HHS Public Access

Author manuscript

Automatica (Oxf). Author manuscript; available in PMC 2019 May 01.

Published in final edited form as:

Automatica (Oxf). 2018 May ; 91: 105–117. doi:10.1016/j.automatica.2018.01.025.

Velocity-weighting & velocity-penalty MPC of an artificial pancreas: Improved safety & performance

Ravi Gondhalekar[a], **Eyal Dassau**[a], and **Francis J. Doyle III**[a]

[a]Harvard John A. Paulson School of Engineering & Applied Sciences, Harvard University, Cambridge, MA, USA

Abstract

A novel Model Predictive Control (MPC) law for the closed-loop operation of an Artificial Pancreas (AP) to treat type 1 diabetes is proposed. The contribution of this paper is to simultaneously enhance both the safety and performance of an AP, by reducing the incidence of controller-induced hypoglycemia, and by promoting assertive hyperglycemia correction. This is achieved by integrating two MPC features separately introduced by the authors previously to independently improve the control performance with respect to these two coupled issues. *Velocity-weighting* MPC reduces the occurrence of controller-induced hypoglycemia. *Velocity-penalty* MPC yields more effective hyperglycemia correction. Benefits of the proposed MPC law over the MPC strategy deployed in the authors' previous clinical trial campaign are demonstrated via a comprehensive in-silico analysis. The proposed MPC law was deployed in four distinct US Food & Drug Administration approved clinical trial campaigns, the most extensive of which involved 29 subjects each spending three months in closed-loop. The paper includes implementation details, an explanation of the state-dependent cost functions required for velocity-weighting and penalties, a discussion of the resulting nonlinear optimization problem, a description of the four clinical trial campaigns, and control-related trial highlights.

Keywords

Model predictive control; Safety-critical control; Nonlinear optimization; State-dependent costs; Artificial pancreas

1 Introduction

People with Type 1 Diabetes Mellitus (T1DM) require exogenous insulin to regulate their Blood-Glucose (BG) concentration and, despite burdensome insulin treatment, tend to suffer difficulty maintaining healthy BG levels. Insufficient insulin causes hyperglycemia, which has few consequences if brief, but if prolonged may beget chronic health problems. Over-insulinization induces hypoglycemia, which has very near-term consequences, including

Corresponding author: Ravi Gondhalekar.

death if severe. The objective of this work is an effective and safe control law for an Artificial Pancreas (AP) (Cobelli et al. 2011, Bequette 2012, Doyle III et al. 2014, Haidar 2016) that performs automated insulin delivery to people with T1DM, with the goals of improving health outcomes and reducing treatment burden. The focus is Model Predictive Control (MPC) (Maciejowski 2002, Rawlings & Mayne 2009), considered in an AP context by others in (Hovorka et al. 2004, Magni et al. 2009, Breton et al. 2012, Bequette 2013, Turksoy et al. 2014, Zavitsanou et al. 2015), and by the authors in (Parker et al. 1999, Grosman et al. 2010, Gondhalekar et al. 2013, Dassau et al. 2015, Gondhalekar et al. 2016, Huyett et al. 2017, Dassau et al. 2017). Of specific interest is an AP for outpatient use, requiring (currently) a Continuous Glucose Monitor (CGM) sensor that provides feedback of BG measurements, and a Continuous Subcutaneous Insulin Infusion (CSII) pump to deliver insulin as commanded by the control law. Subcutaneous sensing and actuation suffer long delays; control action may take an hour to effect BG concentration changes. This major difficulty with CGM and CSII devices may result in controllers that over-deliver insulin during hyperglycemia, inducing subsequent hypoglycemia, and that fail to respond suitably quickly at the start of a BG rise, leading to more severe hyperglycemia. A more conservatively tuned control law may avert controller-induced hypoglycemia, but undesirably promotes hyperglycemia. Analogously, hyperglycemia is better corrected by more aggressive control laws, which abet hypoglycemia. In a traditional MPC framework this tradeoff limits the achievable performance, because safeguarding from hypoglycemia is the overruling priority. Theory on this performance tradeoff, in the context of linear positive systems and T1DM, was developed in (Goodwin et al. 2015). This work focuses on the design of single hormone AP controllers, whose actuation is only insulin infusion, thus possess no control authority to raise glucose levels. The challenge of circumnavigating the aforementioned controls tradeoff may partially be alleviated in bi-hormonal APs by adding glucagon infusion as a further control input (Cobelli et al. 2011, Bequette 2012, Doyle III et al. 2014, Haidar 2016), although research to date does not empower one to reach final conclusions pertaining to the medical efficacy and technological feasibility of bi-hormonal AP systems.

The goal, and main technical contribution, of this work is a novel MPC problem structure that simultaneously elevates the control performance and significantly increases the safety of an AP, compared to the MPC law employed in the authors' previous clinical trial campaign (Dassau et al. 2015, Gondhalekar et al. 2016). These improvements advance the resulting AP towards the best achievable performance, within the envelope of the aforementioned tradeoff (Goodwin et al. 2015), but *without* tradeoff. In (Gondhalekar et al. 2015*a*,*b*) the authors introduced *velocity-weighting* MPC, whereby the MPC cost to penalize the BG deviation is velocity-dependent, i.e., the BG cost parameter itself is a function of the rate of change of the BG output. In a state of hyperglycemia the cost smoothly vanishes with an increasingly negative BG rate of change. The result is that stubborn hyperglycemia is corrected aggressively, but the controller 'backs off' when a hyperglycemic state is in the process of being corrected. In (Gondhalekar et al. 2014*b*) the authors introduced *velocity-penalty* MPC, whereby the actual BG rate of change is penalized using a cost that is state-dependent. Positive BG rates of change are then penalized directly, but only when the BG level is high. This results in a controller that provides an 'extra kick' during the uphill leg of large

hyperglycemic excursions. The algorithmic contribution of this work is to combine these two functionalities, with an enhancement of the velocity-penalty mechanism, and deploy them in tandem. The resulting controller responds with significantly increased insulin delivery at the start of large hyperglycemic excursions, e.g., as experience after an unannounced meal, and delivers very conservatively after the peak of the excursion has been reached. The resulting insulin delivery profile is thus prominently shifted towards the start of the excursion, leading to lower BG peaks, a stark reduction in hypoglycemia risk, and reduced BG variability.

Paper structure

The proposed MPC law is described in Sec. 2. Only few AP specifics are provided there, to highlight the novel MPC problem structure and yield an exposition more accessible to readers unfamiliar with AP control. For completeness, extensive AP related details and parameters, required for algorithm reproducibility (an intended contribution), are provided in the appendices. Velocity-weighting and velocity-penalty MPC notions are discussed in detail in Sec. 3. Details of the required optimization and implementation are presented in Sec. 4. Benefits over, and differences to, the authors' previous MPC strategy of (Dassau et al. 2015, Gondhalekar et al. 2016) are demonstrated in Sec. 5 using the Universities of Virginia (UVA)/Padova United States (US) Food & Drug Administration (FDA) accepted metabolic simulator (Kovatchev et al. 2009, Dalla Man et al. 2014). In the US this simulator is the accepted manner for testing AP controllers prior to deployment in clinical trials, thus the in-silico analysis is comprehensive, and mathematical performance proofs are not presented. The proposed MPC law was deployed in multiple and extensive clinical trials, and some trial highlights are examined in Sec. 6.

2 MPC design

2.1 Preliminaries

The proposed MPC law uses the discrete-time linear time-invariant (LTI) model (1) of insulin-glucose dynamics, with sample-period $T := 5$ [min], insulin (control) input $u \in \mathbb{R}$ [U] ('Units' insulin), BG output $y \in \mathbb{R}$ [mg/dL], and BG velocity output $v \in \mathbb{R}$ [mg/dL/min]. The control input u is the insulin deviation from the subject's basal rate, integrated over sample-period T. The basal rate is a person's baseline insulin requirement and is non-negative. The BG output y is the deviation from a BG setpoint $y_s \in \mathbb{R}_{>0}$. Deviations u_i and y_i may be negative, but absolute insulin inputs and BG concentrations are non-negative (see App. A.1 for definitions). Later figures depict absolute values. The velocity output v is an approximation of the rate of change of BG output y. At each step i we are provided an estimate of model state x_i. For model and estimator details consult App. A.1.

The MPC law penalizes predicted BG excursions from a time-dependent BG target zone with lower and upper bounds $\check{\zeta}$ and $\hat{\zeta}$, respectively. Details of the BG target zone are provided in App. A.2. The zone-excursion function $Z: \mathbb{R} \times \mathbb{N} \to \mathbb{R}$ is defined in (2).

The controller must enforce input constraint (3). The lower bound $\check{\xi}$ enforces that absolute insulin infusion is non-negative, i.e., $\check{\xi}$ is selected to be the negative of the basal insulin

required over sample-period T. A subject's basal rate is generally time-dependent, hence $\underset{\sim}{\xi}$ varies with i. The basal rate is non-negative, hence $\underset{\sim}{\xi}_i \geq 0$ for all i. The upper bound $\hat{\xi}$ is a function of both the time of day the controller is called and also the insulin infusion history. Thus, $\hat{\xi}$ is time-dependent but satisfies $\hat{\xi}_i \geq 0$ for all i. Details of the insulin constraints are provided and discussed in Appendices A.3 & A.4. Note that $u_i = 0$ (i.e., basal rate) is a feasible control input for all i.

$$x_{i+1} = Ax_i + Bu_i, \quad y_i = C_y x_i, \quad v_i = C_v x_i \qquad (1)$$

$$Z(y, i) := \arg\min_{\alpha \in \mathbb{R}} \left\{ \alpha^2 \middle| y + y_s - \alpha \in \left[\underset{\sim}{\zeta}_i, \hat{\zeta}_i \right] \right\} \qquad (2)$$

$$u_i \in \mathbb{U}_i := \left[\underset{\sim}{\xi}_i, \hat{\xi}_i \right] \ni 0 \qquad (3)$$

2.2 MPC problem

Let u, x, y, and v denote the predicted input u, state x, BG output y, and BG velocity output v, respectively. Further denote by $\mathbb{Z}_a^b$ the set of consecutive integers $\{a, \ldots, b\}$, by $N_y \in \mathbb{N}$ the prediction horizon, and by $N_u \in \mathbb{Z}_1^{N_y}$ the control horizon. Then, MPC performs closed-loop control action by applying at each step i the first element u_0^* of the optimal, predicted control input trajectory $\left\{ \mathrm{u}_0^*, \ldots, \mathrm{u}_{N_u - 1}^* \right\}$, characterized as follows.

