AIChE WILEY InterScience DISCOVER SOMETHING GREAT

R&D NOTES

Model Predictive Control for Optimal Oral Anticoagulant Drug Administration

Gabriele Pannocchia and Alessandro Brambilla
Dept. of Chemical Engineering, Industrial Chemistry and Source of Materials,
University of Pisa, Via Diotisalvi, 2 - 56126 Pisa, Italy

DOI 10.1002/aic.10930
Published online July 13, 2006 in Wiley InterScience (www.interscience.wiley.com).

Keywords: oral anticoagulant therapy, model predictive control, identification, optimization

Introduction

Blood must remain fluid within the vasculature and yet clot quickly when exposed to nonendothelial surfaces at sites of a vascular injury. Thrombosis is a pathological process in which a platelet aggregate and/or a fibrin clot occludes a blood vessel. In the normal situation when intravascular thrombi occur, a system of fibrinolysis is activated to restore fluidity. However, in many cases (for example, after heart attacks, installation of heart valves, atrial fibrillation, and so on) patients need to be treated with oral anticoagulant drugs, even life-long. Optimal anticoagulant drug administration is particularly relevant because over-treated patients can experience severe hemorrhagic events, while under-treated patients are at risk of dangerous ischemic necrosis (for example, myocardial infarction due to thrombosis of a coronary artery) or edemas in case of venous thrombosis.

Anticoagulant drugs, as well as many other ones are administrated by physicians using heuristic criteria and practical experience. Nonetheless, in the last decade a significant number of research studies tried to apply automatic control techniques, such as model predictive control, to drug administration in several medical areas, such as anesthesia and diabetes (see, for example, [1-5] and references therein). In most of these applications the drug(s) is administrated continuously (for example, via an external pump or a subcutaneous infuser), and the patient's key parameters (for example, blood pressure, glycemia, and so on) are also measured continuously (that is, at a sampling time equal to or shorter than that used to adjust the drug rate). A distinguishable feature of this application is that the oral anticoagulant drugs are administrated daily, whereas the patient's relevant parameters are measured at a much lower frequency (3-4 weeks or so) and often irregularly. Moreover, due to the specific response of patients to drugs any effective control strategy requires some adaptation steps, such as the development or the identification from data of a personalized model for the patient's key variables response to drugs. In this perspective the patient is seen as a "process" to be controlled where the patient's key parameters are the controlled variables and the drugs are the manipulated ones. Other typical issues arise from the presence of interactions with other drugs which may accelerate or inhibit the effect of the drug administrated, and from the effects of diet and style of life.[6,7] From a control point of view, these can be considered as disturbance variables, which may or may not be measured. In the former case, when a quantitative effect of the interacting drugs, diet and style of life is known, some feed-forward compensation can be implemented.

Model predictive control (MPC) is a control technology widely used in many areas, especially in the process industries.[8] MPC algorithms typically achieve superior performance with respect to other control strategies when manipulated and controlled variables have constraints to meet (possibly with different priority). Moreover, MPC algorithms can efficiently handle multivariable systems, and do not require that the number of controlled and manipulated variables be equal. Thus, they use all manipulated variables to achieve optimization goals, such as: (1) keeping the key variables close to their targets, (2) maintaining all controlled variables within limits, and (3) minimizing the control effort while respecting constraints on the manipulated variables. In order to achieve such objectives, MPC algorithms use a process model, either based on first-principal equations (mass/energy balances, equilibrium equations,. . .), or derived from data. The model is used to compute steady-state targets for manipulated and controlled variables and an "optimal" sequence of future control actions that respect the process constraints and bring the controlled variables to their targets. Feedback from measurements is in-

Correspondence concerning this article should be addressed to G. Pannocchia at g.pannocchia@ing.unipi.it.

corporated by updating the model prediction with a correction term that takes into account the difference between the predicted and the actual measured outputs.[8-10] In this article, a model predictive control algorithm has been adapted to implement a personalized decision support system for oral anticoagulant therapy, characterized by infrequent and irregular measurements.

Oral anticoagulant therapy

Oral anticoagulant drugs, such as Warfarin or coumarin derivatives, are antagonists of vitamin K.[11] Coagulation factors and anticoagulant proteins are synthesized mainly in the liver, and are biologically inactive unless 9 to 12 of the aminoterminal glutamic acid residues are carboxylated. This reaction requires carbon dioxide, molecular oxygen and reduced vitamin K, which must be regenerated from the vitamin K epoxide. Oral anticoagulants block the enzymatic reaction of vitamin K epoxide.

