Experimental Studies on Automated Regulation of Hemodynamic Variables

Ramesh R. Rao†, Cesar C. Palerm‡, Brian Aufderheide† and B. Wayne Bequette† *

Abstract— **Two different control methodologies are developed for automated regulation of hemodynamic variables. These controllers are designed to regulate mean arterial pressure (MAP) and cardiac output (CO) in critical care subjects using inotropic and vasoactive drugs. Both controllers account for inter– and intra–patient variability and handle drug infusion constraints. The first approach is a multiple model predictive controller (MMPC). The algorithm uses a multiple model adaptive approach in a model predictive control framework to account for variability and explicitly handle drug rate constraints. The second approach, a robust direct model reference adaptive controller (DMRAC) is developed for plants with uncertainty in both the time delay elements and in the transfer function coefficients, such as the drug infusion process. Further modifications are introduced to handle drug rate constraints. The controllers are experimentally evaluated on canines that are pharmacologically altered to exhibit symptoms of hypertension and depressed cardiac output.**
Keywords: Biomedical Control Systems, Drug Infusion, Model Predictive Control, Adaptive Control, Dog Experiments.

I. INTRODUCTION

Critical care patients such as those in intensive care or undergoing surgery require close monitoring of all their vital signs. The anesthesiologist or critical care physician must monitor and regulate a wide range of physiological states such as mean arterial pressure, cardiac output, carbon dioxide and oxygen levels, blood acidity, fluid levels, heart contractility, renal function and more. Some physiological states cannot be measured directly and must be inferred by the physician. Physicians maintain patient states within acceptable operating ranges by infusing several drugs and/or intravenous fluids. In addition, during surgical procedures physicians must administer anesthetics and monitor the depth of anesthesia.

Current clinical practice involves manual regulation of drip IV lines to infuse drugs. Programmable pumps are also used to either deliver the drugs at a constant rate or a variable rate to achieve a desired concentration. Such pumps are based on averaged pharmacokinetic data and are essentially open–loop (i.e. there is no feedback of the patient's states), requiring regular intervention by the attending physician or nurse to adjust the drug flow rates. It is desirable to have an automated system that closes the loop on primary variables, but monitors secondary variables and helps the physician perform diagnostics. This would allow the physician to spend more time monitoring the patient for conditions that are not easily measured and assure that the physician is always "in the loop". The physician would use her expertise to diagnose the patient, specify set points or ranges of values for the states to be regulated, choose the drugs best suited to obtain the objective, and mandate permissible infusion rates; this information would then be explicitly used by the controller to automate the regulation of physiological states.

In the case of critical care patients with some degree of congestive heart failure (weak heart), the measured variables that are of primary importance are mean arterial pressure and cardiac output. Secondary variables that are monitored, but not regulated as tightly as the primary variables, include heart rate and pulmonary capillary wedge pressure. In this work we present two control strategies for the regulation of MAP and CO using vasoactive and inotropic drugs and the corresponding experimental results on the evaluation of the controllers on canines.

The vast amount of research on blood pressure control was initiated by Slate et al. [1] who used a PID controller with empirical tuning rules to control MAP using SNP. Since then, more complex control schemes have been used in automation of hemodynamic regulation. Isaka and Sebald [2] provide a comprehensive review of the single–input single–output (SISO) system; they highlight the still unresolved issues of safety and robustness, which are even more important when dealing with multiple drugs to control more than one output. Martin et al. [3] and Kwok et al. [4] have used adaptive control methodologies for blood pressure regulation during surgery. Additional research efforts have been made in the simultaneous regulation of MAP and CO. Serna et al. [5] reported on the simultaneous control of CO and MAP using DPM and SNP. As the CO measurements were available at a rate much slower than MAP, they decoupled the DPM–CO loop from the MAP–SNP loop. One of the more advanced studies in the two–input two–output system was done by Voss et al. [6] on canines. They used the Control Advance Moving Average Controller (CAMAC) which is a class of extended horizon controllers, with a recursive least squares estimate of model parameters. The advantages to this controller are its robustness and ability to handle systems with varying and unknown dead times. Yu et al. [7] used a multiple model adaptive control approach (MMAC) for regulating MAP and CO in canine experiments. Held and Roy [8] developed an expert system that used a fuzzy controller for controlling MAP and CO using SNP and DPM. No clinical work that the authors are aware of has been done on controlling multiple hemodynamic variables simultaneously.

II. SYSTEM OVERVIEW

The control objective is to regulate two hemodynamic variables, mean arterial pressure (MAP) and cardiac output (CO) by the automated infusion of inotropic and vasoactive drugs. Of particular interest is the case of patients with congestive heart failure. Due to the complex, nonlinear behavior of physiological systems, a good controller is difficult to design. In addition, drug responses vary from one patient to another; for the same

† Howard P. Isermann Department of Chemical Engineering, Rensselaer Polytechnic Institute
‡ School of Engineering, Rensselaer Polytechnic Institute
* Corresponding Author, Department of Chemical Engineering, Rensselaer Polytechnic Institute, Troy NY 12180-3590, bequette@rpi.edu

patient, the responses also vary with time. All recent approaches to hemodynamic control have unresolved issues of robustness and safety [2].

Safety constraints translate in part to limits on the infusion rates. Saturation limits must be handled, as drugs can be infused but not removed from the bloodstream. Upper limits are also necessary to prevent overdosing and toxic effects. For example, Sodium Nitroprusside (SNP) breaks down into cyanide, thus a maximum rate of 10 $\mu g/(kg \cdot min)$ is used to prevent accumulation of this toxin. In the case of dopamine (DPM) the inotropic range is from 2 $\mu g/(kg \cdot min)$ to 7 $\mu g/(kg \cdot min)$. Inter– and intra–patient variability is compounded by the tightly coupled dynamics of the variables to be controlled.

Mean arterial pressure is affected by two factors: cardiac output and systemic vascular resistance. An increase in either one, with the other held constant, results in higher blood pressure. For a normal heart, cardiac output follows the Frank–Starling relationship. There are four factors that affect cardiac output: preload (venous return), afterload (mean arterial pressure), myocardial contractility (strength of the heart) and heart rate. The preload is the quantity of blood filling the heart during diastole; the greater the venous return is, the greater the stretching of the myocardium, thus leading to a stronger contraction and higher cardiac output. The afterload is the force that the heart must overcome to pump the blood out. Therefore, the higher the mean arterial pressure the greater the resistance to pumping is, resulting in a lower cardiac output. Heart contractility can be considered as the number of cross–bridges (individual muscle fibers) present in the cardiac muscle during a contraction. An increase in the number of cross–bridges leads to a more powerful contraction and a higher cardiac output.

In the case of congestive heart failure, the heart is unable to pump all the blood in the venous return. An increase in preload only decreases the efficiency of the heart and leads to greater congestion. The venous return is dependent on venous compliance, which is the ease with which the veins can stretch like a balloon. The greater the venous compliance is, the more blood that is retained in the veins and the lower the venous return to the heart. Additional dynamics are present in the form of the baroreceptor reflex, which is the body's own short term controller of mean arterial pressure. It is a part of the autonomous nervous system and it can increase or decrease myocardial contractility, venous compliance and the systemic vascular resistance to regulate blood pressure.

In this study three drugs are used. Sodium Nitroprusside (SNP) is used to lower blood pressure, as it increases the venous compliance and reduces systemic vascular resistance through arterial vasodilation. For congestive heart failure, it is an ideal drug since it decreases both the preload and afterload, so the heart has less blood to pump at a lower resistance. Phenylephrine (PNP) is used to increase blood pressure through arterial vasoconstriction (which increases the systemic vascular resistance). Dopamine (DPM) is used in its inotropic range to enhance cardiac performance by improving heart contractility.

Two different anesthetic agents are used, Isoflurane or Halothane with 50% Nitrous Oxide. Testing with different anesthetics is important, as the dynamics of the system will change depending on the agent used. Both anesthetics decrease the systemic vascular resistance, thus dropping blood pressure. Halothane also decreases the heart contractility. The time constants and time delays of the pharmacodynamics are significantly longer with halothane as well. The baroreceptor reflex is also affected by each anesthetic in different ways. Under isoflurane, the baroreflex remains strong, thus the system is capable of counteracting the controller if the desired MAP is too high or too low compared with normal body function. With halothane, the baroreceptor reflex is affected to the point that it no longer has a significant effect on the blood pressure. Two good sources that cover cardiovascular physiology are the works of Guyton and Hall [9] and Berne and Levy [10].

III. EXPERIMENTAL SETUP

The controllers were evaluated on four mongrel dogs under Institutional Animal Care and Use Committee (IACUC) approved protocol. The experiments were performed on eleven experimental days with 3 to 5 runs each day. A schematic diagram of the experimental setup is shown in figure 1. Following initial anesthetizing with sodium pentobarbithal, the animal was intubated and mechanically ventilated (Siemens–Elena 900C Servo–Ventilator) with Isoflurane or Halothane with 50% Nitrous Oxide anesthetic. An arterial line was placed in the femoral artery to provide continuous arterial pressure tracings on a Mennen Horizon monitor. A Swan–Ganz catheter (Baxter Edwards Swan–Ganz Intellicath CCO/VIP Thermodilution), connected to a Baxter Vigilance monitor, was introduced in the pulmonary arterial tree to provide continuous cardiac output measurement. Control calculations were performed on a Dell Pentium II PC running a custom built Windows–based GUI (figure 2). The pressure and flow measurements were received from the monitors through RS–232 ports. The control loop was closed with rotary infusion pumps (Critikon Simplicity 2100A) modified to accept digital inputs via a digital output card. The sampling time for the controller was set at 30 seconds. The controller computation took about 4 seconds at every sample.

