Published in final edited form as:

Proc Am Control Conf. 2014 June ; 2014: 4224–4230. doi:10.1109/ACC.2014.6859247.

MPC Design for Rapid Pump-Attenuation and Expedited Hyperglycemia Response to Treat T1DM with an Artificial Pancreas

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) for treating Type 1 Diabetes Mellitus (T1DM) is considered in this paper. The contribution of this paper is to propose two changes to the usual structure of the MPC problems typically considered for control of an AP. The first proposed change is to replace the symmetric, quadratic input cost function with an asymmetric, quadratic function, allowing negative control inputs to be penalized less than positive ones. This facilitates rapid pump-suspensions in response to predicted hypoglycemia, while simultaneously permitting the design of a conservative response to hyperglycemia. The second proposed change is to penalize the velocity of the predicted glucose level, where this velocity penalty is based on a cost function that is again asymmetric, but additionally state-dependent. This facilitates the accelerated response to acute, persistent hyperglycemic events, e.g., as induced by unannounced meals. The novel functionality is demonstrated by numerical examples, and the efficacy of the proposed MPC strategy verified using the University of Padova/Virginia metabolic simulator.

I. Introduction

A. Background

Type 1 Diabetes Mellitus (T1DM) is a metabolic autoimmune disease characterized by the destruction of the pancreas' beta cells, and results in the body being incapable of producing insulin, a hormone that serves at least two important functions. The first is to facilitate the absorption of glucose from the blood-stream into many types of cell. The second function is to participate, in conjunction with glucagon (insulin's antagonist), in the endocrine feedback loop that regulates the liver's release/removal of glucose into/from the blood-stream. People with T1DM require the delivery of insulin into their blood-stream from an external source in order to fuel their cells, and furthermore 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

For correspondence contact Francis J. Doyle III.

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 [1,2,3,4,5]. 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) [6]. 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, glycemia controllers based on Model Predictive Control (MPC) [7,8,9,10,11] or proportional-integral-derivative control [12,13] have been proposed. The authors' group has been focusing increasingly on developing zone-MPC strategies [14,15,16,17], whereby blood-glucose levels are controlled with respect to the inclusion within a safe 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 safe zone.

B. Asymmetric input cost functions

Typical MPC formulations employ quadratic functions to penalize output (predicted blood-glucose value) and input (predicted insulin delivery) deviations. The output deviation is the difference between the predicted blood-glucose level and the set-point, or, in the case of zone-MPC, the distance to the safe zone. The input deviation is usually the difference between the absolute insulin infusion rate and the subject's basal infusion rate – the insulin infusion rate that maintains a steady blood-glucose level in a resting and fasting state. Note that a pump-suspension, i.e., a zero delivery in absolute terms, is a control input of the negative basal rate from the point of view of the controller. Positive and negative deviations of equal size are typically penalized equally, i.e, the quadratic cost functions are symmetric. This poses a challenge: A controller tuned to respond conservatively to hyperglycemia, i.e., does not over-deliver in response to high blood-glucose values, has difficulty performing a pump-suspension in the face of hypoglycemia. Conversely, a controller tuned to easily perform pump-attenuations tends to over-correct hyperglycemia, resulting in controller induced hypoglycemia and oscillations following blood-glucose highs. When using symmetric input cost functions the AP's responses to hyperglycemia and hypoglycemia are coupled, and a design trade-off must occur. Implementations of MPC for an AP typically require supervisory control and alarm systems to mitigate the effects of this trade-off and enforce pump-attenuation despite the MPC being tuned for an appropriate response to hyperglycemia.

The first of two contributions of this paper is to propose asymmetric quadratic input cost functions, specifically, a quadratic function for penalizing negative control inputs (pump-

attenuations) with a lower cost then the quadratic function for penalizing positive control inputs (infusion rate increased above basal rate). Thus the design of an AP's responses to hyperglycemia and hypoglycemia is decoupled.

The use of asymmetric cost functions in MPC of an AP was considered in [9,18,19,20]. Those works slightly differ in motivation and implementation, but share two important similarities with one another. First, the asymmetry is applied to the output cost function, not the input cost function. This is in some sense irrelevant, because it is the *ratio* of output-cost vs. input-cost that determines the controller's behavior. However, the aforementioned references attempted to address the asymmetry of the glycemia control problem in terms of risk, in that a glucose deviation below the zone is more risky than one above the zone. In contrast, the objective in this paper is to facilitate swift pump-attenuations by reducing the input cost for negative inputs. Importantly, the level of asymmetry in this paper's proposal is much greater than in the referenced works. An increased risk of a low blood-glucose value by some factor does not imply that the best response is a pump-attenuation by that factor. Because insulin delivery cannot be reversed or corrected for (bar bi-hormonal control and rescue carbohydrates), a hypoglycemic state requires very "aggressive" pump-suspension. The second similarity is that in the aforementioned works the asymmetry was either only considered within a benchmark problem for *in silico* tests, and not meant for implementation, or considered beneficial in *in silico* tests, but the idea abandoned for clinical trials. In contrast, the asymmetric cost functions proposed here have obtained Food and Drug Administration (FDA) approval and are being employed in clinical trials at the time of writing.

C. Velocity penalty MPC

To the authors best knowledge the models employed for MPC of an AP always have a single, scalar output; a blood-glucose value, typically the deviation from some safe level y_s. In this work $y_s := 110$ mg/dL. If the model (and state-estimator) were error-free then a scheme penalizing this output, with sufficiently long prediction horizon, should produce a desirable result. In practice it is difficult to obtain accurate models, and a prediction horizon of 45 min is the longest that yields useful predictive powers employing the model of [16], also employed in this work. Note that this prediction horizon is significantly shorter than the time-scales of the system. For example, in response to a meal the blood-glucose level may take 1–2 hours to peak, and 4–5 hours to descend back to a fasting value. Thus the "proportional action" nature of the typical MPC implementations results in an insulin infusion trajectory that mimics, in shape, the blood-glucose trajectory. This results in an insulin delivery that exceeds the basal-rate long after meal intake, even after the blood-glucose level has peaked and is descending. In contrast, optimal insulin therapy performed manually by subjects involves delivering a large insulin bolus at, or close to, the time of meal ingestion, and delivering the basal-rate thereafter. Such feed-forward action in combination with a feedback-control AP has proven effective. However, the goal of this work is the design of an MPC strategy that performs well based on pure feedback, in order to tackle, e.g., unannounced (i.e., no accompanying bolus) meals.

