Published in final edited form as:

Proc Am Control Conf. 2015 July ; 2015: 1635–1640. doi:10.1109/ACC.2015.7170967.

Velocity-weighting to prevent controller-induced hypoglycemia in MPC of an artificial pancreas to treat T1DM

Ravi Gondhalekar, Eyal Dassau, and Francis J. Doyle III

Department of Chemical Engineering, University of California Santa Barbara (UCSB), USA.

Abstract

The design of a Model Predictive Control (MPC) strategy for the closed-loop operation of an Artificial Pancreas (AP) to treat type 1 diabetes mellitus is considered. The contribution of this paper is to propose a velocity-weighting mechanism, within an MPC problem's cost function, that facilitates penalizing predicted hyperglycemic blood-glucose excursions based on the predicted blood-glucose levels' rates of change. The method provides the control designer some freedom for independently shaping the AP's *uphill* versus *downhill* responses to hyperglycemic excursions; of particular emphasis in this paper is the downhill response. The proposal aims to tackle the dangerous issue of controller-induced hypoglycemia following large hyperglycemic excursions, e.g., after meals, that results in part due to the large delays of subcutaneous glucose sensing and subcutaneous insulin infusion – the case considered here. The efficacy of the proposed approach is demonstrated using the University of Virginia/Padova metabolic simulator with both unannounced and announced meal scenarios.

I. Introduction

Type 1 Diabetes Mellitus (T1DM) is a metabolic auto-immune disease characterized by the destruction of the pancreas' beta cells, and results in the body being incapable of producing insulin, a hormone that facilitates the absorption of glucose from the blood-stream into many types of cell, and that participates in the endocrine feedback loop that regulates the liver's release/removal of glucose into/from the bloodstream. People with T1DM require the delivery of insulin into their blood-stream from an external source, and tend to suffer great difficulty maintaining healthy blood-glucose levels. Hypoglycemia has very near-term consequences and may result in, e.g., dizziness or disorientation if mild, fits or unconsciousness if serious, and coma or death in severe cases. In contrast, a hyperglycemic state has few consequences if it is brief. However, a blood-glucose level that is high on average over long periods of time may result in a variety of health problems, e.g., cardiovascular disease, kidney failure, and retinal damage, possibly many years down the line.

The overall goal of this work is an Artificial Pancreas (AP) for the automated delivery of insulin to people with T1DM [1,2,3]. In particular the subcutaneous-subcutaneous AP

scheme is considered, i.e., an AP where insulin delivery (control input) is performed by a Continuous Subcutaneous Insulin Infusion (CSII) pump, and glucose sensing (output measurement for feedback) is based on a Continuous Glucose Monitor (CGM) [4]. A crucial element of any fully automated AP is a feedback control law that performs algorithmic insulin dosing that is effective and safe. For example, glucose controllers based on Model Predictive Control (MPC) [5,6,7,8,9] or proportional-integral-derivative control [10,11] have been proposed. The authors' group has been focusing increasingly on developing zone-MPC strategies [12,13,14], whereby blood-glucose levels are controlled with respect to the inclusion within a target zone, rather than to track a singular set-point. This has proven effective in *real-life* operation of the controller for two reasons. First, there is generally a significant plant-model mismatch due to the large inter- and intra-subject variability of humans' physiology. Second is that the feedback signals, an estimate of the blood-glucose level provided every 5 minutes by a CGM, suffers large errors and delays, both of which have time-varying properties and have proven difficult to model and correct for. The use of zone-MPC provides robustness against excessively responding to noise in the state estimate when the blood-glucose level is estimated to be within the target zone.