The canines were pharmacologically altered to exhibit hyper– or hypotension and depressed cardiac output. Both SNP and PNP were used to affect the baseline blood pressure, and high levels of halothane were used to significantly depress heart contractility and thus mimic congestive heart failure. The controllers were initially evaluated and tuned in closed–loop simulations using an elaborate nonlinear canine circulatory model [11] as the "patient" before moving to the experimental phase.

IV. MODEL PREDICTIVE CONTROL

Model based predictive control (MPC), which can be implemented quite naturally on constrained multivariable systems, has also been considered for drug delivery (Gopinath et al. [12], Rao et al. [13], [14]). As seen earlier, an important issue in the design of drug infusion systems is the need to impose bounds on infusion rates. Alternatively, the physician may want to specify an operating range of the mean arterial pressure instead of a specific setpoint. While most control strategies handle such constraints in an ad hoc manner, the primary advantage to MPC is its ability to handle constraints explicitly. Its optimization–based framework allows computation of the optimal infusion rates subject to input and output constraints.

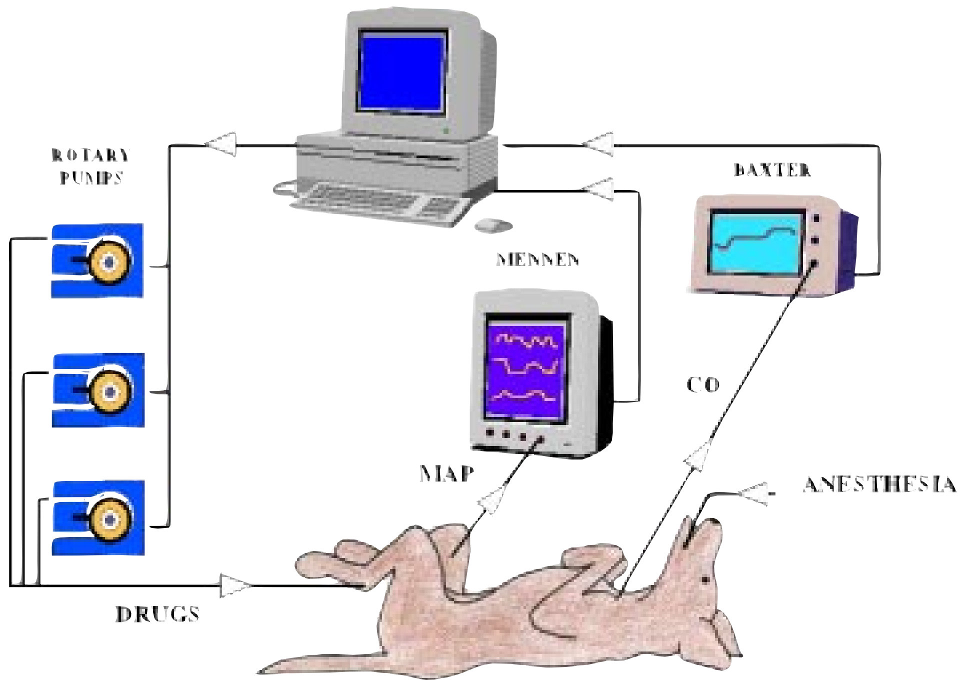

Fig. 1. Schematic diagram of the experimental setup

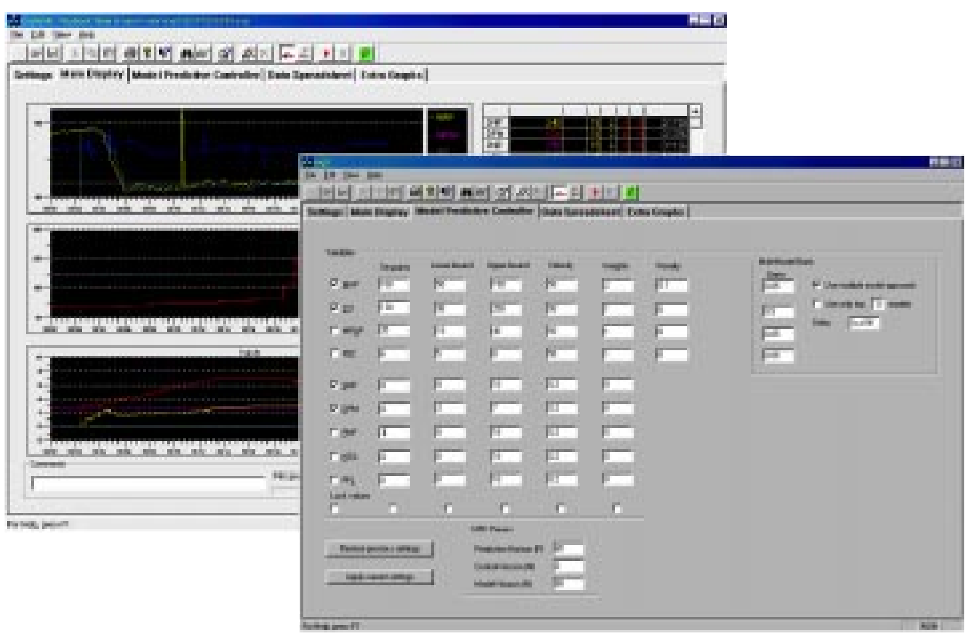

Fig. 2. Computer interface for the controller

MPC uses a model to predict the system response to future control moves and optimizes manipulated variables to minimize the predicted error subject to operating constraints. The basic idea, shown in figure 3, is to select a sequence of M future control moves to minimize an objective function (usually the sum of square of predicted errors) over a prediction horizon of P sample intervals. Using a model, the system response to changes in the manipulated variable is predicted. The M moves of the manipulated variables are selected such that the predicted response has minimal setpoint tracking error. As new measurement information will be available in the next sampling instance, only the first computed change in the manipulated variables is implemented and the optimization is repeated at each sampling interval based on updated measurements of the output variables.

The manipulated variables (drug infusion rates) u are computed to minimize a quadratic objective function

$$\min_{u_k \ldots u_{k+M-1}} J = \sum_{i=k+1}^{k+P} e_i^T \, Q \, e_i \; + \; \sum_{i=k}^{k+M-1} \Delta u_i^T R \Delta u_i \; +$$

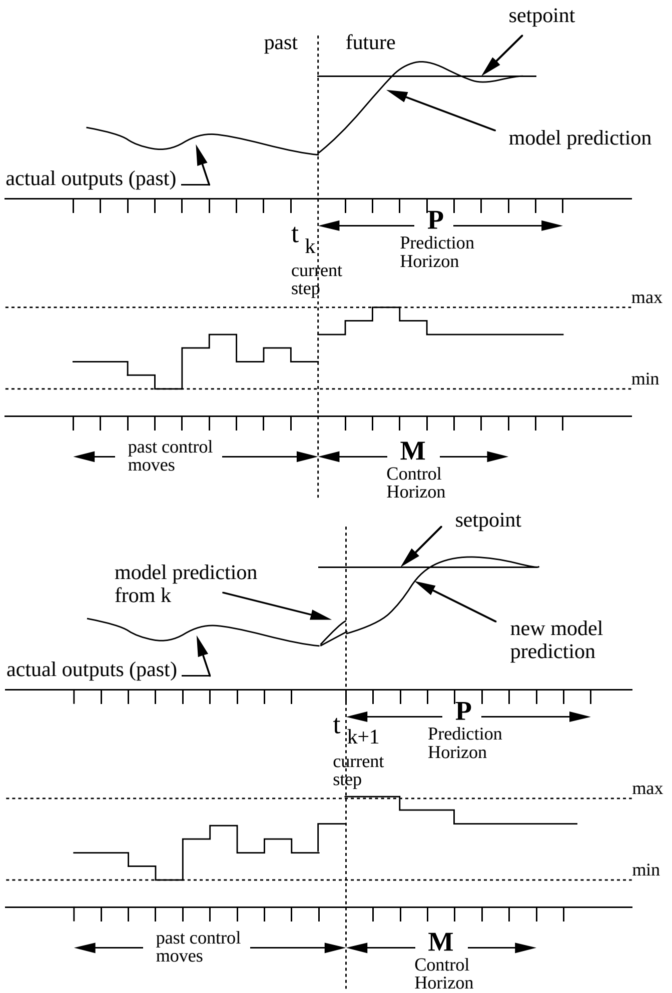

Fig. 3. Model Predictive Control: (a)At the current sampling instance k, a model is used to predict the output behavior of the system P sample intervals into the future based on the past states and M future control moves. The future control moves are optimally estimated to minimize predicted error from setpoint. Feedback is achieved by implementing only the first of the M moves. (b)Based on the actual measurements of the output at the $k+1^{th}$ instance, the model predictions are corrected as an additive disturbance to account for model mismatch and unmeasured disturbances. The optimization procedure is repeated in a receding horizon framework to compute a new set of moves.

$$\sum_{i=k+1}^{k+P} (y_{max} - y_i^c)^T \; S_1 \; (y_{max} - y_i^c) +$$
$$\sum_{i=k+1}^{k+P} (y_i^c - y_{min})^T \; S_2 \; (y_i^c - y_{min}) \qquad (1)$$

subject to absolute and rate constraints on the manipulated variables

$$u_{min} < u_i < u_{max}$$
$$u_{i-1} - \Delta u_{max} \leq u_i \leq u_{i-1} + \Delta u_{max}$$

where, at each sampling instance i, e_i is a vector of model predicted errors $(e_i = r_i - y_i^c)$, y_i^c is a vector of model predicted outputs (MAP, and CO) over a prediction horizon of P, r_i is the desired setpoint, u_i is the vector of manipulated variables (SNP and DPM) over a control horizon M, and Q and R are output and input weighting matrices. The optimization is a quadratic programming (QP) problem and absolute and rate constraints on the manipulated variable are included as linear inequalities. Imposing constraints on the output variables is likely to make the QP problem too restrictive, that is, when computing future moves there may exist no value for which the drug infusion *and* the predicted responses are within the permitted range. The solution is to add a penalty term to the control objective (soft constraints) that penalizes a weighted norm of the output variable violation (last two terms in the objective function). S_1 and S_2 are weighting matrices for the upper and lower soft constraints on the output variables.