The second of two contributions of this paper is to propose an MPC design that attempts to mimic the optimal basal-bolus therapy, by delivering more insulin *before* blood-glucose peaks, and returning to close to the basal-rate when the blood-glucose value is descending. This is achieved by augmenting the insulin-glucose model of [16] with a second output, providing an estimate of the current velocity of the blood-glucose level, and penalizing the velocity by a quadratic function. However, the quadratic function is again asymmetric, because only blood-glucose rises are penalized, not drops. Furthermore, the quadratic function is state-dependent, because blood-glucose rises should only be penalized when the blood-glucose level is high.

II. MPC Design

A. Insulin-glucose transfer function

The insulin-glucose model of [16] is employed in this paper and is summarized as follows. The model is a discrete-time, linear time-invariant (LTI) system with sample-period $T_s = 5$ [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:

$$u_i := u_{\mathrm{IN},i} - u_{\mathrm{BASAL},i}\frac{T_\mathrm{s}}{60\ \mathrm{min}}$$

$$y_i = y_{\mathrm{BG},i} - y_\mathrm{s}.$$

We denote by $\mathcal{Z}^{-1}$ the backwards shift operator, by $\mathscr{Y}(\mathcal{Z}^{-1})$ and $\mathscr{U}(\mathcal{Z}^{-1})$ 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

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

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 [21].

B. Augmented state-space model

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

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

$$y_i=C_{\mathrm{y}}x_i \quad (2b)$$

$$d_i=C_{\mathrm{d}}x_i \quad (2c)$$

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

$$B := \frac{1800\,Fc}{u_{\mathrm{TDI}}}\begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^\top \in \mathbb{R}^3$$

$$C_{\mathrm{y}} := \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{1\times 3}$$

$$C_{\mathrm{d}} := \begin{bmatrix} 0.1 & 0 & -0.1 \end{bmatrix} \in \mathbb{R}^{1\times 3}.$$

The triple (A, B, C_{y}) describes (1). However, for this work a second output was desired, in order to penalize the rate of change of the blood-glucose output y. From the structure of A and C_{y} it can be seen that, without plant-model mismatch, $x_i = [y_{i+2}\ y_{i+1}\ y_i]^\top$. The second output matrix, C_{d}, was chosen such that the auxiliary output d_i of (2c) provides an estimate of the average, over the next $2T_{\mathrm{s}}$ = 10 min, rate of change of blood-glucose level, in units mg/dL/min.

C. State-estimation

An estimate of the state is provided at each step i by a linear state-estimator (see, e.g., [22]). 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. Safe blood-glucose zone

In this work the safe blood-glucose 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 [14,16]. For simplicity the zone is time-invariant, in contrast to [17].

Let $Z: \mathbb{R} \rightarrow \mathbb{R}$ denote the zone-excursion function:

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

E. Insulin delivery constraints

At each step i the controller must enforce the constraint

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

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.

In practice it is common to enforce further constraints on the insulin delivery, e.g., so-called *insulin on board* (IOB) constraints [23], or input constraints that depend on the time of day to reduce the risk of nocturnal hypoglycemia [17]. The lack of IOB constraints in the exposition of this paper renders the proposed method, as presented here, inappropriate for *announced* meals, i.e., scenarios where meal-ingestion is accompanied by a large insulin bolus. In this work supplementary input constraints are dispensed with in order to elucidate the effects of the proposed MPC design without mixing in consequences due to auxiliary mechanisms.

F. MPC problem

For MPC background and fundamentals the reader is referred to [24,25]. We denote by $\mathbb{Z}_+$ the set of positive integers, by $\mathbb{Z}_a^b$ the set of consecutive integers $\{a, \ldots, b\}$, and by u, x, y, d the predicted values of input u, state x, glucose output y, and velocity output d, respectively.

We further denote by $N_{\mathrm{y}} \in \mathbb{Z}_+$ the prediction horizon, by $N_{\mathrm{u}} \in \mathbb{Z}_1^{N_{\mathrm{y}}}$ the control horizon, by $\hat{R}, \check{R} \in \mathbb{R}_{>0}$ weighting factors for the positive and negative control inputs, respectively, and by $D \in \mathbb{R}_{\geq 0}$ a weighting factor for the positive blood-glucose velocity. Then, MPC performs closed-loop control by applying at each step i the first element u_0^* of the optimal, predicted control input trajectory $\{\mathrm{u}_0^*, \ldots, \mathrm{u}_{N_{\mathrm{u}}-1}^*\}$ determined by solving the following problem.

MPC Problem—Determine

$$\{\mathrm{u}_0^*,\ldots,\mathrm{u}_{N_{\mathrm{u}-1}}^*\} := \underset{\{\mathrm{u}_0,\ldots,\mathrm{u}_{N_{\mathrm{u}-1}}\}}{\arg\ \min}\ J\left(x_i,\{\mathrm{u}_0,\ldots,\mathrm{u}_{N_{\mathrm{u}-1}}\}\right)$$

with cost function

$$J(\cdot) := \sum_{k=1}^{N_{\mathrm{y}}} \left(\mathrm{z}_k^2 + \hat{D}\hat{\mathrm{d}}_k^2\right) + \sum_{k=0}^{N_{\mathrm{u}-1}} \left(\hat{R}\hat{u}_k^2 + \check{R}\check{u}_k^2\right) \quad (4)$$

and subject to

$$\mathrm{x}_0 := x_i \quad (5a)$$

$$\mathrm{x}_{k+1} := A\mathrm{x}_k + B\mathrm{u}_k \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{y}-1}} \quad (5b)$$

$$\mathrm{y}_k := C_{\mathrm{y}}\mathrm{x}_k \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{y}}} \quad (5c)$$

$$\mathrm{d}_k := C_{\mathrm{d}}\mathrm{x}_k \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{y}}} \quad (5d)$$

$$0 \leq \mathrm{u}_k + u_{\mathrm{BASAL},(i+k)} \leq u_{\mathrm{MAX}} \quad \forall k \in \mathbb{Z}_0^{N_{\mathrm{u}-1}} \quad (5e)$$

$$\mathrm{u}_k := 0 \qquad \forall k \in \mathbb{Z}_{N_{\mathrm{u}}}^{N_{\mathrm{y}-1}} \quad (5f)$$

$$\mathrm{z}_k := Z(\mathrm{y}_k) \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{y}}} \quad (5g)$$

$$\hat{\mathrm{d}}_k := \max\left(\mathrm{d}_k, 0\right) \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{y}}} \quad (5h)$$

$$\hat{u}_k := \max\left(\mathrm{u}_k, 0\right) \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{u}-1}} \quad (5i)$$

$$\check{u}_k := \min\,(\mathrm{u}_k, 0) \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{u}}-1} \tag{5j}$$

$$\hat{D} := \begin{cases} D & \text{if } y_i + y_{\mathrm{s}} \geq 140\,\mathrm{mg/dL} \\ 0 & \text{otherwise.} \end{cases} \tag{5k}$$

