High-performance Uncertainty Quantification in Large-scale Virtual Clinical Trials of Closed-loop Diabetes Treatment

Asbjørn Thode Reenberg, Tobias K. S. Ritschel, Bernd Dammann, John Bagterp Jørgensen

Abstract— In this paper, we propose a virtual clinical trial for assessing the performance and identifying risks in closed-loop diabetes treatments. Virtual clinical trials enable fast and risk-free tests of many treatment variations for large populations of fictive patients (represented by mathematical models). We use closed-loop Monte Carlo simulation, implemented in high-performance software and hardware, to quantify the uncertainty in treatment performance as well as to compare the performance in different scenarios or of different closed-loop treatments. Our software can be used for testing a wide variety of control strategies ranging from heuristical approaches to nonlinear model predictive control. We present an example of a virtual clinical trial with one million patients over 52 weeks, and we use high-performance software and hardware to conduct the virtual trial in 1 h and 22 min.

I. INTRODUCTION

Clinical trials of medical treatments are crucial to ensuring a high level of safety and efficacy. However, they are also very expensive and time-consuming. Therefore, it is important to assess the treatment performance, identify potential risks, and rigorously compare with state-of-the-art prior to the trials. *Virtual* clinical trials are used for exactly this purpose. In a virtual clinical trial, each patient is represented by a mathematical model and the clinical trial is simulated using high-performance software and hardware. The simulation is carried out for a large population of virtual patients, many different scenarios, and several variations of the treatment. This allows for thorough and *fast* testing of a large variety of different treatment designs.

In this paper, we specifically consider the treatment of type 1 diabetes (T1D). One in eleven adults suffer from diabetes (both types), and in 2019, 10% of the global health expenditure (USD 760 billion) was spent on diabetes [1]. Due to autoimmune destruction of β-cells, people with T1D are unable to produce insulin. Consequently, life-long treatment involving daily injections of insulin is necessary to avoid elevated blood glucose (BG) levels (hyperglycemia), which can lead to several complications and chronic conditions [2]. The BG concentration must be measured in order to determine an appropriate insulin dose. While too little insulin results in hyperglycemia, too much insulin results in hypoglycemia (low BG levels), which can be fatal in very severe cases.

As monitoring the BG and determining the appropriate insulin dose is laborious, there is significant interest in auto-mated closed-loop diabetes treatment systems. Such systems are referred to as artificial pancreases (APs), and they consist of 1) a continuous glucose monitor (CGM), 2) a control algorithm that determines the insulin dose, and 3) a pump which delivers the insulin to the patient (it is possible to use other hormones, e.g., glucagon or amylin, in addition to insulin). Many AP algorithms have been proposed and the majority is based on heuristics [3], proportional-integral-derivative (PID) control [4], fuzzy logic [5], linear model predictive control (MPC) [6]–[8], or nonlinear MPC [9], [10].

All AP algorithms contain hyperparameters, e.g., the gains in a PID controller, and performance assessment is essential to choosing suitable values for these. There exist both process-independent performance measures for this purpose, e.g., setpoint deviation and variance, and process-specific measures [11]–[13]. For diabetes treatment, including closed-loop systems, time-in-range (TIR), Hb1Ac values, and the probability of severe hypoglycemia [14]–[18] are commonly used (process-specific) performance measures. The performance of an AP can vary significantly between T1D patients due to differences in physiology (e.g., in pharmacodynamics and pharmacokinetics). Therefore, it is necessary to evaluate the performance measures for a large population of patients in order to accurately estimate the uncertainty. However, due to computational limitations of standard software, it is common to evaluate the performance using only a small number of patients and over a short time span (days or a few weeks).

In this work, we describe an approach for high-performance uncertainty quantification of the performance of AP algorithms in large-scale long-term virtual clinical trials. The approach involves mathematical models based on stochastic differential equations (SDEs), and we use closed-loop Monte Carlo simulation to quantify the performance uncertainty. Furthermore, we propose multiple ways of 1) visualizing the uncertainty in the performance measures and 2) comparing the performance for different scenarios or AP algorithms. We implement the Monte Carlo simulation and the AP in parallelized high-performance C code, and the computations are carried out on a high-performance computing (HPC) cluster. Finally, we present a numerical example of a virtual clinical trial with one million patients over 52 weeks which can be carried out in 1 h and 22 min.

The remaining part of this paper is organized as follows. In Section II, we describe the virtual clinical trial, and in Section III, we present the approach for uncertainty quantification of AP algorithms. Section IV contains the numerical examples, and conclusions are presented in Section V.

