Published in final edited form as:

Proc IFAC World Congress. 2014 August ; 47(3): 224–230. doi:10.3182/20140824-6-ZA-1003.01085.

State Estimation with Sensor Recalibrations and Asynchronous Measurements for MPC of an Artificial Pancreas to Treat T1DM

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

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

Abstract

A novel state estimation scheme is proposed for use in Model Predictive Control (MPC) of an artificial pancreas based on Continuous Glucose Monitor (CGM) feedback, for treating type 1 diabetes mellitus. The performance of MPC strategies heavily depends on the initial condition of the predictions, typically characterized by a state estimator. Commonly employed Luenberger-observers and Kalman-filters are effective much of the time, but suffer limitations. Three particular limitations are tackled by the proposed approach. First, CGM recalibrations, step changes that cause highly dynamic responses in recursive state estimators, are accommodated in a graceful manner. Second, the proposed strategy is not affected by CGM measurements that are asynchronous, i.e., neither of fixed sample-period, nor of a sample-period that is equal to the controller's. Third, the proposal suffers no offsets due to plant-model mismatches. The proposed approach is based on moving-horizon optimization.

Keywords

State estimation; Model predictive control; Optimization; Artificial pancreas

1. INTRODUCTION

The overall goal of this work is an Artificial Pancreas (AP) for automated insulin delivery to people with Type 1 Diabetes Mellitus (T1DM) (see, e.g., Cobelli et al. [2009], Harvey et al. [2010], Cobelli et al. [2011], Zisser [2011], Doyle III et al. [2014]). In particular, an AP with glucose sensing (measurement for feedback) by a Continuous Glucose Monitor (CGM) (Hovorka [2006]) is considered. A crucial element of an 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) (Parker et al. [1999], Hovorka et al. [2004], Magni et al. [2009], Breton et al. [2012], Turksoy et al. [2013]) have been proposed. The authors' group is focusing increasingly on developing so-called zone-MPC strategies (Grosman et al. [2010, 2011], van Heusden et al. [2012], Gondhalekar et al. [2013, 2014]).

A critical ingredient of every MPC implementation is a mechanism to characterize an initial condition from which to perform predictions. Two main approaches exist. In MPC based on general state-space models, a state estimator is typically employed, e.g., a Luenberger-observer or Kalman-filter (see, e.g., Levine [2011]). Alternatively, when using input-output models, e.g., an Auto-Regressive system with eXogenous inputs (ARX), the initial condition consists trivially of past input and output values (even when using the system's state-space

representation). The state estimator approach is favored by the authors even for ARX model-based predictive control, because it provides simple handles for tuning noise-rejection capabilities. The input-output approach is employed in, e.g., Magni et al. [2007], where it is stated that “The major advantages of this input-output MPC scheme are that an observer is not required”. Both recursive linear state estimators (the class subsumes Luenberger-observers and standard Kalman-filters) and the input-output initialization are straightforward to implement, but have weaknesses. The contribution of this paper is to address three of these weaknesses. A device that initializes MPC predictions is henceforth simply termed a state estimator, regardless of the model class. The proposed state estimator is applicable to both general state-space models as well as input-output models.

The first weakness tackled in this paper is that sensor recalibrations cannot be accommodated well in current state estimators. CGM signals suffer two (at least) types of noise. First, there is high-frequency stochastic noise, the effects of which can, to some extent, be remedied by tuning the gain of a recursive state estimator (Bequette [2004]). Then there is a low-frequency drift, also termed sensor bias, due to slowly undulating characteristics of the CGM sensor gain and changes in the sensor site’s physiology. These low-frequency disturbances are corrected by taking sporadic blood-glucose measurements with a sensor that is more accurate than the CGM, e.g., by a point of care blood-glucose measurement device. The CGM is subsequently “recalibrated” with respect to the reference measurement. Upon receipt of a recalibrated data-point a recursive state estimator could update its state estimate as usual, or possibly employ a higher gain than when updating using CGM data, to reflect the higher confidence. Such an approach was proposed in Kuure-Kinsey et al. [2006] for glucose estimation based on Kalman filtering, and such approaches appear to work well for the purpose of glucose estimation. However, for the purpose of state initialization in MPC the strategy is not ideal, because after a recalibration the state estimator undergoes a period of lively dynamics. These energetic responses may result in meaningless predictions that can lead to serious over-delivery. Thus, in the authors’ controllers, large recalibrations are followed by a period where the insulin infusion rate is constrained to the patient’s basal rate. This seems wasteful, as a recalibration is the introduction of high-fidelity data into the system. Preferably the system could exploit this data and perform *better* after a recalibration, not have to undergo intentional, temporary crippling.

