

A switching control strategy for the attenuation of blood glucose disturbances

Mihalis G. Markakis^{1,*}, Georgios D. Mitsis², George P. Papavassilopoulos³, Petros A. Ioannou⁴
and Vasilis Z. Marmarelis⁵

¹*Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, U.S.A.*

²*Electrical and Computer Engineering, University of Cyprus, Nicosia, Cyprus*

³*Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece*

⁴*Electrical Engineering, University of Southern California, Los Angeles, CA, U.S.A.*

⁵*Biomedical Engineering, University of Southern California, Los Angeles, CA, U.S.A.*

SUMMARY

In this computational study we consider a generalized minimal model structure for the intravenously infused insulin–blood glucose dynamics, which can represent a wide variety of diabetic patients, and augment this model structure with a glucose rate disturbance signal that captures the aggregate effects of various internal and external factors on blood glucose. Then we develop a model-based, switching controller, which attempts to balance between optimal performance, reduced computational complexity and avoidance of dangerous hypoglycaemic events. We evaluate the proposed algorithm relative to the widely studied proportional–derivative controller for the regulation of blood glucose with continuous insulin infusions. The results show that the proposed switching control strategy can regulate blood glucose much better than the proportional–derivative controller for all the different types of diabetic patients examined. This new algorithm is also shown to be remarkably robust in the event of concurrent, unknown variations in critical parameters of the adopted model. Copyright © 2009 John Wiley & Sons, Ltd.

Received 8 September 2008; Revised 15 February 2009; Accepted 21 May 2009

KEY WORDS: diabetes; insulin; glucose disturbance; switching control

1. INTRODUCTION

Diabetes represents a major threat to public health with alarmingly rising trends of incidence and severity in recent years, and numerous detrimental consequences for public health. The mean level of concentration of blood glucose in normal human subjects is about 90 mg/dl and the zone from 70–110 mg/dl is usually defined as the desired state of normoglycaemia. Significant and prolonged deviations from this zone may give rise to numerous pathologies with serious and extensive clinical impact, that is increasingly recognized by current medical practice. When blood glucose concentration falls under 60 mg/dl, we have the acute

*Correspondence to: Mihalis G. Markakis, 77 Massachusetts Ave., Room 32D-740, Cambridge, MA 02139, U.S.A.

†E-mail: mihalis@mit.edu

Contract/grant sponsor: Myronis Foundation

Contract/grant sponsor: European Social Fund

Contract/grant sponsor: National Resources—Operational Program Competitiveness—General Secretariat for Research and Development

Contract/grant sponsor: Empeirikion Foundation of Greece

Contract/grant sponsor: NIH; contract/grant number: P41-EB001978

and very dangerous state of hypoglycaemia that may lead to brain damage or even death if prolonged. On the other hand, when blood glucose concentration rises above 120 mg/dl for prolonged periods of time, we are faced with the detrimental state of hyperglycaemia that may cause a host of long-term health problems (e.g. neuropathies, kidney failure, loss of vision, etc.). The severity of the latter clinical effects is increasingly recognized as medical science advances and the physiological causes of diabetes are better understood, thus revealing it as a major lurking threat to public health with long-term repercussions. Prolonged hyperglycaemia is usually caused by defects in insulin production (Type 1 diabetes), insulin action or both (Type 2 diabetes) [1].

Throughout the last 40 years, many different approaches have been followed for the automated regulation of blood glucose, with the vast majority of publications employing either simple, linear controllers [2–5] or model-based controllers [4, 6–10] due to their conceptual simplicity and effectiveness. However, despite the considerable effort and resources that have been dedicated to this task, no method or approach has been demonstrated yet to produce an effective solution with clinical utility and wide applicability. In our opinion, this is due primarily to the following reasons:

- Most of the computational studies of glucose regulation are restricted to Type 1 diabetes and therefore employ the diabetic minimal model (DMM) [4, 6, 8, 11, 12] to represent the actual system in their simulations. Closer examination of this model indicates that the major function of endogenous insulin production is not represented, which makes the DMM unsuitable for any other diabetic case.
- Most of the publications to date study a controller's ability to regulate blood glucose in the event of a glucose 'disturbance' caused by a single meal or a sequence of meals [6]. However, the form of the actual glucose 'disturbance' that the controller will face is likely to be more complicated, since it will include the aggregate effects of the endocrine system (e.g. systemic secretions of glucagon, cortisol, epinephrine) and of many external factors (e.g. exercise, stress, mental activity, etc.) on blood glucose concentration.

