Technical University of Denmark

DTU

Tuning of Controller for Type 1 Diabetes Treatment with Stochastic Differential Equations

Duun-Henriksen, Anne Katrine; Boiroux, Dimitri; Schmidt, Signe; Skyggebjerg, Ole; Madsbad, Sten; Jensen, Peter Ruhdal; Jørgensen, John Bagterp; Poulsen, Niels Kjølstad; Nørgaard, Kirsten; Madsen, Henrik

Link back to DTU Orbit

DTU Library
Technical Information Center of Denmark

Tuning of Controller for Type 1 Diabetes Treatment with Stochastic Differential Equations [⋆]

Anne Katrine Duun-Henriksen [*] Dimitri Boiroux [*]
Signe Schmidt [**] Ole Skyggebjerg [****] Sten Madsbad [**]
Peter Ruhdal Jensen [***] John Bagterp Jørgensen [*]
Niels Kjølstad Poulsen [*] Kirsten Nørgaard [**] Henrik Madsen [*]

[*] *Department of Informatics and Mathematical Modeling, Technical University of Denmark, 2800 Kgs. Lyngby Denmark (e-mail: akdu@imm.dtu.dk).*
[**] *Department of Endocrinology, Hvidovre University Hospital, 2650 Hvidovre, Denmark*
[***] *Department of Systems Biology, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark*
[****] *Horus Aps, 1671 Copenhagen, Denmark*

Abstract: People with type 1 diabetes need several insulin injections every day to keep their blood glucose level in the normal range and thereby avoiding the acute and long term complications of diabetes. One of the recent treatments consists of a pump injecting insulin into the subcutaneous layer combined with a continuous glucose monitor (CGM) frequently observing the glucose level. Automatic control of the insulin pump based on CGM observations would ease the burden of constant diabetes treatment and management. We have developed a controller designed to keep the blood glucose level in the normal range by adjusting the size of insulin infusions from the pump based on model predictive control (MPC). A clinical pilot study to test the performance of the MPC controller overnight was performed. The conclusion was that the controller relied too much on the local trend of the blood glucose level which is a problem due to the noise corrupted observations from the CGM. In this paper we present a method to estimate the optimal Kalman gain in the controller based on stochastic differential equation modeling. With this model type we could estimate the process noise and observation noise separately based on data from the first clinical pilot study. In doing so we obtained a more robust control algorithm which is less sensitive to fluctuations in the CGM observations and rely more on the global physiological trend of the blood glucose level. Finally, we present the promising results from the second pilot study testing the improved controller.

Keywords: Stochastic differential equations, Model predictive control, Kalman filters, Artificial pancreas, Type 1 diabetes.

1. INTRODUCTION

Type 1 diabetes is a chronic disease characterized by destruction of the insulin producing beta cells of the pancreas. Insulin is crucial for the regulation of the blood glucose level and people with type 1 diabetes are therefore dependent on exogenous insulin supply. The size of the insulin dose must be determined carefully. Underdosing can result in a too high blood glucose level (hyperglycemia) which in the long term can lead to, e.g. kidney failure, blindness or nerve damage. Too much insulin (in case of overdosing) can lead to a too low blood glucose level (hypoglycemia) which can have serious acute effects such as coma or even death.

The insulin is most often delivered subcutaneously either via a pen needle or an insulin pump. The pump delivers a steady basal rate combined with meal boluses (dosages to cover meals) to resemble the normal insulin secretion from a healthy pancreas.