MPC Problem—Determine

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

with cost function

$$J(\cdot) := \sum_{k=1}^{N_y} \left[\underset{\sim}{z}_k^2 + Q(\mathrm{v}_k) \hat{z}_k^2 + \hat{D} \hat{\mathrm{v}}_k^2 \right] + \sum_{k=0}^{N_u - 1} \left[\hat{R} \hat{\mathrm{u}}_k^2 + \underset{\sim}{R}\, \underset{\sim}{\mathrm{u}}_k^2 \right] \qquad (4)$$

subject to

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

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

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

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

$$\mathrm{u}_k \in \mathbb{U}_{i+k} \qquad \forall k \in \mathbb{Z}_0^{N_\mathrm{u}-1} \quad (6)$$

$$\mathrm{u}_k := 0 \qquad \forall k \in \mathbb{Z}_{N_\mathrm{u}}^{N_\mathrm{y}-1} \quad (7)$$

$$\hat{\mathrm{u}}_k := \max\left\{\mathrm{u}_k, 0\right\} \qquad \forall k \in \mathbb{Z}_0^{N_\mathrm{u}-1} \quad (8)$$

$$\breve{\mathrm{u}}_k := \min\left\{\mathrm{u}_k, 0\right\} \qquad \forall k \in \mathbb{Z}_0^{N_\mathrm{u}-1} \quad (9)$$

$$\mathrm{z}_k := Z\left(\mathrm{y}_k, i+k\right) \qquad \forall k \in \mathbb{Z}_0^{N_\mathrm{y}}$$

$$\hat{\mathrm{z}}_k := \max\left\{\mathrm{z}_k, 0\right\} \qquad \forall k \in \mathbb{Z}_0^{N_\mathrm{y}} \quad (10)$$

$$\breve{\mathrm{z}}_k := \min\left\{z_k, 0\right\} \qquad \forall k \in \mathbb{Z}_0^{N_\mathrm{y}} \quad (11)$$

$$\hat{\mathrm{v}}_k := \max\left\{\mathrm{v}_k, 0\right\} \qquad \forall k \in \mathbb{Z}_0^{N_\mathrm{y}} \quad (12)$$

$$\hat{D} := \begin{cases} D & \text{if } y_i + y_\mathrm{s} \in \mathbb{D} \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Eq. (6) enforces input constraint (3) across the control horizon. Eq. (7) implies that beyond the control horizon the basal rate is delivered. Eqs. (8) & (9) provide non-negative and non-positive deviations of input u, and facilitate an asymmetric input cost in (4). Eqs. (10) & (11) provide non-negative and non-positive zone excursions, and facilitate an asymmetric penalization of the zone excursion in (4). Eq. (12) provides the non-negative BG velocity to penalize in (4) using the state-dependent velocity-penalty cost $\hat{D}$ of (13) (see Sec. 3.2). Note that $\hat{D}$ depends on the current output y_i, thus the same $\hat{D}$ is employed at each prediction step k. The set $\mathbb{D} \subset \mathbb{R}_{>0}$ is defined such that the penalty D is only 'activated' when the current BG output y_i indicates mild hyperglycemia.

The proposed MPC law uses $N_{\mathrm{y}} = 9$, $N_{\mathrm{u}} = 5$, $\hat{R} := 6500$, $\check{R} := 100$, $D := 1000$, $\mathbb{D} := [140, 180]$, and $Q(\mathrm{v})$ of (14) with $\varepsilon := 10^{-6}$, depicted in Fig. 1. See (Gondhalekar et al. 2016) for comments on parameter selection. The algorithmic contribution of the proposed MPC law over that of (Gondhalekar et al. 2016) is the use of $Q(\mathrm{v})$ and $\hat{D}\hat{\mathrm{v}}_k^2$ in (4). The MPC cost function of (Gondhalekar et al. 2016) is achieved when $Q(\mathrm{v}) = 1$, $D = 0$, and $\hat{R} = 7000$ (see Table 1); this tuning is henceforth termed 'previous MPC'. Note that the *actual* previous MPC strategy of (Gondhalekar et al. 2016) is different in various ways, e.g., the BG target zone definitions and input constraints (see Appendices A.2–A.4).

$$Q(\mathrm{v}) := \begin{cases} 1 & \text{if } \mathrm{v} \geq 0 \\ \varepsilon & \text{if } \mathrm{v} \leq -1 \\ \frac{1}{2}[\cos(\mathrm{v}\pi)(1-\varepsilon) + (1+\varepsilon)] & \text{otherwise} \end{cases} \quad (14)$$

2.3 Recursive feasibility & stability

The hard constraints are point-wise in time input constraints, and the predicted input trajectory $\mathrm{u}_k = 0 \forall k \in \mathbb{Z}_0^{N_\mathrm{u}-1}$ is a feasible solution for all steps i and states x_i (see (3)). Thus, recursive feasibility is guaranteed. The use of asymmetric costs on input, velocity, and BG target zone excursions affects only the cost-penalization via (4).

A person's BG is never constant, but it is bounded. The purpose of the proposed MPC law is improved safety and performance of the AP system, not to impose closed-loop stability. LTI model (1) is open-loop stable (see App. A.1), thus by the target-zone nature of the proposed MPC law closed-loop stability with respect to model (1) is straightforwardly guaranteed in a

neighborhood of the model's output setpoint y_s. In the US the UVA/Padova US FDA accepted metabolic simulator (Kovatchev et al. 2009, Dalla Man et al. 2014) is the tool to verify AP controllers prior to clinical deployment. It contains a complex, nonlinear model of much higher order than model (1), and simulates sensor artifacts and variations in subject parameters that cannot be measured and explicitly designed for. The focus of this paper is the presentation and discussion of a novel MPC law that has successfully withstood the rigors of both comprehensive in-silico testing using the UVA/Padova simulator, and also multi-month deployments in extensive clinical trials.

3 Velocity-weighting & velocity-penalty MPC

3.1 Velocity-weighting MPC

The **MPC Problem** employs N_y T = 45 min predictions of BG values. This is short compared to the length of hyperglycemic excursions experienced, e.g., after large meals. CGM noise and large plant-model mismatch render long prediction horizons disadvantageous, in the authors' experience. However, short horizons yield controllers that undesirably persist to command insulin delivery surpassing the basal rate when BG predictions exceed the BG target zone's upper bound $\hat{\zeta}$, but are subsiding, i.e., when a state of hyperglycemia is in the process of correcting itself. This supplementary insulin contributes to the risk of controller-induced hypoglycemia. The purpose of the proposed velocity-weighting MPC is to mitigate controller-induced hypoglycemia that is due to insulin over-delivery during waning hyperglycemia. It was first proposed in (Gondhalekar et al. 2015*b*) and further discussed in (Gondhalekar et al. 2015*a*).

If Q(v) = 1 and D = 0 then MPC cost function (4) has the structure of the cost function in (Gondhalekar et al. 2016), where predicted zone excursions z are weighted symmetrically, i.e., non-negative and non-positive excursions $\hat{z}$ and $\check{z}$ are penalized equally. The velocity-weighting function Q(v) of (14) modulates the cost accrued by $\hat{z}$, i.e., modifies the cost only during predicted hyperglycemia. By reducing the zone excursion cost for an increasingly negative BG velocity v the design of Q(v) mildly reduces control action at a low, negative BG velocity, but practically eliminates the cost contribution when BG velocities exceed 1 mg/dL/min downwards, considered steep. Thus, on the downhill leg of a hyperglycemic excursion the proposed controller delivers less insulin while the BG level is descending, but can deliver aggressively during persistent (i.e., v ≥ 0) hyperglycemia. Note that velocity-weighting leaves unchanged the controller's response to predicted BG trajectories that solely rise. Note further that rendering the non-negative zone excursion $\hat{z}$ ineffectual in (4) when v is strongly negative yields u_i = 0, i.e., in absolute terms the basal rate is commanded, not the pump suspended. Due to the open-loop stable nature of the 'plant' and model this has a stabilizing effect and reduces oscillations (see Sec. 5).

Velocity-weighting MPC may be interpreted as realizing a time-varying upper bound $\hat{\zeta}$ of the BG target zone. Fig. 2 depicts an illustration of how the BG target zone's upper bound $\hat{\zeta}$ *would* be varying, *were* the zone excursion cost unity as in (Gondhalekar et al. 2016), in order to accrue equal cost in (4). Importantly, the time-variation is based on CGM *feedback*, because predicted BG velocities v are based on the current state estimate x_i. Time-dependent

BG setpoints[1] were employed by others to reduce controller-induced hypoglycemia. In (Hovorka et al. 2004, Marchetti et al. 2008, Eren-Oruklu et al. 2009, Boiroux et al. 2012) a BG reference trajectory with respect to which predicted BG deviations are penalized is generated at each step. During hyperglycemia the reference trajectory is initialized to the current BG value and has a pre-defined, negative rate of change. Thus, the reference trajectory is based on feedback, but of BG, not its velocity v: The same level of hyperglycemia results in the same reference trajectory regardless of whether the hyperglycemia is persistent (i.e., $v \geq 0$) or is in the process of correcting itself (i.e., $v < 0$). The proposed velocity-weighting MPC approach is able to discriminate between these cases and adapt (4) accordingly. The reference trajectories of the aforementioned works are generated equally irrespective of whether meals are announced. In (Boiroux et al. 2011) meal-announcement triggers the computation of an optimal meal-bolus and associated predicted BG trajectory that is employed as a BG reference trajectory about which to penalize MPC predictions. This approach may suffer weaknesses. First, the reference trajectory is generated based on assumptions, e.g., a model and the meal-size, that may not be correct. Nevertheless it is employed, in disregard of the *true* BG response, without exploiting feedback. In reality models and meal-size estimates are inaccurate, and other factors, e.g., meal-composition, greatly affect the actual BG response. The proposed velocity-weighting MPC is able to modulate its weighting (its 'virtual' reference trajectory) based on feedback. For a given meal-size a high fat content may result in persistent hyperglycemia, which the velocity-weighting would discern and penalize more heavily than had an equi-sized meal with low fat content been consumed. A further weakness of (Boiroux et al. 2011) is that the method is of relevance only when provided by the user with, and during the time following, meal-announcement. In contrast, the proposed velocity-weighting MPC is 'active' continuously, advantageous for both unannounced and announced meals, and is more effective for the former, more challenging situation (see Sec. 5).

3.2 Velocity-penalty MPC

Meal ingestion that is unannounced poses a major challenge for any AP controller, particularly one based on CGM feedback and CSII pumps, due to the large delay in subcutaneous sensing and actuation. Meal consumption inevitably causes a BG rise, and an AP must respond by delivering more insulin. Importantly, earlier delivery is more effective. The objective of velocity-penalty MPC is strong insulin delivery during rapid BG ascents when the subject is *entering* hyperglycemia. Velocity-penalty MPC results in a simple switching strategy that is state-dependent, i.e., is based on feedback, not meal-announcement. Switching controllers for an AP have been proposed before but differently (Marchetti et al. 2008, Messori et al. 2015, Colmegna, Sánchez-Peña, Gondhalekar, Dassau & Doyle III 2016*a*,*b*).

When $y_I + y_S \in \mathbb{D} = [140, 180]$ mg/dL, i.e., a state close to mild hyperglycemia, then $\hat{D} = D > 0$, resulting in the BG velocity v being penalized in (4). The design objective of velocity-penalty MPC is to aggressively drive insulin delivery at the *start* of hyperglycemia. Penalizing only the non-negative BG excursion $\hat{z}$ requires $\hat{z}$ to be sufficiently large, i.e., to

[1]Throughout *this* paragraph the notions of BG setpoints and BG target zones are deemed synonymous.

accrue adequate cost, before the MPC law significantly commands insulin. This induces delay. However, all hyperglycemic excursions start low and rise based on BG velocity. Thus, directly penalizing the non-negative velocity $\hat{v}$ allows the controller to respond sooner. Large, meal-fueled hyperglycemic excursions frequently traverse the BG interval $\mathbb{D}$ with a high velocity, incurring a large MPC cost, leading to abruptly increased insulin delivery. In contrast, excursions that turn out to be small tend to have lower velocities $\hat{v}$, thus are penalized only a little. The velocity-penalty term is quadratic in $\hat{v}$, thus a BG trajectory that undulates nonchalantly (i.e., $\hat{v} \ll 1$) within $\mathbb{D}$ accrues nearly no cost, leading to no significant extra insulin delivery. In the original velocity-penalty proposal of (Gondhalekar et al. 2014*b*) the interval within which the velocity-penalty is activated was [140, ∞) mg/dL. The upper limit of 180 mg/dL was introduced to prevent over-delivery. When the current BG value exceeds this upper bound the usual output cost on the BG excursion $\hat{z}$ results in suitable control action.