As is typical in MPC [15,16], the cost functions usually employed in an AP penalize deviations of the glucose outputs, i.e., differences between predicted glucose levels and the set-point, or, in the case of zone-MPC, distances to the target zone. With no plant-model mismatch and a long prediction horizon that may, as it frequently does in MPC, lead to effective control. However, an AP inevitably operates with large plant-model mismatch, and in practice long prediction horizons do not yield useful predictions. The authors use a 45 minute prediction horizon, which is considerably shorter than a hyperglycemic excursion due to a large meal (up to 8 hours if unannounced). Furthermore, there are long delays associated with CGM sensing and CSII delivery. The result of using a standard quadratic output cost function is an AP that continues to drive insulin delivery even after the measured peak of a hyperglycemic excursion has been reached, when predicted glucose levels may already be steadily converging to the target zone, driven by past insulin delivery. The controller's goal is to accelerate the arrival of the glucose value within the target zone; this is frequently achieved, but often with the undesirable consequence of controller-induced hypoglycemia due to insulin over-delivery. The solution is typically to de-tune the controller, such that insulin delivery is conservative both on the "uphill" and "downhill" leg of a hyperglycemic excursion, resulting in higher glucose peaks and more time spent outside the target zone.

The contribution of this paper is to include a velocity-weighting mechanism in the MPC cost function, such that a predicted zone excursion is penalized taking into account both its value and also the rate of change of the blood-glucose trajectory. The proposed mechanism provides the control designer some new degrees of freedom, and facilitates a somewhat decoupled design of the AP's response during the uphill verses the downhill leg of a hyperglycemic excursion. The goal is to deliver more insulin earlier during an excursion, towards the start of the uphill leg, and allow the controller to deliver less on the downhill leg. In [17] the authors proposed to add a "velocity-penalty" term to the cost function, resulting in more aggressive delivery during uphill travel, directly based on the rate of change of predicted glucose levels. This added term allows an otherwise conservative controller to give

an “extra kick” when the glucose trajectory is rising, but the mechanism does not actively allow the controller to “back off” when the trajectory is falling. The velocity-weighting mechanism proposed here serves this latter purpose; instead of penalizing output values using a fixed cost, and further instead of penalizing the glucose velocity directly, the proposed mechanism penalizes the output *values* based on a weight that is a *function of the velocity*. As the glucose velocity decreases, by choosing the cost to smoothly converge to zero the MPC cost function is increasingly dominated by the input cost, and actively disregards the output cost. The optimized insulin delivery command then converges as desired to the input's set-point, which is the subject's basal infusion rate (see Section II-A).

State-dependent, piece-wise linear/quadratic cost functions are standard practice in hybrid-MPC strategies and result in integer-programming problems. The proposed approach is distinct and smoothly modulates a quadratic (could include linear) cost based on the velocity (i.e., state). Smoothness allows the resulting optimization problem to be solved via a sequence of quadratic programming (QP) problems.

II. MPC Design

A. Insulin-glucose transfer function

The insulin-glucose model of [13] is employed in this paper and is summarized as follows. The model is a discrete-time, linear time-invariant (LTI) system with sample-period $T=5$ [min]. The time step index is denoted by i . The scalar plant input is the administered insulin bolus $u_{IN,i}$ [U] delivered per sample-period, and the scalar plant output is the subject's blood-glucose value $y_{BG,i}$ [mg/dL]. The plant is linearized around a steady-state that is assumed to be achieved by applying the subject-specific, time-dependent basal input rate $u_{BASAL,i}$ [U/h], and is assumed to result in a steady-state blood-glucose output $y_s = 110$ [mg/dL].

The LTI model's input u_i and output y_i are defined as:

$$\begin{aligned} u_i &:= u_{IN,i} - u_{BASAL,i} \frac{T}{60 \text{ min/h}} \\ y_i &:= y_{BG,i} - y_s. \end{aligned}$$

We denote by $\mathcal{Z}^{-1}$ the backwards shift operator, by $\mathcal{Y}(\mathcal{Z}^{-1})$ and $\mathcal{U}(\mathcal{Z}^{-1})$ the z-transform of the time-domain signals of input u_i and output y_i , respectively. The transfer characteristics from u to y are described by

$$\frac{\mathcal{Y}(\mathcal{Z}^{-1})}{\mathcal{U}(\mathcal{Z}^{-1})} = \frac{1800 - Fc}{u_{TDI}} \cdot \frac{\mathcal{Z}^{-3}}{(1 - p_1 \mathcal{Z}^{-1})(1 - p_2 \mathcal{Z}^{-1})^2} \quad (1)$$

with poles $p_1 = 0.98$, $p_2 = 0.965$, a so-called *safety factor* $F := 1.5$ (unitless, can be personalized, but fixed to 1.5 throughout this paper), the subject specific *total daily insulin* amount u_{TDI} [U] (positive scalar), and where the constant

$$c := -60(1 - p_1)(1 - p_2)^2$$

is employed to set the correct gain, and for unit conversion.