Patients are adviced to check their blood glucose level with finger prick measurements several times a day to obtain tight diabetes control [American Diabetes Association (2012)]. The goal is to keep their blood glucose level as close to normal range (4.00-8.00 mmol/L) as possible without compromising safety. Instead of finger prick measurements more and more patients rely on continuous glucose monitors (CGM) to keep track of the blood glucose level. The CGM uses minimally invasive sensors capable of reporting the glucose level every five minutes. The sensor is placed in the subcutaneous layer and thus the observations are delayed in relation to the discrete capillary blood glucose observations that normally are
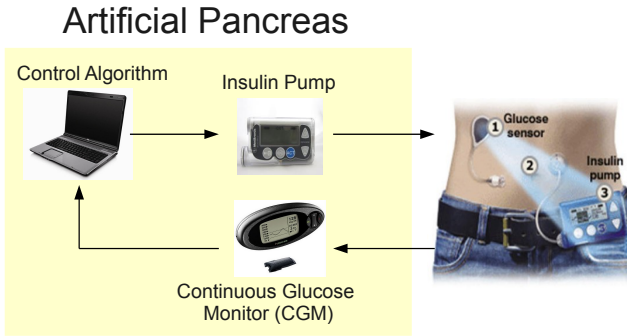
Fig. 1. An overview of the closed loop system consisting of a computer with the control algortihm, a CGM and an insulin pump.
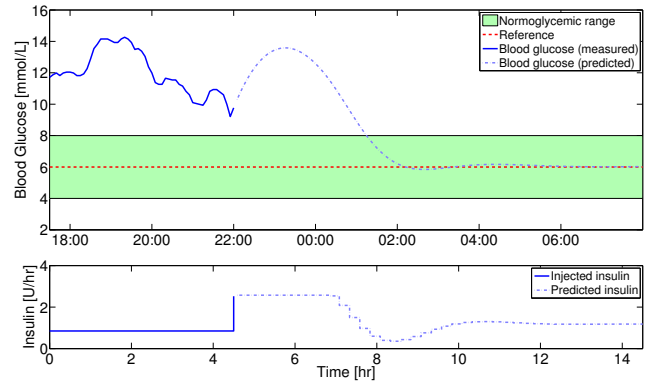

Fig. 2. Top: CGM observations (solid line) and blood glucose prediction at the time of closing the loop (dashed line) from the first pilot study. As seen, the controller relies too much on the local trend. Bottom: Insulin delivered by the pump (solid line) and predicted insulin supply (dashed line).

measured from the finger tips. In addition, the sensors can experience problems with accuracy and precision due to physical, chemical and biological factors [Keenan et al. (2011); Hovorka (2006)].

Even though insulin pump therapy and CGMs can lead to better diabetes control, the treatment still requires constant attention and action from the patient. Furthermore, many patients with type 1 diabetes live in fear of getting severe hypoglycemia during sleep, while they are unable to react.

Several research groups are working on closing the loop with an automatic control algorithm regulating the insulin delivery based on feedback from the CGM [Cobelli et al. (2011); Hovorka (2011)]. The general principle behind this artificial pancreas is illustrated in Fig. 1.

We have developed a control algorithm based on model predictive control (MPC) that predicts future blood glucose values on the basis of the current level. If the control algorithm predicts hyperglycemia, insulin is delivered to bring the blood glucose level down to normal. If the control algorithm predicts hypoglycemia it decreases or shuts off the insulin supply until the blood glucose level increases again.

To test the performance of the controller we did a clinical pilot study on a type 1 diabetes subject. The challenge was to keep the blood glucose level in the normal range overnight. In Fig. 2 the blood glucose level is shown. At the time of starting the closed loop (22:00) the local trend was increasing but the global trend was decreasing after a meal with a corresponding insulin bolus. The MPC prediction of the blood glucose level at this time is also seen in Fig. 2. Due to the local increasing trend this prediction is unrealistically increasing and the controller suggests a too large bolus to compensate for this increase. This lead the subject into hypoglycemia and the clinician had to administrate I.V. glucose to bring the blood glucose level up again immediately.

The main conclusion from this pilot study was that the large amount of noise in the system was preventing the controller from predicting the blood glucose level and thus delivering the correct insulin bolus.

In this paper we address this problem by estimating the process and observation variances separately in the MPC controller using *stochastic differential equations* with data from the pilot study. This paper describes this novel procedure and presents the promising results from the second pilot study testing the improved MPC controller.

2. THE PILOT STUDY SETUP

The data used for tuning the controller came from the first pilot study performed overnight on a type 1 diabetes subject. We chose to do the study overnight because the amount of disturbances to the system, e.g. meals and exercise are minimal at this time of day. In [Boiroux et al. (2012)] details about the setup is described. The subject wore an insulin pump (Medtronic Paradigm Veo, Minneapolis, USA) and two CGMs (Dexcom Seven Plus, San Diego, USA) measuring the glucose level every five minutes. From the two CGMs we had 162 and 158 observations, respectively and both were used for parameter estimation. The first 270 minutes of the pilot study were open loop where the patient was in control of the insulin supply and the last 540 minutes were closed loop.