The most widely used oral anticoagulant drug in Italy and in many other countries (for example, USA, Canada) is Warfarin* (commercial name "Coumadin"). The patient's key parameter on which anticoagulant drug administration is based is the Prothrombin Time (PT), which represents the time required for coagulation *in vitro* of blood plasma after addition of "thromboplastin" (a saline extract of brain that contains tissue factor and phospholipids). It is important to point out that recently an International Normalized Ratio (INR) system of reporting has been adopted in order to standardize testings between different laboratories.[12] Conceptually, INR is the ratio between the patient's PT to a control PT that would have been reported by a standard method using a WHO primary standard (human) thromboplastin. PT measurements are converted to INR measurements by the following equation

$$\mathrm{INR} = \left(\frac{\mathrm{PT}_{\mathrm{pt}}}{\mathrm{PT}_{\mathrm{ref}}}\right)^{\mathrm{ISI}} \tag{1}$$

in which ISI is the International Sensitivity Index (which depends on the thromboplastin used in laboratory). By definition, INR of a patient not treated with oral anticoagulant drugs is around 1, while the therapy typically aims at maintaining INR between 2.0 and 3.0 (in same cases, such as patients with prosthetic heart valves the suggested therapeutic range is 2.5 to 3.5). INR is typically measured on "new" patients every 2-3 days, and then measured every 3-4 weeks, and the daily dose is adjusted consequently. The effectiveness and safety of warfarin therapy are critically dependent on a dose-response relationship, which is extremely individualized and can be influenced by a number of factors.[13] Therefore, patients on oral anticoagulant therapy require a close and careful monitoring of their daily dosage to prevent risks associated to underdosing or overdosing. To assist the control of outpatients on warfarin therapy, a number of computer-based dosage programs have been developed (see for example,[14-16] and references therein). Many of these computer decision support systems (CDSS) are expert systems, which implement a number of (deterministic or probabilistic) rules to adjust the warfarin dosage, whereas other ones, use conventional feedback algorithms.[17]

*Warfarin, whose name is ancronym derived from the name of the patent holder (Wisconsin Alumni Research Foundation) plus the coum*arin*-derived suffix, is a synthetic compound initially produced as a rodenticide.

Method

The objective of this research activity is to propose a predictive-control algorithm that, given the current measurement of INR, past INR measurements and warfarin doses administrated, computes the future warfarin daily dose which keeps INR within the therapeutic limits, and close to the desired target. A key feature of the proposed method is the identification of a personalized response model of the patient's INR to warfarin. In industrial processes a model is obtained from data collected during specific tests traditionally conducted in an "open-loop" fashion. It is clear that the same approach cannot be applied to oral anticoagulant therapy, where it is necessary to bring (and maintain) the patient's INR inside the therapeutic range as quickly as possible. From a control point of view it is important to remark that the sampling time is one day, because warfarin is consumed daily (usually in the afternoon). However, INR is not measured daily (except for inpatients) and its sampling time is usually not regular (3-4 weeks or so). This will raise some issues related to model identification and state estimation, which are discussed in this section. It is also important to clarify that, in some countries (Italy is one of them), warfarin is available in 5 mg pills, which can be divided in quarters, then the daily dose must be a multiple of $u^* = 1.25$ mg.

Patient's response dynamic model: identification and adaptation

In order to develop an MPC algorithm, a dynamic model for the response of the patient's INR to warfarin is required. Studies on pharmacokinetic and pharmacodynamic of warfarin (and other oral anticoagulants) [18-20] reported a stable nonoscillating behavior. From the analysis of retrospective data, it was found that this response can be satisfactorily modeled as a second-order critically damped system, that is, in the Laplace domain

$$y(s) = G(s)u(s) = \frac{K}{(\tau s + 1)^2}\, u(s) \tag{2}$$

in which u is the warfarin dose (input), y is the patient's INR (output, in deviation from the basal value). Given data of the patient's daily doses administrated and of the measured INR, the model parameters, $\theta = [K,\tau]$, can be estimated using an output error criterion, that is, by minimizing the loss function V_{oe} defined as [21]

$$V_{oe} = \frac{1}{N_y} \sum_{k_y=1}^{N_y} (y_{k_y} - \tilde{G}(q|\theta)u_{k_y})^2 \tag{3}$$

in which $\tilde{G}(q|\theta)$ is the discrete transfer function corresponding to (Eq. 2) k_y) are the time-indices at which the output was measured and N_y) is the number of available measurements. It is important to remark that the minimization of the output error-loss function (unlike, for example, the minimization of the prediction error-loss function for the identification of an

