34th Annual International Conference of the IEEE EMBS
San Diego, California USA, 28 August - 1 September, 2012

An approach to controlled drug infusion via tracking of the time-varying dose-response

Nicolò Malagutti[†,*], Arvin Dehghani[‡] and Rodney A. Kennedy[†]

***Abstract*— Automatic administration of medicinal drugs has the potential of delivering benefits over manual practices in terms of reduced costs and improved patient outcomes. Safe and successful substitution of a human operator with a computer algorithm relies, however, on the robustness of the control methodology, the design of which depends, in turn, on available knowledge about the underlying dose-response model. Real-time estimation of a patient's actual response would ensure that the most suitable control algorithm is adopted, but the potentially time-varying nature of model parameters and the limited number of observation signals may cause the estimation problem to be ill-posed, posing a challenge to adaptive control methods. We propose the use of Bayesian inference through a particle filtering approach as a way to overcome these limitations and improve the robustness of automatic drug administration methods. We report on the results of a simulation study modeling the infusion of vasodepressor drug sodium nitroprusside for the control of mean arterial pressure in acute hypertensive patients. The proposed control architecture was able to meet the required performance objectives under challenging operating conditions.**

I. INTRODUCTION

Dynamical systems modeling the dose-response relationship of medicinal drugs can be characterized by a high level of uncertainty and variability in the values of the model parameters (observed both across the patient population, interpatient, and for an individual patient over time, intra-patient) [1]. In current clinical practice, operators administer an initial drug amount, generally according to population statistics, followed by close monitoring and periodic manual adjustment of the dose on the basis of the actual observed clinical outcome. Such an approach is staff-intensive and prone to human error. Clinical studies have demonstrated the potential of delivering superior outcomes when automatic feedback control of administration has been implemented [2]. Ultimate safety of such automatic approaches, however, rests on the robustness of the control methodologies adopted. This remains a challenge [3], since large uncertainty and time-variability demand the use of adaptive control [4], while the availability of only a small number of noisy observation signals may limit the ability to accurately estimate the patient's response and inform adaptation [5].

We consider the specific application of controlled infusion of sodium nitroprusside (SNP). Several approaches for controlled SNP infusion have been proposed in the past (see [6] for a review). SNP is a fast-acting vasodepressor which is administered continuously and intravenously to control mean arterial pressure (MAP) during acute hypertensive episodes, generally exhibited by postoperative patients in critical care settings [7]. It should be acknowledged that SNP is no longer the most commonly used drug for this purpose and that other active ingredients with lower toxicity have been developed in recent years [7]. Nonetheless, SNP is an interesting case study from an engineering perspective as its dose-response relationship has been extensively studied [8] and can be modeled as a single-input-single-output linear parameter-varying system with input time delay and an output offset [9].

We propose that a stochastic approach to patient response estimation could lead to better robustness in automatic drug delivery methods. Particle filtering is a sequential Monte-Carlo method which computes an approximation of the probability distribution of the state of a partially observed dynamical system through Bayesian inference [10], i.e., by iteratively refining an initial distribution as new observations become available. Particle filtering has found biomedical engineering applications in signal processing and, recently, it has been proposed as a tool to assist with differential diagnosis [11]. To the authors' knowledge, its use in supporting control decisions in drug delivery is novel. By including any time-varying parameters among the states of the dose-response model, we seek to estimate and track a patient's actual behaviour as it evolves through time. The statistical features of the probability distribution are then used to inform robust control decisions. The conditionally linear, relatively low-order model of SNP dose-response renders the implementation of a particle filter less computationally taxing and is therefore well suited to investigating the potential of the proposed approach.

The paper is structured as follows: Section II provides a brief description of the problem and the control approach; Section III describes and presents the results of a case study simulation; Section IV discusses the results and outlines possible directions for future work.

