

NIH Public Access

Author Manuscript

Comput Methods Programs Biomed. Author manuscript; available in PMC 2012 May 1

Published in final edited form as:

Comput Methods Programs Biomed. 2011 May ; 102(2): 138-148. doi:10.1016/j.cmpb.2010.04.011.

Anticipating the next meal using meal behavioral profiles: A hybrid model-based stochastic predictive control algorithm for T1DM*

C.S. Hughes^a, S.D. Patek^{a,*}, M. Breton^b, and B.P. Kovatchev^{a,b}

^aDepartment of Systems and Information Engineering, University of Virginia, United States

^bDepartment of Psychiatry and Neurobehavioral Sciences, University of Virginia, United States

Abstract

Automatic control of Type 1 Diabetes Mellitus (T1DM) with subcutaneous (SC) measurement of glucose concentration and subcutaneous (SC) insulin infusion is of great interest within the diabetes technology research community. The main challenge with the so-called "SC–SC" route to control is sensing and actuation delay, which tends to either destabilize the system or inhibit the aggressiveness of the controller in responding to meals and exercise. Model predictive control (MPC) is one strategy for mitigating delay, where optimal insulin infusions can be given in anticipation of future meal disturbances. Unfortunately, exact prior knowledge of meals can only be assured in a clinical environment and uncertainty about when and if meals will arrive could lead to catastrophic outcomes. As a follow-on to our recent paper in the IFAC symposium on Biological and Medical Systems (MCBMS 2009) [1], we develop a control law that can anticipate meals given a probabilistic description of the patient's eating behavior in the form of a random meal (behavioral) profile. Preclinical *in silico* trials using the oral glucose meal model of Dalla Man et al. show that the control strategy provides a convenient means of accounting for uncertain prior knowledge of meals without compromising patient safety, even in the event that anticipated meals are skipped.

Keywords

Diabetes; Artificial pancreas; Behavioral profiles; Linear quadratic Gaussian control; Model predictive control

1. Introduction

Recent national data show that nearly 21 million Americans have diabetes—a lifelong condition that affects people of every age, race, and nationality, and is the leading cause of kidney failure, blindness, and amputations not related to injury. Approximately 1.5 million of these people have Type 1 Diabetes Mellitus (T1DM), in which the immune system

Conflict of interest

^{*}This work was sponsored in part by (1) the Juvenile Diabetes Research Foundation (JDRF) Artificial Pancreas Project and (2) the National Library of Medicine (Award Number T15LM009462). This content is solely the responsibility of the authors and does not necessarily represent the official views of the JDRF, the National Library of Medicine, or the National Institutes of Health.

^{© 2010} Elsevier Ireland Ltd. All rights reserved.

^{*}*Corresponding author at*: Department of Systems and Information Engineering, University of Virginia, 151 Engineers Way, P.O. Box 400747, Charlottesville, VA 22904, United States. Tel.: +1 4349822052. patek@virginia.edu (S.D. Patek).

There are no conflicts of interest for the authors of this article.

destroys the pancreatic beta cells, permanently suppressing insulin secretion. A typical T1DM patient will need 50,000 insulin shots over his/her lifetime, accompanied by testing of blood glucose levels several times a day. Automatic control of blood glucose has been studied for more than three decades and widely different solutions have been proposed. We refer to [2,3] for a review of the early literature, most of which focuses on the case of intravenous (IV) glucose measurement and IV actuation, with both positive (glucose) and negative (insulin) control variables.

As an extension to [1], this paper addresses the challenge to develop effective closed loop systems for regulating blood glucose concentration, via the "SC–SC" route to control, i.e. using commercially available subcutaneous Continuous Glucose Monitoring (CGM) sensors and conventional subcutaneous insulin pump devices. The SC–SC control problem is more challenging than control with intravenous measurement of glucose and insulin infusion, for several reasons. First, because measurement and actuation are achieved via interstitial fluids (and not directly in the bloodstream) there are significant delays associated with this mode of control. State of the art CGM technology relies on sensing glucose in interstitial fluids, which reflects blood glucose with an unknown time lag (between 0 and 45 min [4–6]). Also, the delay introduced by conventional (interstitial) insulin infusion pumps can lead to a limit on control performance, hindering the system's ability to reject major disturbances, specifically meals and exercise.

Model Predictive Control (MPC) has been proposed as a means of overcoming measurement and actuation delay in either IV- or SC-based systems, cf. [2,7,8], where control actions are computed at update intervals based on model predicted optimal future evolution of the system taking into account knowledge about future disturbances. While clinical evaluation of these algorithms is underway, there are clearly many obstacles to automated control of blood glucose *outside of the clinic*, including the need for on-going adaptation to the patient's time-varying physiological parameters and behavior. In this work we address this issue by developing a control methodology that plans insulin delivery given a probabilistic model of meal timing and meal amounts, which can be learned from historical data.

Run-to-run control methods adopted from the chemical process control literature have also been proposed as a means of responding to time-varying aspects of the patient's behavior and physiology. Recent studies have shown that, even with infrequent sampling of blood glucose, run-to-run control can learn bolus and basal rate profiles that are optimal (on average) in the face of time-varying aspects of the patient's daily routine (e.g. timing of meals), cf. [9–11], and circadian fluctuation in insulin resistance (e.g. dawn phenomenon), cf. [12]. The approach developed in this paper differs from run-to-run control in that it employs an explicit probabilistic model of patient meal behavior (including the possibility that the meal will be skipped) and accordingly plans a model-based dynamic response to both measured glucose and anticipated meals.

To outline the sequel, in Section 2, we derive a general control law that optimizes an expected quadratic cost criterion for a linear system subject to a random single shock disturbance, which for this paper represents a potential meal with unknown timing. We set up our Shock-Anticipating Linear Quadratic Regulator (SA-LQR) methodology in Section 3, introducing the notion of a random meal profile characterizing the likelihood, timing, and carbohydrate (CHO) content of the *next* meal and reviewing the glucose–insulin kinetics model used to derive our control law, where state observation is employed. Section 3 also describes a hybrid open-closed loop control strategy, where a closed loop control is used in anticipation of meal arrival, while a meal bolus which is informed by insulin-on-board is delivered at the time of meal arrival. In Section 4, we illustrate the benefits of our control strategy for a breakfast meal scenario. We focus in Section 4 on controller performance in

(2)

cases where the anticipated meal fails to arrive. Throughout, we focus exclusively on the use of insulin for managing blood glucose concentration, ignoring the opportunity to infuse other metabolic hormones such as glucagon, which can provide a so-called "positive" control effect; however, our control methodology can be extended to this setting provided a suitable model of glucagon effect. In Section 5, we compare the SA-LQR strategy to an open loop conventional therapy approach, in which we deliver basal rate continuously and a meal bolus at mealtime. In Section 6, we wrap up the paper with brief conclusions and directions for future work. Finally, we prove the main technical result introduced in Section 2 in the Appendix.

2. Anticipating "shock" disturbances; shock-anticipating LQR (SA-LQR)

In the context of control of T1DM, meals are the primary disturbance. These meal disturbances occur at random times with random magnitudes and can be treated as "shock" disturbances.