The second weakness is that asynchronous CGM data cannot be accommodated in current, recursive state estimators, where “asynchronous” means both that the sample-period of the CGM may not be fixed, and furthermore that the time-instants the CGM and controller perform updates may not be equal. The authors’ controllers (both physical controller and discrete-time prediction model) are based on a $T = 5$ min sample-period. Typical CGMs have the same sample-period, much of the time. However, CGMs may delay their output during times of high uncertainty. Also, communication disruptions between sensor and meter cause delayed measurement updates, only once data-transfer is reestablished. A state estimator based on a fixed sample-period may over-estimate the rate of change of the data if the actual sample-period is elongated, and not compensate for the delay between controller update times and the latest CGM measurements. Both issues cause MPC predictions that are

initialized in such a way that they veer off the CGM trajectory, possibly resulting in inappropriate insulin delivery.

The third weakness is that due to plant-model mismatch, model-based recursive state estimators cannot always achieve offset-free estimates, even in steady-state, when the state is not admissible with respect to the model, input, and measured output. Offsets can be partially remedied by increasing the estimator gain, but this undesirably results in increased responsiveness to high-frequency noise.

The contribution of this paper is to propose a state estimation strategy that tackles the aforementioned three weaknesses. The proposal is based on moving-horizon optimization and is not a recursive estimator. It is inspired by, but *not* equal to, the common notion of moving-horizon estimation (Rawlings and Mayne [2009]). The proposed method performs optimization to fit a continuous-time function to the CGM data. Sensor recalibrations are accommodated straightforwardly by including a discontinuity in the glucose output *value*, but not its derivatives, within the function definition. Importantly, the magnitude of the discontinuity need not be prescribed, but is identified by the optimization. The data fitting exploits the CGM time-stamps *and* controller call time, thus asynchronicity is handled naturally. Crucially, after optimization the fitted function is sampled at *exactly* the controller model's sample-period T , ignoring the recalibration discontinuity, to synthesize an output trajectory. In combination with historical input data, and assuming observability, the current model state is constructed to reflect the fitted output trajectory without offset. To the authors best knowledge the proposed strategy is novel. The proposed strategy can conceivably be combined with a Kalman filter, or other signal processing technique, to pre-treat the CGM data. However, for brevity, the exposition of this paper is based on the use of raw CGM data.

2. PRELIMINARIES

2.1 Linear time-invariant insulin-glucose model

The insulin-glucose model of van Heusden et al. [2012] is employed in this paper and is summarized as follows. The model is a discrete-time, linear time-invariant (LTI) system with sample-period $T = 5$ [min]. The time step index is denoted by i . The scalar plant input is the administered insulin bolus $u_{IN,i}$ [U] delivered per sample-period, and the scalar plant output is the subject's blood-glucose value $y_{BG,i}$ [mg/dL]. The plant is linearized around a steady-state, that is assumed to be achieved by applying the subject-specific, time-dependent basal input rate $u_{BASAL,i}$ [U/h], and is assumed to result in a steady-state blood-glucose output $y_s = 110$ [mg/dL].

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

$$u_i := u_{IN,i} - u_{BASAL,i} \frac{T}{60 \text{ min}}$$

$$y_i := y_{BG,i} - y_s.$$

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

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

with poles $p_1 = 0.98$, $p_2 = 0.965$, a so-called *safety factor* $F := 1.5$ (unitless, personalizable but fixed to 1.5 throughout this paper), the subject specific *total daily insulin* amount $u_{\text{TDI}} \in \mathbb{R}_{>0} [\text{U}]$, and where the constant

$$c := -60(1 - p_1)(1 - p_2)^2 \in \mathbb{R}$$

is employed to set the correct gain, and for unit conversion. The 1800 term stems from the “1800 rule” for estimating blood-glucose decline w.r.t. the delivery of rapid-acting insulin (Walsh and Roberts [2006]).