II. VIRTUAL CLINICAL TRIAL

The virtual clinical trial consists of 1) a population of patients, 2) a protocol containing the trial activities (size and duration of meals, intensity and duration of exercise, etc.), 3) one or more mathematical models of the patients, 4) values of the model parameters, and 5) one or more APs (i.e., control algorithms). Furthermore, it is possible to include both incorrectly announced and unannounced meals and exercise which are some of the key challenges that an AP should be able to address. Finally, the virtual clinical trial allows for *stochastic* mathematical models which can represent unmodeled physiological phenomena, uncertain model parameters, and uncertainty related to meals and exercise, as well as noisy sensor measurements.

A. Patients

The virtual clinical trial contains one million fictive patients. Each patient is represented by the same information that would be available for a real patient. Specifically, each fictive patient is associated with a unique ID and a set of attributes including first and last name, date and place of birth, sex, height, body weight, and resting heart rate. The height, the body weight and the resting heart rate are sampled from normal distributions, and the date of birth is sampled from a uniform distribution.

B. Protocols

A protocol consists of a sequence of model *disturbances*, i.e., uncontrolled inputs to the patient model. Common disturbances are meals and exercise. Each protocol has an ID and for each disturbance, it contains the disturbance type and size as well as time stamps indicating the beginning and end of the disturbance.

Next, to illustrate the concept, we describe a protocol designed to mimic a Northern European lifestyle in terms of meal times, seasons, work weeks, and the number of vacation weeks and public holidays. Furthermore, the protocol involves a high-carb diet (in particular during winter and autumn), which is challenging for APs. We divide the year into 4 seasons each consisting of 13 weeks, and we assume 6 weeks of vacation and 10 public holidays, represented as an additional 2 weeks of vacation. Each season is a different combination of three basis weeks; a *standard week*, an *active week*, and a *vacation week*. Furthermore, each basis week is a different combination of four basis days; a *standard day*, an *active day*, a day with a *movie night*, and a day with a *late night*. Table I shows the compositions of the seasons and the weeks.

The patients are less active and eat more during vacation weeks, they also eat more during winter and autumn, and active weeks contain more active days. Compared to the standard day, 1) the active day has an exercise session, 2) the movie night has an additional snack in the evening, and 3) the late night has two additional snacks in the evening. Fig. 1 shows schematics of the basis days, and Table II shows the meal sizes which depend on the body weight.

TABLE I
COMPOSITIONS OF THE SEASONS AND THE WEEKS

Compositions of the seasons

Season	Standard week	Active week	Vacation week
Winter	6	4	3
Spring	6	6	1
Summer	7	3	3
Autumn	9	3	1

Compositions of the weeks

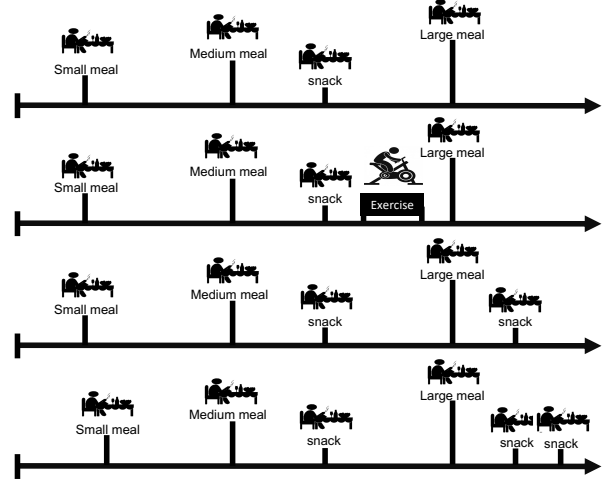| Week type	Standard day	Active day	Movie night	Late night
Standard	4	1	1	1
Active	3	3	1	0
Vacation	5	0	0	2

Fig. 1. Overview of the basis days in winter and autumn. From the top, we show 1) the *standard day*, 2) the *active day*, 3) the *movie night*, and 4) the *late night*. During the summer and spring, the dinner is a *medium meal* and the snack is before lunch.

C. Database

A key component of the virtual clinical trial is a database containing the fictive patients, protocols, model parameters, and simulation results. We use the open-source database system PostgreSQL. Using a database makes it straightforward to share the virtual clinical trial and compare performance results for the exact same patients and protocols. It also allows the user to carry out the clinical trial for a specific demographic, e.g., people with a certain body weight or sex. Furthermore, the database includes several basis days and weeks which the user can combine to form new protocols. Finally, the database can be extended with a graphical user interface in order to 1) visualize the AP performance and characteristics of the patients and protocols and 2) add

TABLE II
WEIGHT-DEPENDENT MEAL SIZES

Meal size	Amount of carbohydrates	For a 70 kg patient
Large meal	1.29 g CHO/kg	90 g CHO
Medium meal	0.86 g CHO/kg	60 g CHO
Small meal	0.57 g CHO/kg	40 g CHO
Snack	0.29 g CHO/kg	20 g CHO

new elements such as patients, protocols, and mathematical models to the database.