3.3 Velocity-weighting + penalty: Demo & comments

Velocity-weighting and velocity-penalty MPC were introduced independently in (Gondhalekar et al. 2015*b*, 2014*b*). A contribution of the proposed MPC approach is their combination and tuning to work in concert. An in-silico comparison of proposed MPC and previous MPC is depicted in Fig. 3. Due to the velocity-penalty, proposed MPC effects a precipitous rise in insulin delivery at the start of the response to the unannounced 90 grams carbohydrate (gCHO) meal at 18:30, reducing the BG peak. However, due to the high BG velocity and rapid BG traversal of the interval $\mathbb{D}$, this insulin peak is short-lived. Due to velocity-weighting, insulin delivery is reduced with respect to previous MPC immediately following the main BG peak. This in-silico subject experiences a double-peak. Proposed MPC resumes elevated insulin commands when hyperglycemia appears persistent, as based on CGM feedback. In contrast the previous MPC law persistently commands elevated insulin throughout the double-peak phenomenon, but commands no substantial peak at the onset of hyperglycemia. No actual hypoglycemia was induced in this example, but with previous MPC the controller was forced to suspend the pump for one hour due to *predicted* hypoglycemia, inducing BG rebound. Proposed MPC yields a less oscillatory, albeit slightly more sluggish, return to the safe zone. Furthermore, proposed MPC commands slightly less total insulin than previous MPC.

Velocity-weighting renders (4) state-dependent with respect to *predicted* state x_k ($v_k = C_v x_k$). Velocity-penalties are state-dependent with respect to *actual* estimated state x_j ($y_j = C_y x_j$). This distinction is important and deliberate, and conforms to the central tenet of zone-MPC; limited intervention, i.e., that remedial action occur when feedback unequivocally indicates its need. Velocity-weighting only attenuates insulin delivery compared to using $Q(v) = 1$, is inherently safer, and is thus ‘active’ at each prediction step. Undesirable outcomes are limited to mild under-correction of hyperglycemia that is already subsiding, but not over-correction, compared to using $Q(v) = 1$. In contrast, velocity-penalties may induce stark insulin infusion increases, and are only safe when BG unambiguously exceeds normoglycemia. During hypoglycemia it is not prudent to command significant insulin based on BG *predictions* entering hyperglycemia, where these predictions are based on an erroneous model and noisy CGM feedback.

The proposed MPC law is strictly an *extension* of (Gondhalekar et al. 2016), and both velocity-weighting and velocity-penalties can be 'tuned out' by selecting $Q(\mathrm{v}) = 1$ and $D = 0$. Thus, importantly, the new, experimental additions can be introduced conservatively to augment the previous MPC law, which performed well in trials (Dassau et al. 2015). Simulator testing indicated a velocity-penalty D higher than selected would work well. However, due to the inherent risk of this new feature it was included with a less assertive tuning.

The reader may wonder whether a velocity-weighting function satisfying $Q(\mathrm{v}) > 1$ for some v is advantageous, or whether such a $Q(\mathrm{v})$ may replace the velocity-penalty. The authors experimented with various (details omitted) more general $Q(\mathrm{v})$ definitions, but concluded the experiments failed to yield a desirable control law.

4 Implementation & computation

4.1 Nonlinear optimization via sequence of QPs

The **MPC Problem** is a nonlinear optimization problem, but is 'nearly' a Quadratic Program (QP), and is solved via a sequence of convex, continuous QPs. Were $Q(\mathrm{v}_k)$ of (14) known $\forall k \in \mathbb{Z}_1^{N_\mathrm{y}}$ then the **MPC Problem** would be a QP: Asymmetric input costs are implemented by explicitly splitting the predicted input u into non-negative/non-positive parts $\hat{\mathrm{u}}/\check{\mathrm{u}}$, resulting in $2N_\mathrm{u}$ decision variables associated with the model input. Asymmetric costs on non-negative/non-positive zone-excursions $\hat{\mathrm{z}}/\check{\mathrm{z}}$, and the non-negative velocity $\hat{\mathrm{v}}$, are implemented using auxiliary decision variables that over/under-bound the excursion z and velocity v. This results in $3N_\mathrm{y}$ decision variables associated with the model outputs. Thus, with $Q(\mathrm{v}_k)\ \forall k \in \mathbb{Z}_1^{N_\mathrm{y}}$ known the **MPC Problem** is a QP with $2N_\mathrm{u} + 3N_\mathrm{y}$ decision variables. This QP *must* be solved subject to hard constraints. (See the appendix of (Gondhalekar et al. 2014*b*) for details.)

However, v_k, thus $Q(\mathrm{v}_k)$, are initially unknown. These are determined iteratively: An initial guess of predicted input sequence $\left\{\mathrm{u}_0, \ldots, \mathrm{u}_{N_\mathrm{u}-1}\right\}$ of zeros, given current state x_i, leads to predicted velocities $\left\{\mathrm{v}_1, \ldots, \mathrm{v}_{N_\mathrm{y}}\right\}$ and associated costs $\left\{Q(\mathrm{v}_1), \ldots, Q(\mathrm{v}_{N_\mathrm{y}})\right\}$. Based on these costs a QP is formulated and solved, yielding an updated predicted input sequence $\left\{\mathrm{u}_0, \ldots, \mathrm{u}_{N_\mathrm{u}-1}\right\}$, and, in turn, an updated set of costs $Q(\cdot)$ and an updated QP. This iterative process is continued until there is no significant change in the costs $\left\{Q(\mathrm{v}_1), \ldots, Q(\mathrm{v}_{N_\mathrm{y}})\right\}$. The function $Q(\cdot)$ is smooth to aid convergence. The solution method is akin to sequential quadratic programming, but simpler, because no numerical differentiation is required.

The v_k values heavily depend on the state estimate x_i, and are affected only little by predicted inputs u_k. Thus, not much is to be gained from a large number of QP iterations. If $\mathrm{v}_k \geq 0 \forall k \in \mathbb{Z}_1^{N_\mathrm{y}}$ then only one iteration is required. The numbers of QP iterations required

for the MPC optimizations of the in-silico study of Sec. 5 and first round of clinical trials of Sec. 6 are tabulated in Table 2. The majority of optimizations require only a single iteration. All cases required ≤ 9 iterations. Testing indicated a cap, e.g., of 4 iterations, would yield closely approximate solutions. In software the number of iterations is hard-limited to 10. Thus, there is no possibility the optimization terminates without a usable solution.

4.2 Matlab for PC, Java for smartphones & PC

The proposed MPC strategy is implemented in two ways. It was coded in Matlab for PCs, to test with the UVA/Padova US FDA accepted metabolic simulator (Kovatchev et al. 2009, Dalla Man et al. 2014) and historic clinical data (see Sec. 5). It was also implemented in Java for deployment on Nexus 5 smartphones within the UVA Diabetes Assistant (DiAs) AP system (Kovatchev et al. 2012, Keith-Hynes et al. 2013) (see Sec. 6). The QP solver employed is quadprog in Matlab, `jOptimizer` (Trivellato et al. n.d.) in Java. When performing simulations on a PC the proposed MPC strategy may be executed in one of three ways: (1) Matlab only; (2) Java only; (3) concurrent Matlab and Java with cross-checking. Option (1) is to develop, test, and debug the control *strategy*; it is just easier in a high-level language. Option (2) is to test and debug the Java *code*. Option (3) checks that the Matlab and Java controllers are functionally identical, by programmatically verifying that each final, and salient intermediate, variable is the same in the Matlab and Java 'spaces'. This setup resulted in high confidence that the Java code is correct, and permitted verification of the Java controller using the Matlab-based UVA/Padova US FDA accepted metabolic simulator.

4.3 Computational burden on Nexus 5

The Nexus 5 has a 2.26 GHz quad-core processor and 2 GB memory, and computes solutions to the nonlinear MPC optimization problem suitably quickly. During the first round of clinical trials (see Sec. 6) the proposed MPC strategy was computed 21, 821 times. Table 3 contains statistics; the difference between the computation-times of the entire controller call (CTRL all) and the MPC optimization (MPC only) is the overhead for non-MPC computations, e.g., data handling, state-estimation, and log-file writing. Controller calls lasted on average 1.2 seconds, where roughly 0.3 seconds is the overhead, which is fairly constant. Table 3 contains the minima, means, and maxima of the computation times of individual QP iterations. Interestingly, higher QP iterations result in a narrower spread of computation times.

5 In-silico performance analysis

The proposed and previous MPC strategies were evaluated using the UVA/Padova FDA accepted metabolic simulator (Kovatchev et al. 2009, Dalla Man et al. 2014) with its entire adult cohort, consisting of 10-subject and 100-subject simulator cohorts. The latter contains a 101st subject 'average', resulting in a combined cohort of 111 subjects. The UVA/Padova simulator's cohort of in-silico subjects has a wide spectrum of parameter values, some of which might be considered at the boundary of physiologic plausibility. Thus, use of the simulator provides a means to test robustness with respect to inter-subject variability. The case-study simulations start at 14:00. Closed-loop control commences at 16:00. Simulations finish at 19:00 next day. This closely mimics the trial protocol of (Dassau et al. 2015), which

stipulates a maximum meal-size of 90 gCHO. Three large, 90 gCHO meals are provided: Dinner at 18:30, breakfast at 07:00 (next day), lunch at 13:00. This scenario is unrealistically challenging, but represents a demanding stress-test for an AP controller. Large unannounced meals are one of the toughest challenges; both unannounced and announced (see App. B) meal scenarios were simulated for each MPC strategy, resulting in four simulation sets. The BG sensor employed is the simulator's CGM, which includes additive, stochastic noise that can be generated based on a user-provided seed. The above scenario was simulated 11 times using all 111 virtual subjects, once with additive CGM noise of zero, and further with non-zero additive CGM noise using random seeds 1 through 10, resulting in a total of 1221 simulations. The multiple CGM realizations yield more significant and insightful statistics.

Simulations employ treatment parameters that are optimal for the in-silico subjects. Prior to seeking FDA approval diverse simulations where performed with parameters mal-adjusted to varying extents, e.g., basal rates and carbohydrate-ratios excessive or insufficient. Other meal-sizes were also simulated. The safety and efficacy of the proposed MPC strategy with respect to such parameter variations was evaluated and deemed satisfactory for clinical trials. Not all simulation results are provided here, but the presented results are indicative of the general trend, and demonstrate the proposed MPC strategy's superiority, and novelty, over that of (Gondhalekar et al. 2016).

Cohort-level trajectories are depicted in Fig. 4. For unannounced meals the differences are easily discernible. Proposed MPC results in BG trajectories that benefit from lower peaks during the meal-response, enjoy higher troughs at the end of the meal-response, descend from peak to trough less steeply, are less oscillatory, and generally display tighter standard-deviation and min-max envelopes. Proposed MPC effects significantly elevated insulin delivery towards the start of the meal-response – a result of velocity-penalty MPC. Following BG peaks velocity-weighting MPC quickly returns insulin delivery to approximately the basal rate; in contrast, previous MPC continues to drive significant insulin delivery in excess of basal beyond the peak, when BG levels are falling. After meal-peaks proposed MPC results in more steady insulin delivery, whereas previous MPC requires much pump-suspension (discussed later), causing a dip in mean insulin delivery, which in turn causes BG rebound. Reducing BG *oscillations* is very important, even at safe BG *values* (Ceriello et al. 2008, Ceriello & Ihnat 2010); BG variability has a detrimental effect on peoples' health, and an AP controller should best mitigate, not exacerbate or induce, BG fluctuations. When meals are announced the proposed and previous MPC strategies' responses are very similar. However, BG responses with proposed MPC are slightly less oscillatory, and the standard-deviation and min-max envelopes are a little tighter. Insulin delivery in the hour following BG peaks is lower on average using the proposed MPC strategy – a result of velocity-weighting. This reduction is of critical benefit for a *small* number of subjects (discussed later).

Numerical results are tabulated in Table 4. The first set of rows lists time-in-range percentages for various BG ranges and thresholds. The second set of rows lists counts of the number of simulations that experience one or more episode of BG beyond the stated thresholds. The third set of rows lists counts of the total number of BG events, i.e., individual episodes of BG beyond the stated thresholds. The fourth set of rows concerns the

BG trajectories' path length, computed by setting 1 mg/dL ≡ 1 min; shorter paths indicate less oscillatory responses. The single row contains the mean total insulin delivery per simulation. The sixth set of rows lists the number of pump-suspensions of various lengths. A Health Monitoring System (HMS) (Harvey et al. 2012), which provides predictive hypoglycemia alarms, was executed during the simulations; the seventh set of rows lists the total number of HMS alarms, as well as the number of simulations that experience one or more alarm. The Low Blood Glucose Index (LBGI) and High Blood Glucose Index (HBGI) were computed according to (Magni et al. 2009). The Control-Variability Grid Analysis (CVGA) conforms to (Magni et al. 2008).