ARX model) is a nonlinear optimization problem, for which no analytical solution is available. However, the choice of an output error criterion is unavoidable due to the fact the output is not sampled daily and regularly.

Using retrospective data of 61 patients, a model for each patient was identified and, after averaging the model parameters, a "preliminary" model was defined with the following parameters: $K = 0.3875\ \text{mg}^{-1}$ (STD: $0.1895\ \text{mg}^{-1}$), $\tau = 65.9$ hr (STD: 24.4 hr). This model is initially used in the MPC algorithm described later in this section. However, due to the significant variability among different patients, the model is updated during the first period of therapy (using at least 3-4 INR measurements) by re-identifying the parameters (K,τ) each time a new INR measurement is available. It is important to point out that the re-identification of the model parameters is carried out in the presence of bound constraints in order to avoid excessive variations of the model parameters, and, hence, to improve the robustness of the control algorithm. After this period, the model is no longer updated and any subsequent mismatch between the actual and predicted output is regarded as a (lumped) disturbance, as described in the next paragraph. Other possible adaptation strategies could be based on a combined state-parameter-disturbance estimation to avoid lumping early disturbances into the model parameters, but were not considered in this application given the limited number of INR measurements available.

The proposed algorithm is based on three main modules: a state and disturbance estimator, a steady-state optimizer and a dynamic optimizer. These modules, described in the next paragraphs, are executed each time an INR measurement is available to update the warfarin dose protocol.

State and disturbance estimation

The algorithm uses a linear discrete time-invariant model in state-space form:

$$\begin{aligned} x_{k+1} &= Ax_k + B(u_k + d_k) \\ d_{k+1} &= d_k \\ y_k &= Cx_k \end{aligned} \tag{4}$$

in which $x \in \mathbb{R}^n$ is the state, and d is a "fictitious" integrating disturbance included to achieve offset-free control.[10] The matrices (A,B,C) are obtained by minimal realization and discretization of Eq. 2 (hence, $n = 2$), and, as discussed, may change during the first period of therapy to match the actual patient's response. Since x_k) and d_k) are not measured, they are estimated from the INR measurement by means of a steady-state Kalman filter

$$\begin{aligned} \hat{x}_{k|k} &= \hat{x}_{k|j} + L_x(y_k - C\hat{x}_{k|j}) \\ \hat{d}_{k|k} &= \hat{d}_{k|j} + L_d(y_k - C\hat{x}_{k|j}) \end{aligned} \tag{5}$$

in which the notation $(\hat{\cdot})_{k|j}$ means the estimate of $(\cdot)_k$ obtained using the information up to time j (when the last INR measurement was available). The gain matrices $L_x \in \mathbb{R}^n$ and $L_d \in \mathbb{R}$ are computed from a steady-state Riccati equation (see for example,[10] for details), appropriately modified to account for the number of sampling times elapsed from the previous to the current INR measurement.

Steady-state dosage optimization

Given the current disturbance estimate and a desired target for INR $\bar{y}$, a module computes the warfarin steady-state daily dose that keeps y as close to $\bar{y}$ as possible while satisfying the constraints. To this aim, the following quadratic program (QP) is solved

$$(\bar{u}_k, \bar{x}_k) = \underset{u_s,\, x_s}{\text{argmin}}\ (\bar{y} - Cx_s)^2 \tag{6a}$$

subject to

$$(I - A)\,x_s = B(u_s + \hat{d}_{k|k}) \tag{6b}$$

$$0 \le u_s \le u_{\max} \tag{6c}$$

in which $u_{\max}$ is the maximum daily dose of warfarin (if any, according to the physician judgment). It is clear that the optimal daily steady-state $\bar{u}_k$ is not necessarily a multiple of u^*. Given the large-time constant ($3 \div 7$ days) it is reasonable to round the weekly dose to a multiple of u^*, rather than the daily dose. More precisely, $\bar{u}_k$ is rounded to the nearest multiple of $u^*/7$, and $\bar{x}_k$ is computed, consequently, from Eq. 6b. Indeed, this allows a finer average daily dose resolution and, hence, a closer INR control.