The prediction horizon P is chosen on the basis of the open-loop settling time. The control horizon M is used to tighten or detune the controller. In general, larger values of M for an input will result in more aggressive action. This yields faster response, but the closed–loop system is less robust to model uncertainty. The output weighting matrix Q is a diagonal matrix used to assign weights to the components of the error function, corresponding to each output in the optimization step. A larger weight for an output will result in tighter control. When soft constraints are used, S_1 and S_2 penalize the error function when the output variables are outside the desired min-max range and Q penalizes deviations from the middle of the min-max range. The values of S_1 and S_2 are chosen to be larger than Q to emphasize tighter control when outside the min-max band. The input penalty matrix R is also a detuning parameter and is used to penalize control action in the objective function. This parameter is especially useful when a large M is used.

A simple form of the prediction model is based on the finite step response (FSR) or the finite impulse response (FIR) convolution model. This is a non-parametric representation of the process and is simply the open-loop response to a unit step or a unit impulse input. The output prediction is computed by convolving the model impulse response with the history of manipulated variable $(u_{k-1}, u_{k-2}, \cdots)$ from the current sampling instance k and given by

$$\hat{y}_k = \sum_{i=1}^{N} H_i \; u_{k-i}$$

where H_i is the ith impulse response coefficient matrix. N is the number of terms in the model, usually chosen to correspond to the settling time of the model. This ensures that we use information about any control move that might have been made in the past until the system settles to the steady state arising from that control move. The corrected predicted output at the j^{th} future point is given by

$$y_{k+j}^c = \sum_{i=1}^{j} H_i \; \Delta u_{k+j-i} \; + \; \sum_{i=j+1}^{N} H_i \; \Delta u_{k+j-i} \; + \; d_k$$

The corrected prediction of output involves three terms on the right hand side. The first term includes the present and all future moves of the manipulated variables which are to be determined so as to solve equation 1. The second term includes the past values of the manipulated variables and is completely known at time k. The third term is a correction (predicted disturbance) which is calculated as the difference between the current measurements and output of the predicted model (i.e. $d_k = y_k - \hat{y}_k$) at the k^{th} sampling instant. This is the 'additive disturbance' which accounts for model mismatch and unmodeled disturbances that enter the system and is assumed to be constant over the prediction horizon due to lack of an explicit means of predicting the mismatch or disturbance.

While adaptive control strategies rely on using a nominal model with on–line adaptation, model predictive approaches depend on the accuracy and availability of a model capable of predicting the patient responses. In all the studies cited above, it is important to note that the motivation to develop advanced strategies is due to the multi–variable, nonlinear behavior of physiological systems. The inherent difficulty arises in the choice of a model, its structure and the associated parameter identification for design and validation of the control system on real subjects. In the next section we present a novel approach combining the MPC and MMAC strategies for regulation of hemodynamic variables. A multiple model strategy is used to provide a prediction model for an MPC framework. This approach has the combined advantage of allowing model adaptation to handle inter– and intra–patient variability and the ability to handle explicit input and output constraint specifications.

A. Multiple Model Predictive Control

Conventional MMAC [15] shown in figure 4, uses a bank of models to capture the possible input–output behavior of patient responses to drug dosages. The control parallel is to use a bank of controllers, to achieve a desired closed–loop performance from a wide variety of possible patients; controller k is designed and tuned based on model k from the model bank. Using a Bayesian approach, the probability of each model representing the patient response is computed and the resultant control action is the probability–weighted average of control moves of each controller. The model probabilities get altered as the drug sensitivities change in inter/intra patient variations. The primary advantage to this approach is that no a priori model identification is necessary during initial stages of drug administration. Although the approach is sub–optimal it allows flexibility to handle a system with large variability such as drug infusion where designing a single nonlinear model is not practical or possible. The emphasis is more on robustness than on performance. However explicit handling of constraints using MPC for each model is computationally unwieldy for a large model bank.

To preserve a multiple model approach that explicitly handles constraints, we use a single constrained MPC controller with a weighted model bank for response predictions, as shown in figure 5. A probability–weighted average of output predictions from a bank of models is used in a MPC framework to calculate drug infusion rates for regulation of mean arterial pressure and cardiac output. This approach has the combined advantage of model adaptation according to patient variations and the ability to handle explicit input and output constraint specifications. The advantage to multiple–model predictive control (MMPC) is that a large number of models can be used and constraints explicitly handled. The issues for MMPC are similar to standard MMAC — determining the choice and number of models to encompass the plant behavior. This may require detuning of MMPC so that it is sufficiently robust in the possible range of prediction models stemming from the model bank.

We used the recursive Bayes theorem for computing the weights assigned for each model in the bank. The theorem calculates the conditional probability of the i^{th} model in the bank being the true model of the plant given this population of models. The probabilities are assumed to be stochastic and Gaussian in nature and thus take a form of the exponential of the negative square of the residuals. The recursive Bayes theorem for the k^{th} step and i^{th} model is

$$p_{i,k} = \frac{exp(-\frac{1}{2}\epsilon_{i,k}^T K \epsilon_{i,k}) p_{i,k-1}}{\sum_{j=1}^{N} exp(-\frac{1}{2}\epsilon_{j,k}^T K \epsilon_{j,k}) p_{j,k-1}} \tag{2}$$

where

$$\epsilon_{i,k} = y_k - y_{i,k}^c \tag{3}$$

is the model residual at the current step. K denotes convergence matrix that can be used to tune the rate of convergence of the probabilities. The recursion is initialized by assigning equal probability to all the models in the bank. The probabilities are computed at every sample "improving" upon the probability computed from the previous sample. The algorithm is computationally inexpensive and ensures that probabilities are bounded between 0 and 1. The models are assigned weights such that

$$W_{i,k} = \begin{cases} \frac{p_{i,k}}{\sum_{j=1}^{N} p_{j,k}} \text{for } p_{i,k} > \delta \\ W_{i,k} = 0 \text{ for } p_{i,k} = \delta \end{cases} \tag{4}$$

Models attaining a probability of zero cannot reenter the recursion and hence an artificial cutoff, δ, is used to keep the models alive. For models with $p < \delta$, the probability is reset to $p = \delta$, and these models are then excluded from being weighted. A weight-averaged prediction model is constructed for MPC to compute optimal drug infusion rates. The blended prediction model is effectively a model of higher order than any of the individual models in the bank. The blended model thus evolves to adapt the patient's drug response and disturbances that occur in the circulatory system.

It is important to realize that the nonlinear system is capable of MIMO dynamics that may be amplified or suppressed compared to the dynamics exhibited when using the drugs individually. The key assumption of our model blending design is: given a bounded set of SISO models, the MIMO interactions will be bounded as well. In other words, as long as the bank of SISO models bounds the entire possible range of input–output behavior, there will exist a linear combination of the MIMO models that will be able to adequately describe the process behavior. The dynamic construction of the prediction model can be considered as a form of online closed loop MIMO identification and, as described in the next section, it is easy to perform relevant updates to the model bank when it is observed that the MIMO bank is not properly bounding the canine's responses.

The recursive computation of probability was observed to attain steady state values of either zero or one (i.e.) converging to a single model in the bank. However, due to the uncertainty involved in such systems, it is unlikely that any single model in the model back would be exactly equivalent to the system under

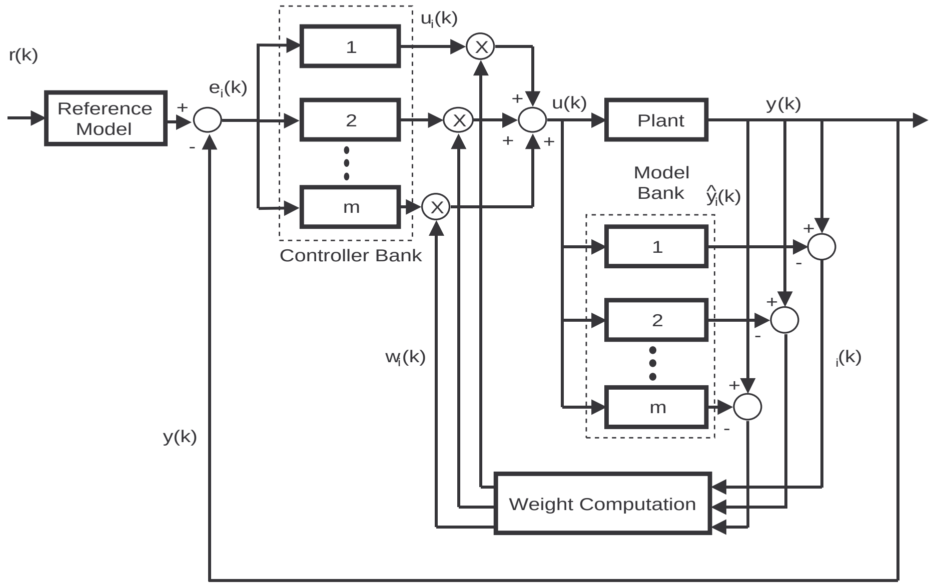

Fig. 4. Schematic of the MMAC strategy.