3. CONTROLLER DESIGN

3.1 Control algorithm

The control algorithm is based on a single input-single output model in continuous time. It is a second order model with the following transfer function:

$$H_s(s) = \frac{M}{(\tau s + 1)^2} \tag{1}$$

$$\tag{2}$$

The output is the blood glucose level in mmol/L and the input is the insulin supply from the pump in mU/min, the standard unit for insulin. τ is the insulin action time defined as the time in minutes it takes to reach the minimum blood glucose level after an insulin injection. M is defined as:

$$M = -\tau \exp(1) I_{SF} \tag{3}$$

I_{SF} is the *insulin sensitivity factor* defined as the maximum decrease in blood glucose level per unit of insulin. M and τ are patient specific and estimated from patient data. The insulin pump injects a bolus every 15 minutes based on the decision of the controller and the glucose level is observed every five minutes with the CGM. The bolus size is computed on the basis of the current glucose level and the blood glucose prediction. For more details on the MPC, see [Boiroux et al. (2012)].

Since the system is influenced by stochastic disturbances we modeled the system by an ARIMAX description:

$$A(q^{-1})y(t) = B(q^{-1})u(t) + \frac{C(q^{-1})}{1 - q^{-1}}\varepsilon(t) \qquad (4)$$

q^{-1} is the backward shift operator, $y(t)$ is the output, $u(t)$ is the input as previously defined. $\varepsilon(t)$ is $N_{iid}(0, \zeta^2)$. A and B are:

$$A(q^{-1}) = 1 + a_1 q^{-1} + a_2 q^{-2} \qquad (5a)$$
$$B(q^{-1}) = b_1 q^{-1} + b_2 q^{-2} \qquad (5b)$$

C is defined as:

$$C(q^{-1}) = 1 - \alpha q^{-1} \qquad (6)$$

α is a tuning parameter. The model in (6) is obtained by considering the noise term, $\varepsilon(t)$ as a sum of white noise and a drift term. In this way the MPC control will ensure off set free tracking. In this paper we used $\alpha = 0.99$. For more details on the choice of α see [Huusom et al. (2010)].

We can realize (4) as a discrete state-space model on innovation form:

$$x_{k+1} = Ax_k + Bu_k + K\varepsilon_k \qquad (7)$$
$$y_k = Cx_k + \varepsilon_k \qquad (8)$$

where x_k is the blood glucose level [mmol/L] above basal level (6 mmol/L), u_k is the insulin input in [mU/min] above a basal rate of 14 mU/min , y_k is the observed blood glucose level above basal level, ε_k is the innovation noise:

$$\varepsilon_k = y_k - Cx_{k|k-1} \qquad (9)$$

A, B, C, and K are matrices in canonical form. For the first pilot study they were defined as:

$$A = \begin{bmatrix} -(a_1 - 1) & 1 & 0 \\ -(a_2 - a_1) & 0 & 1 \\ a_2 & 0 & 0 \end{bmatrix} B = \begin{bmatrix} b_1 \\ b_2 - b_1 \\ -b_2 \end{bmatrix}$$
$$K = \begin{bmatrix} \alpha - (a_1 - 1) \\ -(a_2 - a_1) \\ a_2 \end{bmatrix} C = [1\ 0\ 0] \qquad (10)$$

As seen in Fig. 2 and stated in the section 1 this controller design was too sensitive to noise which resulted in a unrealistic prediction by the controller. For this reason we reformulated the model into a third order model in which
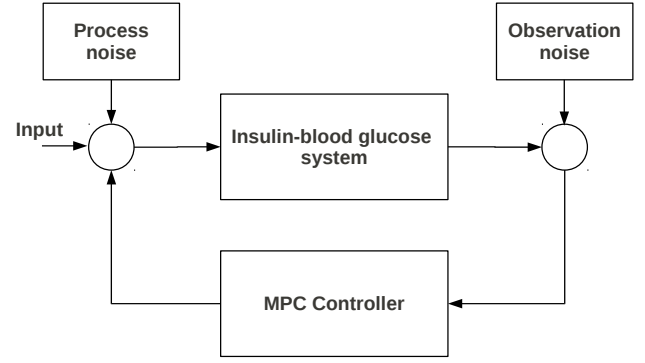

Fig. 3. An illustration of the closed loop system including process and observation noise.