Dynamic dosage optimization

Given the rounded steady-state daily dose $\bar{u}_k$, and the corresponding weekly dose $7\bar{u}_k$, before describing the dynamic optimization module, it is convenient for the to define a "steady-state" sequence of daily doses $\{\bar{u}_{k,i}\}$, over a week period, such that the weekly dose is equal to $7\bar{u}_k$, that is

$$\bar{u}_{k,0} + \bar{u}_{k,1} + \cdots + \bar{u}_{k,6} = 7\bar{u}_k \tag{7}$$

and each $\bar{u}_{k,i}$ is multiple of u^* nearest to $\bar{u}_k$ either from below or from above. The sequence $\{\bar{u}_{k,i}\}$ represents the optimal weekly protocol that should be administrated in stationary conditions, that is, when the disturbance estimate does not change significantly. This situation occurs when the prediction error $e_k = y_k - C\hat{x}_{k|j}$ is small, that is, when the model prediction matches the actual patient's response. It has been observed in practice that when large prediction errors (say greater than 0.5) occur, the steady-state dosage should not be changed too aggressively because the disturbances that affect the system tend to have much stronger effects in the transient than at steady state. This can be effectively achieved by tuning the state and disturbance estimator (Eq. 5) accordingly, that is, by choosing noise covariance matrices (much) larger for the states than for the disturbance. This estimator tuning provides "robustness" to the control algorithm since it avoids large fluctuations of the steady-state weekly dose.

In order to achieve an optimal transient control of the INR, however, it may be necessary to deviate for a few days from the

optimal steady-state weekly protocol. This goal is obtained by solving the following dynamic optimization problem

$$(\{v_j^*\}, \{\overline{\eta}_j^*\}, \{\underline{\eta}_j^*\}) = \underset{\{v_j\},\{\overline{\eta}_j\},\{\underline{\eta}_j\}}{\operatorname{argmin}} \sum_{j=0}^{N-1} \{[C(w_j - \bar{x}_k)]^2 + q(\overline{\eta}_j^2 + \underline{\eta}_j^2) + r(v_j - \bar{u}_{k,j})^2\} + (w_N - \bar{x}_k)^T P(w_N - \bar{x}_k) \quad (8a)$$

subject to

$$w_{j+1} = Aw_j + B(v_j + \hat{d}_{k|k}), \quad w_0 = \hat{x}_{k|k} \quad (8b)$$

$$0 \le v_j \le u_{\max} \quad (8c)$$

$$y_{\min} - \underline{\eta}_j \le Cw_j \le y_{\max} + \overline{\eta}_j \quad (8d)$$

$$\overline{\eta}_j \ge 0, \quad \underline{\eta}_j \ge 0 \quad (8e)$$

in which $N = 7$ is the horizon, q and r are positive scalars, $y_{\min}$ and $y_{\max}$ represent the minimum and maximum therapeutic value for INR, respectively (that is, the therapeutic range) v_j is the warfarin daily dose $\underline{\eta}_j$ and $\overline{\eta}_j$ represent the value of the transient violation of the lower and upper INR bounds, respectively, and P is the positive semidefinite solution of an appropriate Riccati equation for the infinite horizon optimization problem associated to Eq. 8. The rationale behind Eq. 8 is that it finds a dosage for the first week (starting from the current day) that keeps (or brings) the INR within the therapeutic range, close to its target and does not deviate much from the steady-state weekly protocol. Given the optimal warfarin dose sequence $\{v_j^*\}$ each component is rounded to the nearest multiple of u^*, and the protocol is consequently defined by "suggesting" this rounded sequence for the first week and the steady-state sequence for the subsequent days until the next scheduled INR check. Clearly, the physician can overrule the protocol suggested by the algorithm. Moreover, when the current INR deviates significantly from its predicted value, the physician is advised to schedule the next INR check sooner to improve model prediction and ultimately the warfarin dose calculation.