III. CLOSED-LOOP MONTE CARLO SIMULATION

In Monte Carlo simulation, the uncertainty of a quantity of interest (e.g., the TIR for a T1D patient) is estimated by simulating the system (i.e., the clinical trial) with different values of the uncertain quantities (e.g., the model parameters). In this work, we extend the high-performance toolbox for closed-loop Monte Carlo simulation developed by Wahlgreen et al. [19] with a low-memory implementation which circumvents the high storage requirements associated with large numbers of long-term simulations.

A. Mathematical model

The virtual clinical trial can be used for mathematical models of patient physiology in the general stochastic form

$$x(t_0) = x_0, \tag{1a}$$
$$dx(t) = f(t, x(t), u(t), d(t), p)dt$$
$$\qquad + \sigma(t, x(t), u(t), d(t), p)dw(t), \tag{1b}$$
$$z(t) = h(t, x(t), p) \tag{1c}$$
$$y(t_k) = g(t_k, x(t_k), p) + v(t_k). \tag{1d}$$

Here, t is time, and the virtual clinical trial starts at time t_0. The states, x, represent the physiological state of the patient, e.g., the BG concentration and amount of insulin in the body, and x_0 are the initial states. The manipulated inputs, u, are the quantities computed by the AP, e.g., the insulin flow rate. The disturbance variables, d, represent the uncontrolled inputs, e.g., meals and exercise, and p are model parameters. The first term in (1b) is the deterministic drift term and the second term is the stochastic diffusion term.

The AP receives measurements of the observed variables, y, obtained from the measurement function, g, at discrete points in time, t_k, and the outputs, z, obtained from the output function, h, are the quantities relevant to the control objective of the AP. The standard Wiener process w is used to represent uncertainty. Its increment is distributed as $dw(t) \sim N_{iid}(0, Idt)$, and the measurement noise is assumed to be normally distributed: $v(t_k) \sim N_{iid}(0, R(t_k))$. Furthermore, the inputs are assumed to be piecewise constant between sampling times:

$$u(t) = u_k, \qquad t \in [t_k, t_{k+1}[, \tag{2a}$$
$$d(t) = d_k, \qquad t \in [t_k, t_{k+1}[. \tag{2b}$$

Finally, we stress that the form (1b) also includes *deterministic* dynamical systems, i.e., ordinary differential equations (ODEs), where σ is zero.

B. Control algorithm

At time t_k, 1) the control state, x_k^c, is updated, and 2) the AP (i.e., the closed-loop feedback control strategy) computes values of the manipulated inputs based on the previous control state and the measurements, $y_k = y(t_k)$:

$$x_{k+1}^c = \kappa_k(x_k^c, y_k, \bar{u}_k, \bar{y}_k, \hat{d}_k, p_\kappa), \tag{3a}$$
$$u_k = \lambda_k(x_k^c, y_k, \bar{u}_k, \bar{y}_k, \hat{d}_k, p_\mu). \tag{3b}$$

Here, $\bar{u}_k$ and $\bar{y}_k$ are setpoints, and $\hat{d}_k$ are estimates of the disturbances. This form can represent many types of closed-loop control strategies including heuristic strategies based on physiological insight, PID-based strategies, and MPC (including state estimation). Many different values of the hyperparameters p_μ and p_κ can be tested using the virtual clinical trial.

C. Software and hardware

The closed-loop Monte Carlo simulation, the mathematical models, and the AP are implemented using high-performance C code which we parallelize for shared-memory architectures using OpenMP. Whenever a simulation is completed, we immediately compute its contribution to performance indicators such as TIR as well as mean, minimum, and maximum BG concentration as functions of time. Subsequently, the simulation is only stored if it is worse than previous simulations according to some criterion (e.g., lowest BG concentration reached).

We use two AMD EPYC 7542 32-core processors with a clock speed of 2.9 GHz [20]. As the Monte Carlo simulation is highly parallelizable, the speedup in computational performance increases almost linearly with the number of cores. Consequently, the parallel implementation runs almost 64 times faster than a corresponding sequential implementation. Furthermore, it would be computationally infeasible to carry out large-scale long-term virtual clinical trials using a sequential Matlab implementation [19] or similar.