II. MATERIALS AND METHODS

A. The dose-response model

The dose-response for SNP administration is shown in Fig. 1 [9]. This translates into a state-space formulation for the measured MAP output y_{meas} as a function of SNP infusion rate u given by

$$\begin{aligned} \dot{x} &= Ax(t) + Bu(t - T(t)) + Lv(t) \\ y_{meas} &= p_0(t = 0) - Cx(t) + w(t), \end{aligned} \tag{1}$$

with

[†]N. Malagutti and R. A. Kennedy are with the Research School of Engineering, The Australian National University, Canberra, Australia.
[‡]A. Dehghani was formerly with the Department of Electrical and Electronic Engineering, The University of Melbourne, Melbourne, Australia.
[*]corresponding author nicolo.malagutti@anu.edu.au

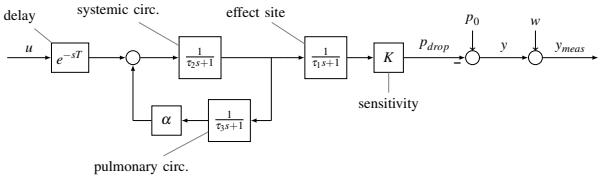

Fig. 1. Model of patient response to SNP [9]. Notation: T is the pure delay parameter; $\tau_1 = 50s$, $\tau_2 = 30s$, $\tau_3 = 10s$ are the time constants of the first-order LTI subsystems; α is the recirculation parameter; K is patient sensitivity; p_0 is the patient's natural MAP; w is output measurement noise.

$$A = \begin{bmatrix} -\left(\frac{1}{\tau_1}+\frac{1}{\tau_2}+\frac{1}{\tau_3}\right) & -\left(\frac{1}{\tau_1\tau_2}+\frac{1-\alpha(t)}{\tau_2\tau_3}+\frac{1}{\tau_1\tau_3}\right) & \frac{\alpha(t)-1}{\tau_1\tau_2\tau_3} & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -0.01 \end{bmatrix} \quad (2a)$$

$$B = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad L = \begin{bmatrix} 0.0625 \\ 0 \\ 0 \\ 1 \end{bmatrix} \quad C = \begin{bmatrix} 0 & \frac{K(t)}{\tau_1\tau_2} & \frac{K(t)}{\tau_1\tau_2\tau_3} & 1 \end{bmatrix}, \quad (2b)$$

where p_0 is a patient's natural MAP in the absence of SNP infusion, and $v(t) \sim N(0,1)$ and $w(t) \sim N(0,2)$ are random noise signals at the input and the output, respectively.

The a-priori-unknown and potentially time-varying parameters for the model are (rate of change per hour hr^{-1}):

$$\begin{aligned} 0.25 < K(t) < 9.5, \quad & \left|\frac{dK}{dt}\right| < 3K(t)hr^{-1} \\ 0.25 < \alpha(t) < 0.75, \quad & \left|\frac{d\alpha}{dt}\right| < 0.5hr^{-1} \qquad \forall t. \\ 10 < T(t) < 50, \quad & \left|\frac{dT}{dt}\right| < 40hr^{-1} \end{aligned} \quad (3)$$

Finally, the offset term $p_0(t)$, whose value can be measured prior to drug infusion ($t = 0$), is also considered to be potentially time-varying for generality. A random, mainly low-frequency behaviour is assumed. To this end,

$$\begin{aligned} p_0(t) &= p_0(t=0) + p_{dist}(t) \\ p_{dist}(t) &: P_{dist}(s) \leq \tfrac{1}{s+0.01}. \end{aligned} \quad (4)$$

Remark: The minimal assumptions adopted make this the most general model formulation ever considered for control design in this application.

It is evident that in light of the presence of a single observation signal $y_{meas}(t)$ and up to four time-varying unknown parameters, the estimation of this dynamical system is an underdetermined problem.