Results

A number of clinical tests have been conducted mainly in collaboration with IFC-CNR (Istituto di Fisiologia Clinica of Consiglio Nazionale delle Ricerche, Pisa-Italy). The method has been tested on 25 patients never treated with warfarin. For each patient, basal INR was measured and, subsequently, INR was measured daily for the first few days, and then measured with a decreasing frequency. As an example a case study of a 65 year-old male patient affected by atrial fibrillation is presented. Patient's INR, warfarin daily dose and weekly target dose are reported in Figure 1 over a period of 140 days in which the patient followed the program's suggested protocol except for four days around the 70th day in which the patient forgot to take warfarin for one day, and the dose for the next three days was "manually" changed. It is clear that for this patient the MPC algorithm was able to keep INR within the therapeutic range. It is interesting to notice that the optimal

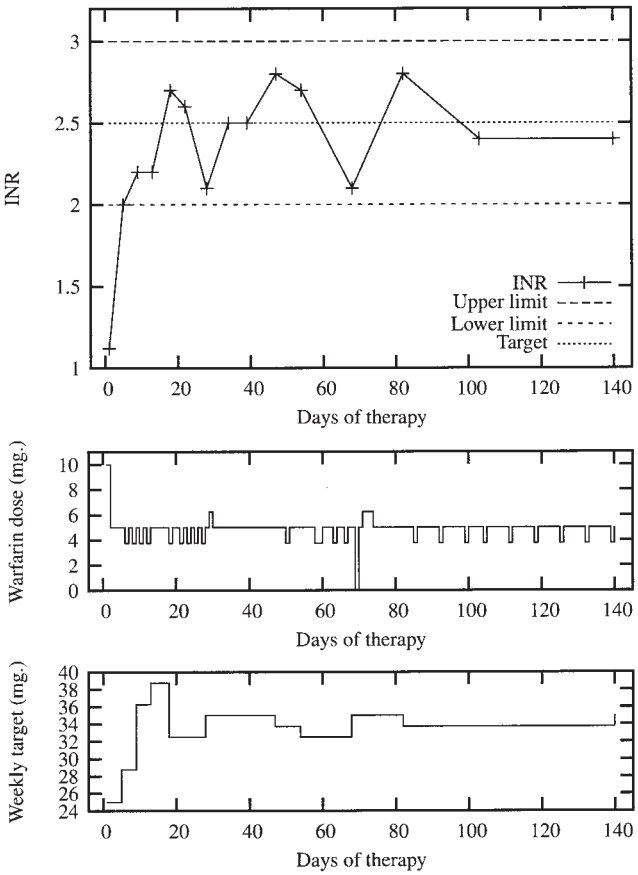

Figure 1. Example of a 65 year old male patient: INR and warfarin dose.

weekly target dose did not remain constant during the five months of therapy (due to initial adaptation of the model to the actual patient response, and because of external disturbances), and the algorithm was able to account for this variability and take appropriate control actions. Using the model of the same patient and starting from steady conditions, rejections of "typical disturbances" were simulated to evaluate the effectiveness of the proposed algorithm related to the INR measurement scheduling. In Figure 2 results of three cases are reported: peak detection (INR is measured when the peak occurs), early detection (INR is measured two days before the peak would have occurred), late detection (INR is measured five days after the peak occurs). In all cases, after the first measurement, INR is then measured every two weeks. It can be noticed that late peak detection could be detrimental to therapy effectiveness. In Figure 3 the results of three cases are reported: open-loop (the weekly dose protocol is not changed), closed-loop 1 (INR is measured and the dose protocol adjusted every week), closed-loop 2 (INR is measured and the dose protocol adjusted every two weeks). In both cases the proposed algorithm adjusts the daily dose correctly to bring INR below the limit of 3.0. It is important to point out that in both cases INR was measured at its peak.

Conclusions and future work

In this article an MPC algorithm was proposed for the optimal oral anticoagulant drug administration. The algorithm

uses a second-order discrete-time model of the response of coagulation key parameter (INR) to warfarin, which is adapted to the actual patient's dynamic response during the first period of therapy. A state and disturbance estimator is used to correct the model prediction with feedback information from INR measurement. Then, the steady-state weekly sequence of warfarin is calculated that keeps INR as close to its desired target as possible. Finally, the warfarin doses for the first week (from the day when INR is measured) are computed by a dynamic optimization module, which assumes that the steady-state weekly sequence is used from the second week. 25 clinical tests have been performed with very promising results. In the future, efforts will be devoted to undertake a large clinical trial, in which the protocol suggested by the proposed algorithm will be (statistically) compared with that decided by expert physicians and other available algorithms for the oral anticoagulant therapy. Other future research activities will regard the development of dynamic models for the response of INR to interacting drugs in order to implement a feed-forward compensation strategy.