2. STATEMENT OF THE PROBLEM

Insulin–glucose dynamics: As mentioned above, most of the computational studies of glucose regulation in diabetic patients adopt the DMM. The implicit assumption of this model is that no endogenous insulin is produced and secreted in plasma. Although this is true for Type 1 cases, it is certainly not true for Type 2 diabetics or other patients in need of continuous insulin therapy (e.g. patients in the ICU—see [13]). In order to incorporate this important phenomenon, we augment the DMM with an additional differential equation describing the endogenous insulin production dynamics. The resulting model, very similar in structure to the original MM [14, 15], operates in 'closed-loop' through the feedback exerted by the endogenously secreted insulin and utilizes intravenous insulin infusions like many similar studies the last years [4, 6, 8–10, 16]:

$$\frac{dI_{\text{ex}}}{dt} = -p_4 I_{\text{ex}}(t) + \frac{U(t)}{V_I}$$

$$\frac{dI_{\text{en}}}{dt} = -\alpha \cdot I_{\text{en}}(t) + \beta \cdot \max[Y(t) - \theta, 0]$$

$$\frac{dX}{dt} = -p_2 X(t) + p_3 [I_{\text{en}}(t) + I_{\text{ex}}(t)]$$

$$\frac{dG}{dt} = -p_1 G(t) - X(t) \cdot Y(t) + D(t)$$

$$Y(t) = G_b + G(t)$$

U is the rate of intravenously infused insulin (deviation from its basal rate), I_{ex} is the concentration of plasma insulin due to exogenous infusion (deviation from its basal value), I_{en} is the concentration of plasma insulin due to endogenous production and secretion (deviation from its basal value), X is the internal variable of 'insulin action', D is the stochastic glucose rate disturbance caused by internal and external factors, G is the deviation of blood glucose concentration from its basal value of $G_b = 90$ mg/dl and Y is the concentration of blood glucose. A similar model structure has been used recently with promising clinical results [10].

We estimate the implicated parameters using nonlinear least squares fitting of the MM to input–output data generated by Sorensen's model [17]. Since

the MM operates in ‘closed-loop’, we follow the technique of artificially ‘opening-the-loop’; in other words every loop of the system (feed-forward or feedback) is identified separately. All input signals used during the system identification procedure are broadband (Gaussian white noise) with carefully selected dynamic ranges, in order to fully excite the dynamics of each loop. The procedure above results in the following parameter values: $p_1 = 0.018 \text{ min}^{-1}$, $p_2 = 0.052 \text{ min}^{-1}$, $p_3 = 1.2 \times 10^{-5} \text{ L/min}^2 \text{ mU}$, $p_4 = 0.28 \text{ min}^{-1}$, $V_I = 3 \text{ L}$, $\alpha = 0.44 \text{ min}^{-1}$, $\beta = 0.108 (\text{mU/L})(\text{mg/dl})^{-1} \text{ min}^{-1}$ and $\theta = 103 \text{ mg/dl}$. However, Sorensen’s model (with the parameter values given in his PhD Thesis) corresponds to healthy and not diabetic subjects, so some parameters of our model have to be properly modified.

- The primary physiological difference between diabetic (either Type 1 or Type 2) and healthy subjects is the delayed and attenuated (compared to the normal) pancreatic responsiveness to hyperglycaemia. In the context of the MM, this implies reducing the values of parameters α and β . For the first parameter we use a nominal value of $\alpha = 0.22 \text{ min}^{-1}$ (double the time constant of endogenous insulin production). Regarding the second (and more important) parameter β , we take $\beta = 0.054$ as reference but various other values are used too, emulating different types of diabetic patients.
- A secondary physiological difference of diabetics is the decreased (insulin-stimulated) tissue glucose uptake. In terms of the MM, this can be translated as reducing the value of parameter p_1 . This statement is in accordance with the conclusions of [12] that many subsequent studies have followed. Throughout this paper we will use $p_1 = 0.009$ (double the time constant), unless stated otherwise.
- Finally, regarding insulin sensitivity, we use a value of $p_3 = 5 \times 10^{-5}$, which is one of the highest values that appear in the literature [8]. Simple analysis can verify that this change increases the magnitude of the bilinear term $X(t) \cdot Y(t)$ and thus, enhances the nonlinear nature of the system. This is in agreement with the authors’ experience from studying real insulin–glucose data.

Closed-loop system: The closed-loop system for the regulation of blood glucose that we implement in the present study can be seen in Figure 1. The MM presented above plays the role of the real system in our simulations. The glucose sensor generates simulated samples of the output Y with a sampling interval T (discretized output of the MM), upon which measurement noise N is added to emulate a realistic situation. A digital controller is used to compute the control input U to the system based on the measured error signal V . The objective of the digital controller is to attenuate the effects of the disturbance signal $D(t)$ and keep the error signal $V(t)$ within bounds, defined as the normoglycaemic region. Usually the targeted value of blood glucose concentration is set equal to the basal value $G_b = 90 \text{ mg/dl}$, and a conservative definition of the normoglycaemic region is from 70 to 110 mg/dl—i.e. $|V(t)| \leq 20 \text{ mg/dl}$.