B. Closed-loop control

1) Performance requirements [12]:

- a settling time of 10min or less;
- a maximum overshoot of 10mmHg during transients;
- during steady state operation, MAP should be contained within ±5mmHg of the desired set-point value;
- under no circumstances should the system display resonant (persistent oscillatory) or unstable behaviour or cause MAP to drop below 60mmHg.

2) Controllers: The proposed feedback control architecture is shown in Fig. 2. A bank of 5 candidate controllers designed using the robust controller design technique of μ synthesis ensures that a stable patient-controller pair capable of delivering the required level of performance exists for all values of the time-varying parameters. The controller design process is described in detail in [4]. Table I summarizes the correspondence between the controllers and the respective subsets of the parameter uncertainty space.

3) Particle filtering: In order to estimate the model parameters (particularly patient sensitivity K) so that the correct controller can be placed in the feedback loop, the estimation problem is recast as a nonlinear tracking problem by including the uncertain parameters in the state, as shown in (5) below. Since the system can be described as linear (with "linear" states x^l) conditionally on the "nonlinear" states x^n, it is possible to use a specialized form of particle filtering called *marginalized particle filtering* [13].

The discrete-time form of the system with augmented state vector using sampling time $Ts = 2$s (one heart beat) is

$$\begin{aligned} x^l(k+1) &= A_d(x^n(k))x^l(k) + B_d u(k - \lfloor \tfrac{T(x^n(k))}{T_s} \rfloor) + w(k) \\ x^n(k+1) &= f(x^n(k)) \\ y(k) &= p_0 - C_d(x^n(k))x^l(k) + v(k), \end{aligned} \quad (5)$$

where $k = \lfloor \frac{t}{T_s} \rfloor$ ($\lfloor \cdot \rfloor$ is the *floor* operator). The subscript d indicates zero-order hold discretisation of (1); i.e.,

$$A_d = e^{AT_s}, \quad B_d = A^{-1}\left(A_d - I\right)B, \quad C_d = C \ . \quad (6)$$

The time update function of the nonlinear states $f(x^n)$ is

$$x^n(k+1) = \begin{bmatrix} K(k+1) \\ T(k+1) \\ \alpha(k+1) \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} x(k) + \chi(k), \quad (7)$$

where $\chi(k)$ is sampled from an array of probability distributions which capture the likely trajectory of these parameters

$$\chi \sim \begin{bmatrix} U(-0.017K(k), 0.017K(k)) \\ U(-0.028, 0.028) \\ U(-0.00028, 0.00028) \end{bmatrix} \quad (8)$$

Uniform distributions (U) are used to capture constraints on the rate of change of parameters expressed in (3) while making no assumptions on possible trends.

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. $p(x(k)|Y(0:k))$, where $Y(0:k) \equiv \{y(i)\}_{i=0}^{k}$. An analytical solution to calculate this function exists for linear Gaussian systems, resulting in the Kalman filter approach. For more general cases, numerical methods must be used. Particle filters are one such method, which approximates the posterior probability with a finite number

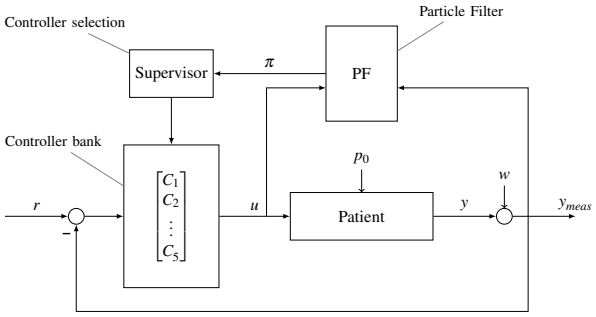

Fig. 2. Control architecture. Notation: r reference signal (desired MAP value); $C_{1,\ldots,5}$ candidate controllers; u control signal (drug infusion rate); p_0 patient's natural MAP; y output MAP; w measurement noise; y_{meas} measured MAP; $\pi = \{\pi_j\}_{j=1,\ldots,5}$ probability of the estimated model parameters belonging to the uncertainty subset for which robust controller C_j has been designed.