Acknowledgments

The authors would like to acknowledge M. Miniati (MD) and S. Monti (MD) of IFC-CNR for their collaboration during the clinical studies, and E. De Candia (MD) of Policlinico Gemelli, Rome, for useful suggestions.

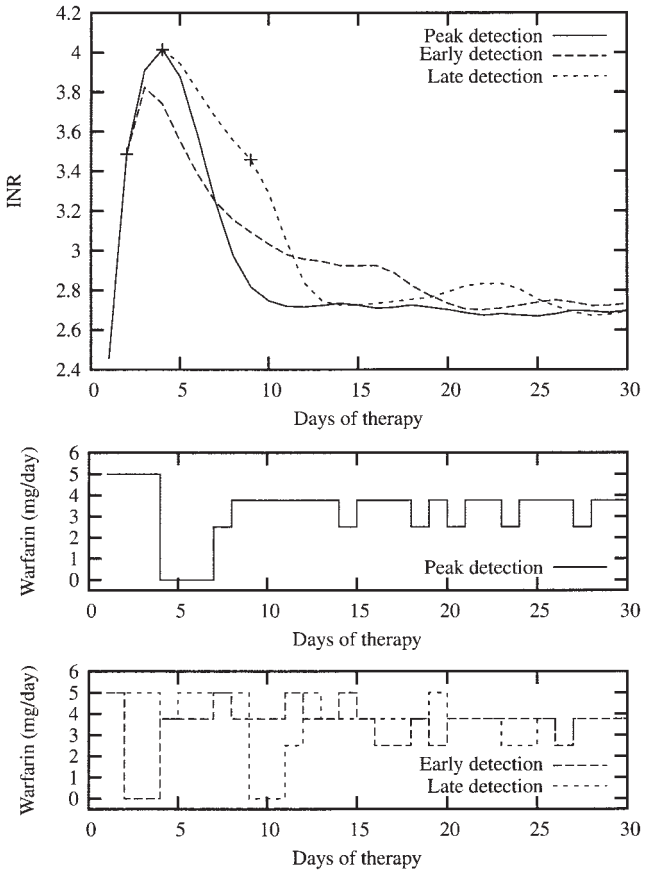

Figure 2. Simulated disturbance rejection: effect of disturbance first detection.

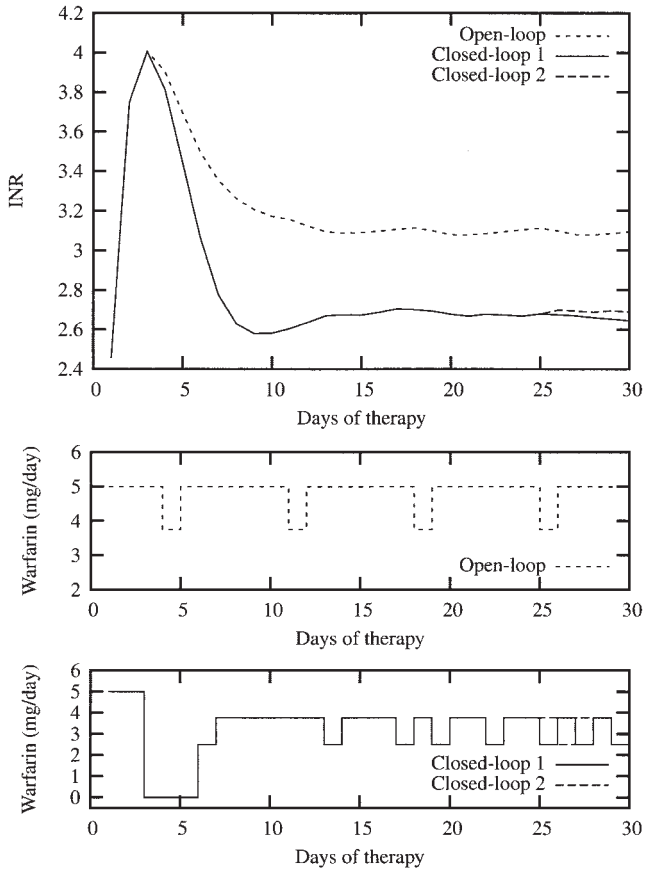

Figure 3. Simulated disturbance rejection: effect of sampling frequency.