Glucose disturbance: A critical set of assumptions in this computational study concerns the design of the glucose rate disturbance signal $D(t)$ that is used in the simulations. These assumptions are based on our best understanding of the physiological factors that affect the blood glucose concentration, independently of the infused insulin. For this purpose, we postulate various deterministic and stochastic components that seek to capture the main factors influencing the blood glucose concentration. Specifically, the following additive terms are incorporated into the disturbance signal $D(t)$:

- Terms of the exponential form $\gamma \cdot \exp(-0.05 \cdot t)$, which represent the standard Fisher meals [11]. Depending on the meal, the parameter γ takes random values within predetermined ranges, corresponding to 15–25 g of an OGTT for breakfast, 40–60 g of an OGTT for lunch and 30–50 g of an OGTT for dinner. The timing of each meal is also selected randomly within a two-hour interval specified for breakfast, lunch and dinner. The effect of this rate disturbance on glucose concentration has the form of a negative gamma-like function that resembles the convolution of this exponential form with the exponential function $\exp(-p_1 \cdot t)$. The peak-time of this gamma-like curve is around 30 min and its peak amplitude varies from

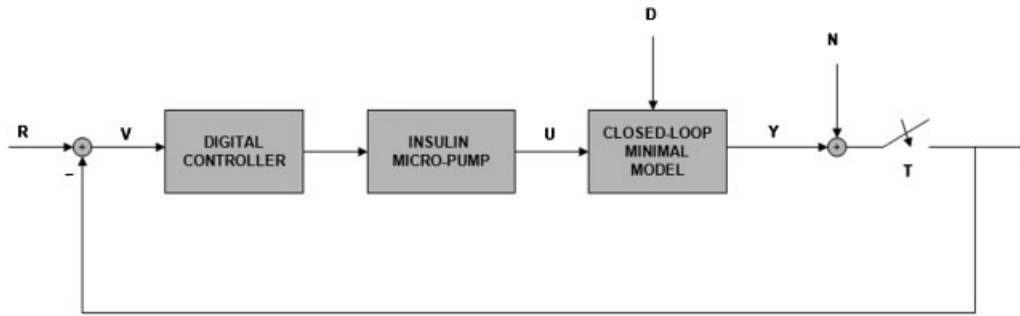


Figure 1. Closed-loop system for the regulation of blood glucose: R is the reference level, which is the target value for the controlled blood glucose (in mg/dl), V is the measured error signal, i.e. the noise-corrupted difference of Y from the constant level R (in mg/dl), N is the measurement noise/errors (in mg/dl) and T is the sampling interval determined by the glucose sensor (in min).

85–310 mg/dl (for nominal parameter values). Fisher meals exhibit much faster dynamics than other, higher-order, meal models, like Lehman–Deutsch meals [18]; this makes them more challenging from a control viewpoint. A simple comparison between Figures 6 and 8 of [6] can justify this argument.

- Terms of the exponential form $\delta \cdot \exp(-0.11 \cdot t)$, which represent small-scale neuro-hormonal effects caused by stress, mental exertion and other external factors affecting the nervous system. The appearance of these terms is modeled with a Poisson process with parameter $\lambda = 0.033 \text{ min}^{-1}$. The parameter δ is uniformly distributed within the range $[-0.07, 0.14]$ in mg/dl/min. The effect of this rate disturbance term on glucose concentration has again the form of a gamma-like function, with peak-time of approximately 25 min and peak amplitude from $[-5, 10]$ mg/dl.
- Terms of the exponential form $\varepsilon \cdot \exp(-0.025 \cdot t)$, which represent larger random effects due to factors such as exercise or strong emotions. The appearance of these terms is modeled again with a Poisson process with parameter $\lambda = 0.016 \text{ min}^{-1}$. The parameter ε is uniformly distributed within the range $[-0.06, 0.135]$ in mg/dl/min. The effect of this rate disturbance term on glucose concentration has again the form of a negative gamma-like function with peak-time of approximately 60 min and peak amplitude from $[-15, 30]$ mg/dl.

- Two sinusoidal terms of the form $\alpha_i \cdot \sin(\omega_i \cdot t + \varphi_i)$ with specified amplitudes and frequencies (α_i and ω_i) and random phase φ_i uniformly distributed within the range $[-\pi/2, \pi/2]$. These terms represent circadian rhythms [19, 20] (endocrine cycles) with periods 8 and 24 h and amplitudes around 10 mg/dl. The effect of these rate disturbance terms on glucose concentration is periodic (with minor harmonics generated by the bilinear term) and their combined amplitude ranges up to 20 mg/dl, depending on the randomly selected phases and the production of endogenous insulin.
- A constant term B , which is uniformly distributed within the range $[0.009, 0.036]$ and represents a random bias of the subject-specific basal glucose from the nominal value of G_b . The effect of this term is a random, steady-state, uniformly distributed offset, ranging from 10 to 40 mg/dl.

Our goal in this paper is to model the disturbance in a way that is consistent with the accumulated qualitative knowledge to date in a realistic context and similar to actual observations in clinical trials (e.g. see the patterns of glucose fluctuations shown in [7, 21, 22]). An illustrative example of the combined effect of these rate disturbance factors on glucose fluctuations is shown in Figure 2.