Here, we derive a generic optimal linear quadratic control law that anticipates a single random "shock" disturbance that, if it arrives at all, will arrive in stages k = 0 through \bar{k} . The setup here is general: we assume that the plant is linear, time-invariant, with $x(k) \in \mathbb{R}^{n_x}$ (where throughout, x(k) is our best estimate of the state $\hat{x}(k)$, which we formulate in 3.1), $u(k) \in \mathbb{R}^{n_u}$, and $\omega(k) \in \mathbb{R}^{n_d}$. Let $\tilde{d} + w \in \mathbb{R}^{n_d}$ denote the "shock disturbance," which arrives (or not) according to the conditional probability of arrival p_k given that the shock has not arrived prior to stage $k \in \{0, ..., \bar{k}\}$, where \tilde{d} is deterministic and w is a zero-mean perturbation.¹ Models of this type have been studied before. Perhaps the closest related work can be found in [13], which addresses finite-horizon, stochastic hybrid systems, subject to Markovian jump parameters. (We pose an infinite horizon problem, where the "shock" disturbance is distinctly non-Markovian.) The control law is designed to minimize expected quadratic cost associated with state and control deviations away from target levels, where the conditional probabilities of meal arrival p_k factor into an anticipatory effect prior to the next meal.

We assume that the controller is aware of the "shock" disturbance when it arrives, and consequently we augment the state vector x(k) of the dynamical system (18) with an additional "disturbance state" m(k) that provides an indication of whether the disturbance is still possible (m(k) = 0 if "yes," m(k) = 1 if "no"). The augmented dynamical equations are as follows:

$$x(k+1) = Ax(k) + Bu(k) + G\omega(k) \tag{1}$$

$$m(k+1) = \max\{m(k), I_k\}$$

where I_k is the indictor variable for the event that the shock arrives in stage k and the initial state of the process is (x(0), m(0)) = (x, 0) for any initial state space vector x. In practice, the transition from m(k) = 0 to m(k + 1) = 1 (i.e. $\omega(k) \neq 0$ or $k = \overline{k}$) would be observed. (In the context of blood glucose control, the patient would either press a "meal now" button or a meal detection algorithm would provide such an indication.) We set as a control objective the expected value of a quadratic functional involving both state errors and control utilization:

¹In the context of control of T1DM, p_k becomes the probability that the meal will "arrive" in stage k, given that the meal has not yet occurred (this is discussed in more detail in Section 3).

Comput Methods Programs Biomed. Author manuscript; available in PMC 2012 May 1.

$$J(x(0), m(0)) = \mathbb{E}\left\{\sum_{t=0}^{\infty} [x(t)'Qx(t) + u(t)'Ru(t)]\right\}$$
(3)

where Q and R are both symmetric and positive semidefinite and positive definite, respectively. Let C denote the "square root" of Q, such that Q = C'C. Also, for notational convenience in the sequel, let

$$J_k^*(x,m) \equiv \min \mathbf{E} \left\{ \sum_{t=k}^{\infty} [x(t)'Qx(t) + u(t)'Ru(t)] \right\}$$
(4)

where the expectation above is conditional on x(k) = x and m(k) = m. For states where m(k) = 1 the shock disturbance is no longer possible, and the problem of minimizing the infinite horizon quadratic cost-to-go amounts to a standard LQR problem. However, for states where m(k) = 0 the shock may arrive at some future time, and the minimization of expected quadratic costs involves an optimal response in *anticipation* of the disturbance. The basic idea is:

 $J_k^*(x, 1)$ =optimal LQR cost-to-go from x, × with no more future disturbances=x'Px

and

 $J_{k}^{*}(x,0) = \text{optimal cost-to-go from } x, \text{ with anticipation of a$ *possible future disturbance* $} = \min_{u} [E[p_{k}J_{k+1}(Ax+Bu+G\tilde{d},1)]] + [[(1-p_{k})J_{k+1}(Ax+Bu,0)]]$

The following proposition establishes the form of both $J_k^*(x, m)$ and the optimal control law in detail.

Proposition 1. *Given that the pair* (*A*, *B*) *is controllable and the pair* (*A*, *C*) *is observable, the optimal control law is uniquely characterized as follows:*

1. At states where m(k) = 1,

$$J_k^*(x,1) = x' P x \tag{5}$$

and the minimizing control action at state (x, 1) in stage k is

$$u_k^*(x,1) = -Kx \tag{6}$$

where P is the unique positive semidefinite solution to the Riccati equation

$$P = A'(P - PB\Psi^{-1}B'P)A + Q \tag{7}$$

with $\Psi = R + B'PB$, and $K = \Psi^{-1} B'PA$.

2. At states where m(k) = 0, the optimal expected cost-to-go function takes the form

$$J_k^*(x,0) = x'Px + \tilde{d}'H_k x + \tilde{d}'L_k \tilde{d} + \Xi_k$$
(8)

and the minimizing control action at state (x, 0) in stage k is

$$u_k^*(x,0) = -Kx - \tilde{K}_k \tilde{d}$$
⁽⁹⁾

where H_k , L_k , Ξ_k , and $\tilde{K_k}$, d erive from the following.

3. Backwards recursion:

a. Stage $k = \overline{k}$:

$$\tilde{K}_{\overline{k}} = p_{\overline{k}} \Psi^{-1} B' P G \tag{10}$$

$$H_{\overline{k}} = G'[2p_{\overline{k}}P - 2p_{\overline{k}}PB\Psi^{-1}B'P]A$$
⁽¹¹⁾

$$L_{\overline{k}} = G'[p_{\overline{k}}P - p_{\overline{k}}^2 P B \Psi^{-1} B' P] G$$
⁽¹²⁾

$$\Xi_{\overline{k}} = p_{\overline{k}} \mathbb{E}_{w}((G\tilde{w})' P(D\tilde{w})) \tag{13}$$

b. Stage $k < \overline{k}$:

$$\tilde{K}_k = \Psi^{-1} B' Q_k \tag{14}$$

$$H_{k} = 2Q_{k}^{'}(I - B\Psi^{-1}B'P)A$$
(15)

$$L_{k} = p_{k}G'PG + q_{k}L_{k+1} - Q_{k}'B\Psi^{-1}B'Q_{k}$$
(16)

$$\Xi_k = p_k \mathbf{E}_w [w'G'PGw] + q_k \Xi_{k+1} \tag{17}$$

where
$$Q_k = (1 - p_k)(H_{k+1})'/2 + p_k PG$$
.

,

The proof of the proposition employs familiar techniques used in the analysis of linearquadratic control models. We refer to the Appendix B for the formal arguments.

3. SA-LQR with behavioral profiles

Our control strategy is motivated from the realization that historical data about a patient's eating behavior can be used to characterize the timing and likelihood of a meal disturbance

within pre-specified meal regimes (e.g. breakfast), along with the magnitude of the disturbance. We refer for example to [14], which outlines a method for establishing a probabilistic meal-pattern profile for individuals based on CGM data. "Profile" information can be used in planning insulin injections prior to the arrival of meals, as long as the uncertainties involved are appropriately considered. In Section 3.4 we also formulate an insulin-informed control action to be taken *at the time* of the disturbance arrival.