Eqs. (5a)–(5d) enforce the prediction dynamics of model (2), initialized to the current state (see Section II-C). Eq. (5e) enforces constraint (3) across the control horizon. Eq. (5f) implies that beyond the control horizon exactly the basal-rate is delivered. Eqs. (5g)–(5h) provide output signals to penalize in (4), where use of max(·) facilitates an asymmetric weighting, penalizing only the positive velocity. Eqs. (5i)–(5j) provide the positive and negative deviations of the input u from the basal-rate, respectively, and again facilitate the implementation of an asymmetric cost function. Eq. (5k) performs state-dependent switching, to penalize $\hat{\mathrm{d}}$ only when the estimated blood-glucose level exceeds the safe zone.

If $\hat{R} = \check{R}$ then the input cost function is symmetric. If $D = 0$ then there exists no velocity penalty. Both conditions being true results in an analogous MPC strategy to the zone-MPC considered in [14,16,17] (except minor details).

Note that at each step i, despite the use of zones and asymmetric cost functions, the above **MPC Problem** can be cast as a continuous, strictly-convex quadratic program (QP). This is explained further in the Appendix.

Eq. (5k) signifies that a positive blood-glucose velocity d is only penalized in (4) if the current, estimated blood-glucose value exceeds 140 mg/dL. Replacing the statement "if y_i …" with "if y_k …" results in state-dependent switching at each prediction step. This could be achieved, e.g., using the mixed logical dynamical hybrid systems approach [26], but would require the solution of mixed-integer QPs. The FDA has in the past approved the use of QPs for clinical trials, but it is not clear whether mixed-integer QPs would receive approval. Hence the approach of Eq. (5k) is favored for now.

III. Demonstrative Examples

In this section one *in silico* adult subject of version 3 of the University of Padova/Virginia FDA accepted metabolic simulator [27] is used to demonstrate the superiority of asymmetric input cost functions, and the usefulness of velocity penalty MPC, in Sections III-A and III-B, respectively. The subject is adult number 1, and has a fairly standard response. For clarity of exposition the feedback signal employed for state-estimation is based on CGM *without* stochastic additive disturbances. However, other CGM dynamics modeled by the simulator, e.g., delays due to interstitial transport and sampling, *are* included. Note that, for demonstration, the controller tunings are highly individualized to the single adult.

A. Asymmetric input cost function

The use of asymmetric input cost functions is demonstrated via an MPC strategy with no velocity penalty ($D = 0$). The closed-loop simulation starts at 8am. At 9am a 2 U bolus is delivered in order to mimic the spontaneous occurrence of a hypoglycemic event. At 2pm a 90 gCHO meal is ingested that is unannounced, i.e., no feed-forward meal-bolus is delivered. Blood-glucose and insulin delivery trajectories are depicted in Fig. 1.

Three tunings are considered: A symmetric high cost, a symmetric low cost, and an asymmetric cost that is high for positive inputs and low for negative inputs. With the symmetric high cost, the pump fails to perform a pump-suspension around 10am despite acute hypoglycemia, but delivers appropriately conservatively w.r.t. the meal, from 2–7pm. Using the symmetric low cost, the pump appropriately suspends delivery around 10am, but catastrophic overdelivery occurs in response to the meal around 3pm. Use of the asymmetric cost function permits the pump to both appropriately suspend in response to hypoglycemia, and also deliver appropriately conservatively w.r.t. the meal.

B. Velocity penalty MPC

Velocity penalty MPC is demonstrated in this section. The closed-loop simulation starts at noon. At 2pm a 90 gCHO, unannounced meal is consumed. Four controller tunings *without* velocity penalty ($D = 0$) are considered: $\hat{R} = 7000$, $\hat{R} = 5000$, $\hat{R} = 3000$, $\hat{R} = 1500$. Note that a smaller value of $\hat{R}$ implies the controller more actively attempts to correct the meal-disturbance. Based on the least conservative tuning with $\hat{R} = 7000$, four further controller tunings *with* velocity penalty are considered: $D = 0$ (as before), $D = 1500$, $D = 3000$, $D = 6000$. Note that, in this case, a *larger* value of D implies the controller more actively attempts to correct the meal-disturbance. For all cases $\check{R} = 100$.

Blood-glucose and insulin delivery trajectories are depicted in Fig. 2. The top subplots are for the cases without velocity penalty, the bottom subplots with. Note that the least conservative response (green) is identical in both upper and lower subplots. The three respective controller tunings that are not the most conservative one were chosen because they result in, respectively, glucose trajectories with similar characteristics after about 7pm. With velocity penalty MPC the blood-glucose peaks are clearly lower, and occur earlier, than in the traditional MPC case. This is caused by a markedly increased insulin delivery towards the start of the meal-disturbance before 3pm.

Plotted in Fig. 3 are the responses of the third most aggressive controllers *without* and *with* velocity penalty, as plotted by a red line in the upper and lower, respectively, subplots of Fig. 2. Also plotted in Fig. 3 are the responses when employing the most conservative MPC scheme with 50% and 100% meal-boluses delivered at mealtime. Clearly the feed-forward action is far superior. Nevertheless, the velocity penalty is beneficial; with velocity penalty the peak is lower, and the minimum at 8pm is higher, than with the analogously tuned traditional MPC scheme.

IV. In Silico Testing – 10 In Silico Adults

In this section the 10 *in silico* adult subjects of the commercial version (v3) of the University of Padova/Virginia FDA accepted metabolic simulator [27] are used to verify the control performance of the proposed MPC strategy. In contrast to Section III, here CGM feedback signals include stochastic disturbances. Closed-loop simulations start at noon and one unannounced 75 gCHO meal is provided at 2pm. Two controller setups are compared: **1)** The "traditional-MPC" scheme employs a symmetric input cost function ($\hat{R} = \check{R} = 11000$) and no velocity penalty ($D = 0$); **2)** the "proposed-MPC" scheme employs an asymmetric input cost function ($\hat{R} = 18000$, $\check{R} = 100$) and includes a velocity penalty ($D = 1800$). These tunings were chosen because they achieve acceptable control performance *on average*. Neither controller may be considered the optimal in its class.

