

RESEARCH

Open Access

Particle filter-based robust adaptive control for closed-loop administration of sodium nitroprusside

Nicolò Malagutti

Correspondence:
dr.nicolo.malagutti@ieee.org
Research School of Engineering,
The Australian National University,
Canberra 0200, Australia

Abstract

Automatic closed-loop administration of medicinal drugs has been the subject of intense research for decades due to its undisputed potential benefits in terms of cost savings and improved patient outcomes. However, concerns still exist about the ultimate safety of engineered feedback controllers. Manual methods remain dominant in clinical practice. In this context, we present a novel feedback control architecture, which combines multiple robust controllers with a particle filter-based method for real-time tracking of a patient's dose-response characteristic. The proposed method is applied to the case of the drug sodium nitroprusside, a vasodepressor used in the treatment of acute hypertension in intensive care and surgery, which is modelled as having a linear-time-varying dose-response characteristic. Our design takes into account the uncertainty in the patient response parameters, as well as potential nonzero-mean disturbances in the baseline arterial pressure and several possible time trends in the variation of the dose-response model. The performance and safety of the new approach are evaluated through an extensive computational simulation campaign. The results show that the proposed system can achieve adequate and safe feedback control of mean arterial pressure, thus validating our analysis and design. Our findings also highlight the fundamental - and possibly clinically overlooked - role of system excitation in ensuring that successful simultaneous identification and control of time-varying drug administration systems can be achieved.

Keywords: Biomedical engineering; Drug delivery; Feedback systems; Adaptive control; Robust control; Closed-loop; Hypertension; Sodium nitroprusside; Particle filtering

Background

In many clinical settings, the administration of suitable medicinal drugs is required to maintain important biological signals within an acceptable range. Unfortunately, significant variability exists in the way different patients respond to the same drug dose (interpatient variability), and variations can even occur in the response of one individual over time (inpatient variability). Appropriate dosing of drugs characterised by a narrow range of therapeutic concentrations and high variability in the response is therefore a challenging task. In these cases, *dose titration* protocols are followed, by which clinical operators administer an initial drug amount, generally according to population statistics

(e.g. age, gender, height, weight,...), followed by close monitoring and periodic manual adjustment of the dose on the basis of the clinical observations. This manual form of closed-loop regulation is prone to human error [1] and can be very time-consuming in clinical environments where timely intervention and/or staff levels can be an issue. In this context, the successful development of safe and effective automatic control methods could bring the significant benefits of improved patient outcomes and better allocation of clinical resources [2].

Over the last 30 years, much research has been devoted to the development of strategies for closed-loop feedback control of drug administration for a variety of applications, including control of haemodynamics, insulin for diabetes management, chemotherapy, anaesthesia and neuromuscular blockade [3-5]. Remarkably, despite promising results in clinical research publications, closed-loop methods for automatic drug administration have not yet been embraced by the medical community, and manual control remains, to date, the standard of care in the clinical setting. Indeed, leading scholars in the field have recently acknowledged that *"Success in the development of [...] closed-loop biomedical devices, will be contingent on the development of robust, verifiable advanced control algorithms"* [6]. This suggests that insufficient guarantees of system robustness may be a key reason - in terms of risk assessment considerations - underpinning clinicians' continued preference for manual methods.

This paper considers the specific problem of intravenous infusion of the drug sodium nitroprusside (SNP). The SNP case study typifies research into closed-loop drug administration. SNP is a fast-acting vasodepressor used in the treatment of acute hypertension in intensive care and surgery patients. The drug is highly effective in reducing mean arterial pressure (MAP); however, incorrect dosing may lead to undesired hypotensive peaks or metabolic toxicity [7], hence the interest in reliable automatic drug dosing. Following the clinical validation of a dose-response model by Slate [8], administration of SNP was the subject of intense research during the 1980s and 1990s, and a wide range of control engineering approaches were proposed, including self-tuning regulators [9], model reference adaptive control [10], multiple model adaptive control [11,12], fuzzy control [13] and model predictive control [14] (see [15] for a review). A commercial device (IVAC titrator; IVAC Medical Systems, Inc., San Diego, CA, USA) developed for SNP titration in intensive care applications reported success in clinical trials [16] but was never successfully marketed (as discussed in [17]). To the best of our knowledge, the automatic management of MAP through SNP and other drugs remains to this day an experimental pursuit.

The work herein presents a novel approach to the control of automatic infusion of SNP with a strong focus on robustness. The proposed methodology is an extension of an earlier approach which featured multiple robust feedback controllers designed with μ synthesis [18] and used Kalman-filter-based estimators [19] to determine controller selection in the context of a robust multiple model adaptive control architecture [20]. While retaining the same robust controllers as the previous work, the new approach uses particle filtering to generate an estimate of the dose-response characteristic in real time and exploits the estimation result to inform appropriate feedback control. The new method has been named Robust Adaptive Control with Particle Filtering (RAC-PF) and we have reported on a nonclinical case study featuring this architecture in [21], where we positively compared the performance of RAC-PF with that of a Kalman-filter-based approach. We deem particle filters to be better suited to the estimation of systems characterised

by time-varying parameters, non-Gaussian disturbances and even nonlinearities, such as pharmacokinetic-pharmacodynamic systems [22]. As a more general estimation tool, particle filtering allows us to remove undesirable *ad hoc* filter adaptations present in our earlier work. The resulting architecture can be readily transposed to other applications as required. This article provides an overview of the main characteristics of RAC-PF and reports on an extensive computational simulation campaign designed to assess the viability of the new methodology and its effectiveness in delivering safe automatic feedback control of SNP infusion for a broad range of response characteristics and disturbances.

The manuscript is structured as follows: the Methods Section comprises of three subsections detailing the SNP dose-response model, the proposed RAC-PF control architecture, and the characteristics of the computational simulations, respectively; the Results Section reports on the outcomes of the simulation campaign; finally, the Discussion and Conclusion Sections comment on the results and the potential of the proposed approach, and outline future research directions.

Methods

Model and problem description

Dose-response model

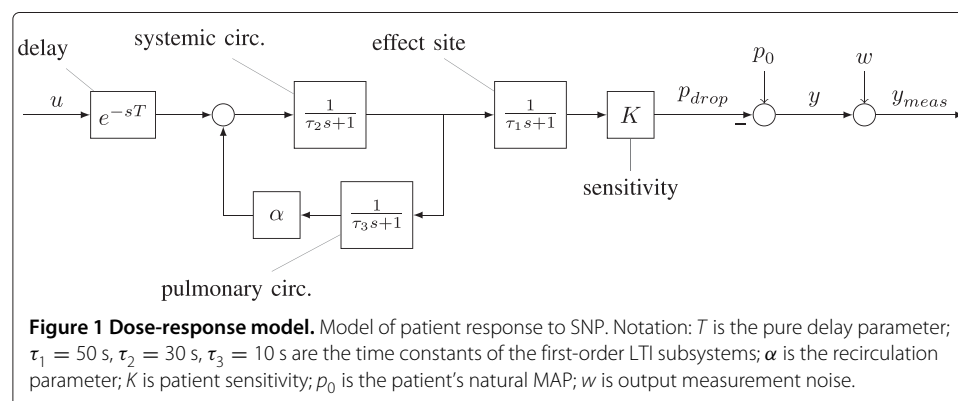
The dose-response model for a patient receiving SNP is shown in Figure 1. The system, taken from [12], is a variation of the experimentally validated model of [8]. Patient response is modelled as three interconnected first-order linear, time-invariant compartments representing the systemic and pulmonary circulation and the drug effect site. In transfer function form, the dynamic behaviour of the system is given by:

$$\frac{P_{\text{drop}}(s)}{U(s)} = e^{-sT} \frac{K(\tau_3 s + 1)}{((\tau_3 s + 1)(\tau_2 s + 1) - \alpha)(\tau_1 s + 1)} \quad (1)$$

and the output is given by the affine transformation

$$y_{\text{meas}}(t) = p_0(t) - p_{\text{drop}}(t) + w(t). \quad (2)$$

The three system parameters are deemed uncertain and *a priori* unknown: K , T and α . These represent the sensitivity in the patient's response, the recirculation fraction and the input time delay, respectively. The allowed ranges for the values and the rates of change of the parameters were chosen so as to render a very general formulation for the problem. Parameter variability has been considered to the best of the knowledge available in the literature, and further details on the assumed parameter ranges and allowable rates of



change are discussed in [19]. A summary is provided below (*note*: the model assumes a drug concentration of 200 µg/ml):

$$\begin{aligned} 0.25 < K(t) < 9.5, \quad \left[\frac{\text{mmHg}}{\text{ml/h}} \right] \quad \left| \frac{dK}{dt} \right| < 3K(t) \text{ h}^{-1} \\ 10 < T(t) < 50, \quad [\text{s}] \quad \left| \frac{dT}{dt} \right| < 40 \text{ h}^{-1} \\ 0.25 < \alpha(t) < 0.75, \quad \left| \frac{d\alpha}{dt} \right| < 0.5 \text{ h}^{-1} \end{aligned} \quad (3)$$