IV. EXAMPLE OF A VIRTUAL CLINICAL TRIAL

In this section, we present an example of a virtual clinical trial involving a million virtual patients following the example protocol described in Section II-B. We briefly describe the mathematical model of the patients' physiology and the AP used in the trial, and we demonstrate how to visualize the uncertainty in the AP performance. Furthermore, we also show how to compare the performance of the AP in two different scenarios; one where the basal rate is correct (trial A) and another where it is underestimated by 50% (trial B). Finally, the computation time is 1 h and 22 min (with a time step size of 30 s in the simulations). Consequently, several virtual clinical trials can be carried out in a single day.

A. Performance measures

We divide the BG concentration into 5 ranges [14] given in mmol/L. Red: severe hypoglycemia (below 3). Light red: hypoglycemia (3–3.9). Green: normoglycemia (3.9–10). Yellow: hyperglycemia (10–13.9). Orange: severe hyperglycemia (above 13.9). Furthermore, we also consider the distributions of the total daily basal and bolus insulin as well as bolus glucagon.

B. Patient model

We use an extension of the model presented by Hovorka et al. [9] to represent the pharmacokinetic and pharmacodynamical responses of the virtual patients to carbohydrate absorption and subcutaneous infusion of insulin and glucagon.

The model is extended with 1) a one-state model of the measured BG concentration, 2) a two-state pharmacokinetic model of subcutaneous glucagon injection [21], and 3) a three-state model of the effect of exercise on the plasma BG concentration [22].

The model parameters related to the measured BG concentration, glucagon infusion, and exercise have the same values for all virtual patients. We sample the remaining model parameters from the distributions presented by Hovorka et al. [23]. We only use parameter sets where the parameter values are nonnegative, the normally distributed parameters are within one standard deviation from the mean, and the insulin basal rate is at least 0.4 U/h.

C. Artificial pancreas

We demonstrate the capabilities of the virtual clinical trial using a dual-hormone AP which switches between an insulin mode and a glucagon mode. Glucagon is used to mitigate hypoglycemia. The insulin mode involves 1) microadjustments of the basal rate, 2) a meal bolus calculator, 3) superboli, and 4) an insulin-to-carb ratio estimator. In the glucagon mode, only microboli are administered. In both modes, a 100 μg glucagon bolus is administered at the beginning of exercise if the blood glucose concentration is below 7 mmol/L. No insulin is administered in the glucagon mode, and no glucagon, apart from the exercise bolus, is administered in the insulin mode. The AP uses filtered estimates of the glucose concentration obtained with a low-pass filter, and several hyperparameters have different values depending on whether the patient is exercising or not.

Fig. 2 shows week 14 of a simulation for a single patient. The AP administers insulin boli at meal times, the basal rate is turned off for a period after the meals, and a small glucagon bolus is administered during exercise, i.e., it is not the 100 μg bolus at the beginning of exercise. As is evident, the TIR is high for the shown period.

D. Performance of the artificial pancreas

Fig. 3 shows the amount of time spent below different BG concentrations. It allows us to inspect the worst-case patient (lowest BG concentration reached), the population average, and the values reached by at least one patient (the span). The worst-case patient is useful for identifying weaknesses in the AP. Here, it seems unlikely that the basal rate is too high because the patient suffers from severe hyperglycemia almost 40% of the time. This can also be seen from the stacked bar chart in the middle of Fig. 4. The population mean can reveal systemic issues. On average, the patients spend almost all of their time above the target of 6 mmol/L. Perhaps too little insulin is administered. Finally, the shaded area can be used to conclude that, e.g., 1) no patient spends more than 6% of their time below 3.9 mmol/L and 2) nobody is above 13.9 mmol/L more than 40% of the time.

Whereas the two stacked bar charts in the left of Fig. 4 provide an intuitive overview, the box plot on the right gives a comprehensive picture of the TIR for the entire population. The red markers show that only a few patients experience
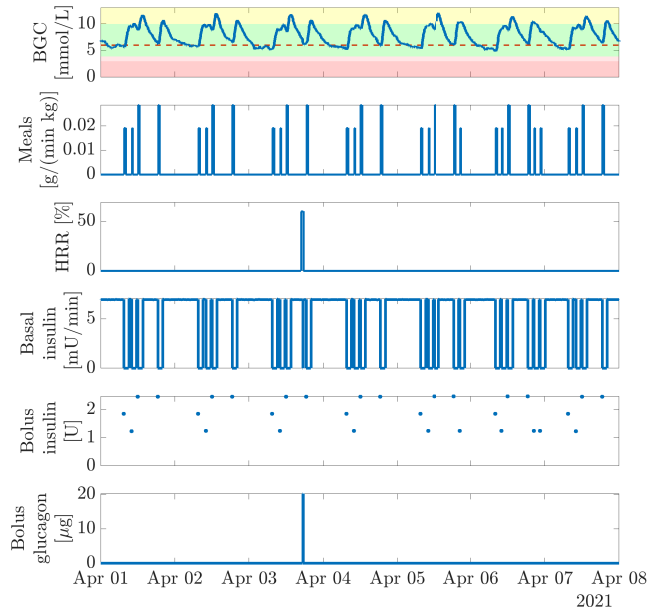