3.1. Random profile for the next meal

As a model for the *next meal* we assume that the patient will experience up to a single meal and that the meal, if it occurs, will involve a random total carbohydrate load of D_{meal} (g CHO) ingested in one sampling period within the discrete interval between stages 0 and \bar{k} . Thus, if a meal is taken at stage k ($k \le \bar{k}$), then $\omega(k) = D_{meal}/\Delta \equiv d_{meaL}$ and $\omega(j) = 0$ for all other stages $j \ne k$. We express d_{meal} as the sum of its expected value d and a zero-mean offset w.

As was introduced in Section 2, p(k) denotes the probability that the meal will "arrive" in stage k, given that the meal has not occurred up to and including stage k - 1. Note that $\omega(k)$ is not a white noise process; it is, in fact, highly correlated in time, since a meal can occur in up to one stage of the process. In practice, the conditional probabilities p_k would come from observation of the patient's eating patterns. For example, a diary of eating behavior can be

used to identify typical meal regimes for individual patients. Note that if $F \equiv \sum_{i=1}^{\overline{k}} f_i < 1$, then the patient occasionally misses the meal and 1 - F is the relative frequency of this happening. (Later, we refer to 1 - F as p(skip), the probability of skipping the meal.) Taking f_i as a model for the probability of the meal being taken in stage *i*, then

 $p_k = f_k / [(1 - F) + \sum_{i=k}^{\overline{k}} f_i]$ for $k \in \{0, ..., \overline{k}\}$. Note that if the meal can only take place in an interval, say $\{k_1, ..., \overline{k}\}$, then we would have $p_k = 0$ for all $k < k_1$. For notational convenience, let $q_k = 1 - p_k$.

3.2. Model of glucose-insulin interaction

The model from which our controller is derived can be traced back to the compartmental "minimal model" of [15]. The model employed here is an 8-state model with extensions which account for subcutaneous oral glucose sensing and actuation, along with oral glucose ingestion. The coefficients of the model reflect population-average glucose–insulin dynamics with a 15 min sample period ($\Delta = 15$). We can express the linearized, discretized implementation of this model in state space form:

$$x(k+1) = Ax(k) + Bu(k) + G\omega(k) \tag{18}$$

where state space matrices *A*, *B*, and *G* are derived from a linearized, discretized form of the model, with model parameters chosen to be representative of an "average" adult patient with T1DM. The matrix values are given in Appendix A.

Additionally, $u(k) = J_{ctrl}(k) - J_{basal}(k)$ (mU/min) is the insulin control signal at stage k, where $J_{ctrl}(k)$ (mU/min) is the current rate of insulin infusion and $J_{basal}(k)$ (mU/min) is the patient's normal basal rate at stage k (as the basal rate can change over time based on the patient's basal rate profile), and $\omega(k) = meal(k) - meal_{ref}$ (mg/min) is the ingested glucose disturbance signal at stage k, where meal(k) (mg/min) is the rate of glucose ingestion at stage k and $meal_{ref}$ (mg/min) is the reference meal input value. We note that this model is employed for use in our meal-to-meal SA-LQR controller and is not intended for more general predictive applications, as it is a population-average, linearized model of glucoseinsulin kinetics.

3.3. Metabolic state observer

Because we cannot measure the state of this system at a given stage k, we estimate the state $\hat{x}(k)$ of the system using a metabolic state observer. The state observer is based on a steady state Kalman filter. Since only CGM measurements are available each minute, it is necessary to compute estimates $\hat{x}(k)$ of x(k) based on the knowledge of infused insulin rate u(k) and measurements y(k), where

$$y(k) = CGM(k) - G_{ref} \tag{19}$$

where CGM(k) (mg/dl) is the readout of the CGM at stage k. We model the measurement signal as

$$y(k) = Cx(k) + v(k) \tag{20}$$

where v(k) (mg/dl) represents CGM signal noise and the matrix C is given by

$$C^{I} = [1, 0, 0, 0, 0, 0, 0, 0]$$
⁽²¹⁾

The metabolic state observer is derived as a Kalman filter based on the state space model of Eq. (18), where, for convenience, we treat $\omega(k)$ (the meal disturbance process) as a zeromean Gaussian process with covariance $R_s = 1$. (We do this despite the fact that in the SA-LQR methodology we treat $\omega(k)$ as a shock disturbance.) We model the sensor noise process as zero-mean Gaussian with covariance $Q_s = .0005$. We point out that, even though meals $\omega(k)$ and sensor noise v(k) are not zero-mean, white, Gaussian processes in reality, the resulting Kalman filter is still a stable state observer. The observer itself can be expressed recursively (as a dynamic process) as

$$\widehat{x}(k|k-1) = A\widehat{x}(k-1|k-1) + Bu(k-1)$$
(22)

$$\widehat{x}(k|k) = \widehat{x}(k|k-1) + L_f(y(k) - C\widehat{x}(k|k-1))$$
(23)

where $\hat{x}(k|k-1)$ refers to the best estimate of $\hat{x}(k)$ using data collected up to stage k - 1, $\hat{x}(k)$ refers to the best estimate of x(k) using data collected up to stage k, the filter gain matrix

$$L_f = AP_f C^T (CP_f C^T + R_s)^{-1}$$
(24)

the estimate update matrix

$$M_f = P_f C^T (C P_f C^T + R_s)^{-1}$$
(25)

the matrix P_f is the unique stabilizing solution to the algebraic Riccati equation

(26)

$$A^{T}PA - A^{T}PG(G^{T}PG + R_{s})^{-1}G^{T}PA + Q_{s} = P$$

In keeping with the discretization used to develop the model, our controller operates by sampling glucose (via CGM) every 15 min and computing a insulin infusion rates via Proposition 1, holding the rate constant between samples. For our control strategy, we apply our best estimate $\hat{x}(k)$ of the state x(k) in Proposition 1, except at the stage of meal arrival, when we deliver an open loop meal bolus as described in Section 3.4 below.

3.4. Hybrid control of diabetes via Proposition 1 and an open loop meal bolus

Section 2 gives us a method for *anticipating* meal arrival. The aggressiveness of this anticipatory effect is driven by the construction of the matrix Q introduced in Eq. (3). In this context, Q is taken to be diagonal, with a weight q > 0 associated with glucose states and a weight of zero associated with meal disturbance states; we set R = 1. The weight parameter q is a tuned patient-specific parameter (the only such parameter in the controller) that determines the aggressiveness with which the controller rejects glucose deviations away from the target operating point (110 mg/dl). At the time of meal arrival, we switch our control to an open loop strategy which is consistent with the action taken by a patient in a conventional therapy setting. At the time of meal arrival, assumed to be known, we deliver an open loop meal bolus which takes into account the insulin that has been delivered in anticipation of the meal, but has yet to act. This insulin is referred to as insulin-on-board (IOB) and we compute the IOB at stage k as follows:

$$IOB(k) = IOB_{curve} \cdot [u(k), u(k-1), u(k-2), \dots, u(k-24)]'$$
(27)

where IOB_{curve} is a vector containing entries which correspond to the proportion of a given insulin injection remaining at stage k (adapted from [16]) and [u(k), u(k-1), u(k-2), ..., u(k-24)] is a vector which contains 6 h of insulin injection history with a sample time of 15 min, where IOB(k) is the assessment of the insulin-on-board above the basal rate at stage k. Using this information, the meal bolus, U(k) (U), given at meal time is computed as:

$$U(k) = d(meal) \cdot CR - IOB(k) \tag{28}$$

where d(meal) is the amount of the meal taken at stage k (in g CHO) and CR is the carbohydrate ratio of the subject (U/g CHO). (We note that the total amount of the meal bolus will be $\Delta \cdot U(k)$, as we hold the injection constant over the Δ -minute interval of stage k. This meal bolus covers a meal of size $D(meal) = d(meal) \cdot \Delta$, where Δ in this formulation is 15 min.) Note that Eq. (28) determines the size of the meal bolus based in part on how much anticipatory insulin has been delivered and when it is delivered. The total amount of insulin can vary from subject to subject and according to different values of p(skip). In the following section, we evaluate the controller performance where we implement our anticipatory closed loop strategy prior to meals and the open loop meal bolus described by Eq. (28) at mealtime.