the C polynomial in (6) was changed into a third order polynomial:

$$C(q^{-1}) = 1 + c_1 q^{-1} + c_2 q^{-2} + c_3 q^{-3} \qquad (11)$$

This change resulted in the following Kalman gain, K, in (10):

$$K = \begin{bmatrix} c_1 - (a_1 - 1) \\ c_2 - (a_2 - a_1) \\ c_3 + a_2 \end{bmatrix} \qquad (12)$$

(the matrices A, B, and C remained unchanged). The coefficients in (11) correspond to the coefficients in:

$$\chi(z) = z^3 + c_1 z^2 + c_2 z + c_3 \qquad (13)$$

which is the characteristic polynomial of $(A - KC)$. By choosing the optimal eigenvalues of $(A - KC)$ i.e. optimal roots of (13) we obtain the best balance between a low sensitivity to noise and a fast decay of the estimation error. If the eigenvalues are close to 1 the controller is less sensitive to noise (and relying more on the global trend) but results in a slow decay of the estimation error. If the eigenvalues are close to 0 the decay of the estimation error is fast but on the other hand the controller is very sensitive to noise and thereby rely too much on the local trend as we saw in the first pilot study.

The optimal placement of the eigenvalues inside the unit circle depends on the amount of noise present in the system. Therefore we decided to use the data from the first pilot study to determine the optimal eigenvalues of $(A - KC)$. We kept $\alpha = 0.99$ as one of the eigenvalues and estimated the remaining two eigenvalues from an stochastic differential model as described in the following section.

3.2 Stochastic differential equations model

Earlier work with stochastic differential equations has mainly dealt with modeling of dynamical systems [Kristensen et al. (2005); Overgaard et al. (2007); Tornøe et al. (2004)]. In this paper we introduce a novel approach for parameter estimation in control algorithms based on stochastic differential equations.

Within the field of biomedical control, stochastic differential equations are highly relevant. Dealing with these often very complex systems, as in this case the insulin-blood glucose system, we have to take into account all the uncertainties not explained by the model in the controller. Especially, when we want to use a very simple model as the one in (1). Uncertainties can be present in the input to the system or occur due to other factors interfering with the system e.g. release of stress hormones or physical exercise.

In a stochastic differential model noise is entering the system in two separate entrances: a process noise and a observation noise as seen in Fig. 3. Using this method we could estimate the process and observation variances separately and thereby obtain a more realistic noise model. From here we could find the optimal roots of the characteristic polynomial in (13).

To estimate the process and observation variances we used CTSM (Continuous Time Stochastic Modelling), a freeware program available on the web. See [DTU Informatics, Technical University of Denmark (2012)] for more information. The control algorithm was first transformed from the continuous time transfer function in (1) into an ordinary differential equation by inverse Laplace:

$$y + 2\tau \frac{dy}{dt} + \tau^2 \frac{d^2 y}{dt^2} = Mu(t) \tag{14}$$

and from here to the stochastic state space model which in the general linear case can be written as:

$$dx_t = (A(\theta)x_t + B(\theta)u_t)dt + \sigma(\theta)d\omega_t \tag{15}$$
$$y_k = C(\theta)x_k + e_k \tag{16}$$

x_t is the state vector, x_k is the discrete state vector, u_t is the input vector, y_k is the output vector in discrete time, θ is a vector of parameters, $A(\cdot)$, $B(\cdot)$, $\sigma(\cdot)$, and $C(\cdot)$ are nonlinear functions, $\{\omega_t\}$ is a standard Wiener process representing the process noise and e_k is a white noise process representing the observation noise for which we assume that $e_k \in N(0, S)$.

In this case, (15) is:

$$\begin{bmatrix} dx_{1t} \\ dx_{2t} \end{bmatrix} = (\begin{bmatrix} 0 & 1 \\ -\frac{1}{\tau^2} & -\frac{2}{\tau} \end{bmatrix} \begin{bmatrix} x_{1t} \\ x_{2t} \end{bmatrix} + \begin{bmatrix} 0 & \frac{1}{\tau_G} \\ \frac{M}{\tau^2} & 0 \end{bmatrix} \begin{bmatrix} u_{1t} \\ u_{2t} \end{bmatrix})dt + \begin{bmatrix} \sigma & 0 \\ 0 & \sigma \end{bmatrix} d\omega_t \tag{17}$$

and the discrete observation equation is:

$$y_k = [1 \ 0] \begin{bmatrix} x_{1k} \\ x_{2k} \end{bmatrix} + e_k \tag{18}$$

x_{1t} is the glucose level [mmol/L] above basal level, x_{2t} is the derivative of the glucose level [mmol/L/min], u_{1t} is the insulin input from the pump above the basal rate, u_{2t} is the intravenously glucose input administered by the clinician in case of severe hypoglycemia [mmol/L]. The parameters are τ, M, τ_G, σ, and S. We decided to fix all parameters except from σ and S because all of them are subject specific and known by the subjects themselves or

the clinician. The fixed parameter values used for variance estimation were $\tau = 300$ min, M = -4.077 min·mmol/L/U and $\tau_G = 1$ min. See [Boiroux et al. (2012)] for more details about the choice of these values.

3.3 Estimation of variance parameters

Parameter estimation was based on the Maximum likelihood criteria, see [Kristensen and Madsen (2003); Kristensen et al. (2004)]. Since the subject wore two CGMs we had two stochastically independent data sets. The likelihood function is the joint probability density of all the observations assuming that the parameters are known:

$$\mathcal{L}(\theta; Y) = p(Y|\theta) \tag{19}$$

where $Y = \begin{bmatrix} \gamma^1_{N_1}, \gamma^2_{N_2} \end{bmatrix}$. θ is now the parameter vector only including the parameters we wish to estimate.

We assume that the densities are Gaussian and implicitly described by the mean and variance. Thus the likelihood function can be expressed as:

$$\mathcal{L}(\theta; Y) = \prod_{i=1}^{L} \left(\prod_{k=1}^{N_i} \frac{\exp\left(-\frac{1}{2}(\epsilon_k^i)^T (R_{k|k-1}^i)^{-1}\epsilon_k^i\right)}{\sqrt{\det R_{k|k-1}^i}(\sqrt{2\pi})} \right)$$
$$\cdot p(y_0^i|\theta) \tag{20}$$

L is the number of datasets, in this case two, y_0 are initial values of Y and:

$$\hat{y}_{k|k-1} = E\{y_k|\gamma_{k-1}, \theta\} \tag{21}$$
$$R_{k|k-1} = Var\{y_k|\gamma_{k-1}, \theta\} \tag{22}$$

are the one-step ahead prediction and variance, respectively. The one-step prediction error is:

$$\epsilon_k = y_k - \hat{y}_{k|k-1} \tag{23}$$

The exact solution is computationally infeasible so an approximative method is used. For a given set of parameters and initial states, x_0, the one-step ahead prediction error and variance are estimated from a continuous-discrete Kalman filter. The output prediction equations are computed as:

$$\hat{y}_{k|k-1} = C\hat{x}_{k|k-1} \tag{24}$$
$$R_{k|k-1} = CP_{k|k-1}C^T + S \tag{25}$$

$P_{k|k-1}$ is the predicted state covariance. The innovation equation is defined as the prediction error in (23). And the Kalman gain is determined by:

$$K_k = P_{k|k-1}C^T R_{k|k-1}^{-1} \tag{26}$$

From here we can update the state and state covariance equations:

Fig. 4. The original data sets were separated into three different data sets with missing observations (NA) replacing those observations present in the other two new data sets.

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k \epsilon_k \tag{27}$$

$$P_{k|k} = P_{k|k-1} - K_k R_{k|k-1} K_k^T \tag{28}$$

And finally, we can predict the state and covariance by solving the differential equations:

$$\frac{dx_{t|k}}{dt} = A\hat{x}_{t|k} + Bu_t, t \in [t_k, t_{k+1}[\tag{29}$$

$$\frac{dP_{t|k}}{dt} = AP_{t|k} + P_{t|k}A^T + \sigma\sigma^T, t \in [t_k, t_{k+1}[\tag{30}$$

Once the approximative likelihood function has been computed, the optimal parameter estimates are found by minimizing the log-likelihood function:

$$\hat{\theta} = \arg\min_{\theta \in \Theta} \{-ln(L(\theta; Y|y_0))\} \tag{31}$$

This parameter estimation is per default based on the one-step ahead prediction as seen above. In our case the step size is 5 minutes corresponding to the time between two CGM observations. Since the controller regulates the insulin infusion every 15 minutes, we found it beneficial to base the parameter estimation on the three-step ahead prediction instead.