The state-space realization of (1) for control synthesis is

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

$$y_i = Cx_i \quad (2b)$$

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

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

$$C := \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{1 \times n}$$

$$n=3.$$

Let $\mathcal{O} := [C^\top (CA)^\top (CA^2)^\top]^\top \in \mathbb{R}^{n \times n}$, and note that $\mathcal{O}$ is equal to the identity matrix flipped top-to-bottom.

Remark 1— $\det(\mathcal{O}) \neq 0$, i.e., (A, C) is observable.

2.2 Nominal model predictive control outline

The reader is referred to Rawlings and Mayne [2009] for MPC background. Let $\mathbb{Z}$ denote the set of integers, $\mathbb{Z}_+$ the set of positive integers, and $\mathbb{Z}_a^b$ the set $\{a, \dots, b\}$ of consecutive integers from a to b . Let $N \in \mathbb{Z}_+$ denote the prediction horizon, and u and x the predicted values of input u and state x . Then, MPC performs closed-loop control by applying, at each step i , the first control input u_0^* of the predicted, optimal control input trajectory $\{u_0^* \dots u_{N-1}^*\}$, characterized by the minimization

$$\{u_0^* \dots u_{N-1}^*\} := \arg \min_{\{u_0, \dots, u_{N-1}\}} J(x_i, \{u_0, \dots, u_{N-1}\}) \quad (3)$$

of a suitable cost function $J(\cdot, \cdot)$ (details omitted for brevity), subject to suitable constraints, and furthermore subject to the predictions performed employing model (2):

$$x_0 := x_i, \quad x_{k+1} := Ax_k + Bu_k \quad \forall k \in \mathbb{Z}_0^{N-1}. \quad (4)$$

The predicted state trajectory is initialized in (4) to the estimated model state, the value of which profoundly affects the performance of the resulting MPC control law. No notational distinction between actual and estimated state is made, because state x of (2) can *only* be estimated.

2.3 Controller timing and input history

The sample-period of (2), and the time interval between controller updates of control input u , are assumed to be the same and equal to T . For simplicity we further assume any controller employing model (2) to have access to the exact control input history, where previous control inputs u_i were applied at time intervals of exactly T . We denote the actual time instants of the controller call by $\tau_i = \tau_{i-1} + T$.

2.4 Sensor timing, sensor recalibration, and output history

Each measurement is defined by a triple $(\tilde{y}_j, t_j, r_j)$, where $j \in \mathbb{Z}_+$ denotes the measurement index that is incremented with each new measurement, $\tilde{y}_j \in \mathbb{R}$ denotes the CGM output, analogous to y of (2b) (i.e., with set-point y_s subtracted), as provided by the CGM at time-instant t_j . The variable $r_j \in \mathbb{Z}$ denotes a recalibration counter, and is incremented each time the sensor is recalibrated ($r_0 = 0$).

The time interval between successive measurements may not be precisely T . However, we suppose that $t_j - t_{j-1} < 2T$ for all j . Analogously, we assume the time interval between a controller call at τ_j and the most recent measurement at t_j to be less than $2T$. If the interval exceeds two sample-periods then, for an interval of a low multiple of sample-periods, a strategy employing open-loop predictions of model (2) to “fill the gap” may be useful. For

simplicity such scenarios are not considered here, although they are in the clinical controller implementations.

Each output measurement $\tilde{y}_i$ at time t_i suffers from errors due to process noise and measurement noise. However, CGM noise has proven difficult to model accurately (Hovorka [2006]), thus in this work we make no assumptions about the measurement errors, and include in the state-estimation scheme no strategy for exploiting perceived knowledge of the noise characteristics. However, we assume that measurements $\tilde{y}_j$ such that $r_j - r_{j-1}$ are exact, because r is incremented when the sensor is recalibrated. The proposed state-estimation strategy achieves rejection of high-frequency disturbances to some (tunable) extent, but even without recent sensor recalibrations, the proposed strategy estimates the state under the assumption that the low-frequency measurement bias is zero. Without further knowledge there appears to be no alternative.

2.5 State-reconstruction based on exact outputs and inputs

We denote by I_a the $a \times a$ identity matrix, by $0_{\{a,b\}}$ the $a \times b$ zero matrix, and by the $\otimes$ Kronecker product.