Fig. 4 shows the mean, and min-max envelopes, of the blood-glucose and insulin delivery trajectories. With the proposed-MPC scheme the means' peak-to-peak distance is shorter, and the return to the safe zone is accellerated. This is achieved by delivering significantly more insulin during the uphill portion of the meal response. Fig. 5 shows the Control-Variability Grid Analysis (CVGA) [28] plot. This demonstrates also that, individually, the blood-glucose maxima are lower, the minima higher, and that the cluster of points is tighter, when using the proposed-MPC scheme.

Except for a simple mechanism to enforces the pump-discretization, the controllers contain no auxiliary functionality over that described in Section II. Thus, no special safety constraints are enforced, and no rescue-carbohydrates given in response to hypoglycemia. This is not how a clinical trial would be performed, but is intended to demonstrate the action of the "pure" control laws. However, the UCSB Health Monitoring System [29] was running during the simulations: The proposed-MPC scheme resulted in no hypoglycemia alarms, whereas the traditional-MPC strategy resulted in four.

V. Conclusion

A novel MPC strategy for controlling an AP that treats T1DM was presented, and its efficacy demonstrated by simulations. The proposed strategy is novel for two reasons.

First, asymmetric input cost functions are employed to facilitate the independent tuning of the responses to hyperglycemia and hypoglycemia, allowing an AP to be both conservative w.r.t. hyperglycemia, and also to enforce "aggressive" pump-attenuations when faced with hypoglycemia, without requiring supervisory control systems. The proposal appears to be a novel implementation of an idea that was previously mooted but abandoned. Importantly, controllers with asymmetric cost functions have obtained FDA approval and are undergoing clinical testing at the time of writing.

The second novel aspect is the state-dependent and asymmetric penalization of the rate of change of the predicted blood-glucose trajectory, facilitating increased insulin delivery during the "uphill" travel of a glucose excursion, while avoiding the continued delivery of excessive insulin when the glucose response is already on its "downhill" travels. This proposal appears to be new, but is still considered experimental and is not yet scheduled for

testing in a clinical trial. One point of concern is the amplified response to CGM noise, and this issue is the subject of future research, both from an MPC as well as state-estimation point of view.

Acknowledgments

The authors gratefully acknowledge funding provided by the National Institutes of Health (NIH): DP3DK094331, R01DK085628.

The authors thank Dr. Howard C. Zisser (Insulet Corp. (Bedford, MA, USA)) for valuable comments.

APPENDIX

The purpose of this appendix is to describe how the MPC problem presented in Section II-F can be formulated as a continuous, convex QP. The outline is as follows: The asymmetric input cost function is implemented by explicitly employing two separate insulin inputs, where one is constrained to be non-negative ($\hat{u} \in \mathbb{R}$, see (5i)), the other non-positive ($\check{u} \in \mathbb{R}$, see (5j)), and penalizing each with an independent, symmetric, quadratic cost function. The cost function penalizing the excursion (see (5g)) of the predicted blood-glucose output $y_k + y_s$ beyond the target zone [80, 140] mg/dL is implemented using two auxiliary optimization variables for each prediction step; the auxiliary variables $\hat{\gamma} \in \mathbb{R}$ and $\check{\gamma} \in \mathbb{R}$ quantify the excursion above and below the zone, respectively. The cost function penalizing the non-negative velocity (see (5h)) is implemented using the auxiliary optimization variable $\upsilon \in \mathbb{R}$ for each prediction step. The three new families of auxiliary variables are penalized using independent, symmetric, quadratic cost functions.

At each step i let

$$\hat{U} := \begin{bmatrix} \hat{u}_0 & \cdots & \hat{u}_{N_{u}-1} \end{bmatrix}^\top \in \mathbb{R}^{N_u}$$

$$\check{U} := \begin{bmatrix} \check{u}_0 & \cdots & \check{u}_{N_{u}-1} \end{bmatrix}^\top \in \mathbb{R}^{N_u}$$

denote the vectors of predicted non-negative and non-positive, respectively, control inputs, over the control horizon, and further let

$$\hat{\Gamma} := \begin{bmatrix} \hat{\gamma}_1 & \cdots & \hat{\gamma}_{N_y} \end{bmatrix}^\top \in \mathbb{R}^{N_y}$$

$$\check{\Gamma} := \begin{bmatrix} \check{\gamma}_1 & \cdots & \check{\gamma}_{N_y} \end{bmatrix}^\top \in \mathbb{R}^{N_y}$$

$$V := \begin{bmatrix} \upsilon_1 & \cdots & \upsilon_{N_y} \end{bmatrix}^\top \in \mathbb{R}^{N_y}$$

denote the vectors of auxiliary optimization variables, over the prediction horizon. At each step i determine the velocity cost $\hat{D}$ of (5k) by querying whether $C_{\mathrm{y}}x_i + y_{\mathrm{s}} \geq 140$.

MPC Problem

Determine

$$\left\{\hat{U}^*, \check{U}^*, \hat{\Gamma}^*, \check{\Gamma}^*, V^*\right\} := \underset{\{\hat{U},\check{U},\hat{\Gamma},\check{\Gamma},V\}}{\arg\min}\ \tilde{J}\left(x_i, \hat{U}, \check{U}, \hat{\Gamma}, \check{\Gamma}, V\right)$$

with cost function

$$\tilde{J}(\cdot) := \sum_{k=1}^{N_{\mathrm{y}}} \left(\hat{\gamma}_k^2 + \check{\gamma}_k^2 + \hat{D}\upsilon_k^2\right) + \sum_{k=0}^{N_{\mathrm{u}-1}} \left(\hat{R}\hat{u}_k^2 + \check{R}\check{u}_k^2\right)$$

and subject to

$$\mathrm{x}_0 := x_i$$

$$\mathrm{x}_{k+1} := A\mathrm{x}_k + B(\hat{u}_k + \check{u}_k) \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{y}-1}}$$

$$\mathrm{y}_k := C_{\mathrm{y}}\mathrm{x}_k \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{y}}}$$

$$\mathrm{d}_k := C_{\mathrm{d}}\mathrm{x}_k \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{y}}}$$

$$\hat{u}_k := \check{u}_k := 0 \qquad \forall k \in \mathbb{Z}_{N_{\mathrm{u}}}^{N_{\mathrm{y}-1}}$$

$$\hat{u}_k \geq 0 \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{u}-1}}$$

$$\hat{u}_k \leq u_{\mathrm{MAX}} - u_{\mathrm{BASAL},(i+k)} \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{u}-1}}$$

$$\check{u}_k \leq 0 \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{u}-1}}$$

$$\check{u}_k \geq -u_{\mathrm{BASAL},(i+k)} \qquad \forall k \in \mathbb{Z}_0^{N_{\mathrm{u}}-1}$$

