

Simulating Insulin Infusion Pump Risks by *In-Silico* Modeling of the Insulin-Glucose Regulatory System.

Sriram Sankaranarayanan¹ and Georgios Fainekos². *

1. University of Colorado, Boulder, CO. Email: `first.lastname@colorado.edu`

2. Arizona State University, Tempe, AZ. Email: `fainekos@asu.edu`

Abstract. We present a case study on the use of robustness-guided and statistical model checking approaches for simulating risks due to insulin infusion pump usage by diabetic patients. Insulin infusion pumps allow for a continuous delivery of insulin with varying rates and delivery profiles to help patients self-regulate their blood glucose levels. However, the use of infusion pumps and continuous glucose monitors can pose risks to the patient including chronically elevated blood glucose levels (hyperglycemia) or dangerously low glucose levels (hypoglycemia). In this paper, we use mathematical models of the basic insulin-glucose regulatory system in a diabetic patient, insulin infusion pumps, and the user's interaction with these pumps defined by commonly used insulin infusion strategies for maintaining normal glucose levels. These strategies include common guidelines taught to patients by physicians and certified diabetes educators and have been implemented in commercially available insulin bolus calculators. Furthermore, we model the failures in the devices themselves along with common errors in the usage of the pump. We compose these models together and analyze them using two related techniques: (a) robustness guided state-space search to explore worst-case scenarios and (b) statistical model checking techniques to assess the probabilities of hyper- and hypoglycemia risks. Our technique can be used to identify the worst-case effects of the combination of many different kinds of failures and place high confidence bounds on their probabilities.

1 Introduction

The goal of this paper is to combine physiological models of the insulin-glucose regulatory system in diabetic patients with medical device models of infusion pumps and continuous glucose meters to perform *in silico* risk assessments. Modern treatments for type-1 and 2 diabetes mellitus require frequent, periodic monitoring of blood glucose levels and the subcutaneous delivery of artificial insulin. Developments in medical device technologies have enabled software-controlled insulin infusion pumps that can deliver precise amounts of insulin in user programmable patterns. Likewise, advances in sensor technologies have enabled continuous glucose monitors (CGM) that can be used to sense the concentration of glucose subcutaneously. These technologies have enabled the development of automatic and manual control strategies, vastly improving the ability of patients to achieve normal glycemic control [12,19,37].

* This work was funded by National Science Foundation (NSF) grants under award numbers CPS-1035845, CNS-1016994 and CNS-1017074. All opinions expressed here are those of the authors and not necessarily of the NSF.

However, the use of these devices and the control strategies are prone to hazards arising from device, software and usage errors. A comprehensive list of these hazards has been compiled by Zhang et al. [43]. These risks primarily arise due to factors such as (a) failures in insulin pumps due to software errors, occlusions, and pump failures; (b) calibration and dropout errors in the glucose monitors; and (c) usage errors including discrepancies between planned and actual meals, incorrect insulin-to-carbs ratios, sensitivity factors, and basal insulin levels. These hazards can expose the patient to significant levels of hypoglycemia (low levels of blood glucose) or hyperglycemia (high levels of glucose), each of which leads to dangerous complications including loss of consciousness for hypoglycemia and ketacidosis for significant hyperglycemia. The long term consequences of elevated post-prandial glucose levels include kidney damage (nephropathy) and eye damage (retinopathy). Given the severity of these risks, a careful study of the various kinds of faults involved in the infusion process and the associated risks is of great importance.

In this paper, we create mathematical models of the overall infusion process by modeling the components involved in an infusion. Our model incorporates physiological models of the gluco-regulatory system [22,11,7,36,9,27,29], models of the various devices involved, the user's infusion strategy [34] and some of the possible faults that can arise during the infusion process [43]. We specify metric temporal logic (MTL) properties [25] for the executions of this model which include (a) absence of hypoglycemia ($G(t) \geq 3$ mmol/L), (b) absence of significant hyperglycemia, ($G(t) \leq 20$ mmol/L), and (c) settling of the blood glucose level to a normal range 3 hours after a meal ($\forall t, t \geq 150 \Rightarrow G(t) \in [4, 10]$ mmol/L). Unfortunately, the resulting models are nonlinear and include discrete mode switches due to the user and device models. Existing symbolic verification tools are inadequate for exhaustively exploring the behaviors of the overall model. Therefore, we adapt two recent approaches based on simulations:

Robustness-Guided Model Checking: We use robustness-guided sampling, assuming that faults are non-deterministic [28,4] to explore the possible worst case scenarios involving a combination of faults. Robustness guided sampling is based on the idea of providing real-valued robustness semantics to formulas in metric temporal logic [30,16,17]. The robustness of a trace w.r.t a given specification can be used as an objective function for a global stochastic optimization approach that seeks to minimize the robustness to falsify a given temporal property. This is a suitable approach in cases where the model is infinite state and non-linear. Such models are generally not amenable to existing symbolic verification techniques. In this paper, we employ this scheme to search for combinations of faults that can cause severe hyperglycemia, hypoglycemia and delayed return to normal glycemc levels following a meal.

Statistical Model Checking: We use statistical model checking by associating probabilities with faults to quantify the risk of hyper- and hypoglycemia with some confidence interval bounds [42,10,44]. Statistical model checking (SMC) repeatedly simulates a stochastic system while evaluating probabilistic temporal logic queries with high confidence. SMC approaches allows us to place bounds on the probability that a formula holds for a given stochastic system.

The models described in this paper are integrate inside the Matlab Simulink/Stateflow(tm) modeling environment. We use our tool S-Taliro which incorporates robustness-guided

state-space exploration using many different global optimization engines including Monte-Carlo search [28], Ant-colony optimization [3], genetic algorithms and cross-entropy sampling [32]¹. Recently, we have extended S-Taliro to support Bayesian Statistical Model Checking (SMC) [23,44]. Using S-Taliro, we examine numerous fault scenarios involving a combination of faults to analyze the worst-case scenarios arising from these situations and to quantify the risk of hyper or hypoglycemia, assuming some prior probabilities for the various faults.

To our knowledge, the use of statistical and robustness-guided model checking to analyze infusion risks in the insulin infusion pump setting is novel. Previously, there have been attempts at quantifying risks involved in model-predictive controllers (MPC) for overnight glycemic control using a simulation environment [40,20]. These simulations derive a risk score for hypoglycemia risk using numerous simulations. However, no confidence intervals are derived for the risk scores. The use of SMC in this paper provides a more systematic and potentially less computationally expensive approach. Jha et al. employ statistical model checking to discover parameters for a PI controller for managing insulin infusion pumps using the Bergman minimal model [24].

Assessing risks in infusion pumps has received increasing attention recently [43,5,39,38]. Our previous work considered the effects of infusion risks in a hospital setting using drug infusion pumps and linear pharmacokinetic models [33]. Therein, we were able to employ bounded-model checking techniques for linear hybrid automata to drive the worst-case search. Currently, the state-of-the-art in symbolic verification techniques are inapplicable to our model which involves non-linear dynamics with switching.