Proposed MPC yields *significantly* reduced hypoglycemia risk. Notably, this is achieved while enforcing considerably *fewer* pump-suspensions. This is important, especially when considering long suspensions. The ability to suspend insulin delivery is crucial to circumnavigate emergencies (see (Gondhalekar et al. 2016)). However, pump-suspensions, especially when long, invariably result in BG rebound, which increases BG variability and oscillations, and tends to elevate hyperglycemia risk. Both proposed and previous MPC employ the asymmetric cost function described in (Gondhalekar et al. 2016) that promotes pump-suspension when required; the strength of proposed MPC is to induce fewer events that require suspension. Proposed MPC results in slightly less total insulin delivery. This alone, but particularly in concert with the reduced need for subjects to consume rescue carbohydrates, indicates that proposed MPC may assist subjects with their body-weight management, an issue people with T1DM struggle with even more than healthy people do (The Diabetes Control and Complications Trial Research Group 1993, Purnell et al. 1998). The risk of severe hyperglycemia is reduced using proposed MPC. However, the time above 180 mg/dL is slightly elevated; the reason is the shallower approach into the safe BG zone after meal-peaks. Nevertheless, due to reduced hypoglycemia risk the time in ranges [80, 140] and [70, 180] mg/dL increases, for both unannounced and announced meals. The reduced path lengths induced by proposed MPC indicate a reduction in BG variability and oscillations. The sizable reduction in HMS alarms and the number of simulations that experience them deserves attention: Many users of AP devices (currently limited to experimental units) bemoan the devices' inordinate alarming. Presumably the alarm-strategies are intended by their designers to be useful; unfortunately, excessive alarms result in subjects reverting to open-loop (standard care) during times they need quiet, e.g., to sleep or work. The fewer alarms using proposed MPC, thus, are not only indicative of a safer and more effective control strategy, but may result in a longer time spent in closed-loop, and more enthusiastic adoption of AP technology.

The CVGA results of Table 4 and Fig. 5 further underline the significantly reduced hypoglycemia risk of proposed MPC. For unannounced meals proposed MPC effects a prominent shift to the left of the cluster of points, meaning that the nadirs of the subjects' BG trajectories have been substantially lifted. Furthermore, the cluster has experienced a slight shift downwards, signifying improved hyperglycemia correction. The standard-deviation of the points (circle radii in Fig. 5) is also smaller using proposed MPC. Note that with large unannounced meals neither MPC law yields a score in the most desirable A zone. With announced meals proposed MPC yields no significant *mean* improvement in CVGA results, because the meal-bolus renders the control law mostly ineffectual, and many points

in the CVGA cluster approximately retain their location. However, a small number of subjects that experience hypoglycemia with previous MPC have their BG trajectory's nadirs notably raised, and hypoglycemia mitigated, thanks to velocity-weighting. With meal-announcement not one subject is 'clipped' at the right edge when using proposed MPC.

In addition to the in-silico and clinical testing outlined in Secs. 5 & 6 proposed MPC was verified using data from the authors' previous clinical trials (Dassau et al. 2015). Extensive historical data were 'fed through' both previous and proposed MPC laws, and differences studied. Outcomes are not presented here, but the proposed MPC law was deemed effective, safe, and beneficial over that of (Gondhalekar et al. 2016).

6 Clinical trials

The proposed MPC law was deployed in four US FDA approved clinical trial campaigns: **(1)** In a pilot-study 6 people were in closed-loop at home (ClinicalTrials.gov ID: NCT02463682). **(2)** Ten adolescents each spent 3 days in closed-loop (ClinicalTrials.gov ID: NCT-02506764). This trial was more challenging than many of the authors' previous trials, because adolescents are generally difficult to control[2], the trial included multiple unannounced exercise periods, and meals were large (frequently > 100 gCHO). Proposed MPC demonstrably outperformed standard, open-loop therapy (Huyett et al. 2017). **(3)** Nineteen adults each spent 14 days in closed-loop at home (ClinicalTrials.gov ID: NCT-02773875). **(4)** Twenty-nine adults each spent 3 months in closed-loop at home (ClinicalTrials.gov ID: NCT-02705053). Trial details and results are discussed elsewhere (e.g., Huyett et al. (2017), Forlenza et al. (2017), Dassau et al. (2017)). Two trial 'snapshots' are discussed in Secs. 6.1 & 6.2. The trials were performed using proposed MPC. They were subsequently simulated using previous MPC, to obtain this alternate control law's command trajectory, had it received *exactly the same* input data. Note that both BG responses and insulin infusion would have been different, had the alternate control law *actually* been employed.

6.1 Clinical trial example 1 of Fig. 6

This clinical trial example demonstrates the effort of velocity-weighting MPC to safeguard from controller-induced hypoglycemia due to excessive insulin commands during waning hyperglycemia. During the half hour preceding 08:00 two feed-forward meal-boluses (see App. B) were delivered to counteract breakfast. These boluses induce a so-called Insulin On Board (IOB) constraint (see App. A.4) to become active. During the ascent to hyperglycemia due to the meal, both proposed MPC and previous MPC are constrained to deliver no more than basal insulin. CGM peaks at about 09:00 and is re-calibrated shortly after. The re-calibration causes the large, discrete step in CGM output. (However, there is no associated step in the actual BG concentration.) Note that the state-estimator (see App. A.1) has a mechanism (details omitted) to accommodate CGM recalibrations and correctly identifies that across the CGM recalibration the CGM trajectory indicates BG is beginning to subside. The jump in CGM value causes the IOB constraint to abruptly rise, providing the controller

[2]Solely referring, indubitably, to glycemic regulation.

more leeway to command insulin in excess of basal. Note that due to the length of time that has elapsed since the meal-boluses the IOB constraint would have loosened soon regardless of the CGM re-calibration. Thus, the point of note is that the IOB constraint relaxes during hyperglycemia, which the control law now is granted authority to remedy. Until 11:15 proposed MPC commands only basal insulin, thanks to velocity-weighting MPC. The subject experiences a hypoglycemic event at noon, *presumed* to be caused by excessive meal-boluses (it is generally difficult for subjects to estimate meal-size). Importantly, previous MPC would have commanded significantly more insulin, and significantly in excess of basal insulin, during most of the CGM descent. It is justifiable to hypothesize that previous MPC could have contributed to calamitous over-insulinization.

6.2 Clinical trial example 2 of Fig. 7

This trial example demonstrates velocity-penalty MPC's effort to assertively correct developing hyperglycemia. In the period 02:00–03:00 the CGM trajectory ascends through the BG interval $\mathbb{D}$ = [140, 180] mg/dL, inducing an increase in insulin delivery. In this case the night-time constraint (see App. A.3) disallows significant insulin delivery. At 12:30 the CGM trajectory very steeply enters the BG interval $\mathbb{D}$ = [140, 180] mg/dL from below, causing a short-lived surge in insulin to be commanded. Note that the insulin surge lasts slightly beyond the CGM peak due to the dynamics of the recursive state-estimator. Around 18:00 the CGM trajectory very slowly roams the BG interval $\mathbb{D}$ = [140, 180] mg/dL, inducing only very slight increase in insulin delivery. The effects of velocity-weighting are evident around 04:00.

7 Conclusion

A novel MPC law was proposed to simultaneously improve both the safety and performance of an AP to treat T1DM. The proposed MPC law allows the control designer to circumvent, to some extent, the well-known waterbed-effect between the abilities to correct hyperglycemia and to minimize hypoglycemia. An extensive in-silico study demonstrated the proposed control law has benefits but suffers no tradeoff. The novel features of the proposed control law are based only on feedback and require no user announcement. The proposed MPC strategy exploits the flexibility of real-time optimization to employ *velocity-weighting* and *velocity-penalties*, two features that appear novel in the context of AP control. Velocity-weighting appears to be an innovative manner of formulating an MPC cost function. The proposed MPC strategy was implemented on smartphones and deployed in US FDA approved clinical trials that demonstrated its safety and efficacy. Spurred by these successes more extensive, long-term trials are planned. In work parallel to that presented in this paper the authors are investigating more extensive personalization (Pinsker et al. 2016, Lee, Dassau, Gondhalekar, Seborg, Pinsker & Doyle III 2016), control-oriented models more elaborate than LTI (Colmegna, Sánchez-Peña, Gondhalekar, Dassau & Doyle III 2016*a*,*b*, Colmegna, Sánchez-Peña & Gondhalekar 2016), enhanced cost functions (Lee, Gondhalekar, Dassau & Doyle III 2016, Rebello et al. 2017), and real-time adaptation of the control strategy (Laguna Sanz et al. 2016, Cao et al. 2017). Planned future work focuses on the design of control laws for people that exercise, for a wider range of users, e.g., young

children and people with disorders in addition to T1DM, and exploiting non-BG feedback by, e.g., activity trackers.

Acknowledgments

The authors thank the following: (a) The trial subjects for their tenacity; (b) Each author of, and person acknowledged in, (Dassau et al. 2015, Huyett et al. 2017, Dassau et al. 2017), for their effort to complete clinical testing of, and trials using, the AP employing the proposed control law; (c) S. A. Deshpande (Harvard University) for making suggestions to improve the manuscript and for making the authors aware of the path-length approach of quantifying BG variability; (d) B. P. Kovatchev (UVA), J. Wiley (William Sansum Diabetes Center), M. Cescon (University of Melbourne), for suggesting references; (e) Dexcom Inc., LifeScan Inc., and Roche AG, for product support; (f) The National Institutes of Health (NIH) for funding: DP3-DK094331, DP3DK104057, R01DK085628. The authors acknowledge that: Access to the complete version of the UVA/Padova metabolic simulator was provided by an agreement with Prof. C. Cobelli (University of Padova) and Prof. B. P. Kovatchev (UVA) for research purposes.

Appendix – AP specific controller details

The proposed control strategy is a modification of that in (Gondhalekar et al. 2016). Details of the feedback MPC law and the feed-forward meal-bolusing strategy are described in Appendices A & B. These details are either similar to, or trivially modified from, those of (Gondhalekar et al. 2016), which the interested reader is encouraged to consult. Presented details focus on *differences* to (Gondhalekar et al. 2016) and are stated to permit discussion, for completeness, and to facilitate algorithm reproducibility when parameter values were revised.

A Feedback MPC law details & settings

A.1 Insulin-BG dynamics: Control-relevant model

LTI system (1) is employed with the following:

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

$$B := \frac{1800K}{u_{\mathrm{TDI}}}[100]^{\top} \in \mathbb{R}^3$$

$$C_{\mathrm{y}} := [001] \in \mathbb{R}^{1\times 3}$$

$$C_{\mathrm{v}} := [0.10 - 0.1] \in \mathbb{R}^{1\times 3}$$

K: = 90 $(p_1 - 1)$ $(p_2 - 1)^2$; p_1: = 0:98; p_2: = 0.965 u_{TDI}> 0: Total daily insulin, subject-specific [U].

The model of insulin-BG dynamics described by A, B, C_y was identified in (van Heusden et al. 2012) and successfully used in many AP trials (Gondhalekar et al. 2013, Dassau et al. 2015, Gondhalekar et al. 2016, Huyett et al. 2017). Model (1) is personalized by the subject-specific u_{TDI}, which one can acquire accurately from patients. The poles p_1 and p_2 are the same for each subject, and were derived to yield a *control-relevant model* (i.e., one that facilitates controller design but not high-fidelity simulations) that leads to a controller that works well for people *on average* but is safe for outlier subjects (see (van Heusden et al. 2012) and Sec. 2.8 of (Gondhalekar et al. 2016)). The poles p_1 and p_2 indicate system (1) is open-loop stable. System (1) is linearized around the subject-specific, time-dependent basal rate input $u_{\mathrm{BASAL},i}$ [U/h], that results in steady-state BG output $y_{\mathrm{s}} := 110$ [mg/dL]. The absolute insulin input and BG output are given by $u_{\mathrm{IN},i} := u_i + u_{\mathrm{BASAL},i} T$ [U] and $y_{\mathrm{BG},i} := y_i + y_{\mathrm{s}}$ [mg/dL].