At each step i , given the *exact* sequence $\{y_k\}_{k=i-n+1}^i$ of past outputs (and present), synchronized to the controller timing $\tau_i = \tau_{i-1} + T$, and further given the *exact* sequence $\{u_k\}_{k=i-n+1}^{i-1}$ of past control inputs, the current state x_i of model (2) may be reconstructed, e.g., as follows. Let

$$U_i := [u_{i-n+1} \ \cdots \ u_{i-1}]^T \in \mathbb{R}^{(n-1)}$$

$$Y_i := [y_{i-n+1} \ \cdots \ y_i]^T \in \mathbb{R}^n$$

$$X_i := [x_{i-n+1}^T \ \cdots \ x_i^T]^T \in \mathbb{R}^{n^2}$$

$$\bar{A} := \begin{bmatrix} I_n & A^T & \cdots & (A^{n-1})^T \end{bmatrix}^T \in \mathbb{R}^{n^2 \times n}$$

$$\hat{A} := \begin{bmatrix} 0 & 0 & \cdots & 0 \\ I_n & 0 & \cdots & 0 \\ A & I_n & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A^{n-2} & A^{n-3} & \cdots & I_n \end{bmatrix} \in \mathbb{R}^{n^2 \times n(n-1)}$$

$$\bar{B} := \hat{A}(I_{n-1} \otimes B) \in \mathbb{R}^{n^2 \times (n-1)}$$

$$\bar{C} := (I_n \otimes C) \in \mathbb{R}^{n \times n^2}$$

$$F := \begin{bmatrix} 0_{\{n, n(n-1)\}} I_n \end{bmatrix} \in \{0, 1\}^{n \times n^2}$$

such that

$$X_i = \bar{A}x_{i-n+1} + \bar{B}U_i \quad (5)$$

$$Y_i = \bar{C}X_i \quad (6)$$

$$x_i = F X_i \quad (7)$$

where at step i all except X_i are known. From (5) and (6):

$$x_{i-n+1} = (\bar{C}\bar{A})^{-1}(Y_i - \bar{C}\bar{B}U_i). \quad (8)$$

The current state x_i is then characterized via (5) and (7). The inverse in (8) exists by Remark 1, because $\bar{C} = \mathcal{O}$.

3. PROPOSAL: STATE-ESTIMATION VIA OUTPUT TRAJECTORY FITTING

3.1 Proposal outline

At each step i the parameter $\theta_i \in \Theta$ defining a continuous-time function $f: \mathbb{R} \times \Theta \rightarrow \mathbb{R}$ is identified such that it closely fits recent data-points. The continuous-time function $f(t, \theta_i)$ is subsequently sampled at time instants $\tau_k, k \in \mathbb{Z}_{i-n+1}^i$ to synthesize a trajectory $\{\hat{y}_k\}_{k=i-n+1}^i$ of synchronous, past (and one present) output values. This manufactured output trajectory is employed, in conjunction with the exact sequence $\{u_k\}_{k=i-n+1}^{i-1}$ of past control inputs, to construct an estimate of the current state x_i by the mechanism described in Section 2.5.

The function fitting is performed using unconstrained least-squares fitting of polynomials. More general cost functions, more general functions $f(\cdot)$, and also constraints, could be considered, but these complexities are dispensed with here to focus on the proposed advantages in terms of timing and sensor recalibrations. The benefit, with regards to timing

and the asynchronous nature of the CGM data-points, is that the function fitting can be performed with data-points that are temporally distributed in an arbitrary way. The important, novel functionality with respect to sensor recalibrations is that due to the optimization-based nature (in contrast to recursive estimators) a discontinuity can be accommodated when a recalibration occurs. Crucially, the discontinuity's size need not be known, but is identified from the data via the optimization. Assuming that at most one recalibration occurred in the near history, the discontinuity is included when fitting data points prior to recalibration, but is not included when fitting more recent data-points. Critically, the discontinuity is not included when sampling $f(\cdot)$ to synthesize the fabricated output trajectory of $\hat{y}$ values.

The optimization penalizes the deviation $\theta_i - \theta_{i-1}$ of the parameter from one step to the next, thus introducing a “viscosity” for rejecting high-frequency disturbances.