2 Overview

We provide a brief overview of the problem and the proposed solutions.

Insulin Infusion Scenario: Consider a commonly occurring scenario of planning a meal for a patient suffering from type-1 (insulin dependent) diabetes. The patient uses an insulin infusion pump to deliver an appropriate bolus dosage of insulin ² before the meal commences. The planning process requires the patient to decide on the following parameters:

1. The insulin bolus amount to be infused through an insulin infusion pump,
2. The timing of the bolus relative to the planned meal time,
3. The width of the bolus,
4. The timings and amounts of any planned corrective dosages to accommodate higher than normal post-prandial blood glucose levels.

Typically, patients suffering from diabetes undergo training by physicians, certified diabetes educators and numerous books on the topic to arrive at suitable strategies for planning meal infusions [34]. A typical calculation that is often automated by an insulin bolus calculator involves the steps detailed below using the planned meal data:

¹ S-Taliro can be downloaded for free from <https://sites.google.com/a/asu.edu/s-taliro/>.

² A bolus dosage is a fixed amount of a drug that is delivered over a relatively short period of time to achieve a near-term effect.

- Divide the amount of carbohydrates in the meal by a personal *insulin-to-carbs ratio* to obtain an appropriate bolus,
- Decide on the timing, shape and width of the bolus based on the Glycemic Index (GI) of the planned meal and the blood glucose reading measured prior to the meal,
- Decide on a correction bolus a few hours after the meal by measuring blood glucose levels and dividing it by a personal *insulin sensitivity factor*.

There are many rules of thumb for deciding upon an appropriate insulin-to-carbs ratio or an insulin sensitivity factor. Often, the patients are required to carefully monitor and adjust these ratios until they can achieve good glycemic control. However, there are numerous risks involved in a typical infusion that can lead to elevated blood glucose levels (hyperglycemia) or very low levels (hypoglycemia). A few commonly occurring faults are summarized below (Zhang et al. provide an exhaustive list of hazards [43]):

1. Software errors in the insulin pump, affecting its ability to deliver insulin of the specified amount and at the specified rates.
2. Calibration errors in the glucose monitors, whose readings are used to compute the correction bolus.
3. Mismatches between the *planned meal* used in the bolus calculations and the *actual meal* consumed.
4. Incorrect timing of the insulin dosage.
5. Incorrect usage of insulin infusion pump (eg., entering a wrong dosage, incorrect bolus shape, unit errors).
6. Failures due to occlusions or pump hardware faults.

The systematic study of the effects of the faults on the overall infusion process is necessary to find and remedy common causes that may result in significant hyperglycemia (elevated blood glucose levels) causing dangerous conditions such as ketacidosis, or hypoglycemia (low blood glucose levels) that may lead to a loss of consciousness or a dangerous coma in the worst case. For instance, it is natural to ask questions such as (a) what are the worst-case effects of a particular single fault or a combination of faults? (b) given probabilities of individual faults, what is the overall probability of a severe hypoglycemia?

While it is possible to predict the qualitative effects of a single fault in isolation, the combined effect of multiple faults are often be hard to predict *quantitatively*. Naturally, a clinical study with real patients using the pump with various controls is the gold standard for providing answers to some of the questions above. However, such studies are expensive, requiring a large set of participants since some of the faults occur infrequently.

An emerging line of research consists of modeling the various components involved in the infusion: the infusion pump, the user's meal planning strategies, physiological models of the insulin-glucose regulation, the glucometer incorporating models of the various faults that may occur [40,27,29]. Such models can then be analyzed for finding worst-cases and their likelihood. This can often point the way towards improving the process to make it safer for patients.

Modeling and Simulation: Figure 1 shows the basic components that are modeled in this scenario and the interactions between these models. Our approach integrates models of the insulin-glucose regulatory system [22,40], meal absorption models [41], a

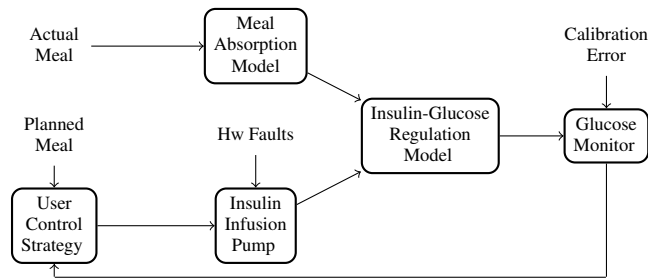


Fig. 1. Key components in modeling meal insulin infusion pump usage scenario.

minimal infusion pump model, glucose meter models [40,15] and a model of the patient’s usage of the pump to cover meals. The latter model is based on an understanding of common pump usage recommendations by physicians and certified diabetes educators [34]. These sub-models are integrated to yield a Matlab Simulink(tm) model.

Figure 2 shows a set of possible blood glucose and insulin levels over time obtained by running simulations with randomly chosen glucose monitor calibration errors and discrepancies between planned and actual meals. The shaded region shows acceptable limits for glucose levels. We note that presence of errors and faults have the effect of potentially causing hyperglycemia as well as hypoglycemia. However, repeated simulations do not suffice to explore worst-case scenarios. If simulations are performed uniformly at random, the number of simulations required to uncover these scenarios is often prohibitively large. The analysis techniques used here explore the worst case outcomes using state-space exploration guided by trace robustness [28,4] and estimate the probability of hypo- and hyperglycemia, given the probabilities for the individual machine faults and user errors [42,23].




Fig. 2. Variation of plasma glucose (left) and plasma insulin (right) concentrations with varying calibration errors and infusion timings. The reference trajectories with no infusion faults is highlighted.

3 Background

We first provide some brief background on diabetes mellitus and its treatment using intensive insulin therapy. More information on topics related to diabetes can be obtained from clinical textbooks on this topic [35].

Insulin-Glucose Regulatory System Diabetes mellitus is the generic name for a class of diseases where critical parts of the natural glyco-control system fail. Type-1 diabetes




Fig. 3. (a,b) Commercial insulin infusion pump models, (c) blood glucose monitor and (d) continuous sub-cutaneous glucose monitors.

results from the loss of pancreatic insulin secretion due to auto-immune destruction of insulin producing β cells in the pancreas. Likewise, type-2 diabetes results from insulin sensitivity, wherein damage to insulin receptors in the cells makes the action of insulin weaker, resulting in the inability of the pancreas to keep up with the demand.

Diabetes is a commonly occurring ailment in the developed world as well as the developing world. A common treatment for chronic diabetes involves the external delivery of artificial insulin (or insulin analogs) directly through a syringe, or sub-cutaneously through an insulin infusion pump. The everyday delivery of insulin is controlled by the patient with advance knowledge of their activities such as diet and exercise. Furthermore, diabetic patients are required to monitor their blood glucose levels intermittently through “finger stick” blood glucose monitors, or recently by continuous glucose monitors (CGMs) that provide a continuous reading of the subcutaneous glucose levels.