The structures of A, B, and C_y indicate that, in the absence of plant-model mismatch and noise, at step i the three state elements $x_{[3]}$, $x_{[2]}$, and $x_{[1]}$ correspond to y_i, y_{i+1}, and y_{i+2}, respectively. The output matrix C_v was chosen such that the velocity output v_i provides an estimate of the average, over the impending $2T = 10$ min, rate of change of BG output y_i, in units mg/dL/min. The structures of A, B, and C_y further indicate that model (1) has a $3T = 15$ min delay from u to BG output y.

An estimate of model state x_i is determined at each step i, based on feedback from a CGM, by the linear recursive estimator (Luenberger observer) described in Sec. 2.2 of (Gondhalekar et al. 2016). This provides a simple tuning handle to reject CGM disturbances that, in practice, are difficult to model or derive realistic stochastic properties of. The authors have experimented in-silico with alternative state-estimators, and certain improvements over the linear recursive estimator were presented in (Gondhalekar et al. 2014*c*,*a*, Lee et al. 2014). However, for clinical trials the current estimator is retained because it is dependable and extensively tested.

A.2 Diurnal BG target zone

The proposed MPC law is a periodic zone-MPC strategy. The BG target zone is the interval [80, 140] mg/dL at times 06:00–22:00, [90, 140] mg/dL at times 24:00–04:00, with two-hour transitions between; see Fig. A.1. Such *diurnal* zones help enforce increased safety from nocturnal hypoglycemia (Gondhalekar et al. 2013, Dassau et al. 2015, Gondhalekar et al. 2016). The current boundary values were modified from previous ones after successful and safe testing. The zone-transition function described in Sec. 2.3 of (Gondhalekar et al. 2016) is used with the updated values.

A.3 Diurnal insulin delivery constraints

At each step i the controller must enforce constraint (A.1) with t_i the time of day at step i. In the period 04:00–22:00, $\bar{u}(t_i) := 1$ [U], i.e., during the day the controller has broad control authority but there exists a safeguard to protect from large, undesirable insulin input commands due to, e.g., data anomalies. At other times $\bar{u}(t_i) := 1.8\, u_{\mathrm{BASAL},i} T$, i.e., at night insulin infusion is constrained to 80% in excess of the subject's basal rate. Note that the basal rate is tailored to the subject and, typically, is time-dependent. Thus, the nighttime

constraint is both personalized and depends on the time of day. This diurnal input constraint is depicted in Fig. A.1 and provides a further safeguard against nocturnal hypoglycemia (Gondhalekar et al. 2013, Dassau et al. 2015, Gondhalekar et al. 2016). Only 14 of 1, 579, 974 MPC commands of the in-silico study of Sec. 5, and only 4 of 21, 821 MPC commands evaluated during the first round of clinical trials of Sec. 6, equal the daytime constraint $\bar{u}(t_i) = 1$. This observation is important: The aggressive hyperglycemia correction effected by velocity-penalty D in (13) is not 'riding the constraint', i.e., tuned to be excessive and reined-in by constraints. The daytime constraint enforces safety from anomalously large control commands, but does not affect the controller during typical operation. It is designed to protect, but permit broad control leeway, based as much as possible on the natural minimization of objective function (4).

$$0 \leq u_i + u_{\mathrm{BASAL},i} T \leq \bar{u}(t_i) \quad \text{(A.1)}$$

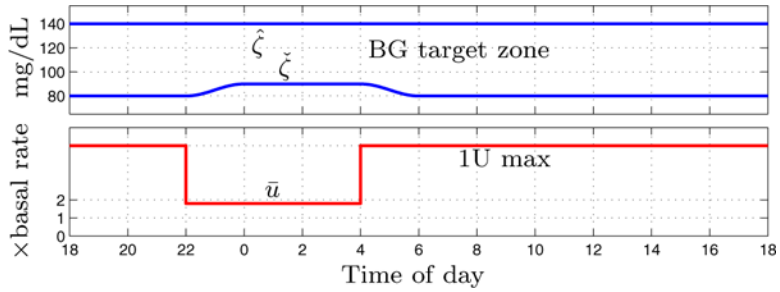

Fig. A.1.

Top: Upper and lower bounds $\hat{\zeta}$ and $\check{\zeta}$ of diurnal BG target zone. **Bottom**: Diurnal, individualized upper-bound $\bar{u}$ on insulin infusion input. See Appendices A.2 & A.3.

A.4 Insulin on board (IOB) constraints

Insulin delivery is subject to an IOB constraint that, based on delivery history, prevents over-delivery when much insulin was recently delivered, e.g., after a meal-bolus (see App. B). The IOB constraint presented here is an important, albeit simple, modification of that in (Gondhalekar et al. 2016). We denote by Θ: $\mathbb{R}_{\geq 0} \times \mathbb{R}_{>0} \rightarrow [0, 1]$ the newly derived IOB decay curve function (A.2), by $\mathcal{T} \in \mathbb{R}_{\geq 0}$ [h] the time since insulin infusion, by $\mathcal{T} \in \mathbb{R}_{\geq 0}$ [h] the curve length; see Fig. A.2. Eq. (A.2) was designed (details omitted) to yield curves closely fitting those of (Gondhalekar et al. 2016), obtained from (Walsh & Roberts 2006, Ellingsen et al. 2009). Introducing (A.2) permits curves of *arbitrary* length $\mathcal{T}$, not only discrete lengths $\mathcal{T} \in \{2, 3, \ldots, 8\}$ depicted in Fig. A.2 as used in (Gondhalekar et al. 2016, Walsh & Roberts 2006, Ellingsen et al. 2009). The length $\mathcal{T}(G) \in [2, 8]$ of the IOB curve employed at current CGM level G [mg/dL] is given by (A.3), a continuous function that replaces the discontinuous switching of (Gondhalekar et al. 2016). At each step i we denote by G_i the most recent CGM value, and by $\theta_i \in \mathbb{R}^{96}$ of (A.4) a vector with values of the IOB decay of length $\mathcal{T}(G_i)$, sampled at $T = 5$ [min] intervals. Analogously let $\theta_{\mathrm{MEAL}} \in \mathbb{R}^{96}$ of (A.5) contain the sampled 4h IOB curve, to be used for meal-boluses.

$$\Theta(\tau,\mathcal{T}) := \begin{cases} \frac{1-b}{1-a}\cos\left(\left(\frac{\tau}{\mathcal{T}}\hat{\tau}\right)^{c}\pi\right)+\frac{b-a}{1-a} & \text{if } \tau\in[0,\mathcal{T}] \\ 0 & \text{otherwise} \end{cases} \quad \text{(A.2)}$$

$$a := -0.995 \qquad c := 0.84$$

$$b := 0.04 \qquad \hat{\tau} := (\operatorname{acos}(a)/\pi)^{1/c}$$

$$\mathcal{T}(G) := \max\{\min\{-G/30+12,8\},2\} \quad \text{(A.3)}$$

$$\theta_{i[j]} := \Theta\left(\frac{5j}{60},\mathcal{T}(G_i)\right) \qquad \forall j\in\{1,2,\ldots,96\} \quad \text{(A.4)}$$

$$\theta_{\mathrm{MEAL}[j]} := \Theta\left(\frac{5j}{60},4\right) \qquad \forall j\in\{1,2,\ldots,96\} \quad \text{(A.5)}$$

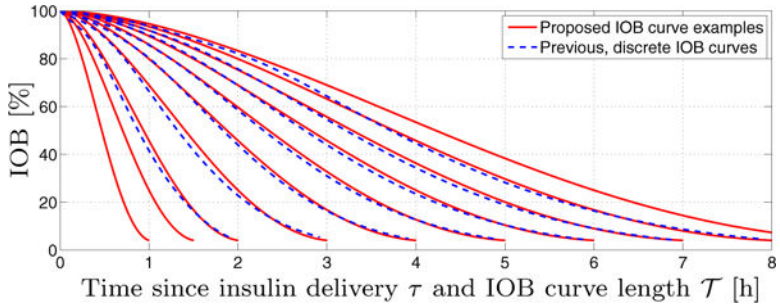

Fig. A.2.
IOB decay curves. Proposed curve (A.2) defined for arbitrary positive length, e.g., $\mathcal{T}\in\{1,1.5,9\}$ as depicted. Previous curves limited to $\mathcal{T}\in\{2,3,\ldots,8\}$ (Gondhalekar et al. 2016, Walsh & Roberts 2006, Ellingsen et al. 2009).

Denote by $\Lambda_{\mathrm{BASAL}}\in\mathbb{R}^{96}$ the 8h history of u_i, the insulin deviation from the basal rate, but setting to zero values at steps a feed-forward meal-bolus was delivered (see App. B). Denote by $\Lambda_{\mathrm{MEAL}}\in\mathbb{R}^{96}$ the 8h history of feed-forward meal-boluses, by $\Phi\in\mathbb{R}$ the estimated IOB, by $\Gamma\in\mathbb{R}$ the required IOB, that $\in$ depends on the current CGM level G_i, and by $C_{\mathrm{F},i}\in\mathbb{R}_{>0}$ [mg/dL/U] the patient's correction-factor at step i. At each step i the IOB upper bound $\bar{u}_{\mathrm{IOB},i}$ is given by (A.6).

$$\bar{u}_{\mathrm{IOB},i} := \max\left\{\Gamma_i - \Phi_i, 0\right\} \quad \text{(A.6)}$$

$$\Phi_i := \theta_i^\top \Lambda_{\mathrm{BASAL},i} + \theta_{\mathrm{MEAL}}^\top \Lambda_{\mathrm{MEAL},i}$$

$$\Gamma_i := \left(G_i - y_S\right)/C_{\mathrm{F},i}$$

The proposed MPC law must enforce IOB constraint (A.7). It holds that $\bar{u}_{\mathrm{IOB},i} \geq 0$, and $\bar{u}_{\mathrm{IOB},i} = 0$ implies the controller delivers no more than $u_{\mathrm{BASAL}} iT$, the basal rate. Thus, after a large bolus insulin delivery is temporarily constrained to the basal rate, but (A.7) cannot constrain insulin delivery to below the basal rate.

$$u_i \leq \bar{u}_{\mathrm{IOB},i} \quad \text{(A.7)}$$

This IOB constraint's improvement over that of (Gondhalekar et al. 2016) is the ability to employ IOB decay curves of arbitrary length, allowing the length function $\mathcal{T}(G)$ of (A.3) to be *continuous*. This has two advantages. First, the IOB constraint no longer displays discrete-steps, which can result in jerky and surprising (potentially alarming) insulin delivery trajectories. Second, it permits more subtle tuning of the IOB constraint. Because velocity-weighting protects from hypoglycemia one can employ slightly shorter IOB decay curves, i.e., be tuned to loosen a little more rapidly, than in (Gondhalekar et al. 2016). The IOB constraint is then not excessively 'hijacked' to protect from over-delivery during post-prandial hyperglycemic excursions. By loosening the IOB constraint earlier following feed-forward meal-boluses the MPC law is sooner provided leeway to exploit feedback provided by the CGM and correct persistent hyperglycemia. However, the velocity-weighting protects from over-delivery in case, at the time the IOB constraint loosens, the BG trajectory is already descending. An example of this is depicted in Fig. 6.

A.5 Pump-discretization

The proposed control strategy was destined for deployment with a CSII pump that has a delivery resolution of 0.1 [U]. After the controller computes the optimal input the final, absolute control input commanded to the pump is characterized according to a so-called *carry-over* scheme; see Sec. 2.7 of (Gondhalekar et al. 2016) for details and discussion. Note that the pump-discretization could not be chosen by the authors, but is dictated to them by the combination of pump and AP platform software. Thus there is no design-related significance to the fact that in (Gondhalekar et al. 2016) a finer pump-discretization of 0.05 [U] was available.