To handle this, we did a reconstruction of the original data file used for the parameter estimation in CTSM. From each original data set we constructed three new data sets starting in t_0, t_1, and t_2, respectively. As illustrated in Fig. 4, we replaced the two following observations with missing observations and repeated this pattern throughout the new data sets.

In the Kalman filter a missing observation is handled by setting the Kalman gain, K_k, to zero and thus the state predictions are computed from the state equations only until a new observation is available.
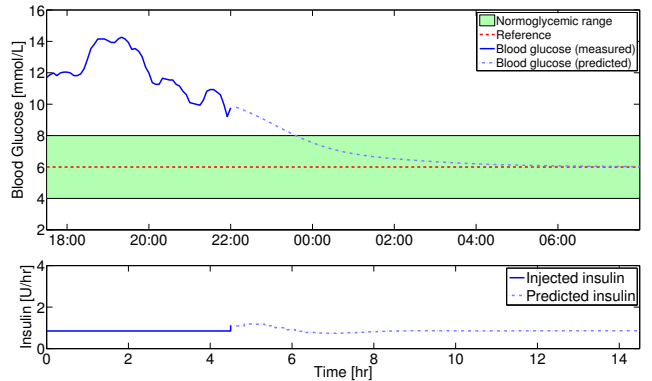

Fig. 5. Top: CGM observations (solid line) and blood glucose prediction (dashed line) from the first pilot study if the tuned Kalman gain had been used in the MPC controller. The normal range is indicated as the green area. Bottom: Insulin delivered by the pump (solid line) and suggested insulin supply (dashed line).

3.4 Tuning of control algorithm

The variance parameters, i.e. the process noise, σ, and the observation noise, S, were estimated to 0.0037 ± 0.00043 and 0.30 ± 0.030 respectively. We used the estimated process and observation variances to compute the Kalman gain, K, by a pole placement method as mentioned in section 3.1. First the continuous state space model in (15) was discretized and from here the two remaining eigenvalues could be computed by solving a Riccati equation. Together with α we computed the roots and thereby the coefficients in the characteristic polynomial in (13) and finally K in (12). The estimated values of the coefficients and K are:

$$c_1 = -2.61 \quad c_2 = 2.28 \quad c_3 = -0.67 \tag{32}$$

$$K = [0.36 \quad -0.66 \quad 0.30] \tag{33}$$

4. RESULTS

We simulated the first pilot study again, this time with the improved Kalman gain, K, in the MPC algorithm. As seen in Fig. 2 the controller originally overestimated the blood glucose prediction at the time of closing the loop. In Fig. 5 the same situation is shown, this time with the new K estimated on the basis of observed data. The controller now relies more on the global trend and the predictions are more realistic. The suggested bolus is likewise smaller and would not cause a hypoglycemic event as the bolus given in the first pilot study.

A second pilot study was performed on the same subject to validate the new algorithm before initiating a larger clinical study. This time the study was performed during day time but without any normal day time disturbances (e.g. meals and exercise) to mimic night time. The result is shown in Fig. 6. The subject had a very low blood glucose level from the beginning of the pilot study. Intravenously glucose was administrated at 10:00 and 12:00 to compensate.
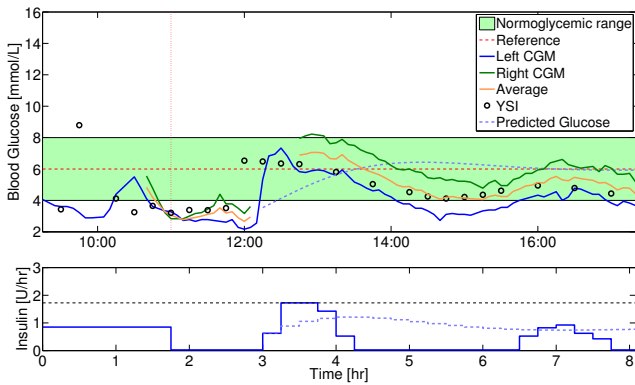
Fig. 6. Top: Results of the second pilot study after the controller had been tuned. The loop is closed at 11:00. The left CGM (solid blue curve) serves as input to the controller. The blood glucose prediction (dashed line) at 12:10 is also shown. YSI observations are blood glucose observations considered as gold standard. Bottom: Insulin delivered by the pump (solid line) and predicted insulin supply (dashed line).