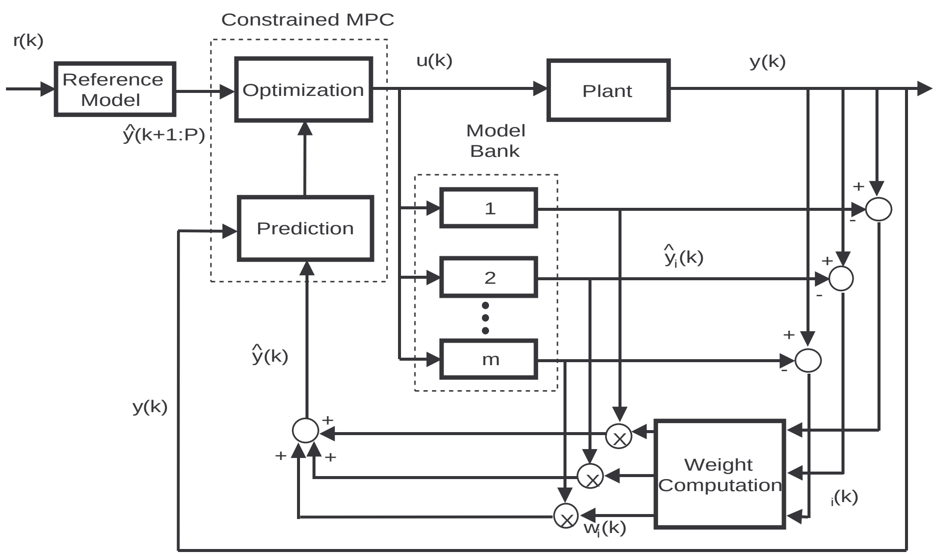

Fig. 5. Schematic of the MMPC strategy.

control and hence blending is preferred. We were able to maintain a reasonable degree of blending by tuning the convergence matrix K (equation 2). In fact, during the experiments, controller performance was consistently good when a single model did not dominate in the weighting.

A.1 Design of the model bank

For this multivariable problem, it was difficult to directly design appropriate multiple-input, multiple-output (MIMO) models. We constructed MIMO models from superposition of single input single output (SISO) models. The SISO models constituted step response coefficients generated using first order plus dead time (FOPDT) transfer functions. Nominal values of model parameters such as the gains, time constants and time delays were obtained from literature and/or evaluated from step response experiments for each drug. Additional models spanned a range around these nominal values. The model gains were spaced by a factor of 1.5. Various combinations of the SISO models were used to form the MIMO model bank. Our experience has been that the observed gain, is not necessarily related to the time constant/time delay of the drug response. Therefore combinations of high time constants with long time delays and short time constants with short time delays were paired with each gain value.

Based on nominal values from literature and prior knowledge of sluggish drug pharmacodynamics under halothane, additional models with longer time constants and time delays (by about 150%) were added to the model bank used in isoflurane experiments. This still proved to be insufficient resulting in deterioration of control performance amongst the first few runs us-

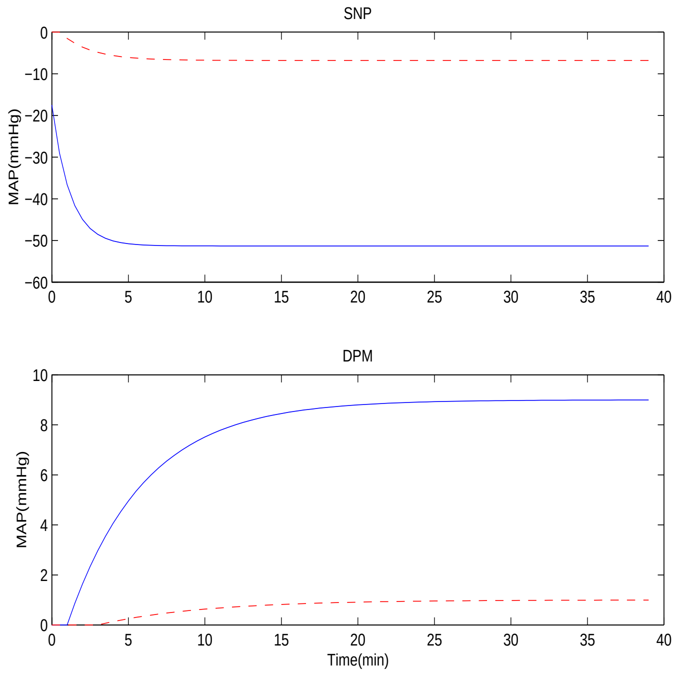

Fig. 6. Step response plots for unit inputs of SNP and DPM vs. MAP to indicate the bounds on the SISO model banks.

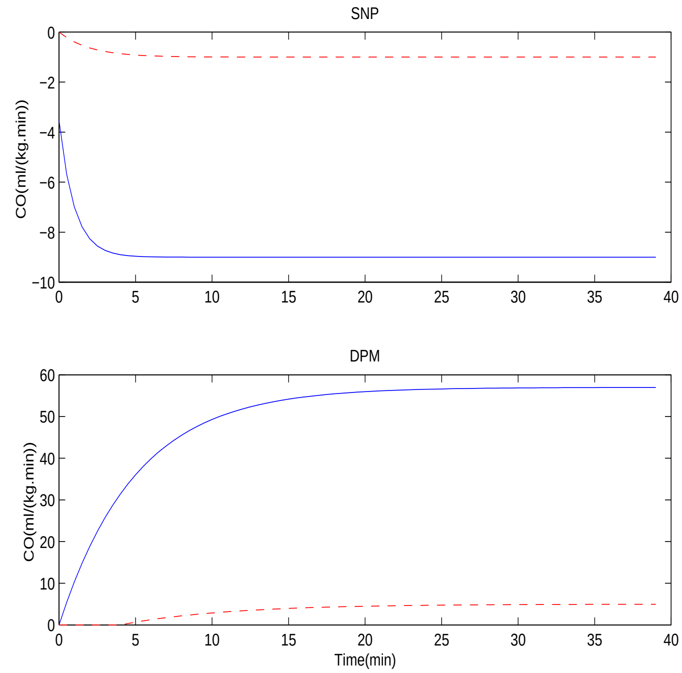

Fig. 7. Step response plots for unit inputs of SNP and DPM vs. CO to indicate the bounds on the SISO model banks.

ing halothane. Post experimental analysis of the weights assigned to the models in the bank indicated that the weighted model mostly constituted an average of models with longer (but still insufficient) time delays and time constants. Models with time constants and time delays (as high as twice those used for isoflurane experiments) were found to be necessary to encompass dynamics exhibited under halothane anesthesia. After this stage, the model bank has required no further changes. The upper and lower bounds of dynamics that the final model bank spanned is shown as step response plots in figures 6 and 7. The FOPDT bounds are made up of the shortest time constant/time delay with largest gain and the longest time constant/delay with smallest gain.

B. MMPC Results

The responses to the drugs administered varied greatly depending on the anesthetic used as well as the canine that received the drugs. Isoflurane lowers systemic vascular resistance yet the contractility of the heart and the baroreceptor reflex remain strong. With isoflurane, the MAP setpoint chosen will determine the baroreceptor response either towards the setpoint or against it strongly perturbing the drug effects on the canine. For halothane with 50% nitrous oxide, systemic vascular resistance, baroreceptor reflex and the heart's contractility are all compromised. As halothane depresses CO values, the pharmacodynamics are very slow, having time delays and time constants in some experiments more than double the values on the same canine under isoflurane. Accordingly, a prediction horizon of $P = 35$ and control horizon of $M = 3$ was used with isoflurane. When using halothane, P was increased to 65.

Case 1: Figure 8 presents experimental results of control of MAP using SNP on a 19 kg female canine anesthetized with isoflurane. Hypertension was induced by infusing PNP at a constant rate of 5 $\mu g/(kg \cdot min)$. Once the MAP stabilized at 120 $mmHg$, the controller was engaged (at $t = 0\ min$) with a setpoint to lower the MAP to 100 $mmHg$. Another setpoint change was made at $t = 15\ min$ to lower the MAP to 80 $mmHg$. A "disturbance", in the form of a change in the hypertension, was created by lowering PNP infusion to 3 $\mu g/(kg \cdot min)$ and back to 5 $\mu g/(kg \cdot min)$ (between $t = 20\ min$ and 30 min). At $t = 38\ min$, the setpoint was further lowered to 75 $mmHg$ and the controller was disengaged at around 50 min. Figure 8 illustrates fairly tight control of MAP within $\pm 5\ mmHg$ with settling times of less than 7 minutes after each setpoint change. The controller also rejected well the disturbance of changing PNP infusion. Note that the lowering of MAP to 80 $mmHg$ and beyond caused the baroreflex to increase the heart rate (HR) and the cardiac output (CO), in an attempt to increase the blood pressure against the influence of the vasoactive drugs. The magnitude of such an unmeasured disturbance can be gauged from the jump in the MAP returning to baseline values after the external drug influences were removed after $t = 50\ min$.