The offset term $p_0(t)$ is also treated as potentially time-varying in the interest of generality. As $p_0(t)$ represents a patient's underlying MAP, it is modelled here as a mainly low-frequency signal, given by the combination of its measurable value at $t = 0$ and an arbitrarily shaped but frequency-domain-bounded component, which we will refer to as *baseline disturbance* $p_{\text{dist}}(t)$:

$$\begin{aligned} p_0(t) &= p_0(t=0) + p_{\text{dist}}(t) \\ p_{\text{dist}}(t) : P_{\text{dist}}(s) &\leq \frac{0.4}{s+0.01} \end{aligned} \quad (4)$$

Thus, $p_0(t=0)$ can be regarded as a known fixed offset at the output, while the disturbance component can be treated as part of the state as shown in the following state space description of the patient response model:

$$\begin{aligned} \dot{x} &= Ax(t) + Bu(t - T(t)) + Lv(t) + Mp_{\text{dist}}(t) \\ y_{\text{meas}} &= p_0(t=0) - Cx(t) + w(t) \end{aligned} \quad (5)$$

with

$$A = \begin{bmatrix} -\left(\frac{1}{\tau_1} + \frac{1}{\tau_2} + \frac{1}{\tau_3}\right) - \left(\frac{1}{\tau_1\tau_2} + \frac{1-\alpha(t)}{\tau_2\tau_3} + \frac{1}{\tau_1\tau_3}\right) \frac{\alpha(t)-1}{\tau_1\tau_2\tau_3} & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 & -0.01 \end{bmatrix} \quad (6a)$$

$$B = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad L = \begin{bmatrix} 0.0625 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad M = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0.4 \end{bmatrix} \quad C = \left[0 \quad \frac{K(t)}{\tau_1\tau_2} \quad \frac{K(t)}{\tau_1\tau_2\tau_3} \quad 1 \right], \quad (6b)$$

where $v(t) \sim N(0, 1)$, $w(t) \sim N(0, 2)$ are normally distributed random noise signals at the input (actuation noise) and the output (measurement noise), respectively. The signal $y_{\text{meas}}(t)$ represents the measurable MAP value.

Control performance requirements

The following performance requirements apply [23]:

- A settling time of preferably 10 min or less, but no more than 15 min
- MAP should be contained within ± 10 mmHg of the desired set-point value at most times
- Under no circumstances should the system display resonant (persistent oscillatory) or unstable behaviour or cause MAP to drop below a pre-determined dangerous threshold (set at 60 mmHg for the purpose of this work)
- To ensure that SNP toxicity is prevented, the infusion rate should not exceed a pre-determined value (set here at 200 ml/h)
- High-frequency dynamics in the control signal should be limited since drug delivery is generally provided through mechanically actuated infusion pumps

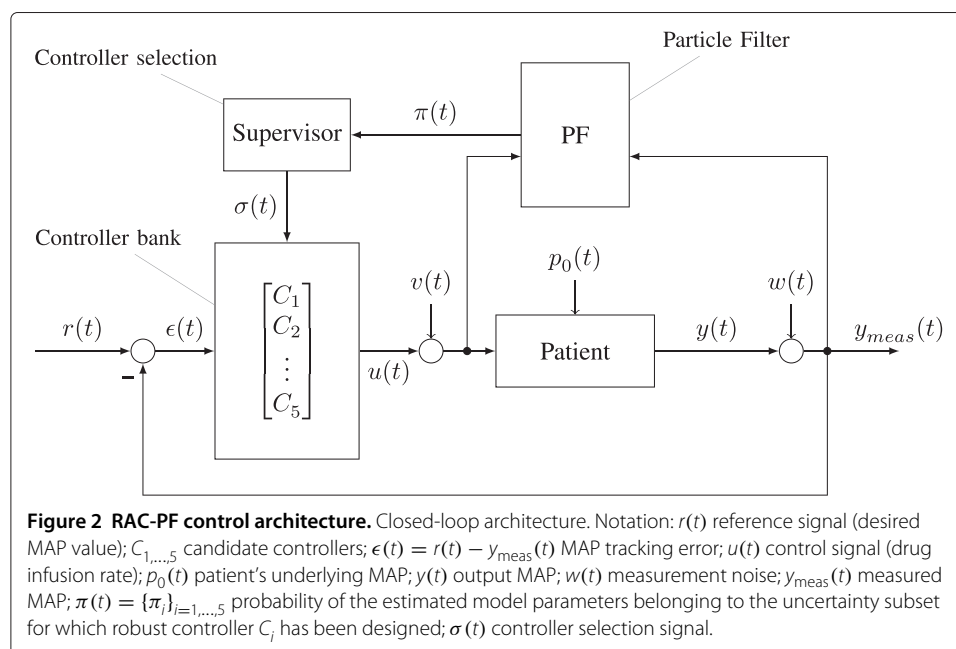
The control approach

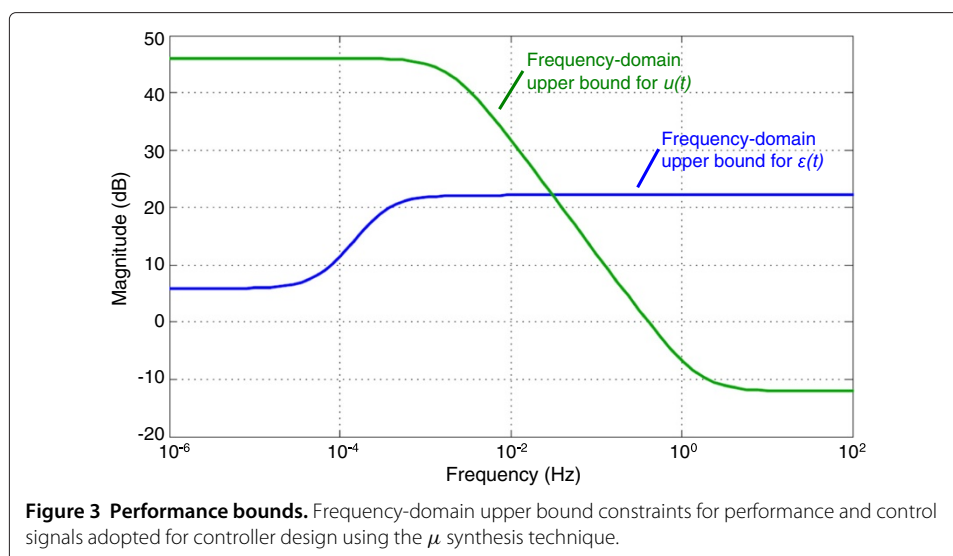
The RAC-PF architecture is shown in Figure 2. In the proposed approach, multiple controllers are required in order to ensure that a controller-plant pair capable of maintaining closed-loop stability and delivering the required level of performance exists for all possible values of the patient parameters. A particle filter tracks the parameters in real time and informs the controller selection algorithm with a probability result, which is in turn used to select the most appropriate controller for insertion in the loop.