3.2 Data fitting with function discontinuity

For consistent interpretation of the value of parameter θ_i as i progresses, the function $f(\cdot)$ is fitted shifting the current time τ_i to the origin. The class of continuous-time functions considered for fitting is the p -order polynomial

$$f(t, \theta_i) := \sum_{k=0}^p a_{(i,k)} (t - \tau_i)^k = [1(t - \tau_i) \dots (t - \tau_i)^p] \theta_i \quad (9)$$

$$a_{(i,k)} \in \mathbb{R} \quad \forall (i, k) \in \mathbb{Z} \times \mathbb{Z}_0^p$$

$$\theta_i := [a_{(i,0)} \dots a_{(i,p)}]^\top \in \mathbb{R}^{p+1},$$

where p is a design parameter. Let the design parameter $M \in \mathbb{Z}_+$ denote a length of measurement history to consider. For each step i , let $c_i \in \mathbb{Z}_+$ denote the index of the most recent measurement, and let $d_i \in \mathbb{Z}_{c_i-M+1}^{c_i}$ denote the index of the most recent measurement that followed a sensor recalibration. The range specified for d implies that a recalibration occurred within the M -length history horizon. The case when the latest recalibration occurred prior to the M -length history horizon is simple and not discussed further. For simplicity we do not discuss the case with multiple recalibrations within the history horizon M , although such cases could be accommodated.

At step i , the measurements employed for state estimation are $(\tilde{y}_j, t_j, r_j)$, $j \in \mathbb{Z}_{c_i-M+1}^{c_i}$. It holds that $r_i = r_j + 1 \forall (i, j) \in \mathbb{Z}_{d_i}^{c_i} \times \mathbb{Z}_{c_i-M+1}^{d_i-1}$. Let $\delta_i \in \mathbb{R}$ denote the (unknown) size of measurement discontinuity resulting from a recalibration, and define the augmented parameter $\bar{\theta}_i := [\delta_i \theta_i^\top]^\top \in \mathbb{R}^{p+2}$. Denote the error, between the discontinuous fitted function and the data, as follows:

$$e_{(i,j)} := \begin{cases} \tilde{y}_j - [0 \ 1 \ (t_j - \tau_i) \ \cdots \ (t_j - \tau_i)^p] \bar{\theta}_i & \text{if } j \in \mathbb{Z}_{d_i}^{c_i} \\ \tilde{y}_j - [1 \ 1 \ (t_j - \tau_i) \ \cdots \ (t_j - \tau_i)^p] \bar{\theta}_i & \text{otherwise.} \end{cases}$$

Let $R_k \in \mathbb{R}_{>0} \forall k \in \mathbb{Z}_1^M$ denote costs to penalize errors $e_{(i,j)}$, time-dependent with respect to relative time the measurement was taken, but not time-dependent with respect to actual time. Further let $Q_i \in \mathbb{R}^{(p+1) \times (p+1)}$, $Q_i \geq 0$ denote a cost for penalizing parameter deviations $\theta_i - \theta_{i-1}$. The optimal augmented parameter $\bar{\theta}_i^*$ is characterized by the solution of the following quadratic program:

$$\bar{\theta}_i^* := \arg \min_{\bar{\theta}_i \in \mathbb{R}^{(p+2)}} (\theta_i - \theta_{i-1}^*)^\top Q_i (\theta_i - \theta_{i-1}^*) + \sum_{k=1}^M R_k e_{(i,c_i-k+1)}^2.$$

The cost R_k should, in general, be chosen such that $R_k \gg R_{k+1}$, i.e., such that more recent measurements influence the optimal parameter estimate θ_i^* more than older ones. The cost matrix Q_i is chosen to penalize parameters $a_{(i,k)}$ of (9), and should generally be diagonal. A higher cost allows to set a “viscosity” on the rate of change of, e.g., the value via $Q_{(1,1)}$, or the velocity via $Q_{(2,2)}$, etc. Note that after a recalibration it is desirable to select $Q_{(1,1)} = 0$ in order to facilitate an instantaneous response to the recalibration step change.

3.3 Output trajectory manufacture

Given the optimized parameter θ_i^* , the synthesized output trajectory, employed for constructing the estimated state via the method described in Section 2.5, is defined by sampling the function $f(t, \theta_i^*)$ at times $t \in \{0, -T, -2T, \dots\}$.