Insulin Infusion Pump and Continuous Glucose Monitors

We will now review some of the basics of insulin infusion pumps and blood glucose monitors. The monograph by Chee and Fernando contains further details [9].

An insulin infusion pump is a device that delivers insulin at a programmable rate over time. Insulin infusion pumps have been shown to deliver insulin accurately even when the requested rate of delivery is very small. This allows the pump to deliver insulin continuously throughout the day at a basal rate to counteract the endogenous production of glucose in the body. Furthermore, pumps also allow for various bolus doses of insulin to be infused before or just after meals to limit the occurrence of hyperglycemia following a meal. The shape, width and amount of the bolus can be fine tuned according to the planned meal. Starchy foods such as rice have a high glycemic index, requiring rapid infusion of (short-acting) insulin while fat and protein-rich foods have a lower glycemic index, requiring an infusion with a spread out peak (eg., square-wave bolus).

A continuous glucose monitor (CGM) provides frequent estimates of the blood glucose level by sensing the amount of glucose subcutaneously in the interstitial fluid. Currently available CGM devices provide readings that can be quite accurate. Furthermore, these devices can communicate wirelessly with a computer or an insulin pump to transmit readings directly. CGMs have proven useful in providing feedback to diabetic patients and their physicians to improve the patient’s ability to achieve normoglycemia.

Model Name	Type	Vars	Remarks
Ackerman	Affine	2	Two compartment linear model [2,1]
Bergman	Nonlinear	3	2 insulin + 1 glucose compartment [7,6]
Cobelli	Nonlinear	~ 11	Comprehensive model including glucagon submodel and renal function model [11,13]
Sorensen	Nonlinear	~ 19	Comprehensive physiological model with compartments for brain, vascular, kidney, renal and peripheral systems [36].
Hovorka	Nonlinear	~ 11	Comprehensive model incorporating endogenous glucose production and renal filtration [22,21,40].

Table 1. Commonly used mathematical models of insulin-glucose regulation.

4 Modeling Insulin Infusion

In this section, we present the overall model developed for insulin infusion. Figure 1 shows the overall model for the insulin infusion setup. The setup consists of an insulin infusion pump used to deliver insulin to the patient to counteract the effect of a meal. The meal itself is modeled based on the meal time, duration and its “glycemic factor” that dictates the time from the start of the meal to the peak in blood glucose. The gut absorption of the meal is modeled using a simple linear two compartment model proposed originally by Worthington [41]. Finally, the model incorporates a user model that attempts to capture the user’s insulin infusion and correction dosages. An ideal user model is first formulated and calibrated based on the best practices advocated in many guide books that are used by patients using infusion pumps (Cf. [34], for instance). We then attempt to model various user mistakes such as discrepancies between the planned meal and the actual meal, calibration errors in glucose readings, mis-timing of the correction dosages, miscalibration of basal insulin levels, insulin-to-carbs ratio and insulin sensitivity factors.

We will now describe the construction of each of the sub-models in detail.

4.1 Glucoregulatory Models

There have been numerous attempts to derive mathematical models of the regulation of glucose by insulin in diabetic patients. In this paper we employ the Hovorka model [22,21,40] originally proposed by Hovorka and co-workers. We note that there are many other models that are widely used. A few of the notable models are summarized in Table 1. We refer the reader to many comprehensive surveys on this topic including Cobelli et al. [12], Hovorka [19] and the monograph by Chee and Fernando [9]. Comparing predictions obtained by various models of the insulin-glucose regulatory system in our risk assessment framework is an important future work.

The Hovorka model refers to a modeling approach that has been used to model the regulatory system based on tracer studies during a standard intravenous glucose tolerance test [22]. The test measures the amount of insulin and the time needed to restore normal plasma glucose concentrations after the direct infusion of an unit of glucose under fasting conditions. Data from this test were fitted to a model that considers the

$$\begin{aligned}
\frac{dQ_1(t)}{dt} &= - \left[\frac{F_{01}^c}{Q_1(t)} + x_1(t) \right] Q_1(t) + k_{12}Q_2(t) - F_R(t) + \text{EGP}_0(1 - x_3(t)) + U_g(t) \\
\frac{dQ_2(t)}{dt} &= x_1(t)Q_1(t) - [k_{12} + x_2(t)]Q_2(t) \\
\frac{dS_1(t)}{dt} &= U_I(t) - \frac{S_1(t)}{t_{max,I}} \\
\frac{dS_2(t)}{dt} &= \frac{1}{t_{max,I}}(S_1(t) - S_2(t)) \\
\frac{dI(t)}{dt} &= \frac{S_2(t)}{t_{max,I}V_I} - k_e I(t) \\
\frac{dx_j(t)}{dt} &= k_{a,j}x_j(t) + k_{b,j}I(t), \quad j = 1, 2, 3 \\
G(t) &= \frac{Q_1(t)}{V_g}
\end{aligned}
\quad
F_{01}^c(t) = \begin{cases} F_{01}, & \text{if } G \geq 4.5 \text{ mmol/L} \\ \frac{F_{01}G(t)}{4.5}, & \text{otherwise} \end{cases}$$

$$F_R(t) = \begin{cases} 0.003(G(t) - 9)V_G, & \text{if } G \geq 9 \text{ mmol/L} \\ 0, & \text{otherwise} \end{cases}$$

Fig. 4. Hovorka’s model for insulin-glucose regulatory system. See [9] for an explanation and comparison with other models. Inputs to the model are $U_I(t)$ the rate of insulin infusion and $U_g(t)$, the rate of plasma glucose infusion. The output is $G(t)$ the blood glucose concentration.

various factors affecting glucose concentration: its uptake by cells, its endogenous production, renal clearance and production due to meal absorption. The complete ODE (with details omitted) is summarized in Figure 4. A detailed explanation is available elsewhere [21,40,9].

The parameter values for a group of “virtual patients” are summarized by Wilinska et al. [40]. These parameter sets capture the observed intra- and inter subject variations seen in real-life patient studies. The Hovorka model has been the basis of a model-predictive controller that has been designed to automatically regulate overnight insulin levels in diabetic through an insulin infusion pump and subcutaneous measurements of glucose concentrations through continuous glucose monitors [21]. The controller has been extensively simulated in-silico to estimate the risk of hypoglycemia [40] and recently has been tested successfully in clinical trials [20].

Meal Sub-Model The meal sub-model is part of the overall glucoregulatory model described by Hovorka et al. to model the rate of absorption of the meal into the bloodstream by the digestive system. We use a two compartment model

$$\frac{dG_1(t)}{dt} = -\frac{G_1(t)}{t_{max,G}} + B \cdot U_D(t) \quad \text{and} \quad \frac{dG_2(t)}{dt} = \frac{1}{t_{max,G}}(G_1(t) - G_2(t)).$$

Here $G_1(t)$, $G_2(t)$ model the amounts of glucose in the two hypothetical compartments, B refers to the bio-availability of the meal (taken to be 0.8 in our simulations), $U_D(t)$ refers to the meal input in terms of millimoles of glucose ingested at time t , and $t_{max,G}$ refers to the time to peak glucose absorption rate. In general, $t_{max,G}$ is a function of the meal glycemic index, wherein meals with high glycemic indices such as starch cause the glucose absorption to peak relatively quickly, while meals with lower glycemic indices such as protein and fat rich meals result in relatively flatter peaks that appear slowly. Throughout our simulation, we will use $t_{max,G}$ as being synonymous with the glycemic index of the chief carbohydrate source in the consumed meal. The input $U_g(t)$ of glucose to the bloodstream resulting from the meal is given by $U_g(t) = \frac{G_2(t)}{t_{max,G}}$.