Controller design

The controllers are designed using μ synthesis. This is an advanced control engineering method used in the design of linear feedback controllers for applications where stability and performance must be guaranteed in the face of uncertainty in the model of the system to be controlled (robustness). In μ synthesis, a dynamical system is considered along with a model or bounding function of any uncertainty it features, whether structured (e.g. parametric) or unstructured (e.g. arising from the presence of delays and/or nonlinearities). Any performance specifications are also expressed in terms of frequency domain bounds. The technique involves a numerical search over the space of stabilising controllers and enables the designer to compute a feedback controller capable of delivering robust performance in the closed-loop configuration, if one such controller exists [24].

Controller design and μ synthesis are not the focus of this paper and we refer the reader to [18,19] for a detailed description of the controller design methods and results as applied to the SNP dose-response model of (1). For the purpose of the presentation hereof, it will suffice to report that the performance specifications were imposed as frequency domain bounds on signals $\epsilon(t)$ and $u(t)$ as shown in Figure 3. These constraints translate into requirements for a small reference tracking error at low frequencies (to ensure proper target following at steady state) and a mainly low-frequency infusion rate signal (to cater for





the slew-rate limitations of the actuator, likely a motorised infusion pump). Our μ synthesis computations have shown that a single linear feedback controller could not deliver the required level of performance due to the large ranges of parametric uncertainty (in particular the sensitivity parameter K). However, satisfactory performance could be achieved by subdividing the uncertainty range of K into multiple subsets and designing a robust controller for each subset. Figure 4 and Table 1 describe the five controllers resulting from our synthesis and the corresponding uncertainty subsets. Note that all controllers cater for the full uncertainty range of parameters α and T , while, with respect to K , each controller is associated with a region of *best performance* and a region of *acceptable performance* based on a minimum required normalised performance index $A = 1$.

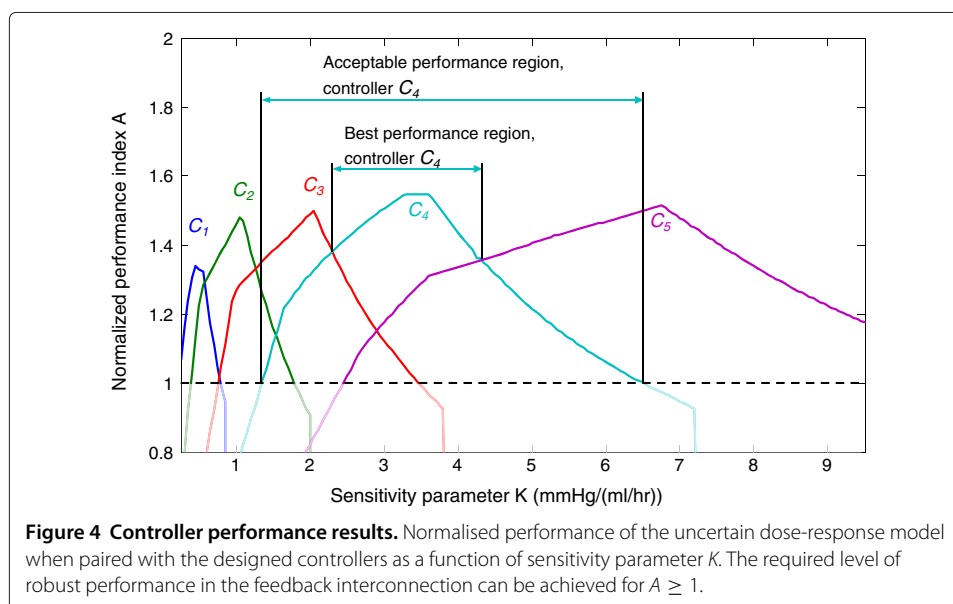


Table 1 Results of controller design

Controller number i	To suit K (best performance) (mmHg/(ml/h))	To suit K (acceptable performance) (mmHg/(ml/h))	To suit α	To suit delay T (s)
1	0.25-0.57	0.25-0.78	0.25-0.75	0-50
2	0.57-1.25	0.37-1.77	0.25-0.75	0-50
3	1.25-2.30	0.75-3.45	0.25-0.75	0-50
4	2.30-4.32	1.33-6.50	0.25-0.75	0-50
5	4.32-9.50	2.45-9.50	0.25-0.75	0-50

Particle filtering

Particle filtering is a sequential Monte Carlo method which iteratively computes the conditional probability distribution of the state of a partially observed dynamical system described by a stochastic state-space model [25]. In this application, the main focus is not to estimate the state but rather to identify the parameters (particularly patient sensitivity K) so that the correct controller can be used in the feedback loop. To this end, the estimation problem is recast as a nonlinear tracking problem to include the uncertain parameters in the state, as shown in (7) below. We refer to the original system states as 'linear' states x^l , as opposed to the 'nonlinear' states x^n representing the parameters. Such a mixed linear/nonlinear formulation is well suited to the application of a specialised form of particle filtering called *marginalised particle filtering* [26].

Particle filtering is a numerical, therefore discrete-time approach. Due to the nature of clinical MAP observations (require averaging over at least 1 cardiac cycle), we consider a sampling time $T_s = 2$ s [19]. The discrete-time description of the system used in the particle filter is given by

$$\begin{aligned} x^l(k+1) &= A_d(x^n(k))x^l(k) + B_d u \left(k - \left\lfloor \frac{T(x^n(k))}{T_s} \right\rfloor \right) + L_d v(k) + M_d p_{\text{dist}}(k) \\ x^n(k+1) &= f(x^n(k)) \\ y_{\text{meas}}(k) &= p_0(0) - C_d(x^n(k))x^l(k) + w(k), \end{aligned} \quad (7)$$

where $k = \left\lfloor \frac{t}{T_s} \right\rfloor$ ($\lfloor \cdot \rfloor$ is the *floor* operator). The subscript d indicates zero-order hold discretisation of continuous-time model (5).

The expression $f(x^n)$ indicates a function of x^n and describes the time update of the nonlinear states. It is given by:

$$x^n(k+1) = \begin{bmatrix} K(k+1) \\ T(k+1) \\ \alpha(k+1) \end{bmatrix} = x^n(k) + \chi(k), \quad (8)$$

where $\chi(k)$ is sampled from an array of probability distributions which are intended to reflect the likely trajectory of the parameters. Uniform distributions (U) are used to capture the constraints on the rate of change of parameters expressed in (3) while making no assumptions on possible trends (tracking by *random walk*), as shown below:

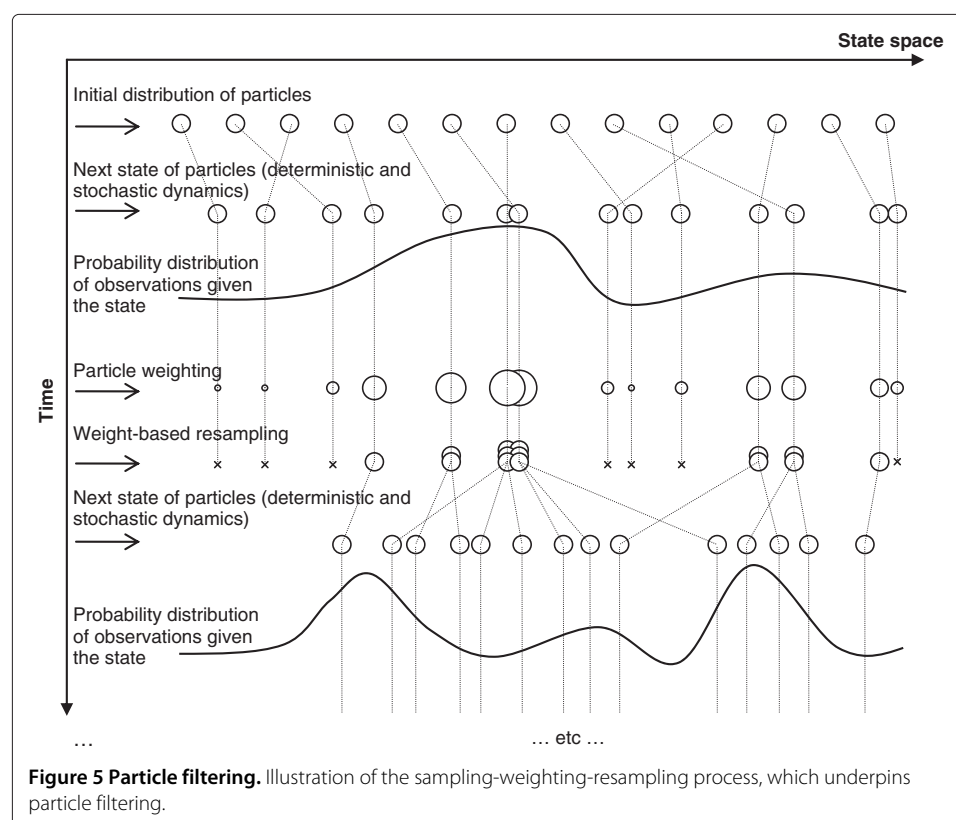
$$\chi(k) \sim \begin{bmatrix} U(-0.017K(k), 0.017K(k)) \\ U(-0.028, 0.028) \\ U(-0.00028, 0.00028) \end{bmatrix}. \quad (9)$$