4. ILLUSTRATIVE EXAMPLES

In this section the behavior and benefits of the proposed state estimation strategy are demonstrated by simple, numerical examples. The parameter choices were made to produce the simplest, within reason, instance of the proposed strategy. The order $p = 1$ of the polynomial of (9) is employed, i.e., we perform a straight-line fit. Note that the number of data-points M employed must equal, or exceed, the degrees of freedom of the function fitting. Thus we select a history horizon $M = 3$, facilitating a straight-line fit with one recalibration. We let $R_1 = R_2 = 1$ and $R_3 = 0.1$, to penalize the error w.r.t. the eldest data-point less than the error associated with the most recent two data-points. Finally, we select $Q = 0_{\{2,2\}}$, i.e., the optimal parameter θ_i^* is independent of the previous step's θ_{i-1}^* .

We compare the responses of the proposed state estimator with a linear state estimator that is based on model (2):

$$\tilde{x}_i = A x_{i-1} + B u_{i-1} \quad , \quad \tilde{y}_i = C \tilde{x}_i \quad (10a)$$

$$x_i = \tilde{x}_i + L(y_i - \tilde{y}_i). \quad (10b)$$

4.1 Rejection of plant-model mismatches

Fig. 1 demonstrates how the gain L of a linear state estimator affects the ability to reject plant-model mismatches. The CGM is constant at 280 mg/dL, and insulin infusion is performed at the basal-rate. A higher gain L rejects mismatches more effectively, but results in elevated responsiveness to noise (not demonstrated). Note that, in this example, the linear state estimator's state is initialized to achieve the CGM value in steady state. Despite this optimal initialization (that is not possible in practice), the linear state estimator's state estimate drifts, inducing a steady state mismatch in estimator output Cx_i and the CGM signal. The reason for this is that model (2) is based on linearization around $y_s = 110$ mg/dL, and that the elevated, steady-state CGM value is not compatible with the basal insulin delivery. In contrast, the proposed state estimator suffers no such mismatches, because the synthesized output trajectory is manufactured based solely on the CGM data, not model (2), and because, by observability, the mechanism of Section 2.5 constructs a state that corresponds exactly to this fabricated output trajectory.

4.2 MPC & CGM synchronized – sample-period incorrect

In this example we demonstrate what happens when the CGM data is transmitted every 9 min instead of $T = 5$ min, under the assumption that the controller updates simultaneously, only every 9 min. This is *not* how the MPC is implemented in practice (see Section 2.3). Nevertheless, it is hoped the example is instructive.

We consider a CGM trajectory that is rising at 1 mg/dL/min, sampled every 9 min. The CGM data are recursively input to estimator (10), that is not able to exploit the data's time-stamps, because model (2) is based on a $T = 5$ min sample-period. The gain L is chosen high, i.e., the estimator is responsive and the output error is rejected well. The result is depicted in Fig. 2. Despite achieving an accurate starting *value* for the output Cx_i , the rate of change is clearly mis-initialized to an over-estimated value, and the MPC predictions veer away from the CGM trajectory. In contrast, the proposed estimator exploits both the controller's call time and also the CGM time-stamps, and accounts for arbitrary timing in an appropriate manner. Based on the proposed estimator's state the MPC predictions are a continuation of the CGM trajectory.

4.3 MPC & CGM sample-instants offset

In this example we demonstrate the ability of the proposed estimator to accommodate delays between the controller update time instants and the CGM. We consider a CGM trajectory with rate of change of 1 mg/dL/min, with a data-point every $T = 5$ min. The controller updates every $T = 5$ min, delayed by 4 mins w.r.t. the latest CGM value. The result is plotted in Fig. 3. Despite the delay the linear estimator causes the MPC predictions to start at the most recent CGM value. In contrast, the proposed estimator initiates the MPC predictions from an extrapolated value lying on a continuation of the CGM data trajectory.

The benefit of the proposed estimator, in regard to delay compensation, is negligible when the CGM's rate of change is low, which is most of the time. However, the CGM signal undergoes rapid change after, e.g., meal ingestion or the commencement of exercise. It is exactly at these challenging times that controller responsiveness is crucial.

4.4 MPC & CGM asynchronous

In this example we consider the case where the controller updates the control input at a sample-period $T = 5$ min, as intended. The CGM value rises at 1 mg/dL/min, but updates its value only every 7 mins. Due to this mismatch in sample-periods, MPC and CGM sometimes update simultaneously, often times there is a delay between them, and other times no CGM update occurred since the previous MPC update.