Simulated glucose sensor and insulin micro-pump: The simulated glucose sensor that is assumed to make

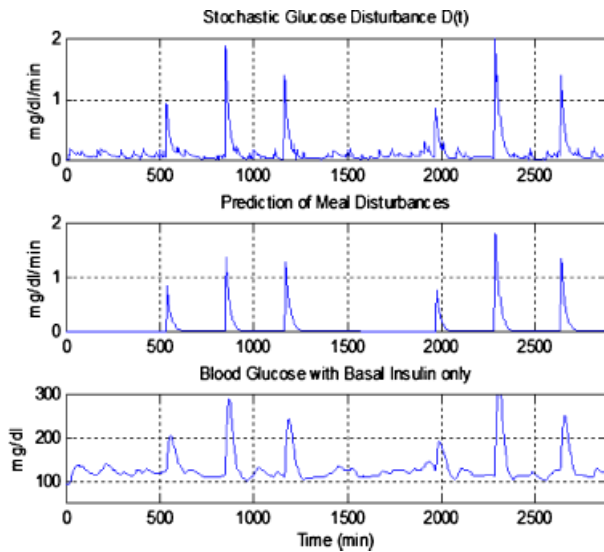


Figure 2. Glucose rate disturbance and its effect on blood glucose: the top panel shows a typical pattern of the glucose rate disturbance signal employed in this study. The large ‘impulsive’ events are caused by meals. The center panel shows the estimate of the meal-related glucose rate disturbances of assumed known exponential form. The bottom panel shows the corresponding blood glucose fluctuations due to this rate disturbance, when only basal insulin is externally delivered ($\beta=0.054$ -nominal values for the other parameters).

the continuous glucose measurements (every 3 min) in this computational study exhibits the shortcomings that have been observed in current continuous glucose sensors: a small time lag between arterial and interstitial glucose (6 min in this study) and underestimation of extreme hyperglycaemic and hypoglycaemic events (e.g. see the figures in [23]). The total ‘sensor noise’ achieving the above is created by two distinct low-pass filters in series—at the output of the second, Gaussian white noise (with high SNR) is added to account for small-scale, random errors during the measurement procedure. Obviously, the presence of noise makes the control task even more challenging.

The use of an insulin micro-pump is simulated with a zero-order hold, which performs digital-to-analog conversion, as well as an upper bound of 200 mU/min on the magnitude of the exogenous insulin rate, although recent advances in the micro-pump

technology have made possible infusion rates of magnitude greater than 300 mU/min [24].

3. METHODS

Proportional-derivative control: First we present briefly the very popular concept of proportional-derivative control (PDC), which will be used as a benchmark to evaluate the performance of our novel algorithm. The discrete-time PDC algorithm utilizes a control signal of the form

$$U(n) = -K_p V(n) - K_d \frac{V(n) - V(n-1)}{T}$$

where n denotes the discrete-time index for samples obtained every $T=3$ min. The values of the controller parameters used in our simulations are $K_p=0.4$ and $K_d=8$ and correspond to the best PDC, as determined by successive trials, using the joint criterion of attenuation of the effects of meal disturbances and avoidance of hypoglycaemic events. The control input is clipped when it exceeds the imposed lower and upper bounds of 0 and 200 mU/min, respectively. A recent study [4] has concluded that a PDC is able to regulate blood glucose and nonlinear, model-based algorithms are in many cases obsolete.

Model predictive control: In this section we outline the concept of model predictive control (MPC), which is at the core of the control algorithm that we developed. Having knowledge of the nonlinear model and of all the past input–output pairs, the goal of the MPC is at every time instant n , to determine the control input value $U(n)$, so that the following cost function is minimized:

$$J(n) = [Y(n+p|n) - R]^T \cdot \Gamma_y \cdot [Y(n+p|n) - R] + \Gamma_U \cdot U(n)^2$$

where $Y(n+p|n)$ is the vector of predicted output values over a future horizon of p steps using the model and the input past values, R is the target reference value for the output, Γ_y is a diagonal matrix of weighting coefficients that assigns greater importance to the near-future predictions and Γ_U a scalar, determining how ‘expensive’ the control input is. We have implicitly

assumed that the so-called ‘move horizon’ is 1, since simulations have shown that only the immediate control input value $U(n)$ is critical. The procedure is repeated at each time step to compute $U(n+1)$ and so on. More details on MPC and relevant control issues can be found in [25].

The MPC algorithm relies critically on our ability to predict the future values of the output. However, to predict these future values accurately we must have a precise estimate of the future values of the rate disturbance (in addition to an accurate model and set of parameters). This is very difficult in the presence of small-scale stochastic disturbances, but feasible in the case of large disturbances caused by meals, when their specific exponential structure (e.g. standardized Fischer meals) can be assumed: using the past input–output pairs and the available model we can reconstruct the glucose rate disturbance. Then, with an empirically determined threshold, we can sort out the large impulsive events caused by meals from the non-meal disturbances. Apart from their exponential form, no other assumptions about the magnitude or the timing of the meals need to be made. Figure 2 illustrates our ability to estimate meal-related glucose disturbances. The observation that MPC can be successfully applied during the time periods of meals but not during small stochastic disturbances is the basic idea behind the switching approach.