The aim of filtering is to obtain the posterior probability density of the state conditioned on the observations up until that time point, i.e. $p(x(k)|Z(0:k))$, where p indicates

probability and $Z(0:k) \equiv \{[u(i) \ y_{\text{meas}}(i)]\}_{i=0}^k$, represents the observations. The posterior probability can be calculated analytically in the case of linear-Gaussian systems using Kalman filters, while, for more general cases, numerical approximation methods must be used [26]. Particle filters are one such method, which approximates the posterior probability with a finite number of samples (particles).

Particles can be understood as realisations of the system to be estimated, sampled from a known or presumed initial probability distribution $p(x(0))$. At every time step, the state of each particle is updated according to the dynamics of the realisation it represents. The output of each particle is then compared with the actual observed output and a weight is generated for the particle. A resampling step follows, during which particles with poorer weights have a higher chance of being discarded and reintroduced as a copy of a better performing realisation. Over a number of iterations, particles will cluster in the state space in a way that approximates the posterior probability density of the state. The process is illustrated graphically in Figure 5. As the number of particles increases, so does the accuracy of the approximation but also the computational burden associated with computing a large number of candidate realisations in parallel. The number of particles required to achieve a given level of approximation grows exponentially with the dimensionality of the target system, rendering the estimation of complex problems computationally onerous.

Marginalised particle filters exploit the fact that a subset of the states can be treated as conditionally linear (x^f here). This methodology involves an estimation of these states using the optimal Kalman filter result (marginalisation), while the other states (x^n) are estimated by the particle filter. As the dimensionality of the numerical problem is thus



reduced, marginalised particle filtering is associated with a comparatively lower computational burden. Here follows a description of our algorithm, which combines a bootstrap particle filter implementation [27] with marginalisation.

Marginalised particle filter algorithm

- (a) *Initialisation.* State $x_i(0)$ for particle $i = 1, \dots, N$ is set as

$$x^l(0) \sim (0, P(0)) \quad x^n(0) \sim \begin{bmatrix} U(0.25, 9.5) \\ U(10, 50) \\ U(0.25, 0.75) \end{bmatrix}$$

with state estimate covariance matrix $P(0) = 0$, as $x^l(0)$ certainly equals 0 since the patient has received no drug for $t < 0$. Also, learn p_0 from observation $y(0)$. Set $k = 0$.

- (b) *Weighting.* For $i = 1, \dots, N$ compute the estimated output from each particle as $\hat{y}_i(k) = p_0 - C_d(x_i^n(k))x_i^l(k)$. Then, evaluate the particles' *normalised importance weights* $\tilde{q}(k)$

$$q_i = p(y(k)|\hat{y}_i(k)) \quad \tilde{q}_i(k) = \frac{q_i(k)}{\sum_{j=1}^N q_j(k)}.$$

- (c) *Resampling.* Resample N particles on the basis of the weights obtained in step (b) using a residual resampling algorithm [28].

- (d) *Time update.* For each particle $i = 1, \dots, N$

- (i) Kalman filter correction of the linear state estimate using the available observation $Z(k)$

$$x_i^l(k|k) = x_i^l(k|k-1) + H_i(k)(y(k) - p_0 - C_{d,i}(k)x_i^l(k|k-1)),$$

with

$$H_i(k) = P_i(k|k-1)C_{d,i}(k)S^{-1}(k)$$

$$S_i(k) = C_{d,i}(k)P_i(k|k-1)C_{d,i}^T(k) + R$$

$$P_i(k|k) = P_i(k|k-1) - H_i(k)C_{d,i}(k)P_i(k|k-1),$$

where H is the Kalman gain, S is the innovation covariance, and R is the variance of noise signal w .

- (ii) Sample χ via (9) and update the nonlinear states via (8), then calculate $A_{d,i}(x_i^n(k))(k)$ and $B_{d,i}(x_i^n(k))(k)$ using (7).

- (iii) Time update of the marginalised states ($x_i^l(k+1|k)$) via (6) and the state estimate covariance matrix using

$$P_i(k+1|k) = A_{d,i}(k)P_i(k|k)A_{d,i}^T(k) + Q, \text{ where } Q \text{ is the covariance matrix of the state/input noise } v.$$

- (e) *Iteration.* Increase $k \rightarrow k+1$ and repeat over from step (b).

Controller selection

Controller selection is carried out by integrating the approximate probability distribution which results from particle filtering. The number of particles n_j associated with each of the *best performance* subsets listed in Table 1 is proportional to the probability of

controller j being the correct one for insertion in the loop. In a weighted approach to controller selection the drug infusion rate u can thus be computed as a weighted sum of the control signals u_j generated by each controller

$$u = \sum_{j=1}^5 \pi_j u_j, \quad \pi_j = \frac{n_j}{N}, \quad (10)$$

where π_j is the probability of the true parameters belonging to subset j and N is the total number of particles.

Theoretical proofs of robustness

The use of the μ synthesis controller design techniques gives a mathematical guarantee of robust stability and performance as long as the true patient sensitivity value is matched by the correct feedback controller (Table 1). Proper operation of the closed-loop system hinges, therefore, on achieving adequately accurate identification of the patient's individual response characteristic.

The inferred parameter probability distribution generated by the particle filter approaches the true $p(x(k)|Z(0:k))$ for $N \rightarrow \infty$ [29] and can be deemed a suitable approximation if a sufficiently large amount of particles is used. Therefore, it is fair to expect for the estimate to converge and, thus, for the system to deliver the required performance, *asymptotically in time*. The presence of a numerical tool in the architecture, however, makes it difficult - if at all possible - for a mathematical proof to be developed. To the best of the author's knowledge, no theoretical proofs of robust performance exist for the several 'robust adaptive control' approaches described in the control literature to date [30], and proofs of asymptotic stability have been developed in restricted cases only (e.g. for linear-time-invariant Gaussian systems in [31]).

In a general formulation such as we consider here, not only would available proofs not apply, but also asymptotic results would be clinically inadequate as the patient must be successfully identified and controlled over a bounded time horizon. In light of the above, while RAC-PF is equipped with conservatively designed controllers and very general estimation tools, its ability to maintain stability and deliver the required level of performance can, at this time, only be evaluated heuristically. We do so, in this work, using computer simulations.

Numerical simulations

A broad computational simulation campaign was conducted in order to evaluate the ability of the proposed adaptive control approach to control MAP through SNP administration in a wide variety of conditions. To reflect the unpredictable nature of blood pressure disturbances and patient parameter variations, a large number of cases were randomly generated and simulated. All simulations involved a control horizon of 10,000 s (approximately 2 h and 45 min), with a target MAP of 100 mmHg for the time period of 0 to 4,000 s and 80 mmHg for 4,000 to 10,000 s.

Two categories, or *simulation streams*, of hypertensive patients were simulated as follows:

- (a) Relatively 'settled' patients, i.e. displaying elevated MAP ($p_0 = 120$ mmHg at $t = 0$) with $p_{\text{dist}}(t)$ modelled as a random, zero-mean, low-intensity additive disturbance (in the range ± 6 mmHg)

- (b) More 'unsettled' patients, i.e. displaying the same initial MAP as (a) but greater intensity of random fluctuations (in the range ± 15 mmHg), as well as two step increases in $p_{\text{dist}}(t)$ (+20 mmHg) at 2,000 and 5,500 s, modelling a worsening hypertensive condition

Scenarios of controlled MAP reduction of 20 to 80 mmHg were deemed quite general while clinically plausible.