Fig. 5. Schematics for insulin delivery profiles supported by most insulin infusion pump models: (left) basal, (middle) spike bolus and (right) square wave bolus.

Recently, Dalla Man et al. consider non-linear models of gut absorption wherein the rate constant of glucose absorption from the gut is itself dependent on the amount of glucose present. This model is shown to fit tracer meal data better than the Worthington model [14,26].

Glucose Monitors The glucose monitor model periodically samples the output of the insulin-glucose regulatory model to simulate readings of the subcutaneous glucose. We assume that the value read by the glucose monitor is subject to a systematic calibration error. Calibration errors in continuous glucose monitors (CGMs) have been studied by Wilinska et al. (ibid.), Castle and Ward [8] and Cobelli and co-workers [15]. CGMs need periodic re-calibration using traditional “finger stick” blood glucose readings. It is conceivable, however, that the user may often delay this process leading to significant calibration error wherein the reading may be off by as much as 40-50%. Our setup models a fixed calibration error parameter that can be set at the start of the simulation. The assumed calibration error for each simulation can occur in either direction. Apart from calibration errors, “dropouts” have been commonly reported wherein the reading from the CGM is attenuated for brief stretches of time. The simulation of “dropouts” due to physical sensor errors is not currently considered in our setup.

Insulin Pump Sub-Model The pump model is itself quite simple: it supports (a) basal delivery of insulin at a fixed rate and upon receiving a command, it provides a spike bolus dosage of a given amount, shape and width. Figure 5 schematically presents the basic modes supported by infusion pumps. Our model has “hardwired” bolus profiles representing a unit bolus amount over a unit time, in the form of lookup tables that summarize fractions of requested amounts against sub intervals. Given a particular amount and time, the values from the lookup table are scaled appropriately and the insulin inputs $U_I(t)$ to the gluco-regulatory models are set.

Insulin pump faults include inaccurate doses delivered due to hardware or software errors, stoppages due to occlusions or pump failure. These can also be modeled during the infusion by using input parameters that specify the times and durations of the various faulty situations.

4.2 User Infusion Control Strategy Model

Typically, the infusion pump is used to deliver a continuous flow of short-acting insulin through the day, and intermittent bolus infusions to cover the glucose level increase

after a meal. Correction doses of insulin are administered to correct for higher than normal levels immediately before a meal and/or a few hours after a meal. This section is based on information available from diabetes education websites and books. We refer the reader to the book by Scheiner for more information [34].

Basal Insulin Requirement Basal insulin refers to a constant flow of insulin delivered by a pump all day to compensate for endogenous glucose production. The basal insulin level I_B requires periodic calibration by the user based on glucose levels observed during the night 3-4 hours after dinner. An appropriately calibrated basal level I_B ensures that the increase in blood glucose level during an extended period of fasting (eg., during the night well after dinner) is as small as possible.

A typical recommendation for basal insulin for “moderately active” adult with type-1 diabetes is $I_B(U/day) = 0.4 \times \text{weight in kilograms}$. Starting from this rule of thumb, the basal levels are adjusted with feedback obtained by frequent monitoring of blood glucose levels to fine tune a basal insulin requirement.

Calculating Pre-Prandial Insulin Bolus In order to adjust for the increase in post-prandial glucose levels, a bolus dose of insulin is delivered through the pump before the meal. The amount and width of a bolus infusion can be directly programmed by the user or calculated by the pump using the planned meal parameters as inputs. The inputs to the calculation include (1) grams of carbohydrates (CHO), (2) glycemic index (GI) of the major CHO source in the food, and (3) personal ratio for insulin-to-carbohydrates.

The bolus size is calculated as follows:

$$\text{pre-meal bolus amt.}(U) = \text{amt. of CHO}(gms) \times \text{insulin-to-carbs ratio}(U/gm).$$

Additionally, based on the current blood glucose reading a correction bolus may also be required. The correction bolus uses the formula

$$\text{correction bolus}(U) = \frac{G(t) - G_{\text{desired}}}{\text{sensitivity}}, \text{ if } G(t) > G_{\text{desired}}.$$

Here $G(t)$ refers to the current blood glucose reading, G_{desired} refers to the desired level, and sensitivity is a parameter that is discovered by calibration during the initial period of pump usage by the patient.

There are many different “rules of thumb” for arriving at an initial estimate of the basal insulin requirement, the insulin-to-carb ratio and the sensitivity factor. A starting guess at the insulin-to-carbs ratio is given by $\frac{850}{\text{weight}(kgs)}$. An initial sensitivity factor is obtained using the formula $\frac{\text{total daily insulin reqd. (mmol/L)}}{94}$. Often starting from these values, the patient is asked to carefully adjust these values over a period of weeks to achieve robust control of their blood sugar levels.

Timing the Bolus: Another key parameter is the timing of the bolus relative to the meal time. The recommended time for the pre-meal bolus depends on the *glycemic index* (GI) of the meal. For instance, high glycemic index meals (starches such as rice, potatoes, white bread) require the bolus infusion 30 – 40 minutes pre-meal while lower glycemic index foods require an infusion that starts with the meal. Another parameter is the shape and the width of the infusion. Typically starch heavy meals are covered by

a spike bolus while a low GI meal rich in fat and proteins is covered by a square wave bolus. Insulin infusion pumps incorporate infusion modes that can support these bolus shapes including combinations for meals that combine various food types.

Overall infusion control model: The infusion control model incorporates a program that calculates the pre-meal bolus requirements using the insulin-to-carbs ratio and the correction insulin using the sensitivity factors. These factors are “calibrated” for the insulin-glucose regulation model using the robustness guided state-space exploration technique described in Section 5. The timing and shape of the bolus are determined by classifying the planned meal GI into three categories high, medium and low [34]. The models developed for this paper and the analysis results will be made available for download as part of the S-Taliro tool.

5 Robustness Guided Search

In this section, we present the basic concepts used in the analyses of models with respect to metric temporal logic properties [4]. We present the details of our analysis methodology at a high level. More details are available from our prior work which deals with the problem of using robustness guided state-space exploration to find falsifying traces for MTL specifications of non-linear hybrid systems [28,17].

5.1 Metric Temporal Properties and Trace Robustness

Table 2. Metric Temporal Logic (MTL) Operators and their formal semantics at time $t = t_0$. $\sigma : [0, T] \mapsto \mathbb{R}^n$ refers to a continuous time signal, \mathcal{I} refers to a real time interval, AP refers to a set of atomic proposition symbols, \mathcal{O} maps each atomic proposition to a subset of \mathbb{R}^n .