4. In silico preclinical trials

In this section, we evaluate the "hybrid" SA-LQR algorithm using a computer simulation environment based on the oral glucose "meal model" of [17,18] equipped with a population of 100 *in silico* adult patients with T1DM, focusing particularly on the controller's ability to anticipate a breakfast meal without endangering the patient if the meal is skipped. Because

the meal model parameters were developed under the assumption of a fixed "mixed" meal composition, we assume that meals are composed of a fixed percentage of protein, fat, and carbohydrates (45% carbohydrate, 15% protein, and 40% fat). The controller that we test is derived from an instance of the model of Eq. (18), representative of an "average" patient with T1DM. In keeping with the discretization used to develop the model, our controller operates by sampling glucose (via CGM) every 15 min and computing insulin infusion rates via Proposition 1, holding the rate constant between samples. In Section 4.3, we show illustrative results for a representative subject with sensor noise. In Section 4.4, we show aggregate results for the population of 100 *in silico* subjects. All simulations are conducted with sensor noise, where we use the CGM model proposed in [19], feeding the corrupted CGM data into the controller. Section 5 compares the SA-LQR strategy to an open loop conventional therapy approach to insulin delivery.

4.1. Experimental scenario

The *in silico* experiments presented here reveal performance characteristics of the controller in a protocol focused on the anticipation of the breakfast meal. The protocol runs from midnight to 12:30 p.m. and includes one meal at time 420 min (7 a.m.). The breakfast meal, if it arrives at all, will take place at minute 420 of the protocol and will involve 55 g CHO.

4.2. Profiles for the breakfast meal

Table 1 shows the meal profile information used in the backwards recursion of Proposition 1 in planning insulin delivery prior to breakfast. Note that the profile for breakfast describes an earliest-possible meal arrival time of 240 min and a latest possible meal time of 525 min. The relative frequency distribution f_k within this interval is derived from a clipped Gaussian distribution centered at the most likely time for breakfast, 420 min. The standard deviation of the Gaussian distribution before clipping is 30 min. The distribution f_k is normalized to achieve different probabilities of skipped breakfast p(skip), as shown in the table, ranging from .0001 (where the meal is almost certain to arrive) to .5 (where the meal arrives or not with equal probability).

The breakfast meal regime starts at minute 0, and at this time uses the backwards recursion of Proposition 1 to compute a feedback gain matrix *K* and set of feedforward gains \tilde{K}_k for the 15 min sampling intervals *k* through either (1) the last possible time of the meal, if the meal has not arrived, or (2) the time of the breakfast meal arrival, at which time the insulin bolus is delivered. For the purpose of comparison, following (1) the last possible time of the meal or (2) the time of meal arrival and delivery of the open loop bolus, we return the subject to basal rate designed to hold the subject at a blood glucose concentration of 112.5 mg/dl. While this scenario only presents the anticipation and arrival of one meal, another meal regime (lunch, in this case), begins either (1) at the stage following last possible time of the breakfast meal regime in the event that breakfast is skipped, or (2) at the stage just following arrival of the breakfast meal.

4.3. Case study results for a representative subject

Here we illustrate the SA-LQR controller as applied to a representative subject within the *in silico* population of adult patients with T1DM, using the experimental scenario and meal profile information above. Of course, performance of the controller will depend strongly on the accuracy of the sensor. Fig. 1 illustrates the case where the meal arrives at minute 420, with the top plot showing the rate of insulin delivery (U/h) and the bottom plot showing BG (mg/dl). Fig. 2 shows the plots of insulin delivery rate and BG for the same subject in the case of a skipped meal, where we focus on the anticipatory insulin effect.

Notice that in either scenario, whether breakfast is taken or skipped, a profile with p(skip) = .5 injects less insulin in anticipation of the meal. Because there is high level of uncertainty associated with the meal occurrence, there is less anticipatory insulin delivery to guard against a skipped meal scenario. Thus in the case when p(skip) = .5 and the meal is in fact skipped, the subject is able to maintain a safe level of blood glucose, while in the case of p(skip) close to zero, the subject experiences a lower glucose minimum and is at greater risk of hypoglycemia. In any case, there is some level of anticipatory insulin delivery which precedes the meal bolus delivery. Table 2 gives the meal bolus amounts for the breakfast meal when it does arrive.

Referring to Table 2, in the case where p(skip) = .5 and the meal does occur, the subject experiences the largest meal bolus delivery at meal time, to compensate for insulin that was not delivered in anticipation of the meal; that is, IOB assessment is smallest for the case of p(skip) = .5 at the time of meal arrival. In the case of p(skip) close to zero, a more aggressive anticipatory insulin delivery results in a larger IOB value at the time of meal arrival, thus yielding the smallest meal bolus at mealtime. Note that the insulin trace of Fig. 2 (for the case the that breakfast meal is skipped) shows a gradual decline at minute 450 lasting to the end of the breakfast meal regime, which can be explained by the fact that the controller recognizes that the most likely time of the meal has passed, and the chance of the breakfast meal occurring after this most likely time gradually decreases until minute 525, the last possible arrival time for breakfast.

The above representative subject results suggest that, by designing the controller around a meal profile with a significant probability that the meal will be skipped, we can obtain control postprandial performance without the threat of severe hypoglycemia when the skipped-meal event actually occurs. In the case of a skipped meal, we are considering the most extreme scenario. We hypothesize that the case of a delayed meal would lead to results which fall somewhere between the skipped meal and the nominal scenario (when the anticipated meal amount is taken at the most likely time). Alternatively, if a meal is taken which is not accounted for in our behavioral profile, the linear quadratic Gaussian controller will inject insulin in reaction, as opposed to anticipation, to an increase in the blood glucose level.

The results of Section 4.4 confirm this observation for an *in silico* population of 100 adult patients with T1DM.

4.4. Population study results

Here we illustrate the SA-LQR controller applied to the full *in silico* population of adults with T1DM, again with sensor noise. Table 3 summarizes the results of the study reporting max and min glucose concentrations from the onset of the protocol to just following the postprandial peak of the breakfast meal (focusing on controller performance leading up to and just after breakfast).