The $p_{\text{dist}}(t)$ signal used in the simulations was generated by filtering the superposition of a Gaussian white noise process and step signals (the latter only for stream b) with a suitable low-pass filter to meet the frequency domain bound assumption of (4). The seeds for the generation of all random processes were also randomised and changed for every simulation: no two simulations in the campaign, therefore, would exhibit the same $p_0(t)$, $w(t)$ and $v(t)$ traces.

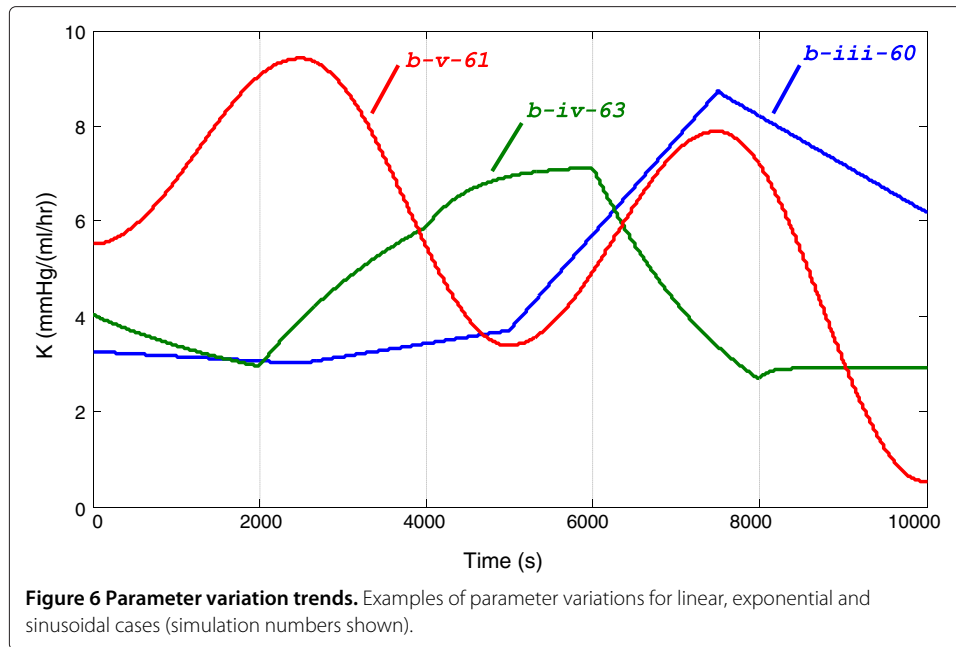
For each simulation stream, the following simulation batches were computed:

- (i) Forty simulations in which mid-range parameter values were chosen, i.e. $K = 4$ mmHg/(ml/h), $\alpha = 0.5$, and $T = 30$ s and held fixed throughout the control horizon
- (ii) Eighty simulations in which the three model parameters K , α , and T were randomly selected from the allowed ranges as per (3) and held *constant* throughout the control horizon
- (iii) Eighty simulations in which the initial values of the parameters were randomly selected as in (ii) with *piece-wise linear variations* throughout the control horizon. A random number of slope changes between 2 and 5 was selected for each run. The slope of the ramp was also selected at random each time ensuring it would not exceed the constraints of (3)
- (iv) Eighty simulations in which the initial values of the parameters were randomly selected as in (ii) with *piece-wise exponential variations* throughout the control horizon. Between 2 and 5 changes were again randomly selected for each run. The exponent was also randomly selected in such a way that the slope of the resulting curve would not exceed the constraints of (3)
- (v) Eighty simulations in which the initial values of the parameters were randomly selected as in (ii) with *piece-wise quarter-sinusoidal variations* throughout the control horizon. Between 2 and 5 changes were again randomly selected for each run. The target value for each change was randomly selected in such a way that the slope of the resulting curve would not exceed the constraints of (3)

The main purpose of batch i was to help verify the ability of the approach to deliver repeatable results. Batches ii to v were intended to test the ability of the control approach to respond to a wide range of parameter variations (for examples of variation trends see Figure 6). The total number of simulations was 720, corresponding to 2,000 *in silico* patient hours.

Due to the large amount of available data, aggregate measures for the control and identification performance were adopted. For each simulation run, the following measures were computed:

1. Measures concerning control performance
 - t_{c_1}, t_{c_2} - the convergence time (10% to 90% of transition) as the two MAP setpoint changes are imposed



- $t_{\pm 5}, t_{\pm 10}, t_{\pm 15}$ - the time (out of 10,000 s) y_{meas} remained within $\pm 5, \pm 10$ and ± 15 mmHg of the setpoint $r(t)$, respectively
- $\epsilon_{\text{max}}, \epsilon_{\text{min}}, \epsilon_{\text{avg}}$ - the positive and negative peaks and average setpoint tracking error recorded over the control horizon
- $t_{\text{pp}}, t_{\text{npp}}$ - the time the RAC-PF controller and the true patient system form a *provably performant* (pp) or *nonprovably performant* (npp) pair as evaluated through μ analysis [24], i.e. a pp closed-loop pair would achieve a performance index $A \geq 1$ in Figure 4 (*note*: due to the inherent conservativeness of μ analysis, a npp pair does not necessarily indicate an unstable or underperforming closed-loop condition, but only that there is no mathematical proof of robust performance. At *npp* times, loop performance should be evaluated from other indices)
- t_{sat} - the time the infusion rate signal remained at the maximum allowed value (controller saturation)

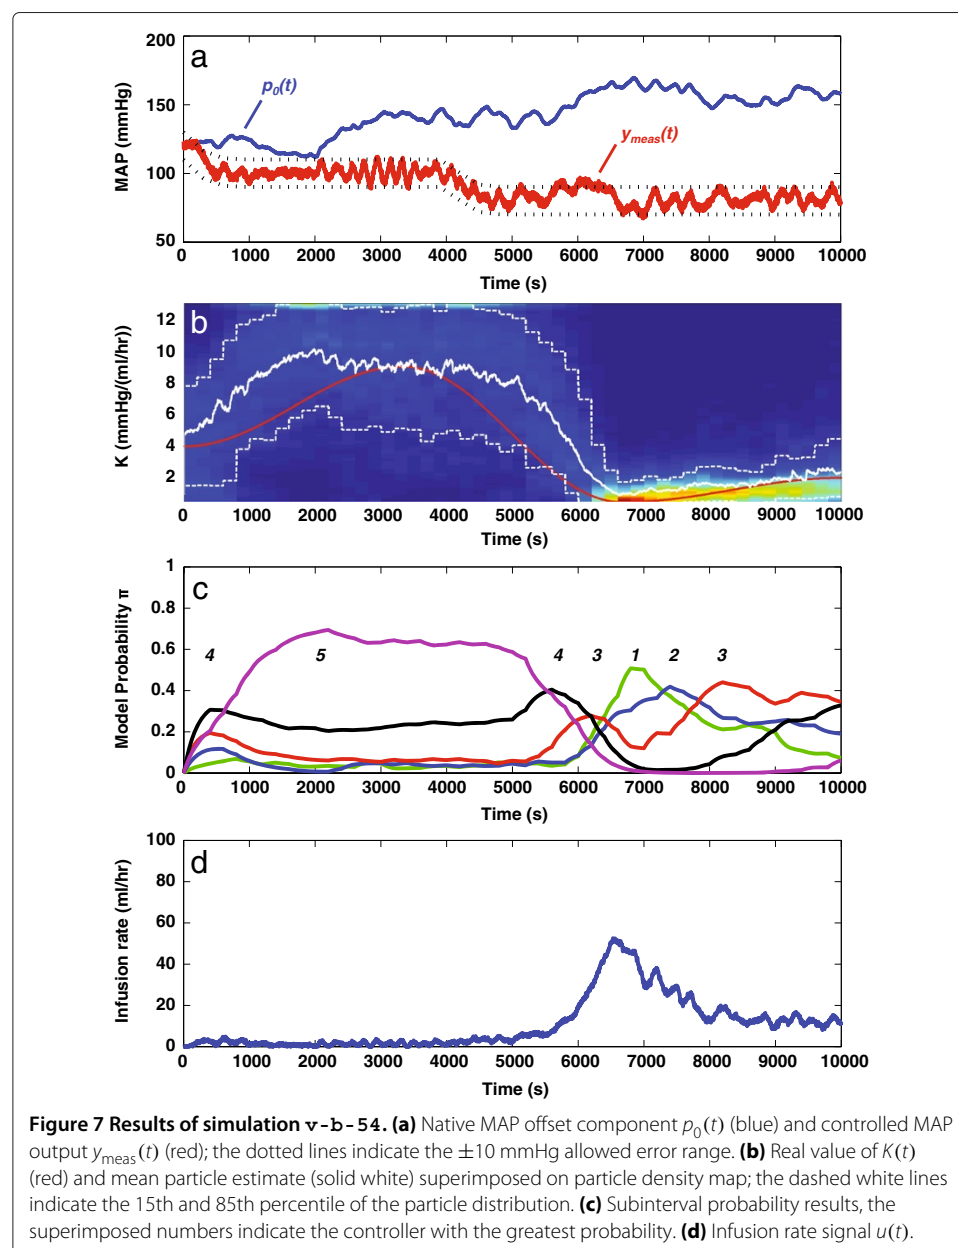
2. Measures concerning estimation of the response characteristic