Even though we consider a system described by a set of ordinary differential equations, our approach is ‘input–output’ based rather than state-space based: we use the known model structure and estimated parameters only to predict the effect of insulin on glucose. Our control design utilizes only the output of the model (Y) and not the whole state vector (i.e. the internal variables I_{ex} , I_{en} , X , G); hence an estimation of the state at all times is not necessary. At the beginning of every simulation run the MPC algorithm assumes that the system is at equilibrium—this assumption might not hold of course. However, since our analysis is not transient but steady-state (simulations of 48 h), the impact of the initial state is very limited and does not affect the results and conclusions of this study.

Switching control strategy: The concept of a switching control strategy (SCS) comes naturally when trying to combine the strengths and mitigate the

weaknesses of the two approaches described above (PDC, MPC), depending on the level of disturbance.

- During meals, when the effects of the bilinear term are significant and disturbance is possible to predict, MPC can be employed since it is known to perform almost optimally and carefully avoid hypoglycaemias. In our simulations we use a prediction horizon of 90 min (30 samples) and no weighting (Γ_y is a unity vector), in order to minimize the risk of hypoglycaemic events. The length of the prediction horizon might seem too long since it is ‘common knowledge’ that the effect of IV insulin on blood glucose has a time constant of 30–40 min. However, this does not imply that there is no causal effect after this period. As another measure of precaution against hypoglycaemia, we also use asymmetric weighting of the predicted output vector, as in [4], whereby we penalize 10 times more the deviations of the vector $Y(n+p|n)$ that are below 80 mg/dl.
- In the presence of non-meal (small-scale) disturbances, when the nonlinear effects are limited and disturbance is hard to predict, PDC is employed. The parameters used here are $K_p=0.8$ and $K_d=16$. A comparison with the parameters of pure PDC shows that we doubled the gain of the controller, or in other words, made the control action more aggressive. Note that using PDC during the long, non-meal periods makes the computational burden of the algorithm much less compared to the pure MPC case.

A switching scheme for the regulation of blood glucose has been proposed in a recent computational study [16]. However, the similarities between the two approaches are otherwise limited: our algorithm, instead of solving Riccati equations and dynamic programming problems, utilizes two widely used controllers for blood glucose with well-defined switching points between them and is tested on a wide variety of diabetic cases.

A note on the issue of stability: the open-loop insulin–glucose model, with the nominal parameter values used in this study, is stable. We can see this, either by linearizing around the operating point and making use of Lyapunov’s linearization method or by

simply noting that the linear part of the model is stable and the nonlinearity (bilinear term) has a stabilizing effect. In fact, one can extend this conclusion to the case when the parameters take values within a wide range around the nominal ones. It is due to these strong stability properties of the open-loop system that we avoid any discussion of closed-loop stability and the consequent technical proofs. Instead, we focus on the issue of performance of the switching control algorithm, which is far more critical in this problem.

4. RESULTS

Mild Type 2 diabetic/ICU case: In this section we try to emulate the case of mild Type 2 diabetics or patients in the ICU with a need of continuous insulin therapy. These subjects have the ability to produce significant amounts of endogenous insulin but of course less than what healthy subjects can. Thus, we assume that $\beta=0.054$, which is half the normal value. The remainder of the parameters retain their nominal values. All parameter values are assumed to be known to the controllers. Note that in the case of the ICU, the MPC controller may be activated only on the relatively rare occasions when the feeding infusion rate is changed (unless there are internal reasons that change significantly the rate-disturbance signal).

Table I reports the averaged performance over 20 independent runs of 48 h of the SCS and PDC. A comparison between the two demonstrates clearly the superiority of the control strategy introduced in this study. Figure 3 presents the SCS in action: the top panel shows the blood glucose without control (dashed line) and after switching control (solid line). The dashed rectangulars on the time axis indicate the time periods when the SCS employs a MPC action. The bottom panel shows the insulin infusions computed by the SCS.

Severe Type 2 diabetic case: In this section we take a look at the case of severe Type 2 diabetics: these subjects have the ability to produce endogenous insulin but significantly less than what is needed to regulate blood glucose. Thus we assume that $\beta=0.01$, which is about one-tenth the normal value. The remainder of the parameters retain their nominal values and are known

to the controllers. Table III reports the average behavior of the two control strategies for 20 independent simulations, of 48 h each. The performance indices of the SCS compare favourably with the PDC case.

Type 1 diabetic case: In this section we try to emulate the case of Type 1 diabetics: these subjects cannot produce any endogenous insulin at all and the regulation of their blood glucose depends solely on exogenously infused insulin. We simulate this case by considering a value of $\beta=0$. The remainder of the parameters retain their nominal values. All parameter values are again known to the controllers. Table IV reports the average behavior of the two control strategies (PDC, SCS) for 20 independent simulations, of 48 h each. The results of Table IV indicate that the SCS can be applied to Type 1 diabetics too, without degradation of performance or risk of patients' safety. Once more, the performance indices of the SCS compare favourably with the PDC case.