The 1800 term is from the “1800 rule” to estimate blood-glucose decrease w.r.t. delivering rapid-acting insulin [18].

B. State-space model

The state-space realization of (1) used in this work is

$$x_{i+1} = Ax_i + Bu_i \quad (2a)$$

$$y_i = C_y x_i \quad (2b)$$

$$v_i = C_v x_i \quad (2c)$$

$$A := \begin{bmatrix} p_1 + 2p_2 & -2p_1p_2 - p_2^2 & p_1p_2^2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \in^{3 \times 3}$$

$$B := \frac{1800}{u_{TDI}} Fc \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^T \in^3$$

$$C_y := \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} \in^{1 \times 3}$$

$$C_v := \begin{bmatrix} 0.1 & 0 & -0.1 \end{bmatrix} \in^{1 \times 3}.$$

The triple (A, B, C_y) describes (1). Note that, without plant-model mismatch, $x_i = [y_{i+2} \ y_{i+1} \ y_i]^T$. The second output matrix, C_v , was chosen such that the output v_i of (2c) provides an estimate of the average, over the next $2T = 10$ min, rate of change of blood-glucose level, in units mg/dL/min. The output v_i is henceforth termed the “velocity” and is employed in the proposed velocity-weighting. This notion of the velocity output was first introduced in [17].

C. State-estimation

An estimate of the state is provided at each step i by a linear state-estimator (see, e.g., [19]). For brevity the state-estimator details are omitted, thus we make the simplifying assumption that the state x_i is available for all i . No notational distinction is made between the actual and estimated state, because the state x of model (2) can *only* be estimated.

D. Blood-glucose target zone

In this work the blood-glucose target zone, i.e., the blood-glucose values for which only the basal-rate is delivered, is the range [80, 140] mg/dL, the same as in [12,13]. For simplicity the zone is time-invariant, in contrast to [14].

The (signed) zone-excursion $Z: \rightarrow$ is defined as:

$$Z(y) := \arg \min_{\alpha \in \mathbb{R}} \left\{ \alpha^2 |y + y_s - \alpha \in [80, 140]| \right\}.$$

E. Insulin delivery constraints

At each step i the controller must enforce the constraint

$$0 \leq u_i + u_{BASAL,i} \leq u_{MAX}, \quad (3)$$

where u_{MAX} denotes the maximum bolus size the CSII pump of choice is capable of delivering. In this work we choose $u_{MAX} = 25$ [U]. Note that this bolus size is so large it is highly unlikely to ever be commanded by the controller.

Insulin delivery is further subject to an Insulin On Board (IOB) constraint – a constraints based on the insulin delivery history, preventing over-delivery when much insulin was recently delivered, e.g., after a meal-bolus. The notion of IOB constraints is taken from [20] and modified slightly; reasons for the modification are not explained here, but the stated IOB constraints have proven themselves effective in clinical trials. Let vectors $\theta_l \in^{96}$ for $l = 2, 4, 6$, and 8 denote the 2, 4, 6, and 8 hour, respectively, decay curves depicted in Fig. 1. The curves are sampled at $T = 5$ min intervals, and each curve is eight hours in duration, padded with trailing zeros when necessary: $8h/T = 96$. The decay curve θ_i applicable at step i depends on the current glucose value $y_{BG,i}$:

$$\theta_i := \begin{cases} \theta_2 & \text{if } y_{BG,i} > 300 \\ \theta_4 & \text{if } y_{BG,i} \in (200, 300] \\ \theta_6 & \text{if } y_{BG,i} \in (140, 200] \\ \theta_8 & \text{otherwise.} \end{cases}$$