$$\hat{\gamma}_k \geq 0 \qquad \forall k \in \mathbb{Z}_1^{N_{\mathrm{y}}}$$

$$\hat{\gamma}_k \geq \mathrm{y}_k - (140 - y_{\mathrm{s}}) \qquad \forall k \in \mathbb{Z}_1^{N_{\mathrm{y}}}$$

$$\check{\gamma}_k \leq 0 \qquad \forall k \in \mathbb{Z}_1^{N_{\mathrm{y}}}$$

$$\check{\gamma}_k \leq \mathrm{y}_k - (80 - y_{\mathrm{s}}) \qquad \forall k \in \mathbb{Z}_1^{N_{\mathrm{y}}}$$

$$\upsilon_k \geq 0 \qquad \forall k \in \mathbb{Z}_1^{N_{\mathrm{y}}}$$

$$\upsilon_k \geq \mathrm{d}_k \qquad \forall k \in \mathbb{Z}_1^{N_{\mathrm{y}}}.$$

The finishing steps of the transformation to a QP can be performed according to standard practice, and are not described further here. Note that the above QP has $3N_{\mathrm{y}} + 2N_{\mathrm{u}}$ optimization variables, whereas the MPC problem as presented in Section II-F is over N_{u} optimization variables.

Given the solution $\{\hat{U}^*, \check{U}^*, \hat{\Gamma}^*, \check{\Gamma}^*, V^*\}$, the insulin delivery command is given by $u_i := \hat{u}_0^* + \check{u}_0^*$.

References

1. Cobelli C, Dalla Man C, Sparacino G, Magni L, De Nicolao G, Kovatchev BP. Diabetes: Models, Signals and Control. IEEE Rev. Biomed. Eng. 2009; 2:54–96. [PubMed: 20936056]
2. Harvey RA, Wang Y, Grosman B, Percival MW, Bevier W, Finan DA, Zisser H, Seborg DE, Jovanovič L, Doyle FJ III, Dassau E. Quest for the Artificial Pancreas: Combining Technology with Treatment. IEEE Eng. Med. Biol. Mag. 2010; 29(2):53–62. [PubMed: 20659841]
3. Cobelli C, Renard E, Kovatchev B. Artificial Pancreas: Past, Present, Future. Diabetes. 2011 Nov. 60:2672–2682. [PubMed: 22025773]
4. Zisser H. Clinical Hurdles and Possible Solutions in the Implementation of Closed-Loop Control in Type 1 Diabetes Mellitus. J. Diabetes Sci. Technol. 2011 Sep.5:1283–1286. [PubMed: 22027329]
5. Doyle FJ III, Huyett LM, Lee JB, Zisser HC, Dassau E. Bench to Clinic Symposia – Closed Loop Artificial Pancreas Systems: Engineering the Algorithms. Diabetes Care. 2014 In press.
6. Hovorka R. Continuous glucose monitoring and closed-loop systems. Diabetic Med. 2006 Jan.23:1–12.

7. Parker RS, Doyle FJ III, Peppas NA. A Model-Based Algorithm for Blood Glucose Control in Type I Diabetic Patients. IEEE Trans. Biomed. Eng. 1999 Feb.46:148–157. [PubMed: 9932336]
8. 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. 2004 Jul.25:905–920. [PubMed: 15382830]
9. 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.
10. 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. 2012 Jun.61:2230–2237. [PubMed: 22688340]
11. 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. 2013 May.15:386–400. [PubMed: 23544672]
12. Steil GM, Rebrin K, Darwin C, Hariri F, Saad MF. Feasibility of Automating Insulin Delivery for the Treatment of Type 1 Diabetes. Diabetes. 2006 Dec.55:3344–3350. [PubMed: 17130478]
13. Marchetti G, Barolo M, Jovanovič L, Zisser H, Seborg DE. A feedforward-feedback glucose control strategy for type 1 diabetes mellitus. J. Process Control. 2008 Feb.18:149–162. [PubMed: 19190726]
14. Grosman B, Dassau E, Zisser HC, Jovanovič L, Doyle FJ III. Zone Model Predictive Control: A Strategy to Minimize Hyper- and Hypoglycemic Events. J. Diabetes Sci. Technol. 2010 Jul.4:961–975. [PubMed: 20663463]
15. Grosman, B., Dassau, E., Zisser, H., Jovanovič, L., Doyle, FJ, III. Multi-Zone-MPC: Clinical Inspired Control Algorithm for the Artificial Pancreas. Proc. 18th IFAC World Congress; Milan, Italy. 2011 Aug.. p. 7120-7125.
16. 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. 2012 Jul. 59:1839–1849. [PubMed: 22127988]
17. 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. 2013 Nov. 7:1446–1460. [PubMed: 24351171]
18. Parker, RS., Gatzke, EP., Doyle, FJ, III. Advanced Model Predictive Control (MPC) for Type I Diabetic Patient Blood Glucose Control. AACC American Control Conf; Chicago, Illinois. 2000 Jun. p. 3483-3487.
19. Hernjak N, Doyle FJ III. Glucose Control Design Using Nonlinearity Assessment Techniques. AIChe J. 2005 Feb.51:544–554.
20. Dua P, Doyle FJ III, Pistikopoulos EN. Multi-objective blood glucose control for type 1 diabetes. Med. Biol. Eng. Comput. 2009 Mar.47:343–352. [PubMed: 19214613]
21. Walsh, J., Roberts, R. Pumping Insulin. 4. San Diego, CA, USA: Torrey Pines Press; 2006.
22. Levine, WS., editor. The Control Handbook. 2. Boca Raton, FL, USA: CRC Press; 2011.
23. Ellingsen C, Dassau E, Zisser H, Grosman B, Percival MW, Jovanovič 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. 2009 May.3:536–544. [PubMed: 20144293]
24. Maciejowski, JM. Predictive Control with Constraints. Harlow, England: Pearson/Prentice Hall; 2002.
25. Rawlings, JB., Mayne, DQ. Model Predictive Control: Theory and Design. Madison, WI, USA: Nob Hill Publishing; 2009 Aug..
26. Bemporad A, Morari M. Control of systems integrating logic, dynamics, and constraints. Automatica. 1999 Mar.35:407–427.