To illustrate how the MPC part of SCS achieves attenuation of meal disturbances in Type 1 diabetics, we present this case in Figure 4. Besides the obvious improvement in performance relative to the case without control, there are two more observations worth making: (a) the SCS is able to fully attenuate the effects of a meal in less than 50 min and this should be satisfactory considering how prolonged the effects of meals can be in Type 1 diabetics; (b) the lowest values of glucose observed in Figure 4 are about 70 mg/dl (very close to the lower bound of our target region), which indicates that the SCS can avoid the hypoglycaemic undershoot that typically follows the attenuation of a meal using PD control. Of course, there is a fundamental trade-off: more conservative tuning of the SCS raises the lowest bound of the glucose values but also prolongs the time-period for attenuation of a meal.

Plant-model mismatch case I: Here we examine the robustness of performance in the face of parametric uncertainty due to inter-patient variability, intra-patient variability, or simply mis-estimation of parameters α , β and p_1 , which determine the production of endogenous insulin and blood glucose clearance. The range of random variations considered for these parameters is $\pm 50\%$ from their nominal values and are not known to the controllers. The effects of the parametric uncertainty

Table I. Averages of 20 independent runs over 48 h each when $\beta=0.054$ (mild Type 2 diabetic/ICU patient).

	No control	PDC	SCS
MV	126.6	106.6	100.8
SD	38.1	26.7	22.5
PTO	68.6	26	14
INSULIN/DAY	0	8.5	10.9
HYPO	0	0	0

Presented are the mean value (MV) and the standard deviation (SD) of glucose fluctuations, the percentage of time glucose is found outside the normoglycaemic region 70–110 mg/dl (PTO), and the average insulin usage (IU) in U/day, as well as the number of hypoglycaemic events.

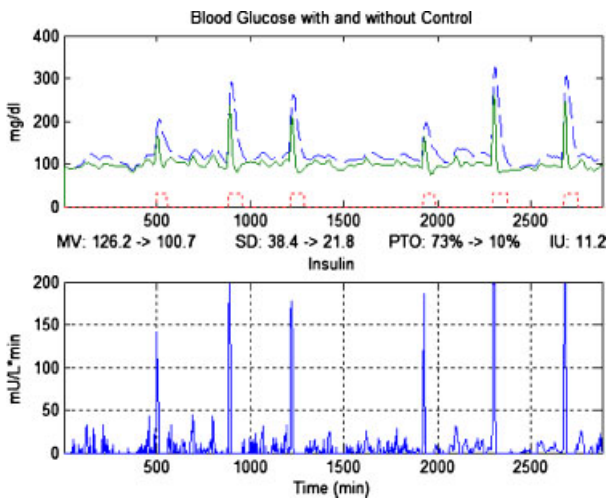


Figure 3. The switching control strategy in action: the top panel shows the blood glucose levels corresponding to the general stochastic disturbance signal, with basal insulin infusion only (dashed line) and after SCS action (solid line). The mean value (MV), standard deviation (SD) and the percentage of time that the glucose is found outside the normoglycaemic region of 70–110 mg/dl (PTO) are reported between the panels for SCS (values to the right of ‘→’) and without control action (values to the left of ‘→’). The dashed rectangular on the time axis determines the periods of time when MPC is switched on. The bottom panel shows the insulin profile determined by the SCS (the average insulin usage is 11.2 U/day in this case).

are very limited for the SCS case, as demonstrated by comparing the results of Table II with Table I. This indicates that SCS is robust in response to variations in the aforementioned parameters. Similar conclusions apply

Table II. Averages of 20 independent runs over 48 h each when $\alpha \in [0.11, 0.33]$, $\beta \in [0.027, 0.081]$ and $p_1 \in [0.0045, 0.0135]$ (random values, uniformly distributed).

	No control	PDC	SCS
MV	125.5	105.2	99.8
SD	38.7	27.1	22.8
PTO	61.8	24.2	13.6
INSULIN/DAY	0	8.2	10.5
HYPO	0	1	0

for the PDC, except for one reported hypoglycaemic event.

Regarding theoretical guarantees of robust performance of the closed-loop system under SCS, the standard H_∞ methodology does not appear promising, since it would require simplifications/assumptions that are not deemed reasonable in this specific application context. As the field of robust MPC of uncertain nonlinear systems advances, this question should be explored—possibly in the direction suggested by a recently published paper [26].

Plant-model mismatch case II: A second case of plant-model mismatch is the widely studied (and of course very important) case of Type 1 diabetes. The setting here is as follows: apart from the fact that we face a Type 1 diabetic patient (hence $\beta=0$), the internal model of MPC has its parameters set to their nominal values. However, we allow the parameter p_3 (which determines insulin sensitivity) of the real system to take random values uniformly over the interval $[5 \times 10^{-6}, 5 \times 10^{-5}]$. These variations are unknown to the

Table III. Averages of 20 independent runs over 48 h each when $\beta=0.01$ (severe Type 2 diabetic patients).

