algorithm to control effect compartment drug concentrations in pharmocokinetic model-driven drug delivery. *IEEE Trans. on Biomed. Eng.*, 40(10):993–999, Oct. 1993.

[19] B. Kabon and A. Kurz. Optimal perioperative oxygen administration. *Current Opinion in Anesth.*, 19:11–18, 2006.

[20] I.S. Kalli. Neuromuscular block monitoring. *Neuromuscular Monitoring in Anesthesia*, (13):442–451, 2003.

[21] A. Kurz and D. Sessler. Opioid-induced bowel dysfunction: pathophysiology and potential new therapies. *Drugs*, 63(7):649–671, 2003.

[22] H. H. Lin, C. L. Beck, and M. Bloom. Multivariable lpv control of anesthesia delivery during surgery. In *ACC*, pages 825–831, 2008.

[23] H. H. Lin, C. L. Beck, and M. J. Bloom. On the use of multivariable piecewise-linear models for predicting human response to anesthesia. *IEEE Trans. on Biomedical Engineering*, 51(11):1876–1887, 2004.

[24] E. Mortier, M. Struys, T. De Smet, L. Versichelen, and G. Rolly. Closed-loop controlled administration of propofol using bispectral analysis. *Anesthesia*, 53(8):749–754, 1998.

[25] KT. Olkkola and H. Schwilden. Adaptive closed-loop feedback control of vecuronium-induced neuromuscalar relaxation. *European Journal of Anaesthesiology*, 8(1):7–12, Jan. 1991.

[26] V. Patel, C. Cao, N. Hovakimyan, K. Wise, and E. Lavretsky. L1 adaptive controller for tailless unstable aircraft in the presence of unknown actuator failures. *International Journal of Control*, pages 705–720, 2009.

[27] M. Ralph. l_1-Adaptive Control for Anesthesia Delivery, 2010.

[28] A.M. Sardesai and R. Griffiths. Monitoring techniques: neuromuscular blockade. *Anesthesia & intensive care medicine*, 6(6):198–200, 2005.

[29] T. W. Schnider, C. F. Minto, and D. R. Stanski. The effect compartment concept in pharmacodynamic modelling. *Anesthetic Pharmacology Review*, 2(3):204–213, 1994.

[30] H. Schwilden and H. Stoeckel. Effective therapeutic infusions produced by closed-loop feedback control of methohexital administration during total intravenous anesthesia with fentanyl. *Anesthesiology*, 73(2):225–229, Aug. 1990.

[31] H. Schwilden and H. Stoeckel. Closed-loop feedback controlled administration of alfentanil during alfentanil-nitrous oxide anaesthesia. *British Journal of Anaesthesia*, pages 389–393, 1993.

[32] H. Schwilden, H. Stoeckel, and J. Schüttler. Closed-loop feedback control of propofol anaesthesia by quantitative EEG analysis in humans. *British Journal of Anaesthesia*, pages 290–296, 1989.

[33] P. Sebel, E. Lang, I. J. Rampil, P. F. White, R. Cork an d M. Jopling, N. T. Smith, P. S. Glass, and P. Manberg. A multicenter study for bispectral electroencephalogram analysis for monitoring anesthetic effect. *Anesthesia and Analgesia*, pages 891–899, 1997.

[34] P. S. Sebel. Bispectral monitoring technology: Clinical applications. In T.H. Stanley and T.D. Egan, editors, *Anesthesia for the New Millennium - Modern Anesthetic Clinical Pharmacology*, pages 99–104. Kluwer Academic Publishers, 1999.

[35] D. R. Wada and D. S. Ward. The hybrid model: A new pharmacokinetic model for computer-controlled infusion pumps. *IEEE Trans. on Biomed. Eng.*, 41(2):134–142, Feb. 1994.

[36] J. Wang, V. Patel, C. Cao, N. Hovakimyan, and E. Lavretsky. Novel l1 adaptive control methodology for aerial refueling with guaranteed transient performance. *AIAA Journal of Guidance, Control and Dynamics*, pages 182–193, 2008.

[37] E. Xargay, N. Hovakimyan, and C. Cao. $\updownarrow_1$ adaptive controller for multi-input multi-output systems in the presence of uncertain nonlinear cross-coupling. In *ACC*, Submitted, 2010.