Case 2: Multivariable control of MAP and CO is a much more difficult problem as the two hemodynamic variables are highly correlated and interdependent. Figure 9 illustrates the results of control MAP and CO using SNP and DPM. A condition of low CO (which mimics congestive heart failure) was created by using high levels of halothane on a 17.8 kg male canine. When the controller was engaged at $t = 0\ min$, a setpoint of lowering MAP to 60 $mmHg$ and increasing CO to 2.3 l/min was sought. The controller achieved both setpoints in about 12 minutes. While MAP was controlled within $\pm 10\ mmHg$, CO had an overshoot at $t = 22\ min$, ten minutes after achieving the setpoint. The overshoot in CO was directly related to the increase in HR (excluding the noise spike at $t = 13\ min$) due to an un-

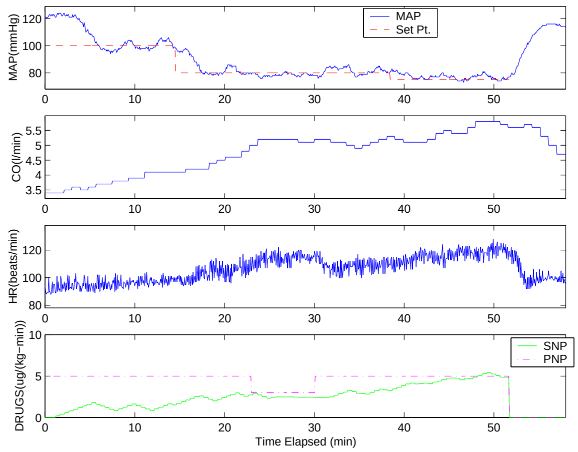

Fig. 8. Case 1: MAP control of hypertensive canine under isoflurane using sodium nitroprusside (SNP). Phenylephrine (PNP) used to induce hypertension.

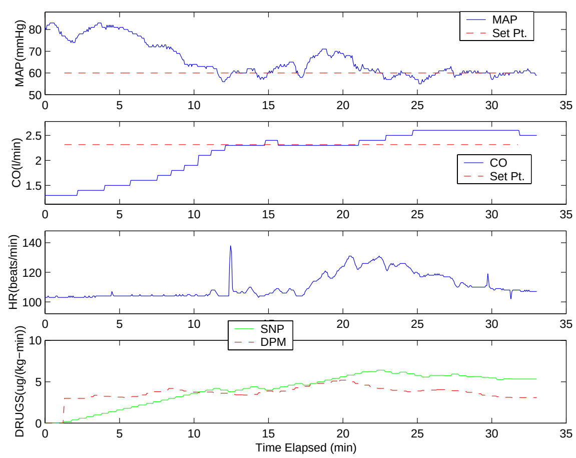

Fig. 9. Case 2: MAP and CO control of depressed CO canine under halothane. Dopamine (DPM) and sodium nitroprusside (SNP) are controlled inputs.

known disturbance. As the HR recedes to its baseline value, CO starts to return to setpoint. Figure 10 presents a comparison of the actual MAP and CO values with the estimates obtained from the weighted model bank. The weighted or blended prediction model was observed to track the hemodynamic variables closely. The success of MMPC strategy relies on the weighted model to provide good estimates and it is imperative that the individual models in the model bank encompass the possible ranges of gains, time constants and time delays.

Case 3: This case presents the results of a multivariable control run to regulate MAP and CO lasting over four hours (figure 11). Instead of a specific setpoint for CO, a desirable operating range was specified. The MPC strategy handles such specifications as output constraints explicitly. The use of output constraints in the controller can easily lead to infeasible solutions in the optimization problems or unstable closed–loop behavior. The output constraints were hence specified as soft constraints. High levels of halothane were used to create a low CO condition

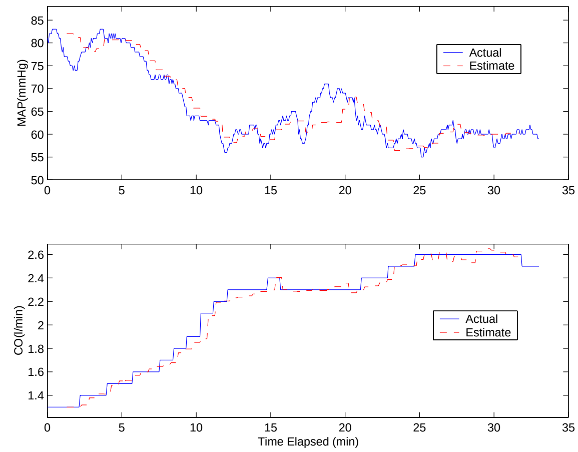

Fig. 10. Case 2: Model bank estimation of MAP and CO versus actual outputs.

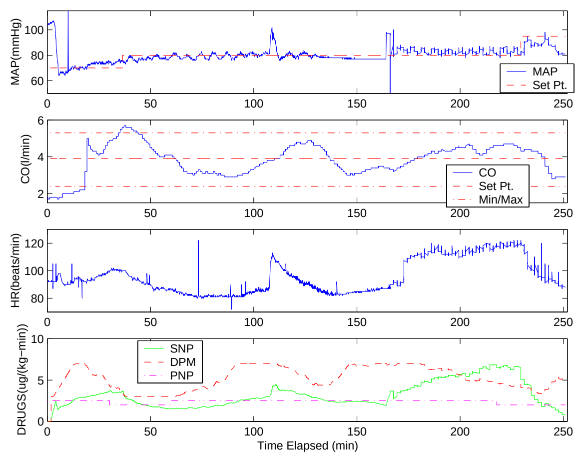

Fig. 11. Case 3: MAP and CO control of depressed CO canine under halothane. Dopamine (DPM) and sodium nitroprusside (SNP) are controlled inputs. PNP was used to induce hypertension.

and PNP was infused at a constant rate of 3 $\mu g/(kg \cdot min)$ to cause hypertension. The control objective was to raise and maintain CO between $2.6 - 5.6\ l/min$ and lower MAP to 70 $mmHg$. The controller infused suitable amounts of SNP and DPM within specified constraints in order to achieve the desired states of the hemodynamic variables with settling times under 15 minutes. The MAP setpoint was changed to $80\ mmHg$ at $t = 40\ min$. An unknown disturbance in the form a sudden increase in HR raises the MAP at around $t = 100\ min$. The controller rejected this disturbance by increasing the SNP infusion and decreasing DPM infusion. A similar but sustained disturbance results in an increase of HR and hence MAP and CO at $t = 170\ min$ and the resulting control action maintained the MAP and CO ranges at desired values. Figure 12 presents a comparison of the MAP and CO values with the corresponding estimates from the weighted model bank.

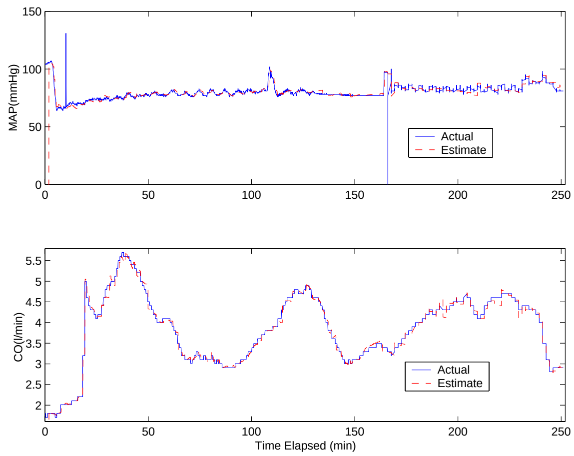

Fig. 12. Case 3: Model bank estimation of MAP and CO versus actual outputs.

C. Discussion

The MMPC control design has demonstrated its efficacy in canine experiments for both hypertensive and depressed cardiac output cases for two different maintenance anesthetics. Inter– and intra–patient variability are the most significant problems to any controller for drug infusion. This variability is compounded by the different anesthetics used in practice and their greatly varying effects on baseline heart contractility, systemic vascular resistance and the strength of the baroreceptor reflex. The success of a drug infusion control system depends on its ability to accommodate such variations. We have presented a multiple–model predictive control approach that does not require a priori specification of model structure and parameters and explicitly handles constraint specifications. The Bayesian weighting scheme used is computationally inexpensive and effective in "building" a higher order prediction model from basic building blocks of first order models. Unlike typical adaptive schemes, the order of the weighted model is not fixed and gets altered dynamically to closely mimic plant behavior. The weighted model bank can account for the variability in the drug infusion system, providing a flexible and bounded prediction model for the online constrained optimization problem. However, it is critical to have the bank of models bound the possible plant behavior. Issues such as the determination of the number of models spanning a given range of system gains and dynamics need to be addressed.

V. DIRECT MODEL REFERENCE ADAPTIVE CONTROL

A. Introduction

Here we present the results of the experimental testing of a modified Direct Model Reference Adaptive Control (DMRAC) algorithm which improves the robustness over previous adaptive control strategies. The simple adaptive control approach to DMRAC of MIMO plants was first proposed by Sobel et al. [16] in 1979. This control structure uses a linear combination of feedforward model states and command inputs and feedback of the error between plant and model outputs. This class of algorithms requires neither full state access nor adaptive observers. Other important properties of this class of algorithms include (1) their applicability to non-minimum phase systems and (2) the fact that the plant (physical system) order may be much higher than the order of the reference model.