Let $\Lambda_{BASAL} \in^{96}$ denote the 8 hour history of u_b , the linearized insulin infusion commands, but setting to zero the value at any step were a manual meal-bolus was delivered. Further denote by $\Lambda_{MEAL} \in^{96}$ the 8 hour history of meal-boluses, by $\Theta \in$ the computed amount of IOB present, by $\Gamma \in$ the required amount of IOB that depends on current blood-glucose

levels, and by C_f [(mg/dL)/U] the patient's so-called *correction-factor*. At each time step i the IOB upper bound u_{IOB} is subsequently given by

$$u_{IOB,i} := \begin{cases} \Theta_i - \Gamma_i & \text{if } \Theta_i > \Gamma_i \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

$$\Theta_i := \theta_i^T \Lambda_{BASAL,i} + \theta_4^T \Lambda_{MEAL,i}$$

$$\Gamma_i := (y_{BG,i} - y_s) / C_f.$$

It holds that $u_{IOB} \geq 0$, and $u_{IOB} = 0$ implies the controller delivers no more than $u_{BASAL,i}$ in absolute terms. Thus, after a large bolus the insulin delivery is temporarily constrained to the basal rate, but note that the IOB constraint cannot constrain insulin delivery to below the basal rate.

F. MPC problem

For MPC background the reader is referred to [15,16]. We denote by $\mathbb{Z}_+$ the set of positive integers, by $\mathbb{Z}_a^b$ the set of consecutive integers, by $\{a, \dots, b\}$ by u, x, y, v the, predicted input u , state x , glucose output y , and velocity output v , respectively, by $N_y \in \mathbb{Z}_+$ the prediction horizon, by $N_u \in \mathbb{Z}_1^{N_y}$ the control horizon, by $\hat{R}, \check{R} \in \mathbb{R}_{>0}$ weighting coefficients for non-negative and non-positive control inputs, respectively, and by $Q: \mathbb{R} \rightarrow \mathbb{R}_{>0}$ a velocity-weighting function (discussed in Section II-H). Then, MPC performs closed-loop control by applying at each step i the first element u_0^* of the optimal, predicted control input trajectory $\{u_0^*, \dots, u_{N_u-1}^*\}$, characterized as follows.

MOC Problem: Determine

$$\{u_0^*, \dots, u_{N_u-1}^*\} := \arg \min_{\{u_0, \dots, u_{N_u-1}\}} J(x_i, \{u_0, \dots, u_{N_u-1}\})$$

with cost function

$$J(\cdot) := \sum_{k=1}^{N_y} \left[Q(v_k) \hat{z}_k^2 + \check{z}_k^2 \right] + \sum_{k=0}^{N_u-1} \left[\hat{R} \hat{u}_k^2 + \check{R} \check{u}_k^2 \right] \quad (5)$$

and subject to

$$\mathbf{x}_0 := \mathbf{x}_i \quad (6a)$$

$$\mathbf{x}_{k+1} := A\mathbf{x}_k + B\mathbf{u}_k \quad \forall k \in \{0, \dots, N_y-1\} \quad (6b)$$

$$\mathbf{y}_k := C_y \mathbf{x}_k \quad \forall k \in \{0, \dots, N_y\} \quad (6c)$$

$$\mathbf{v}_k := C_v \mathbf{x}_k \quad \forall k \in \{0, \dots, N_y\} \quad (6d)$$

$$0 \leq \mathbf{u}_k + u_{BASAL,(i+k)} \leq u_{MAX} \quad \forall k \in \{0, \dots, N_u-1\} \quad (6e)$$

$$\mathbf{u}_k \leq u_{IOB,i} \quad \forall k \in \{0, \dots, N_u-1\} \quad (6f)$$

$$\mathbf{u}_k := 0 \quad \forall k \in \{N_u, \dots, N_y-1\} \quad (6g)$$

$$\hat{\mathbf{z}}_k := \max(Z(\mathbf{y}_k), 0) \quad \forall k \in \{0, \dots, N_y\} \quad (6h)$$

$$\check{\mathbf{z}}_k := \min(Z(\mathbf{y}_k), 0) \quad \forall k \in \{0, \dots, N_y\} \quad (6i)$$

$$\hat{\mathbf{u}}_k := \max(\mathbf{u}_k, 0) \quad \forall k \in \{0, \dots, N_u-1\} \quad (6j)$$

$$\check{\mathbf{u}}_k := \min(\mathbf{u}_k, 0) \quad \forall k \in \{0, \dots, N_u-1\}. \quad (6k)$$