- t_0 - time (out of 10,000 s) for which the K subinterval deemed the most likely by the particle filter, i.e. with the greatest π_i (intervals as per *best performance* in Table 1) contains the true simulated value of K
- t_1 - time for which the K subinterval deemed the most likely by the particle filter is a neighbouring interval to that containing the true value of K
- t_2 - time for which the K subinterval deemed the most likely by the particle filter is 2 intervals away from that containing the true value of K
- $K_{\text{rel}}, T_{\text{rel}}, \alpha_{\text{rel}}$ - the mean relative estimation error for parameters K , T and α , respectively (the mean of the particles is used to determine the estimate)

The simulation environment was programmed using Matlab and Simulink (Math-Works, Natick, MA, USA) and run on a standard desktop computer (Intel Core 2 Quad CPU, 3.0 GHz).

Results

Figure 7 shows the results of simulation v-b-54: a case featuring large variations in the SNP response parameters over time and significant fluctuations in baseline MAP. From the graphs, it can be readily appreciated that the proposed control approach tracks the simulated time-varying response characteristics and meets the required performance specifications.



For completeness and to provide context in the interpretation of the aggregate results for groups of simulations, Table 2 lists the overall results for simulation v-b-54.

The aggregate results for all of the simulations are shown in Table 3. For each measure, the mean value and standard deviation over the corresponding batch are displayed.

The results for batches i-a and i-b confirm the repeatability of results using the proposed method. The standard deviation values are much lower than for other batches, indicating that the system was able to deliver consistent outcomes, as required, when applied multiple times to the same case. The fact that the standard deviation values are not 0 can be explained in terms of the stochastic features present in both the patient model and the particle filter. As no two simulations are identical, some variability in the results exists even with parameters K , α and T being exactly the same.

When considering the other batches (i to v-a/b), the results show that control of MAP was successfully achieved in all cases, with no instability or dangerous pressure drops observed in any of the computed simulations. The controlled MAP trace was maintained within ± 10 mmHg of the desired setpoint for over 94% of the time in any given batch of simulations (considering $M_{t_{\pm 10}} - 2\sigma$, i.e. 9,404.74 s for the worst-performing batch v-b). Similarly, the occurrence of temporary deviations $> \pm 15$ mmHg from the setpoint was prevented for over 98% of the time and peak tracking errors (ϵ_{\max} , ϵ_{\min}) had a magnitude of less than 20 mmHg at all times. The tracking of the setpoint was achieved with negligible bias as shown by the very low value of ϵ_{avg} for all batches. It should be clarified that the above statistics refer to the whole simulation (including the initial condition and the two setpoint changes), thus suggesting that control performance remained entirely adequate throughout the simulation campaign. Furthermore, performance was consistent across simulation batches, indicating that the system can deal with time-varying parameters regardless of the shape of the variations, as long as the assumptions of (3) are met. Transition times $t_{c_{1,2}}$ also met the specifications and were contained below 10 min in almost all cases, with the longest observed transition taking approximately 13 min to settle. As could be expected, due to the simpler nature of the control task (lower disturbance),

Table 2 Aggregate results for simulation v-b-54

Parameter	Result
t_0 (s)	5,924
t_1 (s)	3,442
t_2 (s)	634
t_{pp} (s)	9,716
t_{npp} (s)	284
K_{rel} (%)	81
α_{rel} (%)	-25
T_{rel} (%)	-22
t_{c_1} (s)	186
t_{c_2} (s)	332
t_{sat} (s)	0
$t_{\pm 5}$ (s)	6,558
$t_{\pm 10}$ (s)	9,368
$t_{\pm 15}$ (s)	9,980
ϵ_{\max} (mmHg)	16.7
ϵ_{\min} (mmHg)	-14.07
ϵ_{avg} (mmHg)	1.08

Table 3 Summary of results of the simulation campaign

Batch-stream	t_0 [s]	t_1 [s]	t_2 [s]	t_{pp} [s]	t_{npp} [s]	K_{rel} (%)	α_{rel} (%)	T_{rel} (%)	t_{c1} [s]	t_{c2} [s]	t_{sat} [s]	$t_{\pm 5}$ [s]	$t_{\pm 10}$ [s]	$t_{\pm 15}$ [s]	ϵ_{max} [mmHg]	ϵ_{min} [mmHg]	ϵ_{avg} [mmHg]
i-a	604	9,396	0	10,000	0	69	11	-24	396	360	0	9,442	9,995	10,000	8.43	-8.33	0.16
	(128)	(130)	(0)	(0)	(0)	(5.0)	(3.0)	(5.0)	(46)	(59)	(0)	(104)	(14)	(0)	(0.72)	(0.95)	0.05
i-b	5,986	4,014	0	9,640	360	40	-4.0	-20	397	332	0	8,752	9,968	9,999	11.00	-10.28	0.43
	(788)	(788)	(0)	(512)	(512)	(5.0)	(3.0)	(3.0)	(74)	(72)	(0)	(212)	(147)	(0.6)	(1.46)	(1.18)	(0.06)
ii-a	4,926	3,790	1,280	9,826	167	63	3	25	390	375	0	9,356	9,944	9,994	8.56	-8.43	0.21
	(4,338)	(3,821)	(2,414)	(252)	(252)	(79)	(24)	(54)	(126)	(101)	(0)	(612)	(146)	(34)	(1.98)	(1.98)	(0.22)
ii-b	4,804	4,443	716	9,478	522	42	2	27	392	359	0	8,856	9,907	9,990	11.06	-9.85	0.43
	(3,038)	(2,682)	(1,324)	(576)	(576)	(54)	(22)	(57)	(101)	(75)	(0)	(868)	(87)	(40)	(3.24)	(2.98)	(0.25)
iii-a	2,971	5,358	1,617	9,898	102	85	14	10	421	385	0	9,289	9,941	9,992	9.32	-8.97	0.24
	(2,297)	(2,166)	(1,911)	(204)	(204)	(73)	(19)	(26)	(120)	(110)	(0)	(435)	(114)	(42)	(2.10)	(2.72)	(0.23)
iii-b	4,710	4,460	731	9,396	604	49	2.0	10	428	380	0	8,710	9,898	9,991	12.03	-10.94	0.60
	(1,938)	(1,772)	(882)	(694)	(694)	(39)	(15)	(24)	(124)	(104)	(0)	(612)	(151)	(41)	(2.44)	(2.90)	(0.35)
iv-a	2,152	5,802	1,938	9,860	140	77	12	9	408	384	0	9,405	9,974	10,000	8.78	-8.51	0.21
	(2,096)	(1,945)	(2,135)	(207)	(207)	(59)	(21)	(25)	(89)	(96)	(0)	(398)	(88)	(15)	(1.86)	(1.72)	(0.15)
iv-b	4,701	4,474	729	9,349	651	45	3.0	11	408	366	0	8,860	9,936	9,995	11.44	-10.15	0.50
	(1,730)	(1,744)	(970)	(796)	(796)	(38)	(17)	(25)	(94)	(101)	(0)	(471)	(90)	(31)	(2.15)	(1.95)	(0.21)
v-a	2,834	5,350	1,689	9,898	102	94	13	12	419	404	0	9,164	9,930	9,994	9.95	-8.95	0.37
	(2,159)	(2,030)	(1,344)	(205)	(205)	(62)	(18)	(25)	(110)	(97)	(0)	(117)	(116)	(26)	(2.30)	(1.71)	(0.36)
v-b	4,529	4,568	848	9,298	702	48	0.0	14	390	405	0	8,450	9,846	9,976	13.15	-11.76	0.68
	(1,663)	(1,534)	(882)	(871)	(871)	(31)	(15)	(26)	(107)	(124)	(0)	(674)	(220)	(80)	(3.64)	(3.84)	(0.47)