DMRAC requires that the plant to be controlled satisfy a strictly positive real (SPR) condition (see section VII). That is, for a plant to be controlled, there exists a feedback gain such that the resulting closed-loop system is strictly positive real. Note that the plant satisfying the above condition is called almost strictly positive real (ASPR) (see [17] for a detailed analysis). One way to satisfy this positive real constraint is to design a parallel feedforward compensator.

Since most real systems do not satisfy the ASPR condition, the DMRAC algorithm is extended in [18] and [19] to the class of non-ASPR plants, for which there exists a known dynamic output stabilizing feedback, with transfer matrix $C(s) = H(s)^{-1}$, such that the plant in parallel with $H(s)$ is ASPR. Kaufman and Neat [20] suggested a modification that incorporates part of the supplementary feedforward into the reference model in a manner such that asymptotic tracking of the augmented plant and model outputs implies asymptotic tracking of the original plant and model outputs.

Ozcelik et al. ([21], [22], [23], [24]), have developed and applied systematic design procedures that utilize optimization techniques for both the SISO and MIMO systems with parametric and/or frequency domain uncertainties. However, feedforward compensator design methods were not addressed for plants with uncertain time delay elements, as in the adaptive control of drug infusion [7]. The design of the feedforward compensator used is presented in [25] by Ozcelik et al. They extend the design methodology so that the SPR condition is satisfied in the presence of plant uncertainty which is modeled as variations in both the plant time delay elements and transfer function coefficients.

We have further modified DMRAC by introducing changes to the adaptation law that allow the handling of rate limiting and saturation constraints, without affecting the properties of the algorithm. Formulation of the basic DMRAC algorithm is discussed in Section V-B, the modified algorithm to handle input constraints is presented in Section V-C, and the experimental results for a drug infusion control problem are given in Section V-D. Finally, results are discussed, and conclusions are drawn in Section V-E.

B. Formulation of the DMRAC Algorithm

The linear time invariant model reference adaptive control problem is considered for the linearized plant

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t) \\ y(t) &= Cx(t) \end{aligned} \tag{5}$$

where $x(t)$ is the $(n \times 1)$ state vector, $u(t)$ is the $(m \times 1)$ control vector, $y(t)$ is the $(q \times 1)$ plant output vector, and A, B are matrices with the appropriate dimensions. The range of the plant parameters is assumed to be known and bounded by

$$\begin{aligned} \underline{a}_{ij} \le a(i,j) \le \overline{a}_{ij}, i,j = 1,\cdots,n \\ \underline{b}_{ij} \le b(i,j) \le \overline{b}_{ij}, i,j = 1,\cdots,n \end{aligned} \tag{6}$$

The objective is to find, without explicit knowledge of A and B, the control $u(t)$ such that the plant output vector $y(t)$ follows the reference model

$$\begin{aligned}\dot{x}_m(t) &= A_m x_m(t) + B_m r(t) \\ y_m(t) &= C_m x_m(t)\end{aligned} \tag{7}$$

The model incorporates the desired behaviour of the plant, but its choice is not restricted. In particular, the order of the plant may be much larger than the order of the reference model.

The adaptive control law is given as [17]

$$u(t) = K_e(t)[y_m(t) - y(t)] + K_x(t)x_m(t) + K_r(t)r(t) \tag{8}$$

where $K_e(t)$, $K_x(t)$, and $K_r(t)$ are adaptive gains and concatenated into the matrix $K(t)$ as follows

$$K(t) = [K_e(t) \quad K_x(t) \quad K_r(t)] \tag{9}$$

Defining the vector $v(t)$ as

$$v(t) = \begin{bmatrix} y_m(t) - y(t) \\ x_m(t) \\ r(t) \end{bmatrix} \tag{10}$$

the control $u(t)$ is written in a compact form as

$$u(t) = K(t)v(t) \tag{11}$$

The adaptive gains are obtained as a combination of an integral gain and a proportional gain as shown below [17]

$$\begin{aligned}K(t) &= K_p(t) + K_I(t) \\ K_p(t) &= [y_m(t) - y(t)]v^T(t)\overline{T}, \quad \overline{T} \geq 0 \\ \dot{K}_I(t) &= [y_m(t) - y(t)]v^T(t)T, \quad T > 0\end{aligned} \tag{12}$$

The sufficiency conditions for asymptotic tracking are
1. There exists a solution to the CGT problem (see [17]).
2. The plant is ASPR; that is, there exists a positive definite constant gain matrix K_e, not needed for implementation, such that the closed loop transfer function

$$G(s) = [I + G_p(s)K_e]^{-1}G_p(s) \tag{13}$$

(where $G_p(s) = C(sI - A)^{-1}B$) is strictly positive real (SPR).

The solution to the CGT problem is usually not a problem; but in general, the ASPR condition is not satisfied by most real systems. Therefore, Bar Kana and Kaufman [18] have shown that a non-ASPR plant of the form $G_p(s) = C(sI - A)^{-1}B$ can be augmented with a feedforward compensator $H(s)$ such that the augmented plant transfer function

$$G_a(s) = G_p(s) + H(s) \tag{14}$$

is ASPR. However the resulting adaptive controller will in general result in a model following error that is bounded but not zero in steady state. To eliminate this problem, a modification that incorporates the supplementary feedforward into the reference model output as well as the plant output has been developed by Kaufman and Neat [20]. This configuration is shown in Figure 13.

C. Modified DMRAC Algorithm

One of the main drawbacks of the standard DMRAC algorithm is its inability to handle input constraints, which is critical in the case of drug infusion control for safety reasons. The controller has to be able to handle saturation limits, as no drug can be removed from the bloodstream; an upper bound is also set to avoid overdosing and toxic side effects. The drug infusion rate calculated by the controller depends on the dynamics of the reference model, the magnitude of the error between the desired and actual system responses, and the controller gains. If the manipulated inputs are changing too fast, or if the saturation limits are being exceeded, the simplest way to handle this is to scale back the controller gains. To do this, the adaptation law for the gains is modified as follows

$$\begin{aligned}K(t) &= K_p(t) + K_I(t) \\ K_p(t) &= [y_m(t) - y(t)]r^T(t)\rho(t)\overline{T} \qquad \overline{T} > 0 \\ \dot{K}_I(t) &= [y_m(t) - y(t)]r^T(t)\rho(t)T - \sigma(t)K_I(t)T \geq 0\end{aligned} \tag{15}$$

where $\rho(t)$ is a diagonal matrix with elements $0 < \rho_{ij}(t) \leq 1$ and $\sigma(t)$ is a diagonal matrix with elements $\sigma_{ij}(t) \geq 0$.

From (11) we can see that if the gains increase too fast, this will result in a fast increase of the control commands, possibly violating the command constraints. In this case, the function $\rho(t)$ is used to scale them back and to slow the adaptation as well. If the change in the command is within the prescribed limits, then it remains at 1, but if the change is larger, then $\rho(t)$ decreases in proportion to the magnitude of the violation.

The $\sigma(t)$ term is added to provide for anti–windup of the gains when the saturation constraints are hit. A basic form for this function is

$$\dot{\sigma}(t) = K_{\sigma_1} |u(t) - u_{sat}(u(t))| - K_{\sigma_2}\sigma(t) \tag{16}$$

with

$$u_{sat}(u(t)) = \begin{cases} u_{lb} \text{ for } u(t) < u_{lb} \\ u(t) \text{for } u_{lb} \leq u(t) \leq u_{ub} \\ u_{ub} \text{ for } u(t) > u_{ub} \end{cases}$$

where $K_{\sigma_1} > 0$ and $K_{\sigma_2} > 0$ are tuning parameters, and u_{lb} and u_{ub} are the lower and upper saturation limits. Thus, when the infusion is within the saturation limits $\sigma(t) \to 0$ and increases when the they are violated. It is straightforward to show that this new adaptation law does not change the stability results of the DMRAC algorithm, simply using the same Lyapunov function as used in [17].

D. Results

Previous to the dog experiments, simulations were done using Yu's [7] model of the hemodynamic system. The derivation of the feedforward compensator and the simulation results can be found in [25]. Even though all of the above theory for DMRAC is based on continuous time systems, digital implementation is straightforward. In our case the sample time of the system is 30 sec. The reference model is integrated numerically between the sample times. The same goes for the feedforward dynamics. The integral part of the gain update is also computed at this frequency, using simple Euler type integration. Measurements are kept constant between sampling times.