Eqs. (6a)-(6d) enforce the prediction dynamics of model (2), initialized to the current state (see Section II-C). Eqs. (6e) and (6f) enforce input constraint (3) and (4), respectively, across

the control horizon. Eq. (6g) implies that beyond the control horizon exactly the basal-rate is delivered. Eqs. (6h) and (6i) provide upper and lower zone deviations to penalize in (5), respectively, and facilitate an asymmetric penalization of the zone deviation. Eqs. (6j) and (6k) provide the positive and negative deviations of the input u from the basal-rate, respectively, and again facilitate an asymmetric cost function. Asymmetric input cost functions were proposed in [17] and have proven effective in clinical trials.

G. Pump-discretization

At each step i , after the solution to the **MPC Problem** above is computed, the determined u_0^* is rounded down to the nearest integer multiple of the CSII pump-discretization of 0.05 U. The portion rounded down is added at the next step in a so-called *carry-over* scheme. Other than that, no post-processing of the control command is performed.

H. Velocity-weighting function $Q(\cdot)$

The contribution of this paper is to consider velocity-weighting functions $Q(\cdot)$ that are more general than $Q(v) = 1 \quad \forall v \in \mathbb{R}$. If $Q(v)$ is unity for all velocities v then the output cost function is that of standard (zone) MPC; the cost of a zone deviation is then both invariant and symmetric, and the resulting cost function (5) is that employed in [17], but without the velocity-penalty proposed in [17]. In this work we use the velocity-weighting function of (7), depicted in Fig. 2, with $10^{-6} = \epsilon \ll 1$. It has the desired features that $Q(v) = 1$ if $v \geq 0$ (i.e., we obtain the standard MPC solution for non-negative velocities), $Q(v) = \epsilon \ll 1$ if $v \leq -1$ (i.e., the zone excursion is effectively not penalized at all for steep negative glucose velocities, resulting in the input's set-point being commanded), and with a smooth transition when $-1 < v < 0$, allowing the MPC solution to slowly converge to the input's set-point as the velocity becomes increasingly negative. We require $\epsilon > 0$ to result in a well-posed problem.

$$Q(v) := \begin{cases} 1 & \text{if } v \geq 0 \\ \epsilon & \text{if } v \leq -1 \\ \frac{1}{2} [\cos(v\pi)(1-\epsilon) + (1+\epsilon)] & \text{otherwise} \end{cases} \quad (7)$$

If $Q(\cdot)$ is any positive constant then the **MPC Problem** can be cast as a continuous, strictly-convex QP. If $Q(\cdot)$ is a more general function then the **MPC Problem** is not a QP, but under mild assumptions of continuity and smoothness of $Q(\cdot)$ a solution to the **MPC Problem** can be obtained by sequential quadratic programming (details omitted).

III. Examples and Discussion

A. Simulation setup

In this section the proposed velocity-weighting MPC, utilizing the velocity-weighting function depicted in Fig. 2, is compared to “standard MPC” with $Q(v) = 1$ using the University of Padova/Virginia Food and Drug Administration (FDA) accepted metabolic simulator [21], based on the following scenario. Simulations start at midnight, are 20 hours in duration, and a 90 gram carbohydrate meal is consumed at 04:00. The parameters of both

the simulator and controller are time-invariant, thus the time of day of meal consumption is irrelevant. We consider either *unannounced* meals, where the meal-disturbance is rejected based on feedback only, or *announced* meals, where a meal-bolus is delivered at the start of meal-ingestion, and the bolus size is optimal with respect to the in-silico subject's parameters. The controller parameters are $N_y = 9$ (i.e., 45 minute prediction horizon), $N_u = 5$ (i.e., 25 minute control horizon), $\hat{R} = 7000$, $\hat{Q} = 100$. The feedback signal for control is the simulator's CGM signal, but with no additive noise. CGM noise is omitted in the cases presented in this work in order to demonstrate the velocity-weighting action of the proposed controller; without noise the differences between velocity-weighting MPC and standard MPC occur only *after* the peak of the hyperglycemia excursion. However, cases with CGM noise, and also with more challenging scenarios, were tested, and the overall conclusions of this paper hold for these cases also.