Formula φ	Semantics $(\sigma, t_0, \mathcal{O}) \models \varphi$	Remarks
\top	<i>true</i>	Tautology
$p \in \text{AP}$	$\sigma(t_0) \in \mathcal{O}(p)$	Atomic Proposition holds.
$\varphi_1 \wedge \varphi_2$	$(\sigma, t_0, \mathcal{O}) \models \varphi_1 \wedge (\sigma, t_0, \mathcal{O}) \models \varphi_2$	Conjunction
$\varphi_1 \vee \varphi_2$	$(\sigma, t_0, \mathcal{O}) \models \varphi_1 \vee (\sigma, t_0, \mathcal{O}) \models \varphi_2$	Disjunction
$\neg \varphi$	$(\sigma, t_0, \mathcal{O}) \not\models \varphi$	Negation
$\square_{\mathcal{I}} \varphi$	$(\forall t \in \mathcal{I})((t_0 + t < T) \Rightarrow (\sigma, t_0 + t, \mathcal{O}) \models \varphi)$	φ is Invariant in \mathcal{I}
$\diamond_{\mathcal{I}} \varphi$	$(\exists t \in \mathcal{I})((t_0 + t < T) \wedge (\sigma, t_0 + t, \mathcal{O}) \models \varphi)$	φ eventually holds in \mathcal{I}
$\varphi_1 \mathcal{U}_{\mathcal{I}} \varphi_2$	$(\exists t \in \mathcal{I})((t_0 + t < T) \wedge (\sigma, t_0 + t, \mathcal{O}) \models \varphi_2 \wedge (\forall t' \in [0, t]) (\sigma, t_0 + t', \mathcal{O}) \models \varphi_1)$	φ_1 until φ_2

Metric Temporal Logic (MTL) is a formalism to specify temporal properties of continuous time signals [25]. Table 2 summarizes the syntax and semantics of MTL formulae. MTL formulae can be used to succinctly express key properties of desirable post-prandial blood glucose levels. Let $t = 0$ model the start of a meal and we assume

No Hypoglycemia φ_{hypo}	$\square_{[0,400]}(G \geq 3)$
No significant hyperglycemia φ_{hyper}	$\square_{[0,400]}(G \leq 20)$
Glucose levels settle after digestion φ_{settle}	$\square_{[200,400]}(G \in [3, 10])$

Table 3. MTL specifications for normal post-prandial glycemc control.

$t = 400$ to be end of the simulation period being considered. $G(t)$ is a signal modeling the blood glucose concentration in terms of mmol/L at time t . Table 3 shows the three properties of interest to us along with their descriptions.

Our goal is to find executions of the overall infusion process model that falsify at least one of the properties in Table 3. Here each execution trace corresponds to a different values of planned vs. actual meal data, and calibration error. However, there are potentially infinitely many executions for various values of the input parameters and the models for the insulin-glucose regulatory system are non-linear. Therefore, we use robustness metrics over execution traces to define an objective function over traces to guide us in the search for a falsifying input.

5.2 Trace Robustness

The robustness of signals obtained by simulating hybrid systems is a useful concept that generalizes the standard true/false interpretation of MTL formulae to real valued semantics. Informally, robustness provides a measure of how far away a given trace is from satisfying or violating a property. Real-valued semantics for temporal specifications were considered by Rizk et al. for applications in systems biology [30] and independently by Fainekos and Pappas for testing control systems [16,17]. Figure 6 illustrates the main idea behind the robustness value of a trace σ w.r.t a MTL formula φ . Informally, the robustness value ε indicates the size of the smallest cylinder that can be drawn around σ so that any other trace σ' contained inside this cylinder also has the same valuation for the property as φ . I.e, σ' satisfies φ iff σ does.

Formally, the robustness of a trace σ w.r.t a formula φ , denoted $R(\sigma, \varphi)$, is a real number such that

1. If $R(\sigma, \varphi) > 0$ then $\sigma \models \varphi$. Likewise, if $R(\sigma, \varphi) < 0$ then $\sigma \not\models \varphi$.
2. If $R(\sigma, \varphi) = \varepsilon$, then any trace that lies inside a cylinder of radius $|\varepsilon|$ defined around σ will also have the same outcome for the property φ as σ .

Details on the systematic calculation of robustness values from a sampled continuous-time trace σ and a bounded-time MTL formula φ are available elsewhere [17]. The approach to falsification of a property φ given a model \mathcal{M} is to minimize the objective $R(\sigma, \varphi)$ over all traces σ of the model \mathcal{M} . As noted in our previous work [28], this optimization problem is non-convex for most systems and furthermore, the objective $R(\sigma, \varphi)$ cannot be written down in a closed form. However, it can be evaluated for a given trace σ . Our previous works have explored the use of various global optimization techniques such as Monte-Carlo simulation using Simulated Annealing [28], Ant-Colony Optimization [3], Genetic Algorithms and more recently the Cross-Entropy




Fig. 6. Illustration of robustness of time trajectories. The trajectory is required to lie inside the union of the two rectangles. A “cylindrification” around each trajectory is shown such that any trace in cylindrification has same outcome w.r.t trace as the original trajectory. Leftmost trajectory satisfies property with strictly positive robustness due to larger cylindrification radius, middle trajectory satisfies but with a small robustness value while rightmost trajectory violates property with negative robustness.

Parameter	Ideal	Actual	Remarks
Meal Time (mins)	40	[0, 80]	discrepancy between planned ($t=40$) and actual meal time.
Meal Carbs (gms)	250	[150,350]	discrepancy between planned and actual carbs ingested.
Meal $t_{g,max}$ (mins)	40	[20,80]	meal start time to peak glucose absorption (planned vs. actual GI).
Meal Duration (mins)	30	[10,50]	meal duration planned vs. actual.
Correction Bolus Time	200	[100,300]	time when correction bolus is administered.
Calibration Error	0	[-0.3,0.3]	CGM calibration error.

Table 4. Infusion faults and assumed ranges for minimal robustness search.

Method [32]. These techniques have been implemented in a Matlab toolbox called S-Taliro [4], which supports state-space exploration of Simulink/Stateflow models with MTL specifications.

6 Worst-Case Scenario Search

We fix a scenario consisting of a planned meal at time $t = 40$ minutes after the start of the simulation, a planned duration of 30 minutes, and consisting of 200 grams of CHO (~ 850 calories) with $t_{g,max} = 40$ minutes, indicating a high GI meal (eg., bread, rice or pasta). Our goal here is to explore the risks to the patient arising from these faults in the infusion process. We note that our general framework allows us to explore other meal scenarios and a different set of faults, as well. The overall methodology for this exploration involves the following steps: (a) Formulating MTL properties to falsify (Cf. Table 3); (b) Calibrating the model, so that the MTL properties are satisfied robustly under normal, fault-free situations; (c) Setting up various combinations of faults (Cf. Table 4); (d) Using S-Taliro tool to search for falsifications of the properties in Table 3 in the presence of infusion faults; and (e) Repeating the analysis with various faults disabled to understand the minimal set of faults that are responsible for a given scenario.