Fig. 2. Week 14 of a simulation for 1 patient. From the top we show: 1) The BG concentration (BGC), 2) the meal carbohydrate content (shown as a rate dependent on the body weight), 3) the resting heart rate reserve (HRR), 4) the administered insulin basal rate, 5) the insulin boli, and 6) the glucagon boli.
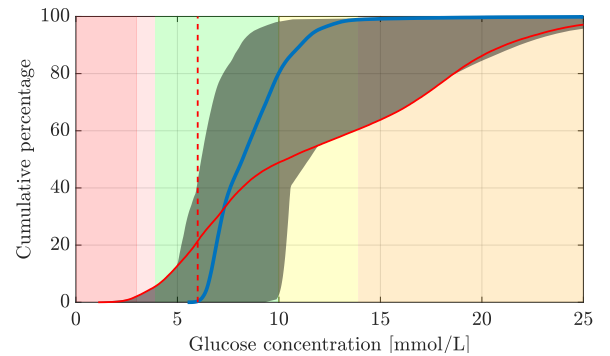

Fig. 3. Cumulative distribution of the BG concentration in a virtual clinical trial with 1 million patients over 52 weeks. Blue solid line: The mean BG concentration. Red solid line: The worst-case patient. Red dashed line: The setpoint. Grey shaded area: The span of all the patients.

hypoglycemia. On average, the patients spend 77% of their time in range, and for most patients, this value is at least 55%. However, a significant part spends between 5% and 40% above 13.9 mmol/L. Finally, Fig. 5 shows distributions of the total insulin and glucagon administered per day. These distributions can be compared with dosage guidelines to see if extreme amounts are administered. For instance, it is positive that for the most part, only small amounts of glucagon are administered here.

E. Comparison of scenarios

When comparing different scenarios or APs, the plots in Fig. 3–5 can be overlaid or combined as shown in Fig. 6–8. It is clear from Fig. 6 that the basal rate is too low in trial B. However, the worst-case patients reach equally low
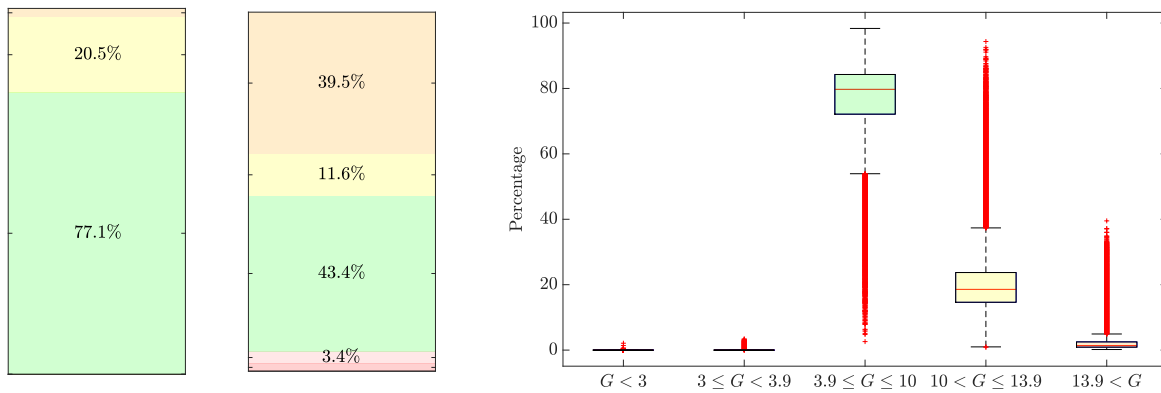
Fig. 4. Distribution of the TIRs for a virtual clinical trial with 1 million patients over 52 weeks. Left: Mean TIRs for all patients. Middle: TIRs for the worst-case patient. Right: Box plot of the TIRs for all patients.
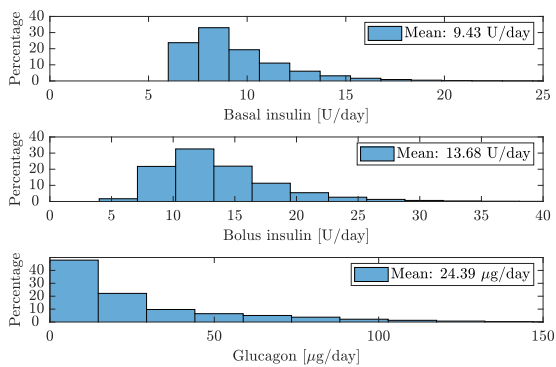