B. Single subject examples

A simulation example of the situation the proposed velocity-weighting MPC is targeted at is depicted in Fig. 3. The use of standard MPC results in significant insulin infusion after the peak, when blood-glucose values are already descending towards the zone, causing hypoglycemia at 09:30. The controller suspends the pump from about 08:30 until 10:00, and after some rebound glucose stabilizes. In contrast, the proposed velocity-weighting MPC reduces insulin infusion to the basal-rate soon after the peak is reached, slowing the descent. No hypoglycemia is induced, and after 06:00 the controller delivers the basal-rate continuously, with no pump-suspension required, and no associated glucose rebound.

In the example of Fig. 4 standard MPC does not induce hypoglycemia, but velocity-weighting MPC again effects an obvious reduction in insulin delivery during the descent. This example demonstrates a possible drawback of the proposed strategy: For some subjects the proposal results in a more sluggish return to the zone, and more time spent in hyperglycemia. However, for such subjects a more aggressively tuned controller, i.e., smaller $\hat{R}$, in combination with the proposed velocity-weighting MPC strategy, delivers more insulin during the ascent, possibly reducing the peak glucose value and time in hyperglycemia. In this paper the value of $\hat{R}$ is fixed in order to elucidate the differences in glucose outcomes induced by only the proposed velocity-weighting.

The examples above had unannounced meals, but the proposed velocity-weighting reduces the risk of hypoglycemia also with announced meals. In the example of Fig. 5 the IOB constraint prevents the controller from delivering more than the basal-rate until 05:45, due to the meal-bolus delivered at 04:00. After 05:45, during the glucose descent, standard MPC delivers in excess of the basal rate, and hypoglycemia is induced at 09:45, with an associated pump-suspension and subsequent glucose rebound. In contrast, the proposed velocity-weighting MPC continuously delivers close to the basal-rate following the meal-bolus, despite the relaxing IOB constraint after 05:45, and no hypoglycemia is induced.

C. Comment on IOB constraint tuning

There is a difficult trade-off to be made when tuning the IOB constraint mechanism based on the decay curves depicted in Fig. 1. On the one hand, a longer decay curve results in less

insulin delivery, because delivered insulin is modeled to be active for longer, and so, following a meal-bolus, it takes longer before the controller is given leeway to attempt correcting hyperglycemia. For standard MPC and the case of Fig. 5, the use of longer decay curves would have prevented controller-induced hypoglycemia. On the other hand, in general it is desirable to allow the controller more leeway to control hyperglycemia. The use of the proposed velocity-weighting MPC partially alleviates this trade-off: One can employ shorter decay curves, giving the controller more leeway to control rising or persistently high glucose values early after a meal-bolus, while enforcing that insulin delivery is not excessive when glucose values are falling due to the delayed action of the meal-bolus.

D. Population results

We consider Control-Variability Grid Analysis (CVGA) [22] of the entire cohort in the University of Padova/Virginia metabolic simulator, consisting of 10-subject and 100-subject simulator cohorts. The latter contains a 101st subject “average”, resulting in a combined cohort of 111 in-silico subjects. The CVGA plots for unannounced and announced meals are depicted in Figs. 6 and 7, respectively. Top sub-plots show the results with standard MPC; bottom sub-plots are for the proposed velocity-weighting MPC. Each in-silico subject's result is denoted by a black dot. The large blue dot is the arithmetic mean of the individual dots, and the blue circle has a radius of the standard deviation of the distances of each dot to the cohort's mean, each in terms of the underlying CVGA space. The circles' radii, and also the CVGA zone inclusion percentages, are tabulated in Table I.