Looking at Table 3, we see that for larger values for p(skip), the controller delivers less insulin in anticipation of the meal on average, a strategy that protects the subject from hypoglycemia in the event that the meal is skipped. This can be seen in the mean glucose minimum of 90.81 mg/dl in Table 3 for p(skip) = .5. Because we compensate for a less aggressive anticipatory insulin delivery by delivering a larger meal bolus at meal time (when the meal arrives) in the case of p(skip) = .5, we avoid a higher postprandial peak that could result from this weaker anticipatory insulin delivery. Alternatively, when we consider the results for p(skip) close to zero, results indicate that when the meal is skipped, a greater risk of hypoglycemia results (a mean glucose min of 62.22 mg/dl). A p(skip) close to zero represents an essentially deterministic meal profile, where we control with near certainty the

timing and the presence of the meal. Because the controller anticipates that a meal will occur with near certainty, this skipped meal poses an unexpected lack of meal disturbance for the controller. From this population study, we can infer that some optimal value of p(skip) exists so as to control both hypoglycemia (in the event that no meal occurs), as well as hyperglycemia (when a meal does occur). By viewing the probability of a skipped meal in this way, p(skip) is a parameter that can be tuned to optimize control. From the perspective of hypoglycemia avoidance and safety, we choose a larger p(skip) when there is uncertainty as to whether or not a subject will eat in a given regime. Thus choosing a larger p(skip)allows for us to control in a safe way while also taking advantage of our ability to anticipate the arrival of a meal to improve postprandial performance.

Fig. 3 presents population outcomes using the Control-Variability Grid Analysis technique of [20]. The plot shows each patient's peak post-breakfast glucose concentration (given that breakfast arrives) plotted against the minimum glucose concentration associated with a skipped breakfast, for p(skip) = .0001 (blue) versus p(skip) = .5 (pink). (Each patient appears twice, once as a pink dot and once as a blue dot.) Despite the uncertainty about whether breakfast will arrive, a significant fraction of the population lands in the A-zone of the plot where peak glucose is below 180 mg/dl (after breakfast) and above 80 mg/dl (when breakfast is skipped).

5. In silico comparison of SA-LQR to conventional therapy

In this section, we compare results obtained when we employ the SA-LQR strategy to those results obtained when a patient is subjected to open loop therapy. We consider again the scenario presented in Section 4, except that in this case, we investigate variability in the standard deviation of the meal arrival time, while keeping p(skip) fixed at a value of p(skip) = .1 (Table 4).

In this section, we consider only the case when the breakfast meal is taken. The purpose of this comparison is to demonstrate that, when we employ anticipatory insulin delivery in a closed loop control strategy, we can improve glucose performance by reducing the postprandial glucose concentration when compared to an open loop strategy. The open loop control strategy applied here is one in which basal insulin is delivered continuously to maintain a steady state glucose concentration of 112.5 mg/dl and at mealtime, a meal bolus is delivered. The amount of this bolus (U) is computed based on the known meal amount (g CHO) and the subject's carbohydrate ratio (U/g CHO). This computation is described by Eq. (28), where, in the case of open loop therapy, IOB = 0 at the time of the meal bolus (as IOB measures insulin-on-board *above* the basal insulin).

5.1. Case study results for a representative subject

Fig. 4 shows a comparison of the SA-LQR strategy to the open loop strategy.

Fig. 4 indicates that, in the case when SA-LQR is employed, there is a "ramping up" in the insulin delivery prior to the meal arrival. This increase in the insulin delivery rate prior to the meal arrival drives the glucose concentration down prior to the arrival of the breakfast meal (for this representative subject, the open loop pre-meal glucose concentration is 112.5 mg/dl, compared to the SA-LQR pre-meal glucose concentration of 86.5 mg/dl). The result of this glucose lowering prior to the meal in the case of the SA-LQR strategy is a lower postprandial peak (as compared to the open loop case). In Section 5.2, we confirm these results for the *in silico* adult population.

5.2. Population study results

For the population comparison, mean minimum and maximum BG concentrations (mg/dl) are computed as mean values for the population of 100 adults, where these values are computed for the time prior to the breakfast meal (to obtain the minimum value) and the postprandial peak following the meal arrival (to obtain the maximum value).

Table 5 shows only a small deviation between the mean min and max glucose concentration in the case of the SA-LQR for the three values of σ . For the three cases of σ , there is a small decrease in the mean min and max BG as the value of σ increases. This is a result of the anticipatory insulin delivery beginning earlier in the case of larger σ , as in this case it is more likely that the meal will occur at the earlier stages of the meal regime than in the case of smaller σ values. This being said, the lower minimum BG in the case of larger σ values is slight; this is a result of the fact that, while the meal may occur at the early stages of the meal regime and we need to prepare for that through insulin delivery, it is also more likely that the meal will arrive at the later stages of the meal regime than in the case of smaller σ values. Thus for larger values of σ , the insulin delivery has a greater spread over the meal regime. In the case of a σ value of 1 min, the meal, if it does occur, will occur with near certainty in minute 420; thus anticipation of the breakfast in this case is concentrated over the most likely breakfast arrival time. In contrast to this, the distinction of mean min and max glucose values between the SA-LQR and the open loop strategy is much greater. This is a result of the anticipatory insulin delivery: prior to the meal arrival, the open loop BG (as we expect) is 109.84 mg/dl, while in the case of SA-LQR, anticipatory insulin drives the BG concentration to approximately 96 mg/dl. In the case of the open loop therapy, delivering basal rate insulin and an open loop bolus at meal time keeps the subject's glucose level at or above 110 mg/dl (the basal steady state level). This translates to a 13 mg/dl difference in the postprandial peaks, where a higher postprandial peak results when no pre-meal insulin is given in the case of open loop therapy. The CVGA results shown in Fig. 5 correspond to the results of Table 5, where we compare a σ of 30 min to the open loop strategy.

6. Conclusions

In [1], we used the notion of a meal profile to develop a model-based stochastic predictive control law for optimal insulin delivery in anticipation of meals. Our *in silico* study shows that being able to account for uncertainty as to *if* and *when* a meal disturbance will occur allows us to achieve reasonable tradeoffs between aggressive anticipatory insulin treatment and the risk of hypoglycemia in the event that meals are skipped. In this paper, we have extended the methodology of [1] to allow the application of a state space model of glucose insulin kinetics with state estimation via a Kalman filter, and we have modified the postprandial action of the controller by introducing an open loop meal bolus that is calibrated to pre-meal insulin "on-board." A linear control law applied to nonlinear system dynamics is shown to perform well in this setting. We have also demonstrated the ability of the control strategy to work in a hybrid environment; anticipating insulin delivery through a closed loop control strategy and delivering an insulin-on-board- informed meal bolus at the time of the meal.

The work presented here uses a model-based stochastic predictive control strategy to anticipate the arrival of meals. In effect, this strategy leads to an offset from the glucose target concentration of 110 mg/dl prior to the meal, with a greater offset associated with increased certainty as to whether not the meal will arrive at all. In this way, the SA-LQR strategy can be viewed as similar to a PID controller with time-varying setpoints chosen to reflect certainty about meal arrival and timing. The "effective setpoint reduction" of SA-LQR is achieved in a systematic fashion based on the patient's behavioral profile. In this sense, the SA-LQR performs optimally under the level of uncertainty we have regarding

future disturbances. Future work hopes to address the way that p(skip) is chosen for each subject, or for each meal regime. We are now considering an implementation of p(skip) in which the value of p(skip) is "learned" for each subject based on the subject's history of past behavior. In this case, the value of p(skip) is also something that can vary over the course of 1 day, or from 1 day to the next.