	No control	PDC	SCS
MV	165.1	110.3	101.7
SD	59.3	28.2	23.2
PTO	93.6	32.7	15.6
INSULIN/DAY	0	9.5	11.9
HYPO	0	0	0

Table IV. Averages of 20 independent runs over 48 h each when $\beta=0$ (Type 1 diabetic patients).

	No control	PDC	SCS
MV	225.1	111.4	101.9
SD	85.4	28.7	23.4
PTO	97.1	35.7	16
INSULIN/DAY	0	9.9	12.1
HYPO	0	0	0

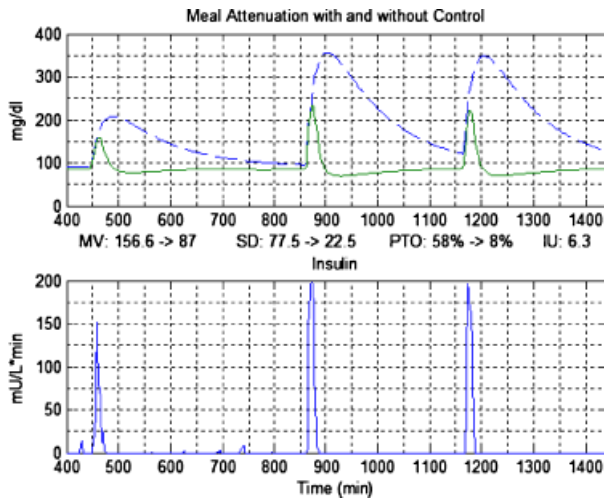


Figure 4. Attenuation of meal disturbances with switching control: the top panel shows the blood glucose levels corresponding to three meals, with basal insulin infusion only (dashed line) and after SCS action (solid line). The bottom panel shows the amounts of infused insulin used by the SCS.

controller. Table V summarizes the results of extensive simulations and shows that, although the performance is not as good as in Table IV (this is to be expected since we are talking about parametric uncertainty of an order of magnitude), SCS remains better than PDC and safe for the patients. Note that we have considered a range of p_3 values that are less than or equal to the nominal, because the nominal value used in this study is the largest one that has appeared in the literature and, thus, there is a higher probability that the controller will face systems with smaller, rather than larger, values of p_3 .

Table V. Averages of 20 independent runs over 48 h each when $\beta=0$ and $p_3 \in [5 \cdot 10^{-6}, 5 \cdot 10^{-5}]$ (random values, uniformly distributed).

	No control	PDC	SCS
MV	221.8	117.4	112.1
SD	84.6	31.8	28.9
PTO	97.3	49.9	35.3
INSULIN/DAY	0	12.1	18.0
HYPO	0	1	0

5. DISCUSSION

The proposed switching control strategy gives very good results and performs much better overall than PDC. Without changing its parameter values, the SCS can deal with Type 1 diabetes, Type 2 diabetes and ICU cases without degradation in performance or risk of the patients' safety. It is also significantly robust to a wide range of variations of crucial model parameters.

An important observation has to do with the performance of the two control algorithms in connection to the value of parameter p_3 : simulation results not presented here show that the performance of the SCS remains unchanged for known variations in p_3 . In contrast, the performance of PDC improves drastically as the value of p_3 decreases (and the effect of the nonlinearity becomes less drastic). When p_3 is sufficiently small then the two algorithms have roughly the same performance (in agreement with [4]).

We observe that the controllers considered in this study usually utilize amounts of insulin from 8–12 U/day, which is less than what appears in the literature (50 U/day). This happens for two reasons: (a) the amount reported does not include the basal

exogenous infusion and (b) we adopted a relatively high value for parameter p_3 , which means that less insulin is needed to reject the same glucose disturbances.

Finally, one might expect that less severe diabetic cases would be in need of less exogenous insulin, since part of the glucose regulation would be carried out by endogenously produced insulin. However, comparing the values of exogenous insulin usage in Tables I–II we see that this is not the case (apart from Table V, the need for exogenous insulin is more or less the same for all diabetic cases examined). The reason behind this rather counter-intuitive result is that endogenous insulin production (which is based on threshold nonlinearity) occurs only when blood glucose exceeds $\theta = 103$ mg/dl and remains quite small as long as blood glucose is in the normoglycaemic region. Thus, the presence of an effective controller minimizes the effect of endogenous insulin on blood glucose.

Future work may extend these results in the non-stationary case that requires adaptive or robust control strategies. Ongoing research examines more closely the case of unknown variations in parameters determining insulin action (p_2 and p_3). Our future efforts should also address the issue of predicting all the stochastic components of the rate disturbance signal (not just meals) and, thus, allowing the application of a pure MPC strategy.

ACKNOWLEDGEMENTS

This work was supported by the Myronis Foundation (Graduate Research Scholarship), by the European Social Fund (75%) and National Resources (25%)—Operational Program Competitiveness—General Secretariat for Research and Development (Program ENTER 04), by a grant from the Empeirikion Foundation of Greece and the NIH Center Grant No. P41-EB001978 to the Biomedical Simulations Resource at the University of Southern California.