Fig. 5. Distributions of the total daily basal and bolus insulin and bolus glucagon for a virtual clinical trial with 1 million patients over 52 weeks. Top: Basal insulin. Middle: Bolus insulin. Bottom: Bolus glucagon.
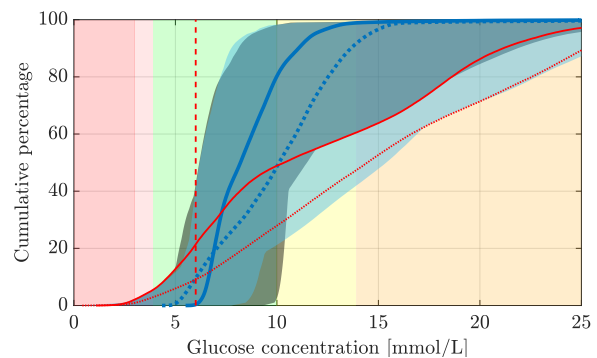

Fig. 6. Cumulative distribution of the BG concentration in two virtual clinical trials of 1 million patients over 52 weeks. Blue solid line: The mean BG concentration of trial A. Blue dotted line: The mean BG concentration of trial B. Red solid line: The patient that reaches the lowest BG concentration in trial A. Red dotted line: The patient that reaches the lowest BG concentration in trial B. Red dashed line: The setpoint. Grey shaded area: The span of all patients in trial A. Light blue shaded area: The span of all patients in trial B.

BG concentrations in both trials. Fig. 7 also clearly shows that far better TIRs are reached with the correct basal rate although the worst-case patient experiences hypoglycemia more often with the correct basal rate. Fig. 8 allows for direct comparison of how much insulin and glucagon that is administered on a daily basis. Obviously, less basal insulin is administered in trial B, and as a result, more bolus insulin is given. Consequently, more glucagon is given in trial B.

V. CONCLUSION

In this paper, we present a virtual clinical trial for assessing the uncertainty in the performance of closed-loop diabetes treatments. We use a high-performance closed-loop Monte Carlo method for quantifying the uncertainty, and we evaluate the performance by examining 1) the distributions of the TIRs for the patient population and 2) the distributions of the total daily doses of basal and bolus insulin as well as bolus glucagon. Furthermore, this approach can be used to compare the performance in different scenarios and for different closed-loop treatments. Finally, we demonstrate that a virtual clinical trial with one million patients over 52 weeks can be completed in 1 h and 22 min by using parallel high-performance software and hardware. The developed software can be used for closed-loop systems with any drug administration device (pump or pen) and measurement device (CGM or self-monitoring of blood glucose (SMBG) device).

REFERENCES

[1] International Diabetes Federation, "IDF diabetes atlas 9th edition," 2019, ISBN: 978-2-930229-87-4.

[2] American Diabetes Association, "Diabetes technology: Standards of medical care in diabetes," *Diabetes Care*, vol. 41, no. 1, 2018.

[3] I. Capel, M. Rigla, G. Garcia-Saez, A. Rodriguez-Herrero, B. Pons, D. Subias, F. Garcia-Garcia, M. Gallach, M. Aguilar, C. Perez-Gandia, E. J. Gomez, A. Caixas, and M. E. Hernando, "Artificial pancreas using a personalized rule-based controller achieves overnight normoglycemia in patients with type 1 diabetes," *Diabetes Technology and Therapeutics*, vol. 16, no. 3, pp. 172–179, 2014.

[4] G. Marchetti, M. Barolo, L. Jovanovic, H. Zisser, and D. E. Seborg, "An improved PID switching control strategy for type 1 diabetes," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 3, pp. 857–865, 2008.