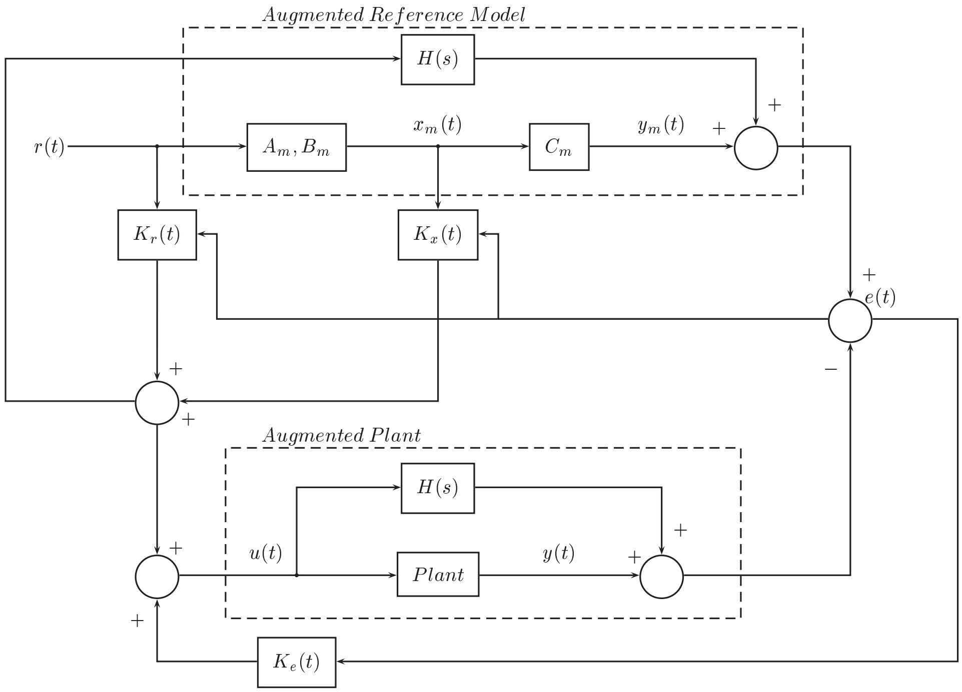

Fig. 13. DMRAC with plant and reference model feedforward

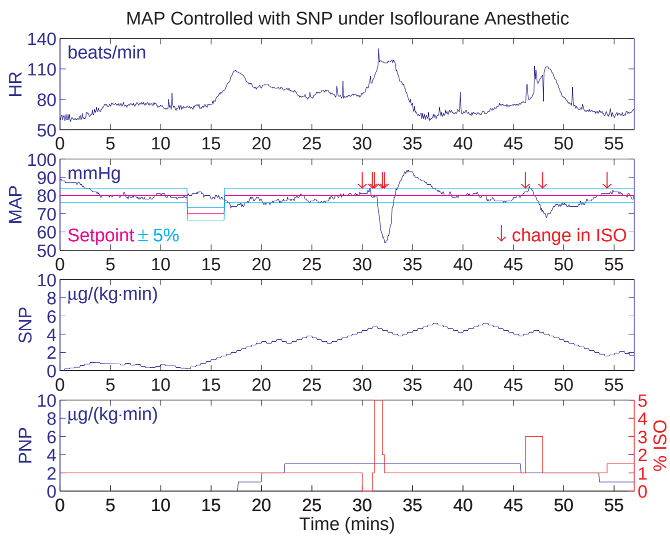

Fig. 14. DMRAC Case 1: MAP control of hypertensive canine using SNP. PNP used to induce hypertension. Significant disturbances due to changes in anesthetic concentration.

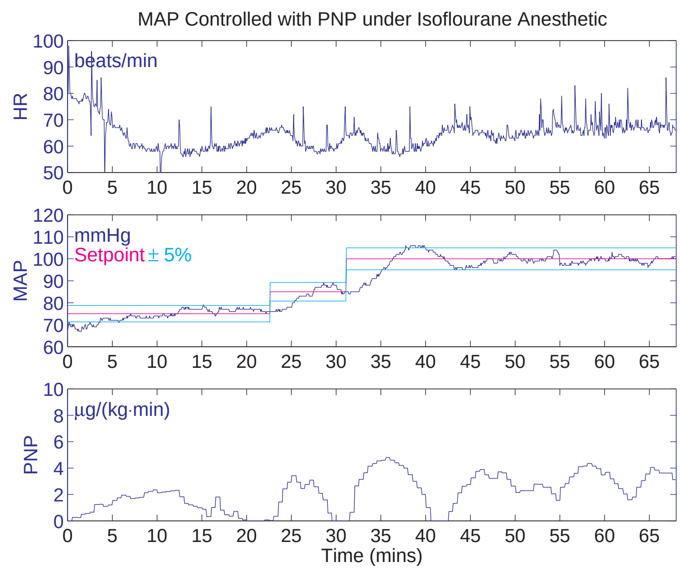

Fig. 15. DMRAC Case 2: MAP control of hypotensive canine using PNP.

The controller was tested on mongrel dogs, using different drug combinations and different anesthetics. For the results presented here, both runs are on the same dog (a 19 kg female) on the same day, using isoflurane as the anesthetic agent. For these runs the $\rho(t)$ parameter is active but $\sigma(t)$ is not, allowing $\rho(t)$ to handle both rate and saturation limits.

Case 1: In figure 14 mean arterial pressure is being regulated by sodium nitroprusside. The cyan lines represent $\pm 5\%$ of the setpoint (magenta line). It can be seen that the controller does an satisfactory job, even when other disturbances are introduced.

Initially the setpoint was set to 80 $mmHg$, and although the controller is doing a good job, the infusion rate of SNP is very small. At $t = 12.7\ min$ the setpoint was dropped to 70 $mmHg$, and the controller responded by increasing the amount of SNP infused. But the baroreflex (which remains strong under isoflurane) also reacted to this, as can be seen by the sudden increase in heart rate. Thus at $t = 16.4\ min$ the setpoint was changed once again to 80 $mmHg$ rather than have the controller fight the baroreflex, which in our experience from other runs would have meant possible SNP infusion rates of close to the upper limit. Instead, the hypertensiveness of the canine was increased by infusing PNP. At $t = 17.7\ min$ the infusion rate for PNP was set to 1 $\mu g/(kg \cdot min)$, at $t = 20\ min$ it was increased again to 1 $\mu g/(kg \cdot min)$ and at $t = 22.3\ min$ it was set at 3 $\mu g/(kg \cdot min)$.

Changes in the concentration of isoflurane administered were another source of disturbances. These changes are indicated by the arrows in the MAP plot. At $t = 30\ min$ the container for the anesthetic had to be refilled, which requires turning off the anesthetic. At $t = 31\ min$, the concentration of isoflurane was set back to 1%. Only moments later the canine started to wake up (as attested by the sharp increase in heart rate). The concentration was then set to 5% from $t = 31.2\ min$ to $t = 32\ min$, at which time it was brought down to 2% and again at $t = 32.2\ min$ to the maintenance level of 1%. The sudden increase in anesthetic concentration to 5% caused the sharp decrease in MAP, at which point the controller started to back down the amount of SNP infused. Once the concentration was brought down to 1% this caused the MAP to swing to the opposite extreme, causing the overshoot around $t = 35\ min$; again, the controller responded by increasing the infusion of SNP. Note that as part of the safety constraints the rate of increase in the infusion rate is limited, and thus large changes in SNP infusion are not possible.

There is one more large transient disturbance, in this case caused by changes in the infusion of PNP as well as the anesthetic. At $t = 45.8\ min$ the infusion rate of PNP was set back to 2 $\mu g/(kg \cdot min)$, practically at the same time the canine started to shiver (an indication that the animal is once again waking up), and thus at $t = 46.2\ min$ the concentration of isoflurane was increased to 3%. At $t = 47.9\ min$ isoflurane was set once more at 1%. The controller once more does a good job in handling the disturbance.

Case 2: The second run shown in figure 15 is for MAP regulation using Phenylephrine, rather than Sodium Nitroprusside as the manipulated input. The canine starts being slightly hypotensive, with a MAP of 67 $mmHg$. The MAP setpoint is started at 75 $mmHg$. The controller's action of infusing PNP actually relaxes the baroreflex which was trying to increase MAP, as can be seen from the decrease in the heart rate. After about 20 minutes the infusion rate of PNP required to maintain the desired setpoint is practically zero, thus at $t = 22.6\ min$ the MAP setpoint is increased to 85 $mmHg$. It is increased once more at $t = 31.2\ min$ to 100 $mmHg$. In all cases it takes the controller about four minutes to bring the MAP to within $\pm 5\%$ of the setpoint. Overall the controller does a good job in regulating MAP. Note that the controller has to keep adjusting the infusion rate of PNP constantly to maintain the MAP at the setpoint, which clearly shows the disadvantage of maintaining hemodynamic variables by manual adjustment.

E. Discussion

The increased robustness of the DMRAC algorithm and the handling of input constraints are a significant improvement over previous adaptive designs. These two features make for a safer controller, with the added factor that it is easily extended to handle multiple–input multiple–output drug infusion systems. At the same time, other possible approaches to command limiting are also being explored, such as modifying the dynamics of the reference model (e.g. the time constant) to tighten or relax the system performance as a function of the command limits.

When controlling cardiac output, together with mean arterial pressure, it is extremely difficult to have good control when specifying a fixed setpoint. CO control is better handled by specifying a range of allowable outputs, or a minimum to maintain. Unfortunately, the DMRAC algorithm can not handle this directly, and thus further modifications are being explored for such cases. Another area of improvement that is under consideration is to incorporate the knowledge of the system time delays into the reference model. Even though on–line estimation of time delay is extremely difficult, it might still prove valuable to include what limited knowledge there might be instead of assuming no delay at all as is done with the basic algorithm.

VI. GENERAL CONCLUSIONS

We have presented two distinct algorithms for the control of hemodynamic variables. The obvious question is which approach is best? Based on current results, MMPC is the recommended approach. Both controllers have their strong and weak points. Multivariable systems with input and output constraints are more easily dealt with using MMPC. DMRAC has a much lower computational overhead than MMPC, but at the same time modifying the algorithm to incorporate variations is not as straightforward. Case in point, the work to allow DMRAC to handle output ranges instead of a fixed setpoint is still in progress, while with MMPC it is quite straightforward. At the same time DMRAC has shown tighter control for the SISO cases, but this might change as both algorithms can still be fine–tuned for the problem at hand. A more detailed comparison is needed in order to make a strong recommendation as to which strategy is better suited.