Fig. 8. Two minimal robustness scenarios for potential hypoglycemia risk. (left) meal too late/infusion too early and (right) meal too early/infusion too late.

Even though our model allows for transient infusion faults, the study performed in this section, does not include such faults. This is primarily due to the lack of available data on the frequency and timings of pump failures.

Model Calibration: First, we calibrate the patient-specific parameters involved in deciding the size of the infusion including the basal insulin I_B , the insulin-to-carbs ratio and the sensitivity factor. After fixing a planned meal as described previously and disabling all faults, we perform this calibration by using S-Taliro to search for parameter values that *maximize* (usually, falsification involves a minimization) the robustness of the correctness property $\psi : \varphi_{hypo} \wedge \varphi_{hyper} \wedge \varphi_{settle}$, as defined in Table 3. By maximizing the robustness, we are, in effect, searching for a trace that (a) satisfies the property ψ , i.e, achieves ideal glycemic control and (b) does so robustly with nearby traces also satisfying the property. The parameter values corresponding to the maximal robustness are chosen as the ideal parameters that achieve the best overall control for a long enough time horizon. Figure 7 shows the output of a fault-free execution using the calibration results to control the infusion.




Fig. 7. Post-meal glucose on “calibration parameters” found using S-Taliro.

The analysis using S-Taliro identified four dangerous scenarios, all characterized by discrepancies between the meal times. The assumed sensor calibration errors of up to 30% in either direction were found to have very little effect. Furthermore, analysis performed here was repeated for many different “optimal” values of basal insulin, insulin-to-carbs ratio, and sensitivity factor found by running the model calibration procedure. Each repetition yields the same qualitative results described in the paper, but with slightly varying robustness values for $G(t)$ (upto 20% variation seen). We will now discuss these scenarios in detail.

6.1 Hypoglycemia Scenarios

We first study the effect of faults to falsify the property $\varphi_{hypo} : \square G \geq 3$. We performed the search inside S-Taliro using our model implemented in Simulink/Stateflow (tm), the




Fig. 9. Glucose concentration for significant hyperglycemia (left) and failure to settle (right). The fault-free output is shown as a dashed line.

calibrated values for the user’s strategy reported previously and the meal timings drawn non-deterministically from the ranges specified in Table 4. The optimization was run using the stochastic optimization algorithms based on Simulated Annealing (SA) and Cross Entropy (CE) roughly twenty five times, each time using different random seeds to produce different minimal robustness scenarios. Each run required approximately 3 – 10 minutes with up to 1000 simulations per run.

Each optimization run discovered property violations that expose potential hypoglycemia. The minimal values of G were around $G = 2.3$. Examining these violations, we found that each scenario falls in one of two distinct categories. Figure 8 shows the blood glucose outputs for both scenarios.

Potential Hypoglycemia Scenario-1: By disabling various faults in turn, we identified three sufficient faults for this scenario:

1. The planned meal time ($t = 40$) is significantly earlier than actual meal time ($t \sim 70$). Alternatively, the insulin bolus is delivered too early.
2. The planned GI ($t_{max,G} = 40$) is lower than its actual GI ($t_{max,G} \sim 20$).
3. The actual amount of CHO ingested in this scenario is *less than* the planned amount of CHO.

Additionally, the actual meal durations were slightly larger than planned.

Potential Hyperglycemia Scenario-2: This scenario is characterized by the following combination of faults:

1. The planned meal time ($t = 40$) is significantly *later* than the actual meal time ($t \sim 5$). Alternatively, the bolus is delivered at or after the start of a high GI meal.
2. The meal’s planned GI is lower than the meal’s actual GI.
3. The actual amount of CHO ingested is lower than the planned CHO.

The actual meal durations were slightly less than the planned duration.

6.2 Scenario Analysis for Significant Hyperglycemia

Next, we consider scenarios for significant hyperglycemia $G \geq 25$ that can lead to dangerous conditions such as ketacidosis. Unlike hypoglycemia, this property is found to be easier to falsify, requiring fewer iterations to find falsifying inputs. The minimal

robustness scenarios found by S-Talro depend chiefly on two faults: (a) discrepancy between planned meal GI ($t_{max,G} = 40$) and actual meal GI $t_{max,G} \sim 20$ and (b) discrepancy between planned meal CHO (= 200) and actual meal CHO (~ 300). Other faults have a minor impact on the maximum value of the blood glucose level $G(t)$. In such a scenario, the bolus of insulin is supplied too late and is insufficient to “cover” the meal.

6.3 Scenario Analysis for Failure to Settle

Finally, we consider minimal robustness scenarios w.r.t failure to settle. Figure 9 (right) shows the output glucose level for this scenario. The scenario depends on a combination of two faults: (a) actual meal time significantly later than the planned time, and (b) the actual CHO is higher than the planned CHO. In this scenario, the peak value of the insulin precedes the peak gut absorption of glucose. Notice the initial dip in $G(t)$ much like the first hypoglycemia scenario followed by a delayed rise in $G(t)$ that fails to settle even after 6 hours.

6.4 Evaluation

We have thus far derived some situations that can cause significant hypoglycemia, hyperglycemia and failure to settle. Can the predictions made by our model be tested? The gold standard evaluation would be to conduct clinical studies of patients to determine if the violations observed can be borne in real life. While we are planning to conduct extensive patient studies with collaborators from medical sciences as part of our ongoing work in this area, such studies require time and significant effort to carry out.

Preliminary evidence is available from web logs maintained by many diabetic patients to check if any of the situations reported by us are also confirmed by diabetic patients. We surveyed many such weblogs and some of the scenarios such as scenario-1 for hypoglycemia and the scenario for failure to settle seem to be well known³. However, barring a few exceptions incident reports by patients focused on the effects and not the root causes.

7 Statistical Model Checking

We have, thus far, used robustness guided state-space explorations to explore the extreme, worst-case scenarios that may happen in the infusion. However, it is equally interesting to find out the probability that the infusion may result in a hypoglycemia, hyperglycemia or a failure to settle (Cf. Table 3). In order to do so, a simple approach is to assume a probability distribution for each of the faults described in Table 4. We assume that the faults are uniformly distributed within their intervals and independent of each other. We then simulate the model by sampling faults from this distribution and find out the fraction of executions that violate each property. However, if the probabilities of

³ For a discussion of meal timing from a patient’s perspective, see <http://thethirstthatchangedmylife.blogspot.com/2010/09/loading-on-carbs.html>

Table 5. Results of SMC for estimating probability of avoiding hypoglycemia (left) and hyperglycemia (right). Each entry shows the posterior probability estimate \hat{p} such that the probability lies within $[\hat{p} - \delta, \hat{p} + \delta]$ with confidence (coverage) indicated by c .