This work presents a method for anticipating meals through the use of a behavioral profile. In future work we plan to more fully evaluate controller performance in the face of uncertainty about meal amounts and meal timing. (While Proposition 1 allows for all three types of uncertainty, the simulation results in this paper only considers uncertainty in meal arrival and meal timing.) The controller developed in this paper only looks ahead to the next (potential) meal, one meal regime at a time. In on-going efforts we are extending the methodology to plan over multiple meal regimes, so that the anticipated response to a lunch meal could be influenced by the anticipated response to breakfast, and vice versa. In conjunction with this effort, we plan to reexamine the form and structure of the meal profile itself, thinking of the model of Section 3.1 as a first step toward a more comprehensive description of patient behavior.

Appendix A. State space matrices for the model of glucose-insulin kinetics

		.9913	-102.7	-1.5×10^{-8}	-2.89×10^{-6}	-4.1×10^{-4}	0	2.01×10^{-6}	4.3×10^{-5}
Α	=	0	.839	5.23×10^{-10}	7.44×10^{-8}	6.84×10^{-6}	0	0	0
		0	0	.9798	0	0	0	0	0
		0	0	.0200	.9798	0	0	0	0
		0	0	1.9×10^{-4}	.0180	.7882	0	0	0
		.0865	-4.667	-2.73×10^{-10}	-6.59×10^{-8}	-1.26×10^{-5}	.9131	6.00×10^{-8}	1.90×10^{-6}
		0	0	0	0	0	0	.9083	0
		0	0	0	0	0	0	.09115	.9891
B^T	$B^{T} = [-3.05 \times 10^{-9}, 1.34 \times 10^{-10}, .9900, .0100, 6.50 \times 10^{-5}, -4.61 \times 10^{-11}, 0, 0]$ $G^{T} = [6.76 \times 10^{-7}, 0, 0, 0, 0, 1.52 \times 10^{-8}, 9534, 0464]$								
G^T									

Appendix B. Proof of Proposition 1

The case where m(k) = 1 is straightforward. Under the assumptions of the proposition, it is clear (e.g. from [DPvol1]), that $J_k^*(x, 1) = x'Px'$ and $u_k^*(x, 1) = -Kx$, where *P* and *K* are as specified. The result for m(k) = 0 follows by mathematical induction, as argued in the following lemmas.

Lemma 2. For $k = \overline{k}$, the expected cost to go function $J_{\overline{k}}^*(x, 0)$ is of the form

$$J_{\overline{k}}^{*}(x,0) = x' P x + \tilde{d}' H_{\overline{k}} x + \tilde{d}' L_{\overline{k}} \tilde{d} + \Xi_{\overline{k}}$$
(B.1)

where

1.
$$H_{\bar{k}} = G'[2p_{\bar{k}}P - 2p_{\bar{k}}PB(R + B'PB)^{-1}B'P]A,$$

2. $L_{\bar{k}} = G'[p_{\bar{k}}P - p_{\bar{k}}^{2}PB(R + B'PB)^{-1}B'P]G,$
3. $\Xi_{\bar{k}} = p_{\bar{k}} \mathbf{E}_{w}((G\tilde{w})'P(D\tilde{w})),$

and the optimal control action at this state is

where $\Psi = R + B'PB$.

Proof. By straightforward dynamic programming:

$$\begin{split} J_{\overline{k}}^{*}(x,0) &= \min_{u} [x'Qx + u'Ru + p_{\overline{k}} \mathbf{E}_{w}(J_{\overline{k+1}}^{*}(Ax + Bu + G(\tilde{d} + w), 1)) + (1 - p_{\overline{k}})J_{\overline{k+1}}^{*}(Ax + Bu, 1)] \\ &= \min_{u} [x'Qx + u'Ru + p_{\overline{k}} \mathbf{E}_{w}((Ax + Bu + G(\tilde{d} + w))' \times P(Ax + Bu + G(\tilde{d} + w))) + (1 - p_{\overline{k}})(Ax + Bu)' \times P(Ax + Bu)] \\ &= \min_{u} [x'Qx + u'Ru + (Ax + Bu)'P(Ax + Bu) + 2p_{\overline{k}} \mathbf{E}_{w}((G(\tilde{d} + w))'P(Ax + Bu)) + p_{\overline{k}} \mathbf{E}_{w}((G(\tilde{d} + w))' \times PG(\tilde{d} + w))] \\ &= \min_{u} [x'Qx + u'Ru + (Ax + Bu)'P(Ax + Bu) + 2p_{\overline{k}}(G\tilde{d})'P(Ax + Bu) + p_{\overline{k}}(G\tilde{d})'P(D\tilde{d}) + 2p_{\overline{k}} \mathbf{E}_{w}((G\tilde{w})'P(D\tilde{d})) + p_{\overline{k}} \mathbf{E}_{w}((G\tilde{w})'P(D\tilde{w}))] \\ &= \min_{u} [x'Qx + u'Ru + (Ax)'P(Ax) + 2(Ax)'P(Bu) + (Bu)'P(Bu) + 2p_{\overline{k}}(G\tilde{d})'P(Ax) + 2p_{\overline{k}}(G\tilde{d})'P(Bu) + p_{\overline{k}}(G\tilde{d})'P(D\tilde{d}) + \Xi_{\overline{k}}]. \end{split}$$

(B.3)

Taking the derivative with respect to u and setting equal to zero, we get the optimal control

$$u_{\overline{k}}^{*}(x,0) = Argmin_{u} [x'Qx + u'Ru + p_{\overline{k}} \mathbb{E}_{w}(J_{\overline{k}+1}^{*}(Ax + Bu + G(\tilde{d}+w), 1)) + (1 - p_{\overline{k}})J_{\overline{k}+1}^{*}(Ax + Bu, 1)] \\ = -(R + B'PB)^{-1}(B'PAx + p_{\overline{k}}B'PG\tilde{d})$$
(B.4)