Results are provided as the mean value across a batch of simulations, with standard deviations shown in parentheses. List of performance indices: t_i , time the model identified with highest probability is the i th neighbour of the true response model; t_{pp} , time spent by the closed-loop system in a provably performant state; t_{npp} , time spent by the closed-loop system in a nonprovably performant state; X_{rel} , average relative error in the estimate of parameter X over a 10,000-s simulation; $t_{c1,2}$, transition time for first and second setpoint change; t_{sat} , time of saturation of control action; $t_{\pm\delta}$, time the controlled MAP remained within δ x mmHg of the target MAP;

$\epsilon_{max,min,avg}$, greatest positive and negative values and average MAP error over a 10,000-s simulation.

the performance results for stream a were better than those for stream b, although only marginally.

The results associated with the identification of patient parameters by the particle filter show that the system is tracked as it evolves through time (with t_0 and t_1 combined representing about 90% of the total simulation time on average). However, the system slightly but consistently overestimates K and delivers estimation errors K_{rel} of about 50% to 70% on average. The large standard deviation values for α_{rel} and T_{rel} , on the other hand, suggest that the latter two parameters are not dependably identified (there is poor convergence of the particle scatter with respect to α and T). Although the estimation errors may seem significant, the precision in identification is adequate in the context of the required performance, since each controller can cater for variations of two- to fourfold in the value of K , and for all possible values of α and T (Table 1). The t_{pp} result confirms this by showing that the system pairs the patient with a suitable controller at most times, with closer inspection of simulation data showing that t_{npp} is accrued mostly in the initial stages when the particle filter has not yet converged. Notably - and perhaps counterintuitively at a first glance - identification of the patient response is more accurate for the more challenging cases of stream b than for stream a as indicated by higher values of t_0 .

Each simulation took approximately 8 min to execute (corresponding to a simulation time-to-real-time ratio of 1:20) using 1,000 particles. This number of particles was found to deliver a reasonable compromise between accuracy in the results and computation time. A number of test simulations conducted ahead of the campaign did not show noticeable improvements in the estimation results when run using 5,000 particles.

Discussion

We have presented RAC-PF, a novel approach for the control of uncertain, time-varying systems, and tailored it to the case study of automatic closed-loop SNP administration for the management of acute hypertension. Having adopted an underlying response model which is, to our knowledge, the most general ever adopted in the literature (with regard to the ranges of variability of parameters, the allowed rates of change and shapes of variations, and the presence of both random and non-zero-mean changes in $p_0(t)$), the system delivered the required performance with no evidence of unsafe behaviour throughout the extensive simulation campaign presented herein. We therefore propose that RAC-PF may represent a viable solution to assist with automated control of hypertension not only in the traditionally considered post-operative setting, but also in the more challenging intraoperative context, where rapid changes in a patient's response characteristic may occur [32].

Only a limited number of authors have considered automatic control of MAP during surgery (although not specifically with SNP). The challenge of on-line adaptation to variations in the patient response has generally been addressed through *ad hoc* rules [33] which mimic physicians' clinical decision process [34]. The RAC-PF method is fundamentally different. By adopting a dose-response model which incorporates uncertainty, μ synthesis allows the designer to conclusively determine a priori the minimum number of controllers required to achieve the desired level of robust performance, while particle filtering tracks the patient's response as it evolves through time to inform correct controller selection. In addressing the argument that methods for automatic drug administration may have been dismissed on grounds of perceived safety in the past, we deem the transparent nature

of RAC-PF, where an explicit relationship is retained between the estimated posterior probability distribution of the patient response and the chosen control action, to be well suited to clinical applications. A further advantage of this work is that the framework is not specific to SNP and the same methodology could be used to design controllers and estimators, given a reference dose-response model, for a different drug delivery problem.

The results also support the viability of particle filtering in a real-time pharmacological application. Particle filters are a very general estimation tool and can be applied to linear as well as nonlinear problems, but are known to be computationally onerous. The fact that simulations ran much faster than real time without requiring particular hardware or software optimisation was therefore a notable positive finding. Particle filter methods have been proposed to assist with the estimation of physiological systems and draw inference from patient data in order to support clinical decision-making processes (see, e.g. [35]). To our knowledge, however, such methods have been generally reserved for off-line computations. Our work highlights the potential of particle filter as a viable tool for real-time closed-loop system identification. As pharmacokinetic/pharmacodynamic models commonly combine linear and nonlinear dynamics [22], we anticipate that marginalisation would be applicable to a broad class of problems in this area.

On the point of identifiability of the system, a number of relevant comments can be made. Perhaps counterintuitively, the results show good control performance, but the identification statistics point to somewhat imprecise parameter estimates. As well as being a result of the underdetermined nature of the estimation problem (multiple uncertain parameters and a single output), this outcome stems from a well-known trade-off between control and identification in feedback systems [36]. The very purpose of a feedback controller is to suppress some of the dynamics from the controlled system, thus the stronger the control action (performance), the lesser input-output information is available to estimate the response characteristic. An extreme situation would be that of a controller capable of forcing a perfect flat-line MAP output at all times: in such a situation no knowledge of the patient characteristic could be gathered (arguably, adaptive control would not be required, either). The necessary imposition of performance constraints in our design, therefore, affects the precision of system identification. The fact that the controllers were designed to cater for all values of α and T and, as a result, the identification of these parameters was inaccurate, is a clear example of this trade-off. Furthermore, accuracy in identification depends on the level of *excitation* of the system, i.e. the amount of energy delivered by exogenous inputs (in this case, $r(t)$ and $p_0(t)$), which elicits observable dynamics at the output. Since $r(t)$ is mostly a constant signal, excitation is provided mainly by the changes in $p_0(t)$. This explains why identification is more accurate (greater t_0) for the more disturbed simulations of stream b. It should be remarked that the trade-off between identifiability and excitation is an inescapable one and affects both traditional manual methods and computer-based solutions equally. A notable question arising in the light of these considerations is whether the commonly stated clinical goal of maintaining a patient in a *settled* state is a desirable one or whether this is actually counterproductive when changes in a patient's response must be timely identified and acted upon. Our analysis suggests that a clinically acceptable, time-varying MAP target would deliver greater system excitation and thus improve ongoing system identification.

Conclusion

Through the extensive simulation campaign presented in this manuscript, we have validated, *in silico*, the methodology underpinning the novel RAC-PF architecture and demonstrated its applicability to the problem of automatic closed-loop administration of SNP for the control of MAP in a model of acute hypertension. These results provide a strong motivation for the approach to be tested further in *in vivo* models (e.g. animal testing).

We have also proven the viability of real-time computation of particle filters on consumer hardware for the relatively simple SNP dose-response system. The analysis of other case studies, including highly nonlinear and multiple-input-multiple-output systems, will be essential to identify any limitations to the practical use of particle filtering in estimating more complex pharmacological response models.

Finally, our analysis has highlighted the known and inescapable trade-off between identification and control in feedback systems and the role of system excitation in adaptive control. While identification proved adequate for control in the SNP case study, this could only be verified heuristically through simulations. In future work we plan to research mathematical methods to quantify the trade-off and design clinically acceptable inputs to guarantee dependable system identification, and thus greater theoretical guarantees of robustness, for this and other automatic drug delivery applications.