[5] T. Biester, J. Nir, K. Remus, A. Farfel, I. Muller, S. Biester, E. Atlas, K. Dovc, N. Bratina, O. Kordonouri, T. Battelino, M. Philip, T. Danne, and R. Nimri, "DREAM5: An open-label, randomized, cross-over study to evaluate the safety and efficacy of day and night closed-loop control by comparing the MD-Logic automated insulin delivery system to sensor augmented pump therapy in patients with type 1 diabetes at home," *Diabetes, Obesity and Metabolism*, vol. 21, no. 4, pp. 822–828, 2019.
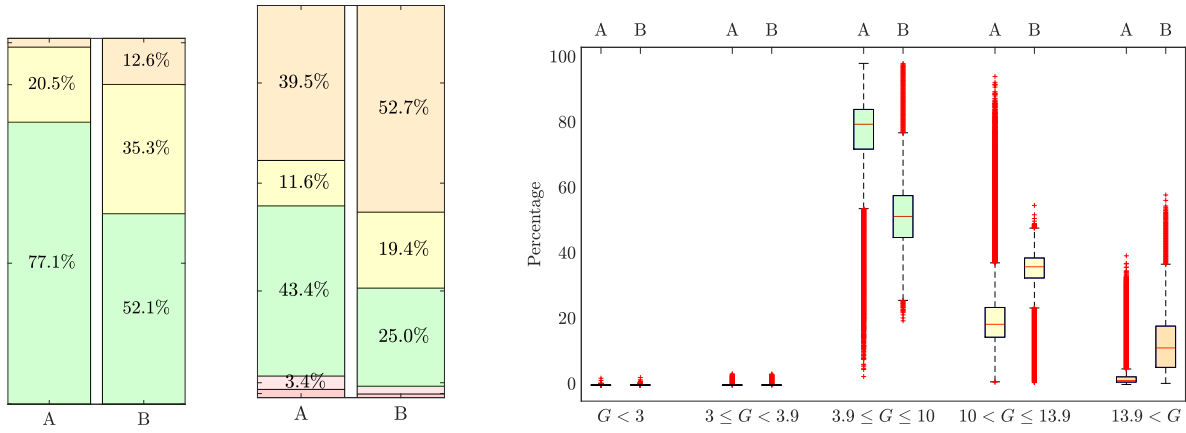
Fig. 7. Distribution of the TIRs for two virtual clinical trials of 1 million patients over 52 weeks. Left: Mean TIRs for all patients in trial A and trial B. Middle: TIRs for the worst-case patients in trial A and trial B. Right: Box plot of the TIRs for all patients in trial A and trial B.
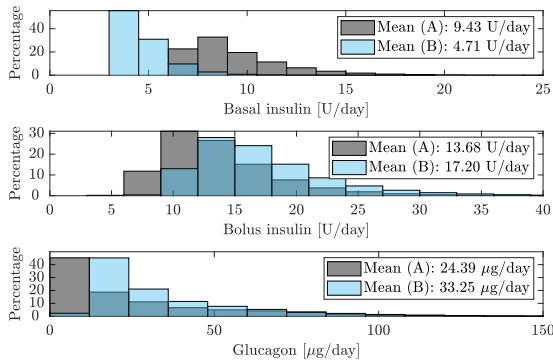

Fig. 8. Distributions of the total daily basal and bolus insulin and bolus glucagon for two virtual clinical trials with 1 million patients over 52 weeks. Top: Basal insulin. Middle: Bolus insulin. Bottom: Bolus glucagon. Grey: Trial A. Light blue: Trial B.

[6] M. Eren-Oruklu, A. Cinar, L. Quinn, and D. Smith, "Adaptive control strategy for regulation of blood glucose levels in patients with type 1 diabetes," *Journal of Process Control*, vol. 19, no. 8, pp. 1333–1346, 2009.

[7] S. Schmidt, D. Boiroux, A. K. Duun-Henriksen, L. Frøssing, O. Skyggebjerg, J. B. Jørgensen, N. K. Poulsen, H. Madsen, S. Madsbad, and K. Nørgaard, "Model-based closed-loop glucose control in type 1 diabetes: The DiaCon experience," *Journal of Diabetes Science and Technology*, vol. 7, no. 5, pp. 1255–1264, 2013.

[8] D. Boiroux, A. K. Duun-Henriksen, S. Schmidt, K. Nørgaard, S. Madsbad, N. K. Poulsen, H. Madsen, and J. B. Jørgensen, "Overnight glucose control in people with type 1 diabetes," *Biomedical Signal Processing and Control*, vol. 39, pp. 503–512, 2018.

[9] R. Hovorka, V. Canonico, L. Chassin, U. Haueter, M. Massi-Benedetti, M. Federici, T. Pieber, H. Schaller, L. Schaupp, T. Vering, and M. Wilinska, "Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes," *Physiological Measurement*, vol. 25, no. 4, pp. 905–920, 2004.