We consider a linear state estimator with high gain, updated with the most recent CGM value at each controller call. The resulting MPC predictions are depicted in the top subplot of Fig. 4. The predictions produce a featherlike spread around the CGM trajectory, where this spread is a result of both an offset in glucose value, as well as mis-initialization of the rate of change. In contrast, plotted in the bottom subplot of Fig. 4 are the MPC predictions when initialized by the proposed estimator. The predictions overlay tightly. They do not overlap perfectly due to the controller tuning; the predictions veer slightly downwards due to the predicted delivery of insulin.

4.5 Recalibration

The initial motivation for the proposed approach was to gracefully accommodate sensor recalibrations – demonstrated next. The controller and CGM are synchronized to the correct sample-period; $T = 5$ min. The CGM reads 200 mg/dL until 14:25, is recalibrated to 250 mg/dL at 14:30, and remains at that reading thereafter.

The response with the linear state estimator is depicted in Fig. 5. Looking at the estimated blood-glucose level (green line), it can be seen that the linear estimator performs admirably in terms of rapid convergence. The linear state estimator has a high gain, leading to a “forceful” correction of the estimator state x to produce an output Cx that equals the CGM value. However, such high gain estimation is inappropriate for initializing MPC predictions, due to the highly dynamic response of the predictions for a protracted period following the recalibration. This response in the state estimate causes a large, undesirable overshoot in the insulin delivery. A low gain linear state estimator may be more desirable for MPC state initialization here, resulting in predictions that have less incline and consequently a more conservative insulin delivery. However, a low gain estimator results in sluggish convergence to the correct glucose level and, depending on the glucose value, an offset due to plant-model mismatch (see Section 4.1).

Fig. 6 shows the response with the proposed estimator. The estimated glucose value instantaneously changes at 14:30 to the recalibrated value, and the state estimate instantaneously switches to a new value that, first, reflects the new CGM value, and, second, reflects the rate of change of the CGM trajectory in recent history. The MPC predictions beyond the recalibration are therefore nearly not visible. The resulting insulin delivery undergoes a step change upwards at the recalibration time instant. The two short-term

deviations from steady-state delivery are due to the pump's discretization and carryover scheme. Both before and after the insulin step change a delivery in excess of the basal rate is desirable, due to the hyperglycemia. Thus, a safety mechanism that enforces basal delivery during the rapid transients of the state estimator would pose an obstacle to effective glycemia control.

5. CONCLUSION

A novel state estimation scheme, based on moving-horizon optimization, was proposed to tackle problems associated with recursive state estimators for initializing MPC optimizations based on CGM data. The mechanics and benefit of the proposed strategy were demonstrated using simple, synthetic examples. The proposed method was tested via the University of Padova/Virginia Food and Drug Administration (FDA) accepted metabolic simulator (Kovatchev et al. [2009]) and behaves comparably to a responsively tuned linear state estimator in “normal” circumstances, i.e., when not dealing with the problem instances that motivated the proposal. In future work more challenging scenarios will be simulated, to challenge and verify the proposed estimator's performance with CGM data that more closely resembles that obtained in clinical trials.

The proposed scheme offers a exible foundation for extensions. (a) The proposed method can be combined with, e.g., a Kalman filter, for tackling high-frequency noise, when far from a recalibration. (b) The proposed method was described using polynomials as the fitting function, but in Miller and Strange [2007] it was suggested that Fourier series are effective for fitting to CGM data. (c) Future CGMs may provide data richer than only blood-glucose estimates, e.g., with accompanying estimates of confidence bounds. The optimization based approach may offer an avenue to exploit such auxiliary information. (d) The notion of bias-control – the ability to manipulate the state estimate in a well-defined manner based on further sensors or user input – may be facilitated, potentially leading to improved safety after detecting, e.g., a meal, exercise, a pump failure, or a sudden loss of CGM sensitivity. These avenues will be investigated in future research.