Due to the lack of CGM noise and fixed $\hat{R}$ the glucose peak (vertical position) for each subject is equal for both MPC strategies. The MPC strategy affects only the depth of the nadir (horizontal position). For unannounced meals most subjects undergo a shift to the left, i.e., the nadir is higher, meaning the hypoglycemia risk is reduced. The average is significantly shifted to the left, and the spread is smaller. For announced meals the peaks are significantly lower, due to the feed-forward nature of the meal announcement. For many subjects there is no difference between standard MPC and the proposed velocity-weighting MPC. However, there is an overall shift to the left, the spread is less, and for a small number of subjects the use of velocity-weighting clearly alleviates a significant risk of controller-induced hypoglycemia.

IV. Conclusion

A velocity-weighting mechanism was proposed for MPC of an AP, such that deviations of predicted blood-glucose values are penalized by a cost-function that takes into account both the size of the deviation and also the velocity of the trajectory the predicted glucose value is an element of. Thus the MPC cost-function is made more “velocity aware”, and the AP's response during the uphill and downhill leg of a hyperglycemic excursion can be tailored somewhat independently. By simulation examples it was demonstrated that the risk of controller-induced hypoglycemia is partially alleviated by the proposed strategy, as per design objective.

Acknowledgments

Access to the complete version of the University of Virginia/Padova metabolic simulator was provided by an agreement with Prof. C. Cobelli (University of Padova) and Prof. B. P. Kovatchev (University of Virginia) for research purposes. The authors acknowledge funding provided by the National Institutes of Health: DP3DK094331, R01DK085628.

References

1. Zisser H. Clinical Hurdles and Possible Solutions in the Implementation of Closed-Loop Control in Type 1 Diabetes Mellitus. *J Diabetes Sci Technol*. Sep.2011 5:1283–1286. [PubMed: 22027329]
2. Cobelli C, Renard E, Kovatchev B. Artificial Pancreas: Past, Present, Future. *Diabetes*. Nov.2011 60:2672–2682. [PubMed: 22025773]
3. Doyle FJ III, Huyett LM, Lee JB, Zisser HC, Dassau E. Closed Loop Artificial Pancreas Systems: Engineering the Algorithms. *Diabetes Care*. May.2014 37:1191–1197. [PubMed: 24757226]
4. Hovorka R. Continuous glucose monitoring and closed-loop systems. *Diabetic Med*. Jan.2006 23:1–12.
5. Parker RS, Doyle FJ III, Peppas NA. A Model-Based Algorithm for Blood Glucose Control in Type I Diabetic Patients. *IEEE Trans Biomed Eng*. Feb.1999 46:148–157. [PubMed: 9932336]
6. Hovorka R, Canonico V, Chassin LJ, Haueter U, Massi-Benedetti M, Federici MO, Pieber TR, Schaller HC, Schaupp L, Vering T, Wilinska ME. Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. *Physiol Meas*. Jul.2004 25:905–920. [PubMed: 15382830]
7. Magni L, Raimondo DM, Dalla Man C, De Nicolao G, Kovatchev B, Cobelli C. Model predictive control of glucose concentration in type 1 diabetic patients: An in silico trial. *Biomed Signal Process Control*. 2009; 4(4):338–346.
8. Breton M, Farret A, Bruttomesso D, Anderson S, Magni L, Patek S, Dalla Man C, Place J, Demartini S, Del Favero S, Toffanin C, Hughes-Karvetski C, Dassau E, Zisser H, Doyle FJ III, De Nicolao G, Avogaro A, Cobelli C, Renard E, Kovatchev B. Fully Integrated Artificial Pancreas in Type 1 Diabetes: Modular Closed-Loop Glucose Control Maintains Near Normoglycemia. *Diabetes*. Jun.2012 61:2230–2237. [PubMed: 22688340]
9. Turksoy K, Bayrak ES, Quinn L, Littlejohn E, Cinar A. Multivariable Adaptive Closed-Loop Control of an Artificial Pancreas Without Meal and Activity Announcement. *Diabetes Technol Ther*. May.2013 15:386–400. [PubMed: 23544672]
10. Steil GM, Rebrin K, Darwin C, Hariri F, Saad MF. Feasibility of Automating Insulin Delivery for the Treatment of Type 1 Diabetes. *Diabetes*. Dec.2006 55:3344–3350. [PubMed: 17130478]
11. Marchetti G, Barolo M, Jovanovi L, Zisser H, Seborg DE. A feedforward-feedback glucose control strategy for type 1 diabetes mellitus. *J Process Control*. Feb.2008 18:149–162. [PubMed: 19190726]
12. Grosman B, Dassau E, Zisser HC, Jovanovi L, Doyle FJ III. Zone Model Predictive Control: A Strategy to Minimize Hyper- and Hypoglycemic Events. *J Diabetes Sci Technol*. Jul.2010 4:961–975. [PubMed: 20663463]
13. van Heusden K, Dassau E, Zisser HC, Seborg DE, Doyle FJ III. Control-Relevant Models for Glucose Control Using A Priori Patient Characteristics. *IEEE Trans Biomed Eng*. Jul.2012 59:1839–1849. [PubMed: 22127988]
14. Gondhalekar R, Dassau E, Zisser HC, Doyle FJ III. Periodic-Zone Model Predictive Control for Diurnal Closed-loop Operation of an Artificial Pancreas. *J Diabetes Sci Technol*. Nov.2013 7:1446–1460. [PubMed: 24351171]
15. Maciejowski, JM. Predictive Control with Constraints. Pearson/Prentice Hall; Harlow, England: 2002.
16. Rawlings, JB., Mayne, DQ. Model Predictive Control: Theory and Design. Nob Hill Publishing; Madison, WI, USA: Aug. 2009