Plugging back into the expression above, we get

$$\begin{split} J^*_{\overline{k}}(x,0) &= [x'Qx + x'A'PAx + 2p_{\overline{k}}\tilde{d}'G'PAx + p_{\overline{k}}\tilde{d}'G'PG\tilde{d} + u^*_{\overline{k}}(x,0)'(R+B'PB)u^*_{\overline{k}}(x,0) + 2(x'A'PB + p_{\overline{k}}\tilde{d}'G'PB)u^*_{\overline{k}}(x,0) + \Xi_{\overline{k}}] \\ &= [x'Qx + x'A'PAx + 2p_{\overline{k}}\tilde{d}'G'PAx + p_{\overline{k}}\tilde{d}'G'PG\tilde{d} - (x'A'PB + p_{\overline{k}}\tilde{d}'G'PB)(R+B'PB)^{-1} \times (B'PAx + p_{\overline{k}}B'PG\tilde{d}) + \Xi_{\overline{k}}] \\ &= [x'Qx + x'A'PAx - x'A'PB(R+B'PB)^{-1}B'PAx + 2p_{\overline{k}}\tilde{d}'G'PAx - 2p_{\overline{k}}\tilde{d}'G'PB(R+B'PB)^{-1}B'PAx + p_{\overline{k}}d'G'PG\tilde{d} - p^2_{\overline{k}}\tilde{d}'G'PB(R+B'PB)^{-1}B'PAx + p_{\overline{k}}d'G'PG\tilde{d} - p^2_{\overline{k}}\tilde{d}'G'PB(R+B'PB)^{-1}B'PAx + p_{\overline{k}}d'G'PG\tilde{d} - p^2_{\overline{k}}\tilde{d}'G'PB(R+B'PB)^{-1}B'P]Ax + \tilde{d}'G'[2p_{\overline{k}}P - 2p_{\overline{k}}PB(R+B'PB)^{-1}B'P]Ax + \tilde{d}'G'[p_{\overline{k}}P - p^2_{\overline{k}}PB(R+B'PB)^{-1}B'P]G\tilde{d} + \Xi \\ &= x'Px + \tilde{d}'H_{\overline{k}}x + \tilde{d}'L_{\overline{k}}\tilde{d} + \Xi_{\overline{k}}, \end{split}$$

(B.5)

where we have used the fact that $P = Q + A'[P - PB(R + B'PB)^{-1}B'P]A$.

Lemma 3. For $k < \overline{k}$, the expected cost to go function $J_k^*(x, 0)$ is also of the form:

$$J_k^*(x,0) = x' P x + \tilde{d}' H_k x + \tilde{d}' L_k \tilde{d} + \Xi_k$$
(B.6)

where

1.
$$H_k = 2Q'_k (I - B\Psi^{-1}B'P)A$$

Comput Methods Programs Biomed. Author manuscript; available in PMC 2012 May 1.

(B.2)

2.
$$L_k = p_k G' P G + (1 - p_k) L_{k+1} - Q'_k B \Psi^{-1} B' Q_k (symmetric)$$

3. $\Xi_k = p_k \mathbf{E}_w [w'G' P G w] + (1 - p_k) \Xi_{k+1}.$
and
4. $Q_k = (1 - p_k) (H_{k+1})'/2 + p_k P G.$

The optimal policy is

$$u_k^*(x,0) = -\Psi^{-1}B'[PAx + Q_k\tilde{d}]$$
(B.7)

Proof. Using the induction hypothesis:

$$J_{k}^{*}(x,0) = \min_{u} \left[x'Qx + u'Ru + p_{\overline{k}} \mathbf{E}_{w}(J_{k+1}^{*}(Ax + Bu + G(\overline{d} + w), 1)) + (1 - p_{\overline{k}})J_{k+1}^{*}(Ax + Bu, 0) \right]$$
(B.8)

 $J_{k}^{*}(x,0) = \min_{u} [x'Qx+u'Ru+p_{\overline{k}}\mathbf{E}_{w}[(Ax+Bu+G(\tilde{d}+w))' \times P(Ax+Bu+G(\tilde{d}+w))] + (1-p_{\overline{k}}) \times [(Ax+Bu)'P(Ax+Bu)+\tilde{d}'H_{k+1}(Ax+Bu)+\tilde{d}'1] \\ = \min_{u} [x'Qx+u'Ru+p_{\overline{k}}[(Ax+Bu)'P(Ax+Bu)+2\mathbf{E}_{w}((G(\tilde{d}+w))'P(Ax+Bu)) + \mathbf{E}_{w}((G(\tilde{d}+w))'PG(\tilde{d}+w))] + (1-p_{\overline{k}})[(Ax+Bu)'P(Ax+Bu)'P(Ax+Bu)] \\ = \min_{u} [x'Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d} + \mathbf{E}_{w}((Gw)'PGw)] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+u'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(G\tilde{d})'P(Ax+Bu) + (G\tilde{d})'PG\tilde{d}] + (1-p_{\overline{k}})[\tilde{d}'H_{k+1}(Ax+Bu) + \tilde{d}'L_{k+1}] \\ = \min_{u} [x'(Qx+U'Ru+(Ax+Bu)'P(Ax+Bu) + p_{\overline{k}}[2(Q'A'P)] + (1-p_{\overline{k}})] \\ = \min_{u} [x'(Qx+U'Ru+(Ax+Bu) + p_{\overline{k}}[2(Q'A'P)] + (1-p_{\overline{k}})] \\ = \min_{u} [x'(Qx+U'Ru+(Ax+Bu) + p_{\overline{k}}[2(Q'A'P)] + (1-p_{\overline{k}})] \\ = \min_{u} [x'(Qx+U'Ru+(Ax+Bu) + p_{\overline{k}}] \\ = \min_{u} [x'(Qx+U'Ru+(Ax+Bu) + p_{\overline{k}}] \\ = \min_{u} [x'(Qx+U'Ru+(Ax+Bu) + p_{\overline{k})] \\ = \min_{$

(B.9)

Taking the derivative with respect to *u* and setting equal to zero, we get

 $2\Psi u + 2B' [PAx + Q_k \tilde{d}] = 0$

so that $u_k^*(x, 0) = -\Psi^{-1}B'[PAx+Q_k\tilde{d}]$. Plugging into the expression above, we get

$$\begin{split} J_{k}^{*}(x,0) &= [x'(Q+A'PA)x + (x'A'P+\tilde{d}'Q'_{k})B\Psi^{-1}\Psi\Psi^{-1}B'(PAx+Q_{k}\tilde{d}) - (2x'A'P+2\tilde{d}'Q'_{k})B\Psi^{-1}B'(PAx+Q_{k}\tilde{d}) + 2\tilde{d}'Q'_{k}Ax + \tilde{d}'(p_{k}G'PG+(1-p_{k})L_{k+1})\tilde{d} + \Xi_{k}] \\ &= [x'(Q+A'PA)x - (x'A'P+\tilde{d}'Q'_{k})B\Psi^{-1}B'(PAx+Q_{k}\tilde{d}) + 2\tilde{d}'Q'_{k}Ax + \tilde{d}'(p_{k}G'PG+(1-p_{k})L_{k+1})\tilde{d} + \Xi_{k}] \\ &= [x'(Q+A'PA)x - x'A'PB\Psi^{-1}B'PAx + 2\tilde{d}'Q'_{k}Ax - 2x'A'PB\Psi^{-1}B'Q_{k}\tilde{d} + \tilde{d}'(p_{k}G'PG+(1-p_{k})L_{k+1})\tilde{d} - \tilde{d}'Q'_{k}B\Psi^{-1}B'Q_{k}\tilde{d} + \Xi_{k}] \\ &= [x'[Q+A'(P-PB\Psi^{-1}B'P)A]x + \tilde{d}'[2Q'_{k}(I-B\Psi^{-1}B'P)A]x + \tilde{d}'[p_{k}G'PG+(1-p_{k})L_{k+1} - Q'_{k}B\Psi^{-1}B'Q_{k}\tilde{d} + \Xi_{k}] \\ &= x'Px + \tilde{d}'H_{k}x + \tilde{d}'L_{k}\tilde{d} + \Xi_{k} \end{split}$$

(B.10)

where again we have used the fact that $P = Q + A'[P - PB(R + B'PB)^{-1}B'P]A$.