Coverage (c)	$\delta = 0.05$			$\delta = 0.01$			Coverage(c)	$\delta = 0.05$			$\delta = 0.01$		
	.95	.99	.999	.95	.99	.999		.95	.99	.999	.95	.99	.999
$\square G \geq 3$	0.75	0.7	0.74	0.73	0.72	0.72	$\square G \leq 25$	0.97	0.97	0.96	0.97	0.97	0.98
$\square G \geq 2.8$	0.87	0.88	0.88	0.89	0.88	0.88	$\square G \leq 30$	0.97	0.98	0.98	0.99	0.99	0.99
$\square G \geq 2.7$	0.91	0.92	0.93	0.93	0.93	0.93	$\square G \leq 35$	0.97	0.98	0.99	0.99	0.99	0.99

these bad outcomes are tiny to begin with, the number of executions required become prohibitively large. Recent advances in *Statistical Model-Checking* (SMC) have given rise to techniques that can estimate these probabilities efficiently while running as few simulations as possible. SMC was originally formulated by Younes and Simmons [42]. A promising extension involving the use of Bayesian reasoning, has been used to applied to problems in systems biology [23] and control systems verification [44]. Rather than estimating the probabilities of hypoglycemia or hyperglycemia empirically, these techniques use repeated simulations to bound the required probability inside an interval of given half-width δ with a given confidence c . In addition to worst-case search, we have extended S-Taliro to support Bayesian SMC. Existing techniques implemented inside S-Taliro such as the Cross-Entropy Method can be directly used to find a suitable prior distribution to reduce the number of simulations required.

Table 5 shows the probability estimates for significant hypoglycemia with various confidence levels and intervals around the probability. We find that hypoglycemia ($G < 3$) has roughly 30% chance of occurrence, while significant hypoglycemia ($G < 2.7$) has an estimated 7% chance of occurrence with very high confidence. The probabilities for hyperglycemia are also presented in Table 5. Once again, we estimate a 2.2% chance of significant hyperglycemia $G > 25$ and 0.3% chance of finding $G > 35$ (which exposes the patient to dangerous ketacidosis).

The findings of this section are qualitatively borne out by an informal analysis of incidents reported by diabetic patients and some clinical studies [18]. The majority of the infusion faults due to discrepancies between planned and actual meals result in hypoglycemia. Hyperglycemia risks commonly occur due to silent pump failures that were not modeled in this study.

8 Threats to Validity

In this section, we discuss some of the threats to validity and address remedial steps taken to ensure that the results in this work are applicable to real-life situations.

With any result involving *in silico* simulations, there is a risk that we are observing modeling quirks that are not reflective of what happens in reality. However, the models used here have been extensively evaluated against studies on real patients [22,20], providing evidence for their validity.

Another concern is that we assume that the food ingested has a single carbohydrate source with fixed (high) GI. While this can be a good approximation in some cases (eg.,

a meal consisting mostly of pasta or a CHO heavy drink), such meals are not advised for diabetic patients. Insulin pumps provide combination boluses to offset for different types of foods with varying GIs. We plan to investigate these effects as part of our ongoing research. Mixed meal simulations have been considered in the past by Della Man et al. [14,27]. The effects of exercise are also a factor. However, modeling physical activity and its effect on the blood glucose regulation is an active area of research with few established models [31].

A shortcoming of assigning probabilities on the occurrence of faults is that there is no available mathematical evidence that the distribution of planned meal times vs. actual meal times are uniformly distributed in the interval of interest. Another assumption is that of independence of the various faults. It is conceivable that a larger discrepancy between planned mealtimes and actual mealtimes indicates a larger discrepancy between the planned and actual meal CHO or GI. Building fault models based on observations of insulin infusion pump usage of real patients are critical to construct these models. The probabilities reported by SMC are likely to change if we included more types of faults such as pump failures in our study.

9 Conclusions

We have presented an *in-silico* evaluation of the risks involved in the infusion process. Our approach has been two-fold: (a) using robustness-guided model checking to search for potential worst-case scenarios. Here, we report on some scenarios causing hypoglycemia, hyperglycemia and failure of the blood glucose to settle to normoglycemia. Some of our scenarios are borne out by patient reports reported online. (b) We use Statistical Model Checking to place small bounds on the risks with very high confidence.

Our future direction is to consider *individualized risk studies*. Here, we seek to develop models and risk analysis fitted to individual patients. This can yield lifestyle analysis tools that can help advise patients on the best pump calibration parameters to maintain normal glucose levels.

Acknowledgments: We gratefully acknowledge helpful discussions with Dr. Phillip Dzwonczyk, NY State Veterans Home, Oxford, NY and Prof. Clayton Lewis, University of Colorado, Boulder. We thank the anonymous reviewers for their comments.

References

1. E. Ackerman, L. Gatewood, J. Rosevear, and G. Molnar. Blood glucose regulation and diabetes. In F. Heinmets, editor, *Concepts and Models of Biomathematics*, pages 131–156. Marcel Dekker, 1969.
2. E. Ackerman, J. Rosevear, and W. McGuckin. A mathematical model of the insulin-glucose tolerance test. *Physics in Medicine and Biology*, 9:202–213, 1964.
3. Y. S. R. Annapureddy and G. E. Fainekos. Ant colonies for temporal logic falsification of hybrid systems. In *Proceedings of the 36th Annual Conference of IEEE Industrial Electronics*, pages 91 – 96, 2010.
4. Y. S. R. Annapureddy, C. Liu, G. E. Fainekos, and S. Sankaranarayanan. S-taliro: A tool for temporal logic falsification for hybrid systems. In *Tools and algorithms for the construction and analysis of systems*, volume 6605 of *LNCS*, pages 254–257. Springer, 2011.