[10] D. Boiroux and J. B. Jørgensen, "Nonlinear model predictive control and artificial pancreas technologies," in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 284–290.

[11] K. J. Åström, "Assessment of achievable performance of simple feedback loops," Department of Automatic Control, Lund Institute of Technology, Tech. Rep. TFRT-7411, 1989.

[12] T. J. Harris, "Assessment of control loop performance," *The Canadian Journal of Chemical Engineering*, vol. 67, pp. 856–861, 1989.

[13] M. Jelali, "An overview of control performance assessment technology and industrial applications," *Control Engineering Practice*, vol. 14, pp. 441–466, 2006.

[14] R. I. G. Holt, J. H. DeVries, A. Hess-Fischl, I. B. Hirsch, M. S. Kirkman, T. Klupa, B. Ludwig, K. Nørgaard, J. Pettus, E. Renard, J. S. Skyler, F. J. Snoek, R. S. Weinstock, and A. L. Peters, "The management of type 1 diabetes in adults. A consensus report by the American Diabetes Association (ADA) and the European Association for the Study of Diabetes (EASD)," *Diabetes Care*, vol. 44, pp. 2589–2625, 2021.

[15] H. Calsbeek, N. A. B. M. Ketelaar, M. J. Faber, M. Wensing, and J. Braspenning, "Performance measurements in diabetes care: The complex task of selecting quality indicators," *International Journal for Quality in Health Care*, vol. 25, no. 6, pp. 704–709, 2013.

[16] R. A. Lal, C. L. Maikawa, D. Lewis, S. W. Baker, A. A. A. Smith, G. A. Roth, E. C. Gale, L. M. Stapleton, J. L. Mann, A. C. Yu, S. Correa, A. K. Grosskopf, C. S. Liong, C. M. Meis, D. Chan, J. P. Garner, D. M. Maahs, B. A. Buckingham, and E. A. Appel, "Full closed loop open-source algorithm performance comparison in pigs with diabetes," *Clinical and Translational Medicine*, vol. 11, p. e387, 2021.

[17] I. Hajizadeh, S. Samadi, M. Sevil, M. Rashid, and A. Cinar, "Performance assessment and modification of an adaptive model predictive control for automated insulin delivery by a multivariable artificial pancreas," *Industrial & Engineering Chemistry Research*, vol. 58, pp. 11 506–11 520, 2019.

[18] M. Sejersen, D. Boiroux, S. E. Engell, T. K. S. Ritschel, A. T. Reenberg, and J. B. Jørgensen, "Initial titration for people with type 1 diabetes using an artificial pancreas," *IFAC PapersOnLine*, vol. 54, no. 15, pp. 484–489, 2021.

[19] M. R. Wahlgreen, A. T. Reenberg, M. K. Nielsen, A. Rydahl, T. K. S. Ritschel, B. Dammann, and J. B. Jørgensen, "A high-performance Monte Carlo simulation toolbox for uncertainty quantification of closed-loop systems," in *2021 IEEE Conference on Decision and Control (CDC)*, 2021, pp. 6755–6761.

[20] DTU Computing Center, "DTU Computing Center resources," 2021, Technical University of Denmark. [Online]. Available: https://doi.org/10.48714/DTU.HPC.0001

[21] S. L. Wendt, J. K. Møller, C. B. Knudsen, H. Madsen, A. Haidar, and J. B. Jørgensen, "PK/PD modelling of glucose-insulin-glucagon dynamics in healthy dogs after a subcutaneous bolus administration of native glucagon or a novel glucagon analogue," Technical University of Denmark, Tech. Rep. 2, 2016.

[22] M. Rashid, S. Samadi, M. Sevil, I. Hajizadeh, P. Kolodziej, N. Hobbs, Z. Maloney, R. Brandt, J. Feng, M. Park, L. Quinn, and A. Cinar, "Simulation software for assessment of nonlinear and adaptive multivariable control algorithms: Glucose–insulin dynamics in type 1 diabetes," *Computers & Chemical Engineering*, vol. 130, p. 106565, 2019.

[23] R. Hovorka, F. Shojaee-Moradie, P. V. Carroll, L. J. Chassin, I. J. Gowrie, N. C. Jackson, R. S. Tudor, A. M. Umpleby, and R. H. Jones, "Partitioning glucose distribution/transport, disposal, and endogenous production during IVGTT," *American Journal of Physiology-Endocrinology and Metabolism*, vol. 282, pp. E992–E1007, 2002.