B Feed-forward meal-bolusing strategy

Subjects may opt to announce meals to the controller, resulting in feed-forward meal-boluses. The bolus size computation is a variation of that in Sec. 3 of (Gondhalekar et al. 2016). Let $M \in \mathbb{R}_{>0}$ [gCHO] denote the meal-size, and $C_R \in \mathbb{R}_{>0}$ [gCHO/U], $\in C_F \in \mathbb{R}_{>0}$ [mg/dL/U], and $S \in \mathbb{R}$ the subject's carbohydrate-ratio, correction-factor, and self-monitoring BG measurement, respectively, specific to the time of day of meal-announcement. Following a meal-announcement the MPC law of Sec. 2 is by-passed and the insulin input $u_{IN,i}$ is characterized according to (B.1). As a precaution the basic bolus M/C_R is reduced by 20% if $S < 120$ mg/dL. If insufficient the controller can increase insulin delivery later, but removing excessive insulin is not possible. A correction is added to the basic bolus when $S > 150$ mg/dL, but only if no such correction was administered within the preceding two hours; this lockout period prevents a sequence of announced meals causing a sequence of corrections to address the *same* BG excursion. The correction is limited to 2 U. The switching thresholds, bound, lockout period, and reduction ratio were selected after discussions with endocrinologists. Without meal-announcement the controller counters meal-responses based solely on CGM feedback and the MPC law of Sec. 2.

$$u_{\mathrm{IN},i} := u_{\mathrm{MEAL}} + \min\left\{u_{\mathrm{COR}}, 2\right\} \quad \text{(B.1)}$$

$$u_{\mathrm{MEAL}} := \begin{cases} M/C_R & \text{if } S \geq 120 \\ 0.80\,M/C_R & \text{otherwise} \end{cases}$$

$$u_{\mathrm{COR}} := \begin{cases} (S-150)/C_F & \text{if } (S \geq 150) \wedge (\Delta > 120) \\ 0 & \text{otherwise} \end{cases}$$

$$\Delta : \text{Time}\left[\min\right] \text{since preceding } u_{\mathrm{COR}} > 0$$

Biographies

Ravi Gondhalekar is a research engineer at the Charles Stark Draper Laboratory, USA, where he investigates and develops guidance and control algorithms for spacecraft and

atmospheric flight vehicles. Ravi received BA and MEng degrees in engineering in 2002 from the University of Cambridge, UK, and a PhD degree in informatics in 2008 from the Tokyo Institute of Technology, Japan. From 2008 to 2012 Ravi was an assistant professor at Osaka University, Japan, where he performed research into constrained control theory and model predictive control methods. From 2012 to 2017 Ravi was at the University of California Santa Barbara, USA, and Harvard University, USA, where he investigated applying constrained model predictive control techniques to an artificial pancreas that performs automated delivery of insulin to people with type 1 diabetes. Ravi held shortterm positions at the Massachusetts Institute of Technology, USA, the University of Cambridge, UK, Princeton University, USA, Pi Technology, UK, the Rutherford Appleton Laboratory, UK, and the United Kingdom Atomic Energy Authority, UK.

Eyal Dassau is a Senior Research Fellow in BiomedicalÊEngineering in the Harvard John A.ÊPaulson School of Engineering andÊApplied Sciences. He is also an Adjunct Senior Investigator with the Sansum Diabetes Research Institute, Santa Barbara, CA and an Adjunct Faculty at the Joslin Diabetes Center, Boston, MA. Prior to that he was an Associate Research Engineer, senior investigator and diabetes research manager in the Chemical Engineering Department and the Institute for Collaborative Biotechnologies at the University of California, Santa Barbara (UCSB). He received his B.Sc., M.Sc. and Ph.D. degrees in Chemical Engineering from the Technion Israel Institute of Technology, Haifa, Israel in 1999, 2002 and 2006, respectively. Dr. DassauÕs research is focused on the development of an artificial pancreas for people with type 1 diabetes mellitus. Dr. Dassau is a senior member of the American Institute of Chemical Engineers (AIChE) and the Institute of Electrical and Electronics Engineers (IEEE) Engineering in Medicine and Biology Society, a member of the American Diabetes Association (ADA), European Association for the study of Diabetes (EASD) and the American Association for the Advancement of Science (AAAS). Dr. Dassau published more than 120 per-review papers on the artificial pancreas and has submitted more than 20 IDEs to the FDA to support artificial pancreas studies. He is a member of the editorial board of the journal of diabetes science and technology and diabetes yearbook published by Mary Ann Liebert, Inc.

Frank Doyle is the John A. Paulson Dean of the Harvard John A. Paulson School of Engineering and Applied Sciences at Harvard University, where he also is the John A. & Elizabeth S. Armstrong Professor. He received a B.S.E. degree from Princeton, C.P.G.S.

from Cambridge, and Ph.D. from Caltech, all in Chemical Engineering. He has held faculty appointments at Purdue University, the University of Delaware, and UCSB. In addition he has held visiting positions at DuPont, Weyerhaeuser, and Stuttgart University. He has been recognized as a Fellow of multiple professional organizations including: IEEE, IFAC, AIMBE, and the AAAS. He was the President for the IEEE Control Systems Society in 2015, and was the Vice President of the International Federation of Automatic Control from 2014 to 2017. In 2005, he was awarded the Computing in Chemical Engineering Award from the AIChE for his innovative work in systems biology, and in 2015 received the Control Engineering Practice Award from the American Automatic Control Council for his development of the artificial pancreas. In 2016, he was inducted as a Fellow into the National Academy of Medicine for his work on biomedical control. His research interests are in systems biology, network science, modeling and analysis of circadian rhythms, and drug delivery for diabetes.

References

Bequette BW. Challenges and recent progress in the development of a closed-loop artificial pancreas. Annu Rev Control. 2012; 36(2):255–26. [PubMed: 23175620]

Bequette BW. Algorithms for a Closed-Loop Artificial Pancreas: The Case for Model Predictive Control. J Diabetes Sci Technol. 2013; 7(6):1632–1643. [PubMed: 24351190]

Boiroux D, Duun-Henriksen AK, Schmidt S, Nørgaard K, Madsbad S, Skyggebjerg O, Jensen PR, Poulsen NK, Madsen H, Jørgensen JB. Overnight Control of Blood Glucose in People with Type 1 Diabetes. IFAC Symp Biol & Med Systems. 2012:73–78.

Boiroux D, Finan DA, Jørgensen JB, Poulsen NK, Madsen H. Strategies for Glucose Control in People with Type 1 Diabetes. IFAC World Congress. 2011:3765–3770.

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; 61(9):2230–2237. [PubMed: 22688340]

Cao Z, Dassau E, Gondhalekar R, Doyle FJ III. Extremum Seeking Control Based Zone Adaptation for Zone Model Predictive Control in Type 1 Diabetes. IFAC World Congress. 2017:15639–15644.

Ceriello A, Esposito K, Piconi L, Ihnat MA, Thorpe JE, Testa R, Boemi M, Giugliano D. Oscillating Glucose Is More Deleterious to Endothelial Function and Oxidative Stress Than Mean Glucose in Normal and Type 2 Diabetic Patients. Diabetes. 2008; 57:1349–1354. [PubMed: 18299315]

Ceriello A, Ihnat MA. 'Glycaemic variability': a new therapeutic challenge in diabetes and the critical care setting. Diabetic Medicine. 2010; 27:862–867. [PubMed: 20653741]

Cobelli C, Renard E, Kovatchev B. Artificial Pancreas: Past, Present, Future. Diabetes. 2011; 60(11): 2672–2682. [PubMed: 22025773]

Colmegna P, Sánchez-Peña RS, Gondhalekar R. Control-Oriented Linear Parameter-Varying Model for Glucose Control in Type 1 Diabetes. IEEE Multi-Conf Systems & Control. 2016:410–415.

Colmegna P, Sánchez-Peña RS, Gondhalekar R, Dassau E, Doyle FJ III. Reducing Glucose Variability Due to Meals and Post-prandial Exercise in T1DM Using Switched LPV Control: In Silico Studies. J Diabetes Sci Technol. 2016a; 10(3):744–753. [PubMed: 27022097]

Colmegna P, Sánchez-Peña RS, Gondhalekar R, Dassau E, Doyle FJ III. Switched LPV Glucose Control in Type 1 Diabetes. IEEE Trans Biomed Eng. 2016b; 63(6):1192–1200. [PubMed: 26452196]

Dalla Man C, Micheletto F, Lv D, Breton M, Ko-vatchev B, Cobelli C. The UVA/PADOVA Type 1 Diabetes Simulator: New Features. J Diabetes Sci Technol. 2014; 8(1):26–34. [PubMed: 24876534]

Dassau E, Brown SA, Basu A, Pinsker JE, Kudva YC, Gondhalekar R, Patek S, Lv D, Schiavon M, Lee JB, Dalla Man C, Hinshaw L, Castorino K, Mallad A, Dadlani V, McCrady-Spitzer SK, McElwee-Malloy M, Wakeman CA, Bevier WC, Bradley PK, Kovatchev B, Cobelli C, Zisser HC, Doyle FJ III. Adjustment of Open-Loop Settings to Improve Closed-Loop Results in Type 1 Diabetes: A Multicenter Randomized Trial. J Clin Endocr Metab. 2015; 100(10):3878–3886. [PubMed: 26204135]

Dassau E, Pinsker JE, Kudva YC, Brown SA, Gondhalekar R, Dalla Man C, Patek S, Schiavon M, Dadlani V, Dasanayake I, Church MM, Carter RE, Bevier WC, Huyett LM, Hughes J, Anderson S, Lv D, Schertz E, Emory E, McCrady- Spitzer SK, Jean T, Bradley PK, Hinshaw L, Laguna Sanz AJ, Basu A, Kovatchev B, Cobelli C, Doyle FJ III. 12 Week 24/7 Ambulatory Artificial Pancreas with Weekly Adaptation of Insulin Delivery Settings: Effect on Hemoglobin A1c and Hypoglycemia. Diabetes Care. 2017; 40(12):1719–1726. [PubMed: 29030383]

Doyle FJ III, Huyett LM, Lee JB, Zisser HC, Dassau E. Closed Loop Artificial Pancreas Systems: Engineering the Algorithms. Diabetes Care. 2014; 37(5):1191–1197. [PubMed: 24757226]

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; 3(3):536–544. [PubMed: 20144293]

Eren-Oruklu M, Cinar A, Quinn L, Smith D. Adaptive control strategy for regulation of blood glucose levels in patients with type 1 diabetes. J Process Control. 2009; 19(8):1333–1346.

Forlenza GP, Deshpande S, Ly TT, Howsmon DP, Cameron F, Baysal N, Mauritzen E, Marcal T, Towers L, Bequette BW, Huyett LM, Pinsker JE, Gondhalekar R, Doyle FJ III, Maahs DM, Buckingham BA, Dassau E. Application of Zone Model Predictive Control Artificial Pancreas during Extended Use of Infusion Set and Sensor: A Randomized Crossover-Controlled Home-Use Trial. Diabetes Care. 2017; 40(8):1096–1102. [PubMed: 28584075]

Gondhalekar R, Dassau E, Doyle FJ III. Moving-horizon-like state estimation via continuous glucose monitor feedback in MPC of an artificial pancreas for type 1 diabetes. IEEE Conf Decision & Control. 2014a:310–315.

Gondhalekar R, Dassau E, Doyle FJ III. MPC Design for Rapid Pump-Attenuation and Expedited Hyperglycemia Response to Treat T1DM with an Artificial Pancreas. AACC American Control Conf. 2014b:4224–4230.

Gondhalekar R, Dassau E, Doyle FJ III. State Estimation with Sensor Recalibrations and Asynchronous Measurements for MPC of an Artificial Pancreas to Treat T1DM. IFAC World Congress. 2014c:224–230.