5. D. E. Arney, R. Jetley, P. Jones, I. Lee, A. Ray, O. Sokolsky, and Y. Zhang. Generic infusion pump hazard analysis and safety requirements: Version 1.0, 2009. CIS Technical Report, University of Pennsylvania. Available on-line: http://repository.upenn.edu/cis_reports/893, Accessed May 2011.
6. R. N. Bergman. Minimal model: Perspective from 2005. *Hormone Research*, pages 8–15, 2005.
7. R. N. Bergman and J. Urquhart. The pilot gland approach to the study of insulin secretory dynamics. *Recent Progress in Hormone Research*, 27:583–605, 1971.
8. J. Castle and K. Ward. Amperometric glucose sensors: Sources of error and potential benefit of redundancy. *J. Diabetes Sci. and Tech.*, 4(1), January 2010.
9. F. Chee and T. Fernando. *Closed-Loop Control of Blood Glucose*. Springer, 2007.
10. E. Clarke, A. Donze, and A. Legay. Statistical model checking of analog mixed-signal circuits with an application to a third order $\delta - \sigma$ modulator. In *Hardware and Software: Verification and Testing*, volume 5394/2009 of LNCS, pages 149–163, 2009.
11. C. Cobelli, G. Federspil, G. Pacini, A. Salvan, and C. Scandellari. An integrated mathematical model of the dynamics of blood glucose and its hormonal control. *Mathematical Biosciences*, 58:27–60, 1982.
12. C. Cobelli, C. D. Man, G. Sparacino, L. Magni, G. D. Nicolao, and B. P. Kovatchev. Diabetes: Models, signals and control (methodological review). *IEEE reviews in biomedical engineering*, 2:54–95, 2009.
13. C. Cobelli and A. Mari. Control of diabetes with artificial systems for insulin delivery — algorithm independent limitations revealed by a modeling study. *IEEE Trans. on Biomed. Engg.*, BME-32(10), Oct. 1985.
14. C. Dalla Man, R. A. Rizza, and C. Cobelli. Meal simulation model of the glucose-insulin system. *IEEE Transactions on Biomedical Engineering*, 1(10):1740–1749, 2006.
15. A. Facchinetti, G. Sparacino, and C. Cobelli. Modeling the error of continuous glucose monitoring sensor data: Critical aspects discussed through simulation studies. *J. Diabetes Sci. and Tech.*, 4(1), January 2010.
16. G. Fainekos and G. J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science*, 410:4262–4291, 2009.
17. G. E. Fainekos. *Robustness of Temporal Logic Specifications*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, 2008.
18. L. Fox, L. Buckloh, S. D. Smith, T. Wysocki, and N. Mauras. A randomized controlled trial of insulin pump therapy in young children with type 1 diabetes. *Diabetes Care*, 28(6), June 2005.
19. R. Hovorka. Continuous glucose monitoring and closed-loop systems. *Diabetic Medicine*, 23(1):1–12, 2005.
20. R. Hovorka, J. M. Allen, D. Elleri, L. J. Chassin, J. Harris, D. Xing, C. Kollman, T. Hovorka, A. M. Larsen, M. Nodale, A. D. Palma, M. Wilinska, C. Acerini, and D. Dunger. Manual closed-loop delivery in children and adolescents with type 1 diabetes: a phase 2 randomised crossover trial. *Lancet*, 375:743–751, February 2010.
21. R. Hovorka, V. Canonico, L. Chassin, U. Haueter, M. Massi-Benedetti, M. Federici, T. Pieber, H. Shaller, L. Schaupp, T. Vering, and M. Wilinska. Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. *Physiological Measurement*, 25:905–920, 2004.
22. R. Hovorka, F. Shojaee-Moradie, P. Carroll, L. Chassin, I. Gowrie, N. Jackson, R. Tudor, A. Umpleby, and R. Hones. Partitioning glucose distribution/transport, disposal and endogenous production during IVGTT. *Am. J. Physiol. Endocrinol. Metab.*, 282:992–1007, 2002.
23. S. K. Jha, E. M. Clarke, C. J. Langmead, A. Legay, A. Platzter, and P. Zuliani. A bayesian approach to model checking biological systems. In *CMSB*, volume 5688 of *Lecture Notes in Computer Science*, pages 218–234. Springer, 2009.

24. S. K. Jha, R. Datta, C. Langmead, S. Jha, and E. Sassano. Synthesis of insulin pump controllers from safety specifications using bayesian model validation. In *Proceedings of 10th Asia Pacific Bioinformatics Conference, (APBC)*, 2012.
25. R. Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.
26. C. Man, M. Camilleri, and C. Cobelli. A system model of oral glucose absorption: Validation on gold standard data. *Biomedical Engineering, IEEE Transactions on*, 53(12):2472–2478, dec. 2006.
27. C. D. Man, D. M. Raimondo, R. A. Rizza, and C. Cobelli. GIM, simulation software of meal glucose-insulin model. *J. Diabetes Sci. and Tech.*, 1(3), May 2007.
28. T. Nghiem, S. Sankaranarayanan, G. E. Fainekos, F. Ivančić, A. Gupta, and G. J. Pappas. Monte-carlo techniques for falsification of temporal properties of non-linear hybrid systems. In *Hybrid Systems: Computation and Control*, pages 211–220. ACM Press, 2010.
29. S. Patek, B. Bequette, M. Breton, B. Buckingham, E. Dassau, F. Doyle III, J. Lum, L. Magni, and H. Zisser. In silico preclinical trials: methodology and engineering guide to closed-loop control in type 1 diabetes mellitus. *J Diabetes Sci Technol.*, 3(2):269–82, 2009.
30. A. Rizk, G. Batt, F. Fages, and S. Soliman. On a continuous degree of satisfaction of temporal logic formulae with applications to systems biology. In *6th International Conference on Computational Methods in Systems Biology*, number 5307 in LNCS, pages 251–268. Springer, 2008.
31. A. Roy and R. Parker. Dynamic modeling of exercise effects on plasma glucose and insulin levels. *J. Diabetes Sci. and Tech.*, 1(3):338–347, May 2007.
32. S. Sankaranarayanan and G. E. Fainekos. Falsification of temporal properties of hybrid systems using the cross-entropy method. In *HSCC*, pages 125–134. ACM, 2012.
33. S. Sankaranarayanan, H. Homaei, and C. Lewis. Model-based dependability analysis of programmable drug infusion pumps. In *FORMATS*, volume 6919 of *Lecture Notes in Computer Science*, pages 317–334, 2011.
34. G. Scheiner. *Think like a pancreas: A Practical guide to managing diabetes with insulin*. Da Capo Press, 2011.
35. J. S. Skyler (editor). *Atlas of Diabetes: Fourth Edition*. Springer Science+Business Media, 2012.
36. J. Sorensen. *A Physiological Model of Glucos Metabolism in Man and its use to Design and Access Improved Insulin Therapies for Diabetes*. PhD thesis, Massachusetts Inst. of Technology (MIT), 1985.
37. R. E. Teixeira and S. Malin. The next generation of artificial pancreas control algorithms. *J. Diabetes Sci. and Tech.*, 2:105–112, Jan 2008.
38. H. Thimbleby. Ignorance of interaction programming is killing people. *ACM Interactions*, pages 52–57, 2008.
39. H. Thimbleby. Is it a dangerous prescription? *BCS Interfaces*, 84:5–10, 2010.
40. M. Wilinska, L. Chassin, C. L. Acerini, J. M. Allen, D. Dunber, and R. Hovorka. Simulation environment to evaluate closed-loop insulin delivery systems in type 1 diabetes. *J. Diabetes Science and Technology*, 4, January 2010.
41. D. Worthington. Minimal model of food absorption in the gut. *Medical Informatics (London)*, 22(1):35–45, March 1997.
42. H. L. S. Younes and R. G. Simmons. Statistical probabilistic model checking with a focus on time-bounded properties. *Information & Computation*, 204(9):1368–1409, 2006.
43. Y. Zhang, P. L. Jones, and R. Jetley. A hazard analysis for a generic insulin infusion pump. *J. Diabetes Sci. and Tech.*, 4(2):263–282, March 2010.
44. P. Zuliani, A. Platzer, and E. M. Clarke. Bayesian statistical model checking with application to simulink/stateflow verification. In *HSCC*, pages 243–252. ACM, 2010.