of samples (particles). Particles can be interpreted as realisations of the system to be estimated. At every time step, each particle updates its state according to the dynamics of the realisation it represents, and produces an estimated output. This is compared with the actual observation and a weight for the particle is generated. Particles with poorer weights have a higher chance of being discarded during a resampling step and reintroduced as a copy of a better performing realisation. Over a number of iterations, particles cluster in the state space in a way that approximates the posterior probability density of the state. As the number of particles increases, so does the accuracy of the approximation but also the computational burden. Also, models with more states require a larger number of particles to be approximated.

Marginalized particle filters exploit the fact that a subset of the states can be treated as conditionally linear and can be estimated using the optimal Kalman filter result (marginalisation), while the other (nonlinear) states are estimated by the particle filter. A formal description of the method can be found in [10]. Only an algorithmic description will be given here. Due to the lower dimension of the numerical approximation problem, marginalized particle filters have a lower computational cost than a standard particle filter.

Marginalized Particle Filter algorithm

(a) *Initialisation*. State $x_i(0)$ for particle $i = 1,\ldots,N$ 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 equal 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 $i = 1,\ldots,N$ 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' *normalized 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 [14].

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

(i) Kalman filter correction of the linear state estimate using the available observation $y(k)$
$x_i^l(k|k) = x_i^l(k|k-1) + H_i(k)(y(k) - p_0 - C_{d,i}(k)x_i^l(k|k-1))$,
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 (8) and update the nonlinear states via (7), then calculate $A_{d,i}(x_i^n(k))(k)$ and $B_{d,i}(x_i^n(k))(k)$ using (6).

(iii) Time update of the marginalized states $(x_i^l(k+1|k))$ via (5) and the state estimate covariance matrix using $P_i(k+1|k) = A_{d,i}(k)P_i(k|k)A_{d,i}^T(k) + Q$, where Q is the covariance matrix of the state/input noise v.

(e) *Iteration*. Increase $k \to k+1$ and repeat over from Step (b).

4) 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 subsets listed in Table I 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

$$u = \sum_{j=1}^{5} \pi_j u_j \quad \pi_j = \frac{n_j}{N}, \tag{9}$$

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

TABLE I

LIST OF CONTROLLERS AND CORRESPONDING PARAMETER RANGES

Controller number j	To suit K (mmHg/(ml/hr))	To suit α	To suit delay T (s)
1	0.25 – 0.71	0.25 – 0.75	0 – 50
2	0.71 – 1.63	0.25 – 0.75	0 – 50
3	1.63 – 3.48	0.25 – 0.75	0 – 50
4	3.48 – 6.72	0.25 – 0.75	0 – 50
5	6.72 – 9.50	0.25 – 0.75	0 – 50

III. SIMULATION EXAMPLE AND RESULTS

We present a simulation case study in which the proposed approach is to control MAP in a virtual hypertensive patient with a baseline of 120mmHg. The clinical goal (reference) is to lower this first to 100mmHg and then to 80mmHg (at $t = 4000$s). The total control period is 10000s. The control task is challenging, with all patient parameters set to be time-varying, including K. Large variations in the offset $p_0(t)$ occur, with random fluctuations in the range $\pm$20mmHg, and step increases of 10mmHg (low-pass filtered) introduced at times $t = 3500$s, $t = 6100$s and $t = 9500$s simulating a worsening hypertensive condition. It is expected that in such highly-varying conditions the system should be able to maintain MAP within $\pm$10mmHg of the required set point.

The simulation was implemented in Matlab and Simulink. We evaluated the adherence of the controlled MAP to the specifications and the quality of the filter estimate in terms of particle scatter within the space of uncertain parameters.