27. 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. 2009 Jan.3:44–55. [PubMed: 19444330]

28. 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. 2008 Jul.2:630–635. [PubMed: 19885239]

29. Harvey RA, Dassau E, Zisser H, Seborg DE, Jovanovič L, Doyle FJ III. Design of the Health Monitoring System for the Artificial Pancreas: Low Glucose Prediction Module. J. Diabetes Sci. Technol. 2012 Mar.6:1345–1354. [PubMed: 23294779]

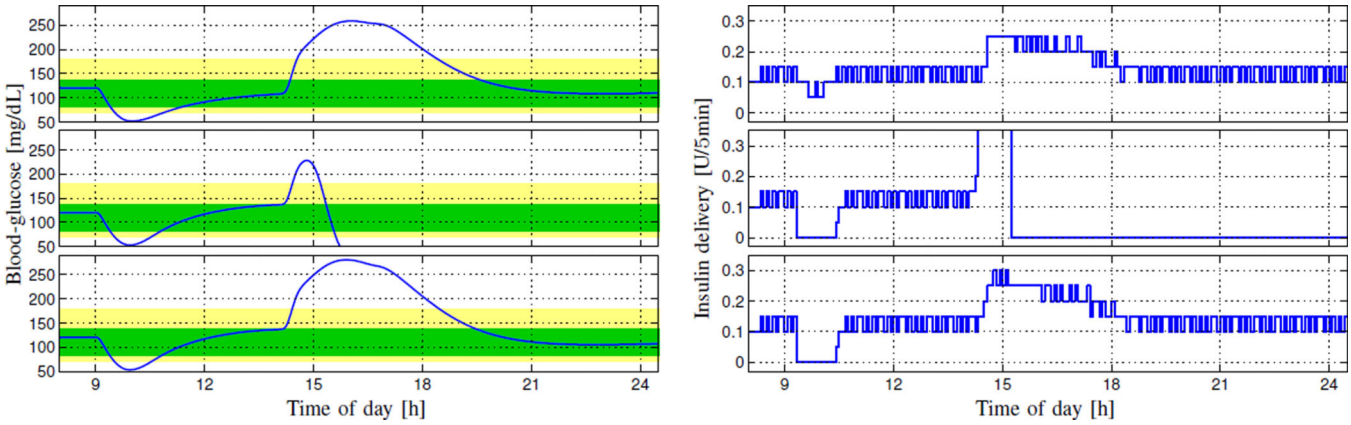

Fig. 1.

Demonstration of asymmetric input cost function. **Top**: Symmetric, high cost, $\hat{R} = \check{R} = 7000$. **Middle**: Symmetric, low cost, $\hat{R} = \check{R} = 100$. **Bottom**: Asymmetric, $\hat{R} = 7000$, $\check{R} = 100$. Green area: [80, 140] mg/dL. Yellow area [70, 180] mg/dL.

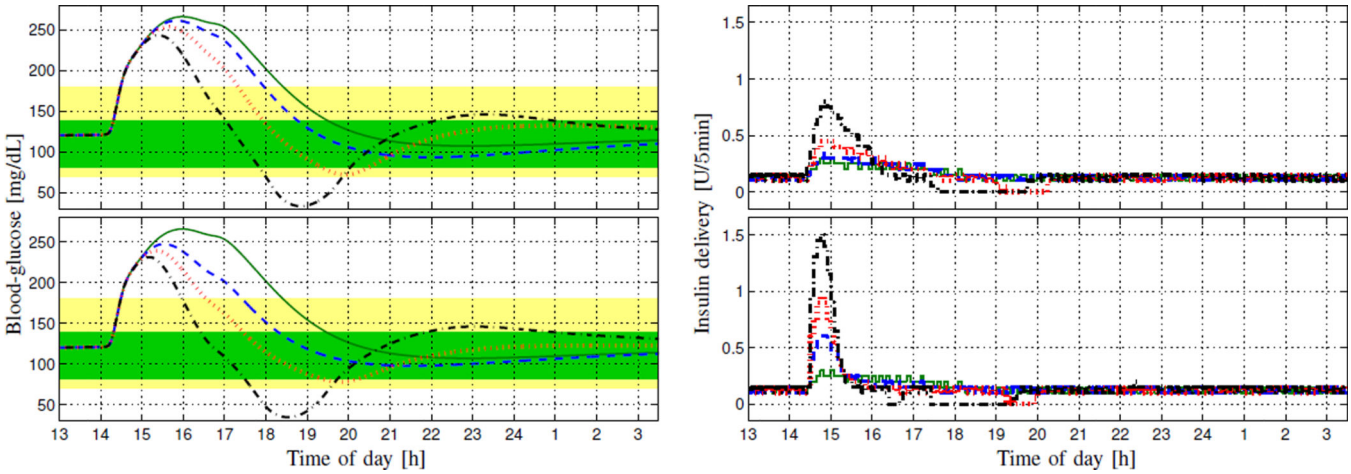

Fig. 2.

Demonstration of velocity penalty zone-MPC. **Top**: No velocity penalty ($D = 0$), i.e., standard zone-MPC: $\hat{R} = 7000$ (green solid), 5000 (blue dashed), 3000 (red dotted), 1500 (black dash-dotted). **Bottom**: Velocity penalty with $\hat{R} = 7000$: $D = 0$ (green solid), 1500 (blue dashed), 3000 (red dotted), 6000 (black dash-dotted). Green area: [80, 140] mg/dL. Yellow area [70, 180] mg/dL.