REFERENCES

- Patek, S.; Hughes, C.; Breton, M.; Kovatchev, B. Anticipating meals with behavioral profiles: towards stochastic model predictive control of T1DM; Proceedings of the 7th IFAC Symposium on Modelling and Control in Biomedical Systems; 2009. p. 37-42. Available in IFAC Papers-OnLine: http://www.ifac-papersonline.net/Detailed/39984.html
- Parker RS, Doyle FJ, Peppas NA III. A model-based algorithm for blood glucose control in type 1 diabetic patients. IEEE Trans. Biomed. Eng. 1999; BME-46(2):148–157. [PubMed: 9932336]
- 3. Bequette BW. A critical assessment of algorithms and challenges in the development of a closed-loop artificial pancreas. Diab. Technol. Ther. 2005; 7:28–47.
- Boyne M, Silver DM, Kaplan J, Saudek C. Timing of changes in interstitial and venous blood glucose measured with a continuous subcutaneous glucose sensor. Diabetes. 2003; 52:2790–2794. [PubMed: 14578298]
- Rebrin K, Steil G, van Antwerp W, Mastrototaro J. Subcutaneous glucose predicts plasma glucose independent of insulin: implications for continuous monitoring. Am. J. Physiol. Endocrinol. Metab. 1999; 277:561–571.
- Kulcu E, Tamada J, Reach G, Potts R, Lesho M. Physiological differences between interstitial glucose and blood glucose measured in human subjects. Diabetes Care. 2003; 26:2405–2409. [PubMed: 12882870]
- Hovorka R, Canonico V, Chassin LJ, Haueter U, Massi-Benedetti M, Federici MO, Pieber T, Schaller H, Schaupp L, Vering T, Wilinska ME. Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. Physiol. Meas. 2004; 25:905–920. [PubMed: 15382830]
- Magni, L.; Raimondo, DM.; Dalla Man, C.; Nicolao, GD.; Kovatchev, B.; Cobelli, C. Model predictive control of glucose concentration in subjects with type 1 diabetes: an in silico trial; 17th IFAC World Congress; 2008. p. 4246-4251.
- Zisser H, Jovanovic L, Doyle FJ, Ospina P, Owens C III. Run-to-run control of meal-related insulin dosing. Diab. Technol. Ther. 2005; 7:48–57.
- Owens C, Zisser H, Jovanovic L, Srinivasan B, Bonvin D, Doyle FJ III. Run-to-run control of blood glucose concentrations for people with type 1 diabetes mellitus. IEEE Trans. Biomed. Eng. 2006; 53:996–1005. [PubMed: 16761826]
- Palerm C, Zisser H, Bevier W, Jovanovic L, Doyle FJ III. Prandial insulin dosing using run-to-run control: application of clinical data and medical expertise to define a suitable performance metric. Diabetes Care. 2007; 30:1131–1136. [PubMed: 17303792]
- Palerm C, Zisser H, Jovanovic L, Doyle FJ III. A run-to-run control strategy to adjust basal insulin infusion rates in type 1 diabetes. J. Process Control. 2008; 18:258–265. [PubMed: 18709180]
- Campo L, Bar-Shalom Y. Control of discrete-time hybrid stochastic systems. IEEE Trans. Autom. Control. 1992; 37(10):1522–1527.
- Patek, S.; Breton, M.; Cobelli, C.; Dalla Man, C.; Kovatchev, B. Adaptive meal detection algorithm enabling closed-loop control in type 1 diabetes; Proceedings of the 2007 Diabetes Technology Meeting (Abstract); 2007.
- Bergman R, Ider Y, Bowden C, Cobelli C. Quantitative estimation of insulin sensitivity. Am. J. Physiol. 1979; 236:E667–E677. [PubMed: 443421]
- 16. Swan KL, Dziura JD, Steil GM, Voskanyan GR, Sikes KA, Steffen AT, Martin ML, Tamborlane WV, Weinzimer SA. Effect of age of infusion site and type of rapid-acting analog on pharmacodynamic parameters of insulin boluses in youth with type 1 diabetes receiving insulin pump therapy. Diabetes Care. 2009; 32:240–244. [PubMed: 19017777]
- Dalla Man C, Rizza RA, Cobelli C. Meal simulation model of the glucose–insulin system. IEEE Trans. Biomed. Eng. 2007; 54(10):1740–1749. [PubMed: 17926672]
- Dalla Man C, Raimondo DM, Rizza RA, Cobelli C. GIM, simulation software of meal glucose– insulin model. J. Diabetes Sci. Technol. 2007; 1(3):323–330. [PubMed: 19885087]
- Breton M, Kovatchev B. Analysis, modeling, and simulation of the accuracy of continuous glucose sensors. J. Diabetes Sci. Technol. 2008; 2(5):1–10. [PubMed: 19885170]

 Magni L, Raimondo DM, Dalla Man C, Breton M, Patek SD, Nicolao GD, Cobelli C, Kovatchev BP. Evaluating the efficacy of closed-loop glucose regulation via control-variability grid analysis. J. Diabetes Sci. Technol. 2008; 2(4):630–635. [PubMed: 19885239]



Fig. 1. Representative subject: meal taken.



Fig. 2. Representative subject: meal skipped.

NIH-PA Author Manuscript

Hughes et al.



Fig. 3.

CVGA results for 100 adults comparing p(skip) = .0001 (blue) to p(skip) = .5 (pink); for p(skip) = .0001: (B) 22%, (C) 74% and (D) 4%; for p(skip) = .5: (A) 46% and (B) 54%. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)



Fig. 4.

Representative subject: comparing various values of σ in SA-LQR to open loop therapy.

Hughes et al.



Fig. 5.

CVGA comparing SA-LQR, Std. Dev. 30 min (blue) to open loop (pink); for SA-LQR: (A) 75% and (B) 25%; for open loop: (A) 88% and (B) 12%. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Table 1

Meal profile data for the experiments.

	Breakfast
Meal size (g CHO)	55
Earliest meal time (min)	240
Most likely meal time (min)	420
Latest meal time (min)	525
Std. Dev. of meal time	30
p(skip)	{.0001,.01,.1,.5}

Table 2

Meal bolus insulin amounts (U).

p(skip)	Meal bolus (U)
.0001	8.51
.01	8.53
.1	8.75
.5	9.71

Table 3

Results of the population study

	Meal skipped mean pre-breakfast min (std deviation)	Meal taken mean post-breakfast max (std deviation)
Noisy CGM		
p(skip) = .0001	62.22 (10.38)	143.90 (26.75)
p(skip) = .01	63.70 (10.34)	144.35 (27.09)
p(skip) = .1	73.56 (9.22)	144.86 (27.10)
p(skip) = .5	90.81 (6.58)	149.53 (27.78)

Table 4

Meal profile data for the experiments.

	Breakfast
Meal size (g CHO)	55
Earliest meal time (min)	240
Most likely meal time (min)	420
Latest meal time (min)	525
Std. Dev. of meal time (min)	1, 30, 60
p(skip)	.1

Table 5

Results for population comparison.

	Std. Dev. 1	30	60	Open loop
Noisy CGM				
Min BG	96.72	96.25	94.77	109.84
Max BG	145.23	144.87	144.22	156.98