REFERENCES

1. Available from: http://en.wikipedia.org/wiki/Diabetes_mellitus.
2. Chee F, Fernando T, Savkin A *et al.* Expert PID control system for blood glucose control in critically ill patients. *IEEE Transactions on Information Technology in Biomedicine* 2003; **7**:419–425.
3. Fischer U, Salzsieder E, Freyse E *et al.* Experimental validation of a glucose insulin control model to simulate patterns in glucose-turnover. *Computational Methods and Programs in Biomedicine* 1990; **32**:249–258.
4. Hernjak N, Doyle F. Glucose control design using nonlinearity assessment techniques. *American Journal of Chemical Engineers* 2005; **51**:544–554.
5. Salzsieder E, Albrecht G, Fischer U *et al.* Kinetic modeling of the glucoregulatory system to improve insulin therapy. *IEEE Transactions on Biomedical Engineering* 1985; **32**:846–855.
6. Dua P, Doyle F, Pistikopoulos E. Model-based blood glucose control for Type 1 diabetes via parametric programming. *IEEE Transactions on Biomedical Engineering* 2006; **53**:1478–1491.
7. Hovorka R, Canonico V, Chassin L *et al.* Nonlinear model predictive control of glucose concentration in subjects with Type 1 diabetes. *Physiological Measurement* 2004; **25**: 905–920.
8. Lynch S, Bequette B. Model predictive control of blood glucose in Type 1 diabetics using subcutaneous glucose measurements. *Proceedings of American Control Conference*, Anchorage, AK, 2002; 4039–4043.
9. Parker R, Doyle F, Peppas N. A model-based algorithm for blood glucose control in Type I diabetic patients. *IEEE Transactions on Biomedical Engineering* 1999; **46**:148–157.
10. Van Herpe T, Haverbeke N, Pluyms B *et al.* The application of model predictive control to normalize glycemia of critically ill patients. *Proceedings of European Control Conference*, Kos, Greece, 3116–3123.
11. Fisher M. A semiclosed-loop algorithm for the control of blood glucose levels in diabetics. *IEEE Transactions on Biomedical Engineering* 1991; **38**:57–61.
12. Furler S, Kraegen E, Smallwood R *et al.* Blood glucose control by intermittent loop closure in the basal model: computer simulation studies with a diabetic model. *Diabetes Care* 1985; **8**:553–561.
13. Van den Berghe G, Wilmers A, Hermans G *et al.* Intensive insulin therapy in the medical ICU. *New England Journal of Medicine* 2006; **354**:449–461.
14. Bergman R, Phillips L, Cobelli C. Physiologic evaluation of factors controlling glucose tolerance in man. *Journal of Clinical Investigations* 1981; **68**:1456–1467.
15. Toffolo G, Bergman R, Finegood D *et al.* Quantitative estimation of beta cell sensitivity to glucose in the intact organism: a minimal model of insulin kinetics in the dog. *Diabetes* 1980; **29**:979–990.
16. Chee F, Savkin A, Fernando T *et al.* Optimal H_∞ insulin injection control for blood glucose regulation in diabetic patients. *IEEE Transactions on Biomedical Engineering* 2005; **52**:1625–1631.
17. Sorensen J. A physiological model of glucose metabolism in man and its use to design and assess improved insulin therapies for diabetes. *Ph.D. Thesis*, MIT, 1985.
18. Lehmann E, Deutsch T. A physiological model of glucose-insulin interaction in Type 1 diabetes mellitus. *Journal of Biomedical Engineering* 1992; **14**:235–242.

19. Van Cauter E, Shapiro E, Tillil H *et al.* Circadian modulation of glucose and insulin responses to meals: relationship to cortisol rhythm. *American Journal of Physiology* 1992; **262**:467–475.
20. Lee A, Ader M, Bray G *et al.* Diurnal variation in glucose tolerance. *Diabetes* 1992; **41**:750–759.
21. Chee F, Fernando T, Van Heerden V. Closed-loop glucose control in critically ill patients using continuous glucose monitoring system (CGMS) in real time. *IEEE Transactions on Information Technology in Biomedicine* 2003; **7**:43–53.
22. Dudde R, Vering T, Piechotta G *et al.* Computer-aided continuous drug infusion: setup and test of a mobile closed-loop system for the continuous automated infusion of insulin. *IEEE Transactions on Information Technology in Biomedicine* 2006; **10**:395–402.
23. Freckmann G, Kalatz B, Pfeiffer B *et al.* Recent advances in continuous glucose monitoring. *Experimental and Clinical Endocrinology & Diabetes* 2001; **109**:347–357.
24. Moore K, Saudek C, Greene A *et al.* Implantable insulin pump therapy: an unusual presentation of a catheter-related complication. *Diabetes Technology and Therapeutics* 2006; **8**:397–401.
25. Bertsekas D. *Dynamic Programming and Optimal Control*, vol. 1. Athena Scientific: Belmont, MA, 2005; 281–400.
26. Mhaskar P, El-Farra N, Christofides P. Robust hybrid predictive control of nonlinear systems. *Automatica* 2005; **41**:209–217.