Literature Cited

1. Yu C, Roy RJ, Kaufman H, Bequette BW. Multiple-model adaptive predictive control of mean arterial pressure and cardiac output. *IEEE Trans Biom Eng.* 1992;39:765-778.
2. Gentilini A, Schaniel C, Morari M, Bieniok C, Wymann R, Schnider T. A new paradigm for closed-loop intraoperative administration of analgesics in humans. *IEEE Trans Biom Eng.* 2002;49:289-299.
3. Parker R, Doyle FJ, Peppas NA. The intravenous route to blood glucose control. *IEEE Eng in Medicine and Biology.* 2001;20:65-73.
4. Hovorka R, Canonico V, Chassin LJ, et al. Nonlinear model predicitive control of glucose concentration in subjects with type I diabetes. *Physiol Meas.* 2004;25:905-920.
5. Bequette BW. A critical assessment of algorithms and challanges in the development of a closed-loop artificial pancreas. *Diab Technol & Therapeutics.* 2005;7:28-47.
6. Wells PS, Holbrook AM, Crowther NR, Hirsh J. Interactions of warfarin with drugs and food. *Ann Intern Med..* 1994;121:676-683.
7. Wittkowsky AK. Drug interactions update: drugs, herbs, and oral anticoagulant. *J. Thrombosis and Thrombolysis.* 2002;12:67-71.
8. Qin SJ, Badgwell TA. A survey of industrial model predictive control technology. *Cont Eng Pract.* 2003;11:733-764.
9. Mayne DQ, Rawlings JB, Rao CV, Scokaert POM. Constrained model predictive control: stability and optimality. *Automatica.* 2000;36:789-814.
10. Pannocchia G, Rawlings JB. Disturbance models for offset-free model predictive control. *AIChE J.* 2003;49:426-437.
11. Goodman A, Gilman A. *The Pharmacological Basis of Therapeutics*. McGraw-Hill: New York; 1992.
12. Kitchen S, Preston FE. Standardization of prothrombin time for laboratory control of oral anticoagulant therapy. *Seminars in Thrombosis and Hemostasis.* 1999;25:17-25.

13. O'Reilly RA, Aggeler PM. Determinants of the response to oral anticoagulant drug in man. *Pharmacol Rev.* 1970;22:35-96.
14. Chatellier G, Colombet I, Degoulet P. An overview of the effect of computer-assisted managment of anticoagulant therapy on the quality of anticoagulation. *Intl J of Medical Informatics.* 1998;49:311-320.
15. Duxbury BM, Poller L. The oral anticoagulant saga: past, present and future. *Clinical and Applied Thrombosis-Hemostasis.* 2001;7:269-275.
16. Manotti C, Pattacini C, Quintavalla R, Tagliaferri A, Lombardi M, Tassoni M. Computer assisted anticoagulant therapy. *Pathophysiology of Haemostasis and Thrombosis.* 2003;33:366-371.
17. Vadher BD, Patterson DLH, Leaning MS. Validation of an algorithm for oral anticoagulant dosing and appointment scheduling. *Clinical and Laboratory Haematology.* 1995;17:339-345.
18. Haustein KO. Pharmacokinetic and pharmacodynamic properties of oral anticoagulants, especially phenprocoumon. *Semin Thrombosis and Hemostasis.* 1999;25:5-11.
19. Keller C, Matzdorff AC, Kemkes-Matthes B. Pharmacology of warfarin and clinical implications. *Semin Thrombosis and Hemostasis.* 1999;25:13-16.
20. Hirsh J, Fuster V, Ansell J, Halperin JL. Foundation guide to warfarin therapy. *Circulation.* 2003;107:1692-1711.
21. Ljung L. *System Identification: Theory for the User*. 2nd ed. Upper Saddle River, New Jersey: Prentice Hall Inc; 1999.

Manuscript received Jun. 16, 2005, and revision received May 23, 2006.