Competing interests

The author declares no competing interest.

Authors' information

NM received the B.E. degree in Biomedical Engineering in 2005 and the M.E. degree in Industrial Engineering (Biomedical) in 2007 from Politecnico di Milano, Milan, Italy. He is currently a lecturer in Digital Systems at The Australian National University, Canberra, Australia, where he has recently completed a Ph.D. degree with a focus on drug delivery, robust adaptive control and cardiovascular modelling.

Acknowledgements

The Author would like to thank Dr. Jochen Trumpf for his precious advice in structuring this paper.

Received: 1 August 2013 Accepted: 3 February 2014

Published: 25 March 2014

References

1. Cosgrove DM, Petre JH, Waller JL, Roth JV, Shepherd C, Cohn LH: **Automated control of postoperative hypertension: a prospective, randomized multicenter study.** *Ann Thorac Surg* 1989, **47**(5):678–683.
2. Bailey JM, Haddad WM: **Drug dosing control in clinical pharmacology.** *IEEE Control Syst Mag* 2005, **25**(2):35–51.
3. Parker RS, Doyle FJ: **Control-relevant modeling in drug delivery.** *Adv Drug Deliv Rev* 2001, **48**:211–228.
4. Haddad WM, Bailey JM: **Closed-loop control for intensive care unit sedation.** *Best Pract Res Clin Anaesthesiol* 2009, **23**:95–114.
5. Mendonca T, Lemos J, Magalhaes H, Rocha P, Esteves S: **Drug delivery for neuromuscular blockade with supervised multimodel adaptive control.** *IEEE Trans Control Syst Technol* 2009, **17**(6):1237–1244.
6. Doyle FJ, Bequette BW, Middleton R, Ogunnaike B, Paden B, Parker RS, Vidyasagar M: **Control in biological systems.** In *The Impact of Control Technology*. Edited by Samad T, Annaswamy AM. Piscataway: IEEE Control Systems Society; 2011.
7. Varon J, Marik PE: **Perioperative hypertension management.** *J Vasc Health Risk Manage* 2008, **4**(3):615–627.
8. Slate JB, Sheppard LC, Rideout VC, Blackstone EH: **A model for design of a blood pressure controller for hypertensive patients.** In *Proceedings of the 1st Annual Conference, IEEE Engineering in Medicine and Biology Society, Denver, Colorado*. Piscataway: IEEE; 1979.
9. Mansour NE, Linkens DA: **Pole-assignment self-tuning control of blood pressure in postoperative patients: a simulation study.** *IEEE Proc Control Theory Appl* 1989, **136**:1–11.
10. Pajunen GA, Steinmetz M, Shankar R: **Model reference adaptive control with constraints for postoperative blood pressure management.** *IEEE Trans Biomed Eng* 1990, **37**(7):679–687.
11. He WG, Kaufman H, Roy R: **Multiple model adaptive control procedure for blood pressure control.** *IEEE Trans Biomed Eng* 1986, **33**:10–19.
12. Martin JF, Schneider AM, Smith NT: **Multiple-model adaptive control of blood pressure using sodium nitroprusside.** *IEEE Trans Biomed Eng* 1987, **34**(8):603–611.
13. Ying H, McEachern M, Eddleman DW, Sheppard LC: **Fuzzy control of mean arterial pressure in postsurgical patients with sodium nitroprusside infusion.** *IEEE Trans Biomed Eng* 1992, **39**(10):1060–1070.

14. Yu C, Roy RJ, Kaufman H, Bequette BW: **Multiple-model adaptive predictive control of mean arterial pressure and cardiac output.** *IEEE Trans Biomed Eng* 1992, **39**(8):765–778.
15. Yu YC: **Blood pressure, automatic control of.** In *Encyclopedia of Medical Devices and Instrumentation*. 2nd edition. Edited by Webster JG. New York: Wiley; 2006:490–500.
16. Chitwood WR, Cosgrove DM, Lust RM: **Multicenter trial of automated nitroprusside infusion for postoperative hypertension.** *Ann Thorac Surg* 1992, **54**(3):517–522.
17. Bequette BW: **Modeling and control of drug infusion in critical care.** *J Process Control* 2007, **17**(7):582–586.
18. Malagutti N, Dehghani A, Kennedy RA: **Improved robust performance in a system for automatic administration of a vasoactive drug.** In *Proceedings of BIOSTEC Biosignals Conference, Vilamoura, Portugal*. Montreal: SciTePress; 2012:282–290.
19. Malagutti N, Dehghani A, Kennedy RA: **Robust control design for automatic regulation of blood pressure.** *IET Control Theory Appl* 2013, **7**(9):387–396.
20. Fekri S, Athans M, Pascoal A: **Issues, progress and new results in robust adaptive control.** *Int J Adaptive Control Signal Process* 2006, **20**(10):519–579.
21. Malagutti N, Hassani V, Dehghani A: **An adaptive multi-controller architecture using particle filtering.** In *Proceedings of the 2012 Australian Control Conference, Sydney, Australia*. Barton: Engineers Australia; 2012:392–398.
22. Bonate P: *Pharmacokinetic-Pharmacodynamic Modeling and Simulation*. Berlin: Springer; 2011.
23. Malagutti N, Dehghani A, Kennedy RA: **A robust multiple-model adaptive control system for automatic delivery of a vasoactive drug.** In *Proceedings of 9th IASTED Conference on Biomedical Engineering*. Innsbruck, Austria, 15–17 February 2012:74–79.
24. Skogestad S, Postlethwaite I: **Robust stability and performance analysis for MIMO systems.** In *Multivariable Feedback Control, Analysis and Design*. New York: Wiley; 2006:309–368.
25. Cappe O, Godsill SJ, Moulines E: **An overview of existing methods and recent advances in sequential Monte Carlo.** *Proc IEEE* 2007, **95**(5):899–924.
26. Schon T, Gustafsson F, Nordlund PJ: **Marginalized particle filters for mixed linear/nonlinear state-space models.** *IEEE Trans Signal Process* 2005, **53**(7):2279–2289.
27. Doucet A, de Freitas N, Gordon N: **An introduction to sequential Monte Carlo methods.** In *Sequential Monte Carlo Methods in Practice*. New York: Springer; 2001:3–13.
28. Douc R, Cappe O: **Comparison of resampling schemes for particle filtering.** In *Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis: Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia*. 2005:64–69.
29. Doucet A, Johansen AM: *A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later, Volume 12*. Oxford: Oxford University Press; 2009.
30. Ioannou PA: **Robust adaptive control: the search for the Holy Grail.** In *Proceedings of the 47th IEEE Conference on Decision and Control*. Piscataway: IEEE; 2008:12–13.
31. Hassani V, Hespanha J, Athans M, Pascoal AM: **Stability analysis of robust multiple model adaptive control.** In *Proceedings of The 18th IFAC World Congress*. Milan: IFAC; 2011.
32. Meijers RH, Schmartz D, Cantraine FR, Barvais L, d'Hollander AA, Blom JA: **Clinical evaluation of an automatic blood pressure controller during cardiac surgery.** *J Clin Monit Comput* 1997, **13**(4):261–268.
33. Furutani E, Araki M, Kan S, Aung T, Onodera H, Imamura M, Shirakami G, Maetani S: **An automatic control system of the blood pressure of patients under surgical operation.** *Int J Control, Automation, Syst* 2004, **2**:39–54.
34. Hoeksels SA, Blom JA, Jansen JR, Maessen JG, Schreuder JJ: **Automated infusion of vasoactive and inotropic drugs to control arterial and pulmonary pressures during cardiac surgery.** *Crit Care Med* 1999, **27**(12):2792–2798.
35. Zenker S, Rubin J, Clermont G: **From inverse problems in mathematical physiology to quantitative differential diagnoses.** *PLoS Comput Biol* 2007, **3**(11):e204.
36. Forssell U, Forssell U, Ljung L, Ljung L: **Closed-loop identification revisited.** *Automatica* 1997, **35**:1215–1241.

doi:10.1186/2194-3990-1-8

Cite this article as: Malagutti: Particle filter-based robust adaptive control for closed-loop administration of sodium nitroprusside. *Journal of Computational Surgery* 2014 **1**:8.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com