The results are presented in Fig. 3. The system is able to provide the required level of performance, with MAP converging to the reference setpoints in under 10 minutes (600s) and remaining within the allowed error range for over 95% of the time; no persistent oscillations or dangerous downward spikes were observed. Tracking of K was accurate, resulting in correct controller allocation. Larger grey areas in the density maps for T and α, however, indicate that under the challenging simulation conditions chosen the values of these two parameters could not be accurately resolved.

The processing time for the 10000s simulation was 45min (2700s, approximately 3.7$\times$ faster than real time) on a standard desktop computer (Intel®Core™2 Quad CPU 3.00GHz) using 3000 particles. This number of particles was deemed to be a reasonable compromise between accuracy in the results and computational time. The same simulation did not deliver noticeable improvements in the estimation results when run using 5000 particles.

IV. DISCUSSIONS AND FUTURE WORK

We have introduced a new approach for the control of highly uncertain, time-varying systems, and applied it to the drug delivery example of controlled SNP administration. The

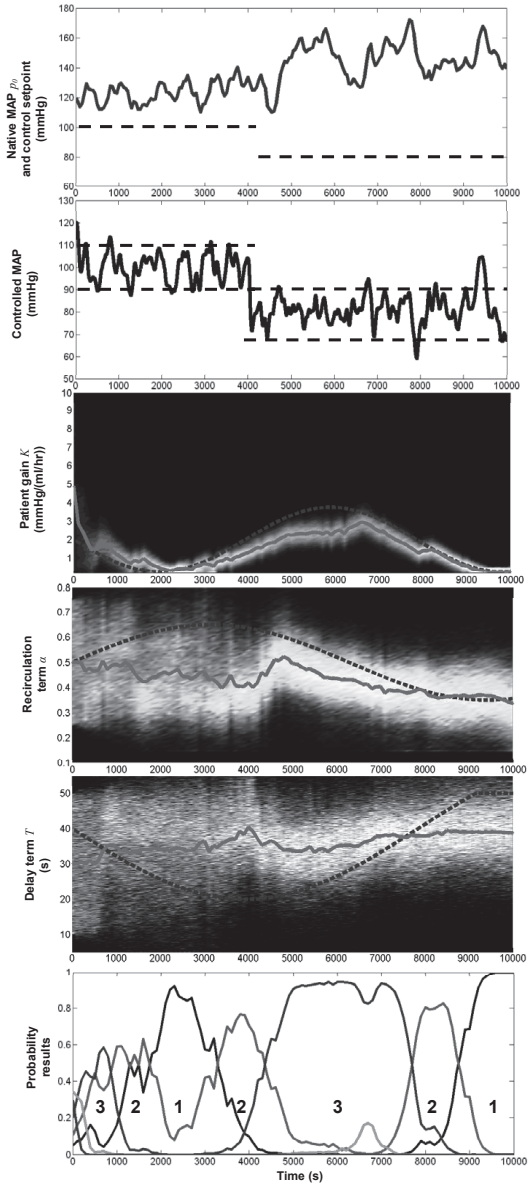

Fig. 3. Results of simulation. The dashed lines in the Controlled MAP plot show the $\pm$10mmHg allowed error range. The plots of the time-varying parameters K, α and T feature the real value (dashed line) and the particle filter estimate (mean of the distribution, solid line); the background of the plots are density maps showing the distribution of the particles in the uncertainty space (lighter areas indicate greater density). The numbers in the controller probability plot indicate the controller with the greatest probability.

proposed architecture features multiple controllers and uses the results of a particle filter for controller selection. Under challenging simulation conditions, the approach delivered promising results which are comparable with those obtained in our earlier work with multiple-model adaptive control [12] and arguably superior to those of past control architectures which did not consider the challenge of a time-varying output offset in their modeling. Future work will involve thorough experimentation to confirm these positive results.

The fact that the computational burden associated with particle filtering was compatible with real-time operation without requiring special computational hardware, albeit for the simple model chosen here, is also a notable positive finding. We plan to examine other case studies, including nonlinear and multiple-input-multiple-output systems, to identify any limitations to the practical use of this technique in addressing more complex problems.