Acknowledgments

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

REFERENCES

- Bequette, BW. AACC American Control Conf. Boston, MA, USA: 2004 Jun. Optimal Estimation Applications to Continuous Glucose Monitoring; p. 958-962.
- 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(9):2230–2237. [PubMed: 22688340]
- 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]
- Cobelli C, Renard E, Kovatchev B. Artificial Pancreas: Past, Present, Future. *Diabetes*. 2011 Nov; 60(11):2672–2682. [PubMed: 22025773]

- 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.
- 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(6): 1446–1460. [PubMed: 24351171]
- Gondhalekar, R., Dassau, E., Doyle, FJ, III. AACC American Control Conf. Portland, OR, USA: 2014 Jun. MPC Design for Rapid Pump-Attenuation and Expedited Hyperglycemia Response to Treat T1DM with an Artificial Pancreas.
- Grosman B, Dassau E, Zisser HC, Jovanovi L, Doyle FJ III. Zone Model Predictive Control: A Strategy to Minimize Hyper- and Hypoglycemic Events. *J. Diabetes. Sci. Technol.* 2010 Jul; 4(4): 961–975. [PubMed: 20663463]
- Grosman, B., Dassau, E., Zisser, H., Jovanovi L., Doyle, FJ, III. Proc. 18th IFAC World Congress. Milan, Italy: 2011 Aug. Multi-Zone-MPC: Clinical Inspired Control Algorithm for the Artificial Pancreas; p. 7120-7125.
- Harvey RA, Wang Y, Grosman B, Percival MW, Bevier W, Finan DA, Zisser H, Seborg DE, Jovanovi 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]
- Hovorka R. Continuous glucose monitoring and closed-loop systems. *Diabetic Med.* 2006 Jan; 23(1): 1–12.
- 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]
- 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(1):44–55. [PubMed: 19444330]
- Kuure-Kinsey, M., Palerm, CC., Bequette, BW. EMBS Annual Int. Conf. New York City, NY, USA: 2006 Aug. A Dual-Rate Kalman Filter for Continuous Glucose Monitoring; p. 63-66.
- Levine, WS., editor. The Control Handbook. 2. Boca Raton, FL, USA: CRC Press; 2011.
- 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.
- Magni L, Raimondo DM, Bossi L, Dalla Man C, De Nicolao G, Kovatchev B, Cobelli C. Model Predictive Control of Type 1 Diabetes: An *in Silico* Trial. *J. Diabetes Sci. Technol.* 2007 Nov; 1(6):804–812. [PubMed: 19885152]
- Miller M, Strange P. Use of Fourier Models for Analysis and Interpretation of Continuous Glucose Monitoring Glucose Profiles. *J. Diabetes Sci. Technol.* 2007 Sep; 1(5):630–638. [PubMed: 19885131]
- 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(2):148–157. [PubMed: 9932336]
- Rawlings, JB., Mayne, DQ. Model Predictive Control: Theory and Design. Madison, WI, USA: Nob Hill Publishing; 2009 Aug.
- 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(5):386–400. [PubMed: 23544672]
- 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(7):1839–1849. [PubMed: 22127988]
- Walsh, J., Roberts, R. Pumping Insulin. 4. San Diego, CA, USA: Torrey Pines Press; 2006.
- 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(5):1283–1286. [PubMed: 22027329]

Fig. 1. Demonstration of plant-model mismatches. CGM = 280 mg/dL (●). Low-gain linear estimator (▴). High-gain linear estimator (■). Proposed estimator (—).

Fig. 2. Demonstration of elongated sample-period. CGM and MPC synchronized. CGM (●) rate of increase: 1 mg/dL/min. Linear estimator-based MPC predictions (▲). Proposed estimator-based MPC predictions (●).

Fig. 3.

Demonstration of 4 min delay between MPC update and CGM. CGM (●) rate of increase: 1 mg/dL/min. Linear estimator-based MPC predictions (▲). Proposed estimator-based MPC predictions (—).

Fig. 4.

Demonstration with MPC and CGM asynchronous. MPC sample-period: 5 min. CGM sample-period: 7 min. CGM (●) rate of increase: 1 mg/dL/min. **Top:** Linear estimator, MPC predictions (►). **Bottom:** Proposed estimator, MPC predictions (—).

Fig. 5.

Demonstration of recalibration response with linear state estimator and *no* safety features. MPC and CGM synchronous: $T = 5$ min. CGM (●). MPC predictions (▲). Estimated blood-glucose value (■).

Fig. 6.

Demonstration of recalibration response with pro-posed state estimator and *no* safety features. MPC and CGM synchronous: $T = 5$ min. CGM (●). MPC predictions (▲). Estimated blood-glucose value (■).