Note that the insulin delivery is updated every 15 minutes as new CGM measurements are available. Thus the insulin delivery profile differs from the predicted profile.

The blood glucose prediction from the improved control algorithm at 12:10 is also seen in Fig. 6. At this point in time the CGM measurements are rapidly increasing due to the glucose administrated intravenously at 12:00. As seen, the controller predicts a slower and more realistic increase in blood glucose due to the improved performance. The rapid increase is an artificial situation and does not occur in daily life. The controller is thereby acting as expected and characterizes the local increase as noise.

Even though tuning has improved the prediction performance, the left CGM measurements decrease below the normal range after 14:00. This could indicate that despite the improvements, the intravenously glucose administration still causes the controller to overestimate the needed bolus size.

In general, the second pilot showed that the controller is able to keep the blood glucose level in the normal range after 12:00 despite the prior disturbances. However, since the second pilot study further improvements have been implemented to increase the robustness.

5. CONCLUSION

This work illustrates how MPC can benefit from stochastic differential equations modeling. By estimating the process and observation noise in two separate terms the Kalman gain was calibrated to the actual system with a rather simple and effective method using stochastic differential equations. Additionally, we based the parameter estimation on the three-step ahead prediction instead of the one-step ahead prediction which was in accordance with the time between controller inputs. It is important to note that this method is not only restricted to this specific application but can be used within many areas dealing with control of complex and stochastic systems.

Currently, a clinical study testing our artificial pancreas controller is ongoing including several subjects and different scenarios.

REFERENCES

American Diabetes Association (2012). Tight diabetes control. URL www.diabetes.org/living-with-diabetes/ treatment-and-care/blood-glucose-control/ tight-diabetes-control.html. Accessed February 27th 2012.

Boiroux, D., Schmidt, S., Duun-Henriksen, A., Frøssing, L., Nørgaard, K., Madsbad, S., Skyggebjerg, O., Poulsen, N., Madsen, H., and Jørgensen, J. (2012). Control of blood glucose for people with type 1 diabetes: an in vivo study. In *Proceedings of the 17th Nordic Process Control Workshop*.

Cobelli, C., Renard, E., and Kovatchev, B. (2011). Artificial pancreas: Past, present, future. *Diabetes*, 60(11), 2672–2682.

DTU Informatics, Technical University of Denmark (2012). Ctsm - continuous time stochastic modelling. URL www2.imm.dtu.dk/ ctsm/. Accessed February 22nd 2012.

Hovorka, R. (2006). Continuous glucose monitoring and closed-loop systems. *Diabetic medicine*, 23(1), 1–12.

Hovorka, R. (2011). Closed-loop insulin delivery: from bench to clinical practice. *Nature Reviews Endocrinology*, 7(7), 385–395.

Huusom, J.K., Poulsen, N.K., Jørgensen, S.B., and Jørgensen, J.B. (2010). Tuning of methods for offset free MPC based on ARX model representations. In *2010 American Control Conference (ACC)*, 2355–2360. Baltimore, MD, USA.

Keenan, D., Grosman, B., Clark, H., Roy, A., Weinzimer, S., Shah, R., and Mastrototaro, J. (2011). Continuous glucose monitoring considerations for the development of a closed-loop artificial pancreas system. *Journal of diabetes science and technology*, 5(6), 1327.

Kristensen, N. and Madsen, H. (2003). *Continuous time stochastic modelling, CTSM 2.3, Mathematics guide*.

Kristensen, N., Madsen, H., and Ingwersen, S. (2005). Using stochastic differential equations for pk/pd model development. *Journal of pharmacokinetics and pharmacodynamics*, 32(1), 109–141.

Kristensen, N., Madsen, H., and Jørgensen, S. (2004). Parameter estimation in stochastic grey-box models. *Automatica*, 40(2), 225–237.

Overgaard, R., Holford, N., Rytved, K., and Madsen, H. (2007). Pkpd model of interleukin-21 effects on thermoregulation in monkeys—application and evaluation of stochastic differential equations. *Pharmaceutical research*, 24(2), 298–309.

Tornøe, C., Jacobsen, J., and Madsen, H. (2004). Grey-box pharmacokinetic/pharmacodynamic modelling of a euglycaemic clamp study. *Journal of mathematical biology*, 48(6), 591–604.