Finally, we have shown that particle filtering can give an indication as to what extent the estimation of an underdetermined system can be carried out successfully through the resulting approximate probability distribution (in this example, K was a sufficiently sensitive parameter to be estimated, unlike T and α). While stability was not affected here, from a general perspective the distribution may be used to develop alternative supervisory methods which reduce the risk of pairing the patient with a destabilizing controller below an acceptable threshold (e.g., by locating a specific percentile). We intend to investigate whether this concept can contribute to the design of safer and more dependable automatic drug delivery solutions.

REFERENCES

[1] J. M. Bailey and W. M. Haddad, Drug dosing control in clinical pharmacology, IEEE Contr. Syst. Mag., vol. 25(2), pp. 35–51, 2005.

[2] R. A. de Asla, A. M. Benis, R. A. Jurado, and R. S. Litwak, Management of postcardiotomy hypertension by microcomputer-controlled administration of sodium nitroprusside, J. Thorac. Cardiovasc. Surg., vol. 89, no. 1, pp. 115–20, 1985.

[3] F. J. Doyle, B. W. Bequette, R. Middleton, B. Ogunnaike, B. Paden, R. S. Parker, and M. Vidyasagar, Control in biological systems, in T. Samad and A. M. Annaswamy, eds., The Impact of Control Technology, IEEE Control Systems Society, 2011.

[4] N. Malagutti, A. Dehghani, and R. A. Kennedy, Improved robust performance in a system for automatic administration of a vasoactive drug, in Proc. of BIOSTEC Biosignals Conf., pp. 282–290, Vilamoura, Portugal, 2012.

[5] N. Malagutti, A. Dehghani, and R. A. Kennedy, Safety issues in adaptive control systems for automatic administration of vasoactive drugs, in Proc. of 18th IFAC World Congress, Milan, Italy, 2011.

[6] Y. C. Yu, Blood pressure, automatic control of, in J. G. Webster, ed., Encyclopedia of Medical Devices and Instrumentation, 2nd ed., pp. 490–500, John Wiley & Sons, Inc., 2006.

[7] J. Varon and P. E. Marik, Perioperative hypertension management, Vasc. Health. Risk. Manag., vol. 4, no. 3, pp. 615–627, 2008.

[8] 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 Proc. of the 1st Annu. Conf., IEEE Engineering in Medicine and Biology Society, Denver, Colorado, 1979.

[9] J. F. Martin, A. M. Schneider, and N. T. Smith, Multiple-model adaptive control of blood pressure using sodium nitroprusside, IEEE Trans. Biomed. Eng., vol. 34, no. 8, pp. 603–611, 1987.

[10] O. Cappe, S. J. Godsill, and E. Moulines, An overview of existing methods and recent advances in sequential monte carlo, Proc. IEEE, vol. 95, no. 5, pp. 899–924, 2007.

[11] S. Zenker, J. Rubin, and G. Clermont, From inverse problems in mathematical physiology to quantitative differential diagnoses, PLoS Comput. Biol., vol. 3, no. 11, p. e204, 2007.

[12] N. Malagutti, A. Dehghani, and R. A. Kennedy, A robust multiple-model adaptive control system for automatic delivery of a vasoactive drug, in Proc. of 9th IASTED Conference on Biomedical Engineering, pp. 74–79, Innsbruck, Austria, 2012.

[13] T. Schon, F. Gustafsson, and P.-J. Nordlund, Marginalized particle filters for mixed linear/nonlinear state-space models, IEEE Trans. Sig. Proc., vol. 53, no. 7, pp. 2279–2289, 2005.

[14] R. Douc and O. Cappe, Comparison of resampling schemes for particle filtering, in Proc. of the 4th Int. Symposium on Image and Signal Processing and Analysis, pp. 64–69, Zagreb, Croatia, 2005.