There are several areas where this work is still in progress. There is a need for further canine experiments, with the aim of taking the technology to clinical trials and thus a step closer to clinical practice. Related to this will be the need to show clearly that indeed such a controller can perform better than a skilled cardiac nurse. To mimic congestive heart failure, halothane was used to create a condition of depressed cardiac output. We are exploring other methods of inducing congestive heart failure. We are also integrating microdialysis techniques, using fast–liquid–chromatography, to obtain real time blood drug concentrations. This information will be used to improve the nonlinear simulation model, and to improve control by directly relating drug concentrations to MAP and CO.

ACKNOWLEDGMENTS

This work was supported by grants from the Whitaker Foundation and National Science Foundation (BES 9522639). C. Palerm's work is supported in part by the Consejo Nacional de Ciencia y Tecnología (CONACYT) of México. Our thanks to Whitney Hartz and Millie Smith for veterinary support, and Melaku Ayele for modifying the Critikon pump for digital control. We are grateful to Professor Jonathan Newell for his time, support and advice.

VII. APPENDIX

The basic definition for strictly positive real (SPR) is:

Definition 1: A $p \times p$ proper rational function matrix $Z(s)$ is called positive real if

• All elements of $Z(s)$ are analytic for $Re[s] > 0$
• Any pure imaginary pole of any element of $Z(s)$ is a simple pole and the associated residue matrix of $Z(s)$ is positive semidefinite Hermitian, and
• for all real ω for which $j\omega$ is not a pole of any element of $Z(s)$, the matrix $Z(j\omega) + Z^T(-j\omega)$ is positive semidefinite.

The transfer function $Z(s)$ is called strictly positive real if $Z(s-\epsilon)$ is positive real for some $\epsilon > 0$.

References

[1] J B Slate, L C Sheppard, V C Rideout, and E H Blackstone, "A model for design of a blood pressure controller for hypertensive patients," in *IEEE Engineering in Medicine and Biology Conference*, Denver, October 1979, pp. 867–872.
[2] S Isaka and A V Sebald, "Control strategies for arterial blood pressure regulation," *IEEE Trans. Biomed. Eng.*, vol. 40, no. 4, pp. 353–363, 1993.
[3] J F Martin, N T Smith, M L Quinn, and A M Schneider, "Supervisory adaptive control of arterial pressure during cardiac surgery," *IEEE Transactions on Biomedical Engineering*, vol. 39, no. 4, pp. 389–393, 1992.
[4] K E Kwok, S L Shah, B A Finegan, and G K Kwong, "An observational trial of a computerized drug delivery system on two patients," *IEEE Transactions on Control Systems Technology*, vol. 5, no. 4, pp. 385–393, 1997.
[5] V Serna, R Roy, and H Kaufman, "Adaptive control of multiple drug infusion," in *1983 American Control Conference*, San Francisco, CA, 1983, pp. 22–26.
[6] G I Voss, P G Katona, and H J Chizeck, "Adaptive multivariable drug delivery control of arterial pressureand cardiac output in anesthetized dogs," *IEEE Trans. Biomed. Eng.*, vol. 34, pp. 617–623, 1987.
[7] C L Yu, C L, R J Roy, H Kaufman, and B W Bequette, "Multiple–model adaptive predictive control of mean arterial pressure and cardiac output," *IEEE Trans. Biomed. Eng.*, vol. 39, no. 8, pp. 765–778, 1992.
[8] C M Held and R J Roy, "Multiple drug hemodynamic control by means of a supervisory–fuzzy rule based adaptive control system: Validation on a model," *IEEE Trans. Biomed. Eng.*, vol. 42, no. 4, pp. 371–385, 1995.
[9] A. C. Guyton and J. E. Hall, *Medical Physiology*, W. B. Sannders Company, 9th edition, 1996.
[10] R. M. Berne and M. N. Levy, *Cardiovascular Physiology*, The C. V. Mosby Company, 7th edition, 1997.
[11] C. Yu, R. Roy, and H. Kaufman, "A circulatory model for combined nitroprusside–dopamine therapy in acute heart failure," *Medical Progress Through Technology*, vol. 16, pp. 77–88, May 1990.
[12] R S Gopinath, B W Bequette, R J Roy, H Kaufman, and C Yu, "Issues in the design of a multirate model–based controller for a nonlinear drug infusion system," *Biotech. Prog.*, vol. 11, no. 3, pp. 318–332, 1995.
[13] R R Rao, J W Huang, B W Bequette, H Kaufman, and R J Roy, "Control of a nonsquare drug infusion system — a simulation study," *Biotechnology Progress*, vol. 15, no. 3, pp. 556–564, 1999.
[14] R R Rao, B Aufderheide, and B W Bequette, "Multiple model predictive control of hemodynamic variables — an experimental study," in *Proceedings of 1999 American Control Conference*, San Diego, 1999, pp. 1253–1257.
[15] K D Schott and B W Bequette, "Multiple model adaptive control," in *Multiple Model Approaches to Modelling and Control*, R Murray-Smith and T A Johanson, Eds., chapter 11, pp. 269–291. Taylor & Francis, London, UK, 1997.
[16] K. Sobel, H. Kaufman, and L. Mabius, "Model reference output adaptive control systems without parameter identification," in *18th IEEE Conference on Decision and Control*, Ft. Lauderdale, FL, 1979, pp. 347–351.
[17] H. Kaufman, I. Bar-Kana, and K. Sobel, *Direct Adaptive Control Algorithms*, Springer Verlag, 2nd edition, 1998.
[18] I. Bar-Kana and H. Kaufman, "Global stability and performance of a simplified adaptive algorithm," *Int. J. Control*, vol. 42, no. 6, pp. 1491–1505, December 1985.
[19] I. Bar-Kana and H. Kaufman, "Robust simplified adaptive control for a class of multivariable continuous systems," in *24th IEEE Conference on Decision and Control*, Ft. Lauderdale, FL, 1985, pp. 141–146.
[20] H. Kaufman and G. Neat, "Asymptotically stable multiple–input multiple–output direct model reference adaptive controller for processes not necessarily satisfying a positive real constraint," *Int. J. Control*, vol. 58, no. 5, pp. 1011–1031, November 1993.
[21] S. Ozcelik and H. Kaufman, "Robust direct model reference adaptive controllers," in *34th IEEE Conference on Decision and Control*, New Orleans, LA, 1995, IEEE, pp. 3955–3960.
[22] S. Ozcelik, *Design of Robust Feedforward Compensators for Direct Model Reference Adaptive Controllers*, Ph.D. thesis, Rensselaer Polytechnic Institute, Troy, NY, 1996.
[23] S. Ozcelik and H. Kaufman, "Frequency domain design conditions for robust direct model reference adaptive controllers," in *Proceedings American Control Conference*, Albuquerque, NM, 1997, IEEE, pp. 1652–1656.
[24] S. Ozcelik and H. Kaufman, "Design of MIMO robust direct model reference adaptive controllers," in *36th IEEE Conference on Decision and Control*, San Diego, CA, 1997, IEEE, pp. 1890–1895.
[25] Selahattin Ozcelik, Cesar C. Palerm, and Howard Kaufman, "Multi–drug infusion control using a robust direct adaptive controller for plants with time delays," in *Proceedings Mediterranean Control Conference*, 1999.

About the Authors (seated left to right)

Ramesh R. Rao received a B.E. in Chemical Engineering in 1991 from Annamalai University, India. He received his M.Tech. from the Indian Institute of Technology, Kanpur in 1993. Before joining Rensselaer Polytechnic Institute to pursue a Ph.D., he spent two years as a research engineer at the Tata Research Design & Development Centre, Pune. His current research interests are in model predictive control and gain scheduling in chemical and biomedical engineering applications.

Smooch is a 27kg female mixed hound and was the first canine in our experiments. She also was the first dog of five to be successfully adopted after our experiments, initiating a new adoption policy at Albany Medical Center's Animal Research Facility. She is currently a cherished member of the Aufderheide family.

Brian Aufderheide received his B.S. in Chemical Engineering from Lehigh University in 1987. After teaching math, chemistry and engineering at New Hampton School, he became a process engineer and quality control supervisor at Adell Compounding, Inc. Two and a half years later, Brian went back to school to earn a M.S. in Chemical Engineering from LSU. He then spent two years as Peace Corps volunteer in Honduras doing hillside agriculture. He is currently studying for a Ph.D at Rensselaer Polytechnic Institute. His research interests lie in biomedical control and modeling.

Cesar C. Palerm obtained his B.Sc. in Electrical Engineering in 1992 from ITESO University, Mexico. He received his M.Sc. in Electrical Engineering in 1996 and a M.Sc. in Mathematics in 1997, both from Rensselaer Polytechnic Institute, where he is currently pursuing his Ph.D. His research interests are in adaptive control and its biomedical applications. He is a member of the IEEE.

B. Wayne Bequette is a Professor of Chemical Engineering at Rensselaer Polytechnic Institute, where he conducts research on chemical process and biomedical systems control topics. He is an Associate Editor for Automatica, and is Program Chair for the 2001 American Control Conference and General Chair for the 2003 American Control Conference. He holds a B.S.Ch.E. from the University of Arkansas and a Ph.D. from the University of Texas. He was a process engineer with American Petrofina, a postdoctoral research associate at the University of Texas, and a visiting lecturer at the University of California at Davis. He has been on the faculty at Rensselaer Polytechnic Institute since 1988, and has spent sabbatical leaves at Merck and Co. and Northwestern University. He actively consults on process development and control problems in the pharmaceutical industry.