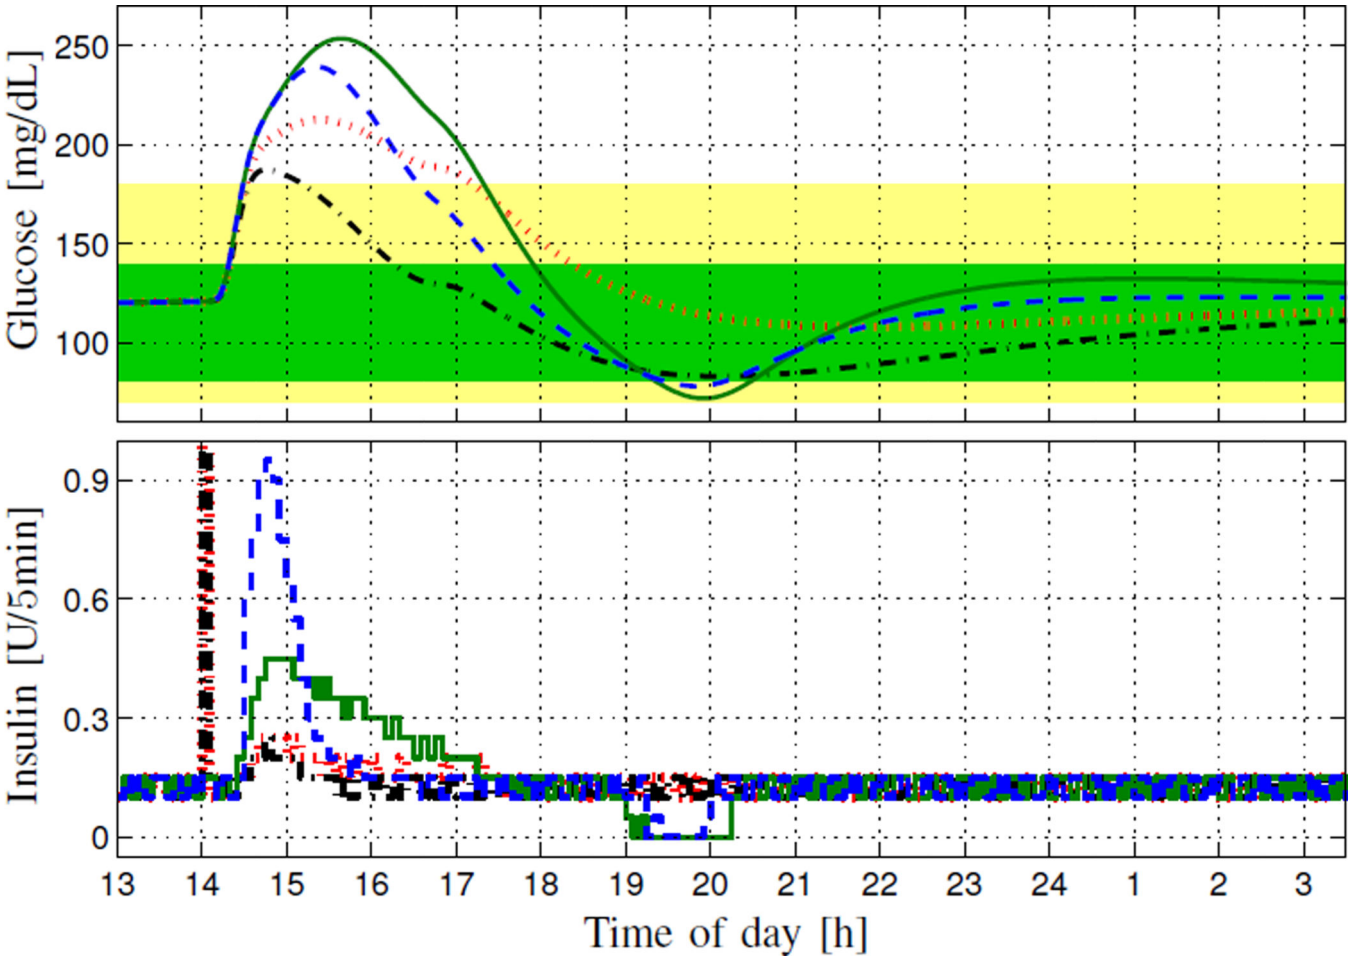

Fig. 3.

Green solid: Third most aggressive MPC *without* velocity penalty ($\hat{R} = 5000$, $D = 0$, red-dotted in upper subplots of Fig. 2). **Blue dashed**: Third most aggressive MPC *with* velocity penalty ($\hat{R} = 7000$, $D = 3000$, red-dotted in lower subplots of Fig. 2). **Red dotted**: 50% optimal meal-bolus delivered at mealtime, with most conservative MPC *without* velocity penalty ($\hat{R} = 7000$, $D = 0$, green-solid in Fig. 2). **Black dash-dotted**: 100% optimal meal-bolus delivered at mealtime, with most conservative MPC *without* velocity penalty ($\hat{R} = 7000$, $D = 0$, green-solid in Fig. 2). Green area: [80, 140] mg/dL. Yellow area [70, 180] mg/dL.