Gondhalekar R, Dassau E, Doyle FJ III. Tackling problem nonlinearities & delays via asymmetric, state-dependent objective costs in MPC of an artificial pancreas. IFAC Conf Nonlinear MPC. 2015a:154–159.

Gondhalekar R, Dassau E, Doyle FJ III. Velocity-weighting to prevent controller-induced hypoglycemia in MPC of an artificial pancreas to treat T1DM. AACC American Control Conf. 2015b:1635–1640.

Gondhalekar R, Dassau E, Doyle FJ III. Periodic zone-MPC with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes. Automatica. 2016; 71:237–246. [PubMed: 27695131]

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; 7(6):1446–1460. [PubMed: 24351171]

Goodwin GC, Medioli AM, Carrasco DS, King BR, Fu Y. A fundamental control limitation for linear positive systems with application to type 1 diabetes treatment. Automatica. 2015; 55:73–77.

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; 4(4):961–975. [PubMed: 20663463]

Haidar A. The Artificial Pancreas: How closed-loop control is revolutionizing diabetes. IEEE Control Systems. 2016; 36(5):28–47.

Harvey RA, Dassau E, Zisser H, Seborg DE, Jo-vanovič L, Doyle FJ III. Design of the Health Monitoring System for the Artificial Pancreas: Low Glucose Prediction Module. J Diabetes Sci Technol. 2012; 6(2):1345–1354. [PubMed: 23294779]

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; 25:905–920. [PubMed: 15382830]

Huyett LM, Ly TT, Forlenza GP, Reuschel-DiVirgilio S, Messer LH, Wadwa RP, Gondhalekar R, Doyle FJ III, Pinsker JE, Maahs DM, Buckingham BA, Dassau E. Outpatient Closed-Loop Control with Unannounced Moderate Exercise in Adolescents using Zone Model Predictive Control. Diabetes Technol Ther. 2017; 19(6):331–339. [PubMed: 28459617]

Keith-Hynes P, Guerlain S, Mize B, Hughes-Karvetski C, Khan M, McElwee-Malloy M, Kovatchev BP. DiAs User Interface: A Patient-Centric Interface for Mobile Artificial Pancreas Systems. J Diabetes Sci Technol. 2013; 7(6):1416–1426. [PubMed: 24351168]

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; 3(1):44–55. [PubMed: 19444330]

Kovatchev BP, Keith-Hynes PT, Breton MD, Patek SD. The Diabetes Assistant (DiAs) – Unified Platform for Monitoring and Control of Blood Glucose Levels in Diabetic Patients. 2012 PCT/US12/43910.

Laguna Sanz AJ, Doyle FJ III, Dassau E. An Enhanced Model Predictive Control for the Artificial Pancreas Using a Confidence Index Based on Residual Analysis of Past Predictions. J Diabetes Sci Technol. 2016; 3(11):537–544.

Lee JB, Dassau E, Gondhalekar R, Seborg DE, Pinsker J, Doyle FJ III. An Enhanced MPC (eMPC) Strategy for Automated Glucose Control. Ind Eng Chem Res. 2016; 55(46):11857–11868. [PubMed: 27942106]

Lee JB, Gondhalekar R, Dassau E, Doyle FJ III. Shaping the MPC Cost Function for Superior Automated Glucose Control. IFAC Symp Dynamics & Control of Process Systems. 2016:779–784.

Lee JJ, Gondhalekar R, Doyle FJ III. Design of an Artificial Pancreas using Zone Model Predictive Control with a Moving Horizon State Estimator. IEEE Conf Decision & Control. 2014:6975–6980.

Maciejowski JM. Predictive Control with Constraints Pearson/Prentice Hall; Harlow, England: 2002

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; 2(4):630–635. [PubMed: 19885239]

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.

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; 18(2):149–162. [PubMed: 19190726]

Messori M, Ellis M, Cobelli C, Christofides PD, Magni L. Improved Postprandial Glucose Control with a Customized Model Predictive Controller. AACC American Control Conf. 2015:5108–5115.

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; 46(2):148–157. [PubMed: 9932336]

Pinsker JE, Lee JB, Dassau E, Seborg DE, Castorino K, Gondhalekar R, Bevier WC, Bradley PK, Zisser HC, Doyle FJ III. A Randomized Crossover Clinical Comparison of Personalized MPC and PID Artificial Pancreas Control Algorithms. Diabetes Care. 2016; 39(7):1135–1142. [PubMed: 27289127]

Purnell JQ, Hokanson JE, Marcovina SM, Steffes MW, Cleary PA, Brunzell JD. Effect of Excessive Weight Gain With Intensive Therapy of Type 1 Diabetes on Lipid Levels and Blood Pressure: Results from the DCCT. J Am Med Assoc. 1998; 280(2):140–146.

Rawlings JB, , Mayne DQ. Model Predictive Control: Theory and Design Nob Hill Publishing; Madison, WI, USA: 2009

Rebello G, Gondhalekar R, Dassau E, Doyle FJ III. Mixed Linear-Quadratic Cost Function Design for MPC of an Artificial Pancreas — Improved Treatment & Safety for a Broad Range of Meal Sizes. IFAC World Congress. 2017:7990–7996.

The Diabetes Control and Complications Trial Research Group. The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus. N Engl J Med. 1993; 329(14):977–86. [PubMed: 8366922]

Trivellato A, , Moriconi C, , Squillantini E. (n.d.), 'www.joptimizer.com'. [Accessed Dec 10, 2017]

Turksoy K, Quinn L, Littlejohn E, Cinar A. Multivariable Adaptive Identification and Control for Artificial Pancreas Systems. IEEE Trans Biomed Eng. 2014; 61(3):883–891. [PubMed: 24557689]

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; 59(7):1839–1849. [PubMed: 22127988]

Walsh J, , Roberts R. Pumping Insulin 4. Torrey Pines Press; San Diego, CA, USA: 2006

Zavitsanou S, Mantalaris A, Georgiadis MC, Pistikopoulos EN. In Silico Closed-Loop Control Validation Studies for Optimal Insulin Delivery in Type 1 Diabetes. IEEE Trans Biomed Eng. 2015; 6(10):2369–2378.

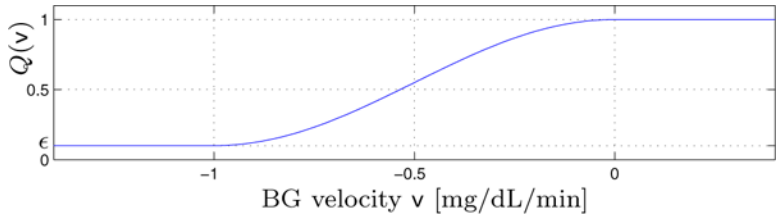

Fig. 1.
Schematic of velocity-weighting function $Q(\mathrm{v})$ of (14).

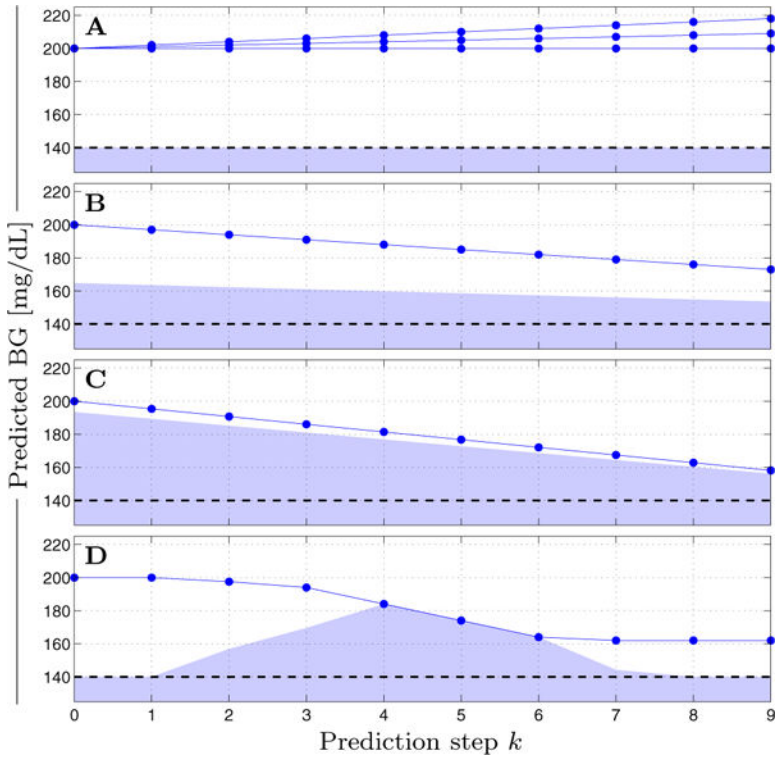

Fig. 2.

Schematic demonstration of velocity-weighting MPC. Actual BG target zone upper bound $\hat{\zeta}$; black dashed line. 'Effective' zone; blue fill. Predicted BG values; blue dots. Predicted BG velocities: Non-negative (**A**); negative (**B**, $\mathbf{v}_k = -0.6$ mg/dL/min); more negative (**C**, $\mathbf{v}_k = -0.93$ mg/dL/min); variable negative (**D**, $\mathbf{v}_k \in [-2, 0]$ mg/dL/min).

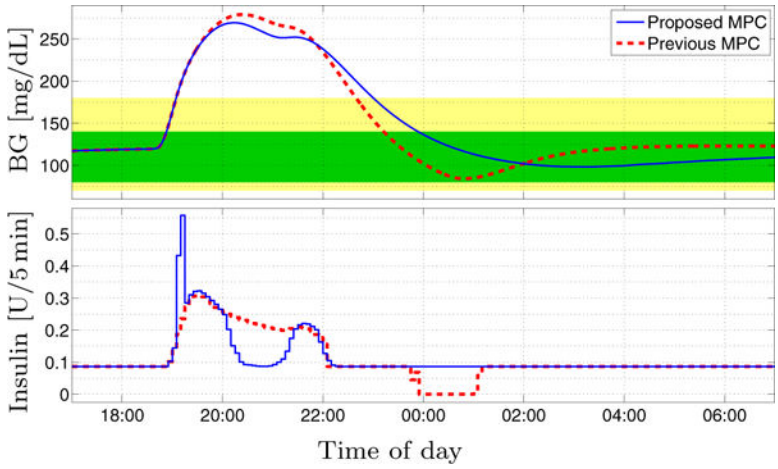

Fig. 3.

Proposed MPC vs. previous MPC demonstrated using UVA/Padova simulator subject 59. Green area: [80, 140] mg/dL. Yellow area: [70, 180] mg/dL. Unannounced 90 gCHO meal ingested at 18:30; see protocol of Sec. 5. Total insulin (18:00–06:00): 16.44 U (proposed), 16.76 U (previous).

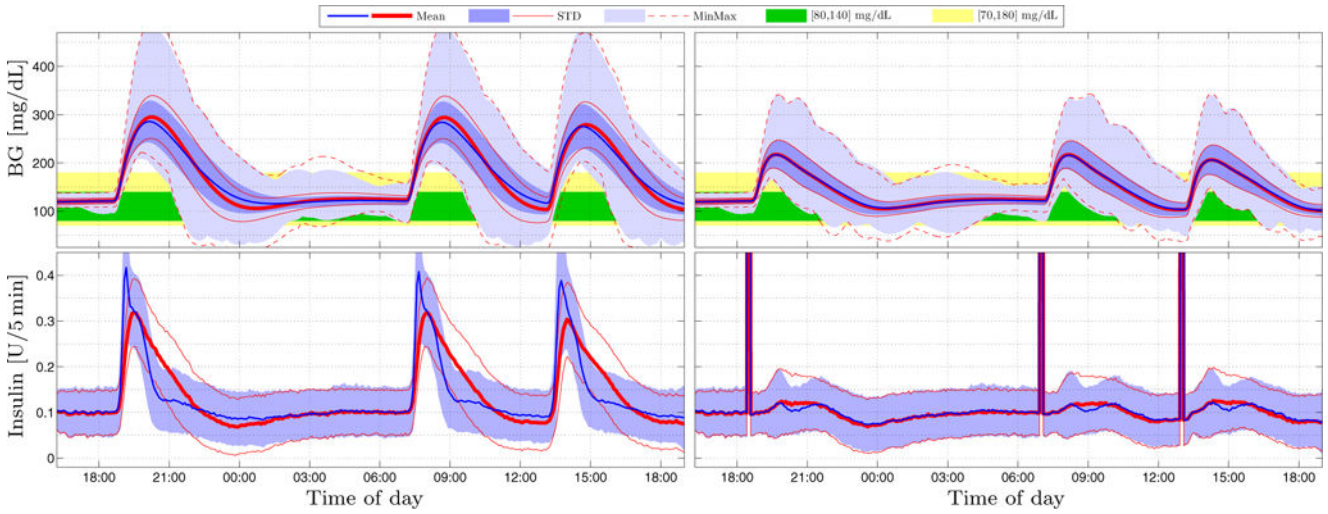

Fig. 4.