17. Gondhalekar, R., Dassau, E., Doyle, FJ, III. MPC Design for Rapid Pump-Attenuation and Expedited Hyperglycemia Response to Treat T1DM with an Artificial Pancreas. Proc AACC American Control Conf; Portland, OR, USA. Jun 2014; p. 4224-4230.
18. Walsh, J., Roberts, R. Pumping Insulin. 4 ed.. Torrey Pines Press; San Diego, CA, USA: 2006.
19. Levine, WS., editor. The Control Handbook. 2 ed.. CRC Press; Boca Raton, FL, USA: 2011.
20. Ellingsen C, Dassau E, Zisser H, Grosman B, Percival MW, Jovanovi L, Doyle FJ III. Safety Constraints in an Artificial Pancreatic β Cell: An Implementation of Model Predictive Control with Insulin on Board. J Diabetes Sci Technol. May.2009 3:536–544. [PubMed: 20144293]
21. Kovatchev BP, Breton M, Dalla Man C, Cobelli C. *In Silico* Preclinical Trials: A Proof of Concept in Closed-Loop Control of Type 1 Diabetes. J Diabetes Sci Technol. Jan.2009 3:44–55. [PubMed: 19444330]
22. Magni L, Raimondo DM, Dalla Man C, Breton M, Patek S, De Nicolao G, Cobelli C, Kovatchev B. Evaluating the Efficacy of Closed-Loop Glucose Regulation via Control-Variability Grid Analysis. J Diabetes Sci Technol. Jul.2008 2:630–635. [PubMed: 19885239]

Fig. 1.
IOB decay curves from [18] and used in IOB constraints in [20].

Fig. 2.
Velocity-weighting function (7) used in Section III.

Fig. 3.
Unannounced meal. In-silico subject #12 of 100.

Fig. 4.
Unannounced meal. In-silico subject #8 of 100.

Fig. 5.
Announced meal. In-silico subject #13 of 100.

Fig. 6. Unannounced meal. **Top:** Standard MPC. **Bottom:** Proposed velocity-weighting MPC. For zone counts and cluster circle radii see Table I.

Fig. 7. Announced meal. **Top:** Standard MPC. **Bottom:** Proposed velocity-weighting MPC. For zone counts and cluster circle radii see Table I.

TABLE I

CVGA zone (A,...,E) inclusion percentages, and cluster circle radii [mg/dL], of the four cases depicted in Figs. 6 and 7.

Meal-type	MPC	A	B	C	D	E	Radius
Unannounced	Standard	0%	64%	13%	15%	8%	9.3
	Proposed	0%	64%	23%	11%	2%	7.4
Announced	Standard	7%	89%	0%	3%	1%	6.6
	Proposed	7%	91%	0%	2%	0%	4.7