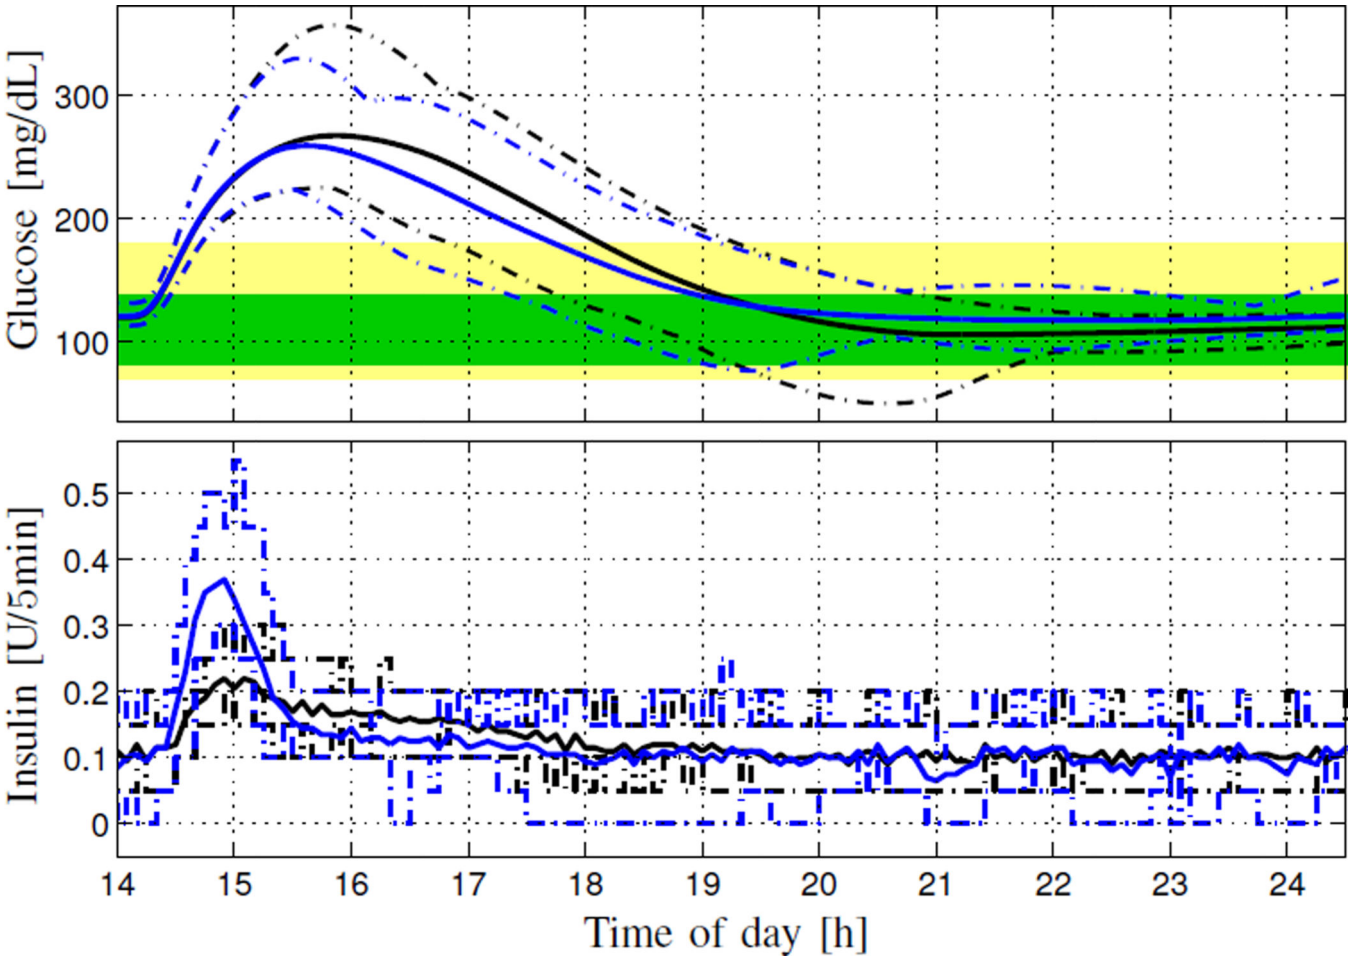

Fig. 4.

10 *in silico* adults: Traditional-MPC (black), proposed-MPC (blue). Mean (solid) and min-max envelopes (dash-dotted). Green area: [80, 140] mg/dL. Yellow area [70, 180] mg/dL.

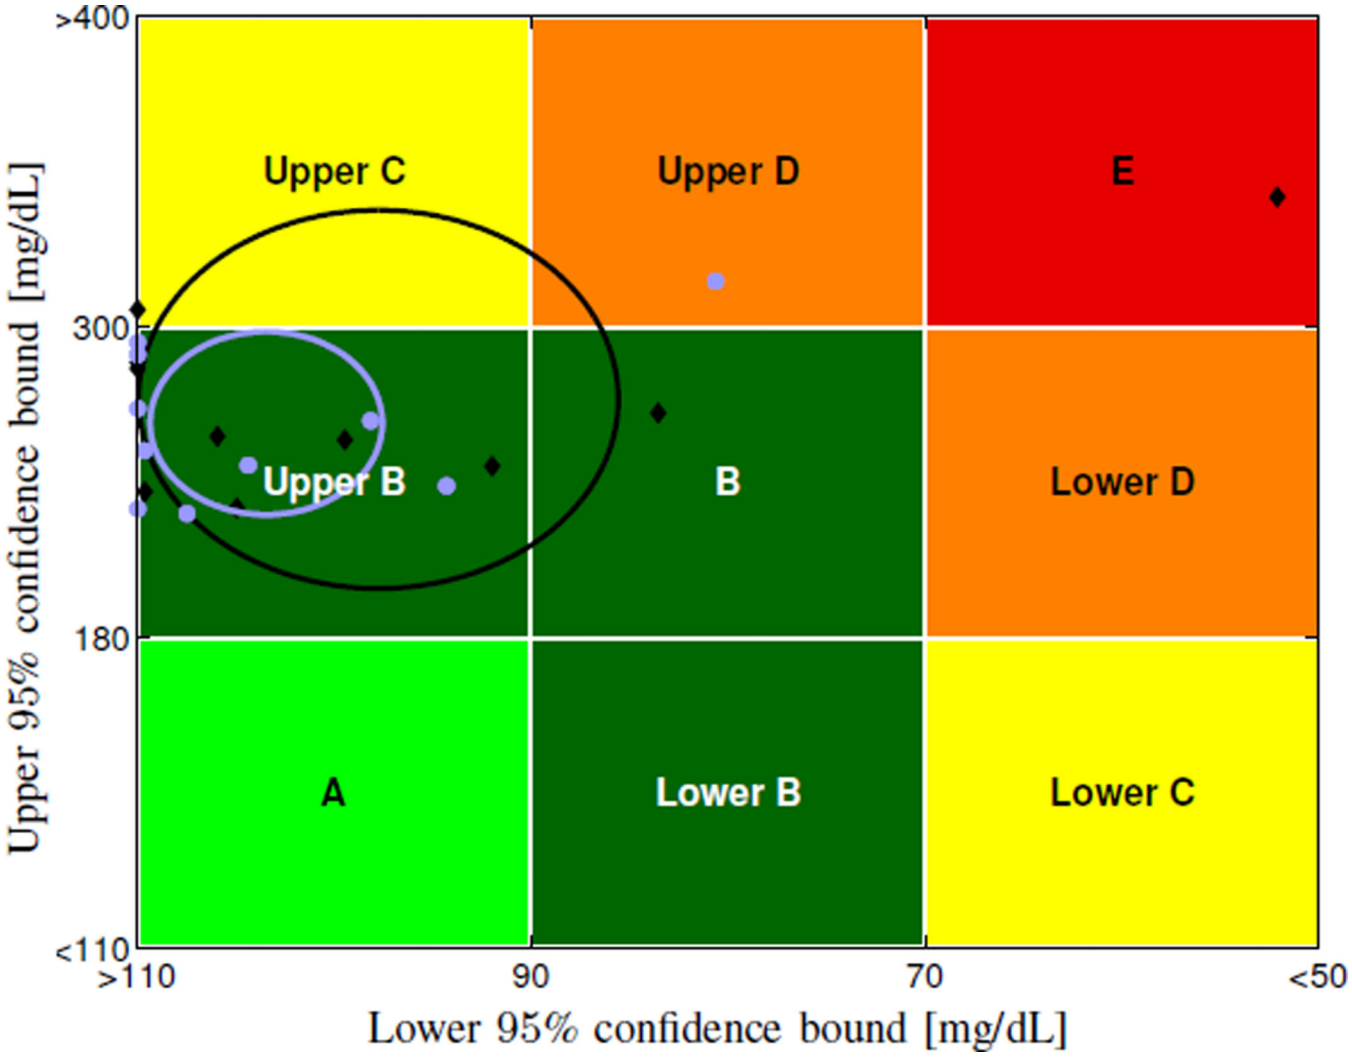

Fig. 5.

CVGA plot for 10 *in silico* adults. Traditional-MPC (black), proposed-MPC (blue). Circles: Centered on mean, standard deviation radius. (A,B,C,D,E) %-age zone-membership: Traditional-MPC (0,80,10,0,10), proposed-MPC (0,90,0,10,0).