In-silico performance analysis: Mean, standard-deviation (STD), and min-max envelopes (only BG). Proposed (blue) vs. previous (red) MPC strategy. Unannounced (**left**) vs. announced (**right**) meals. See Table 4 for statistics.

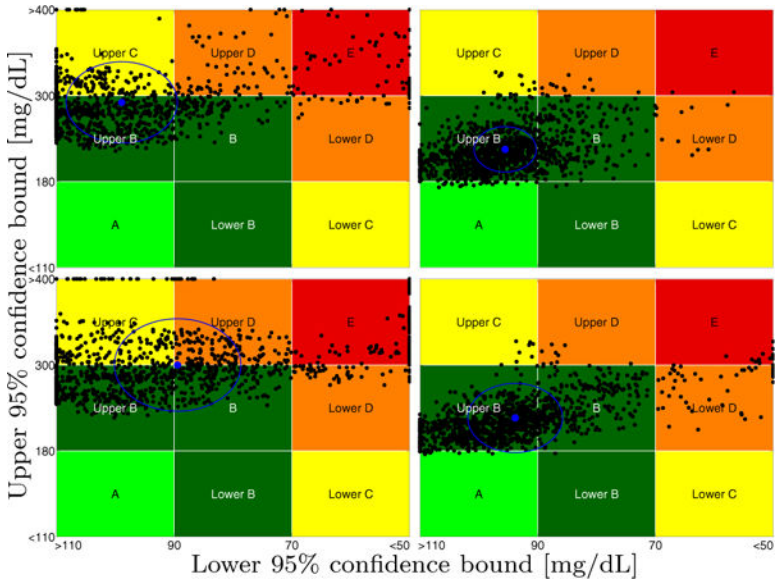

Fig. 5.

In-silico performance analysis: CVGA plots. Proposed (**top**) vs. previous (**bottom**) MPC strategy. Unannounced (**left**) vs. announced (**right**) meals. Blue dot and circle depict mean and standard-deviation. See Table 4 for statistics.

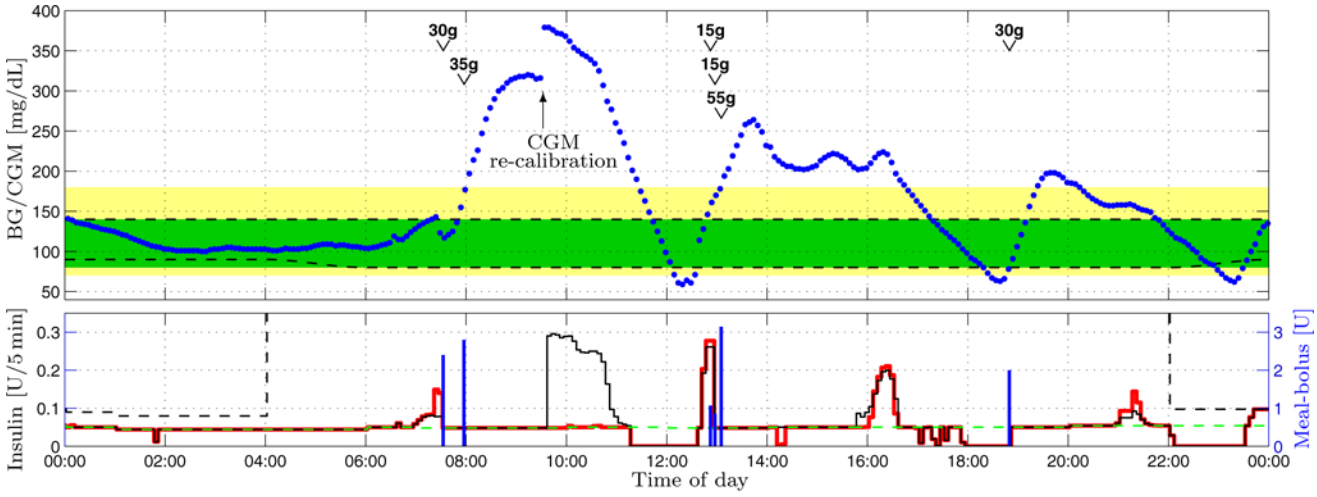

Fig. 6.

Clinical trial example 1; see Sec. 6.1. **Top:** CGM (blue dots), [80, 140] mg/dL (green area), [70, 180] mg/dL (yellow area), diurnal BG target zone (dashed black lines), meal-sizes [gCHO]. **Bottom:** Insulin delivery of trial using proposed MPC (thick red line), insulin delivery as *would have been commanded* by previous MPC (thin black line), feed-forward meal-boluses (blue bars, right scale), basal rate (dashed green line), diurnal safety constraint (dashed black line).

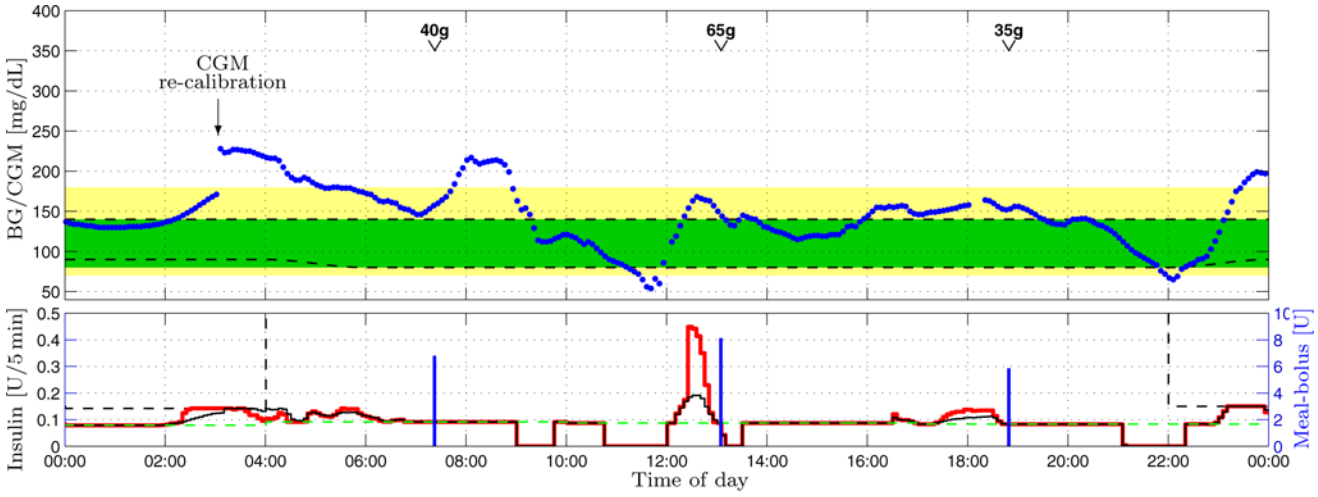

Fig. 7.

Closed-loop clinical trial example 2; see Sec. 6.2. Key same as Fig. 6, but note the different insulin axis scale.

Table 1

The proposed MPC and previous MPC tunings compared.

MPC	$\widehat{R}$	D	$Q(v)$
Proposed	6500	1000	Eq. (14)
Previous	7000	0	$Q(v) = 1$

Table 2

QP iteration numbers required for optimization convergence. Line 1: In-silico study of Sec. 5, total 1, 579, 974 solutions. Line 2: First clinical trials of Sec. 6, total 21, 821 solutions.

Its.	1	2	3	4	5	6	7	8	9
	1, 393, 091	112, 150	36, 263	26, 625	9, 875	1, 678	285	7	0
	16, 039	4, 526	794	364	63	16	12	4	3

Table 3

Computation-time [ms] on Nexus 5. Total 21, 821 controller calls. Entire controller call (CTRL all), MPC optimization only, individual QP iterations.

			QP iterations							
	CTRL all	**MPC only**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Min	518	262	262	278	357	428	444	558	551	574
Ave.	1, 224	920	681	664	713	731	752	817	791	865
Max	8, 337	8, 000	2, 429	1, 686	1, 629	1, 499	1,158	1,140	1, 077	1, 097

Table 4

Results of the four cases depicted in Figs. 4 & 5.

Meals		**Unannounced**		**Announced**	
MPC strategy		**Proposed**	**Previous**	**Proposed**	**Previous**
BG [mg/dL] % time	∈ [80, 140]	48.43	48.26	62.37	61.82
	∈ [70, 180]	61.86	61.00	80.78	80.49
	< 80	0.98	3.05	0.28	0.93
	< 70	0.51	2.02	0.07	0.42
	< 60	0.27	1.45	0.02	0.20
	< 50	0.12	0.99	0.00	0.06
	< 40	0.03	0.61	0.00	0.00
	> 180	37.62	36.97	19.15	19.08
	> 250	16.09	17.85	1.80	1.83
	> 300	4.27	5.50	0.21	0.22
	> 350	1.00	1.19	0.00	0.00
	> 400	0.41	0.51	0.00	0.00
#Sims.w.BG [mg/dL]	< 80	165	388	106	184
	< 70	100	243	31	89
	< 60	69	180	10	48
	< 50	43	146	1	30
	< 40	19	114	0	4
	> 180	1221	1221	1219	1221
	> 250	1056	1117	212	218
	> 300	450	573	41	45
	> 350	95	141	0	0
	> 400	52	66	0	0
#Events BG [mg/dL]	< 80	342	827	140	335
	< 70	196	553	37	158
	< 60	124	438	10	89
	< 50	70	340	1	40

Meals		Unannounced		Announced	
MPC strategy		**Proposed**	**Previous**	**Proposed**	**Previous**
	< 40	24	253	0	4
	> 180	3661	3664	3532	3578
	> 250	2900	3068	498	512
	> 300	1132	1400	92	96
	> 350	243	312	0	0
	> 400	116	140	0	0
BG path: Mean [min]		2155	2240	1923	1932
BG path: Min [min]		1857	1882	1763	1763
BG path: Max [min]		3029	3230	2484	2512
Mean insulin/sim [U]		42.77	43.63	48.86	48.88
#Pump-suspensions	≥ 15 min	3722	4446	5006	5001
	≥ 30 min	1439	2176	1831	1992
	≥ 60 min	318	761	220	349
	≥ 90 min	130	449	27	101
	≥ 120 min	37	282	0	27
	≥ 150 min	5	88	0	11
	≥ 180 min	1	26	0	4
	≥ 210 min	1	14	0	0
	≥ 240 min	0	6	0	0
#HMS alarms		507	1469	235	523
#Sims. with alarms		181	366	176	241
Mean LBGI		0.18	0.77	0.12	0.20
Mean HBGI		8.98	9.29	3.63	3.63
CVGA	Zone count: A	0	0	20	13
	Zone count: B	802	670	1160	1132
	Zone count: C	275	207	16	8
	Zone count: D	86	191	23	55

Meals		Unannounced		Announced	
MPC strategy		Proposed	Previous	Proposed	Previous
	Zone count: E	58	153	2	13
	Circle radius	9.45	10.75	5.27	8.05