Model-Based Falsification of an Artificial Pancreas Control System.

Sriram Sankaranarayanan[1], Suhas Akshar Kumar[1], Faye Cameron[2], B. Wayne Bequette[2], Georgios Fainekos[3] and David M. Maahs[4]

1. University of Colorado, Boulder, CO, USA.
2. Rensselaer Polytechnic Institute, Troy, NY, USA.
3. Arizona State University, Tempe, AZ, USA.
4. Barbara Davis Center for Childhood Diabetes, University of Colorado, Denver, CO, USA.

ABSTRACT

We present a model-based falsification scheme for artificial pancreas controllers. Our approach performs a closed-loop simulation of the control software using models of the human insulin-glucose regulatory system. Our work focuses on testing properties of an overnight control system for hypoglycemia/hyperglycemia minimization in patients with type-1 diabetes. This control system is currently the subject of extensive phase II clinical trials.

We describe how the overall closed loop simulator is constructed, and formulate properties to be tested. Significantly, the closed loop simulation incorporates the control software, as is, without any abstractions. Next, we demonstrate the use of a simulation-based falsification approach to find potential property violations in the resulting control system. We formulate a series of properties about the controller behavior and examine the violations obtained. Using these violations, we propose modifications to the controller software to improve its performance under these adverse (corner-case) scenarios. We also illustrate the effectiveness of robustness as a metric for identifying interesting property violations. Finally, we identify important open problems for future work.

1. INTRODUCTION

This paper presents a case-study on the use of robustness-guided falsification techniques for analyzing properties of a closed-loop artificial pancreas system. The systematic testing of closed-loop medical devices is an important step towards ensuring the safety of their users. To this end, staged clinical trials have served as the gold standard for evaluating the safety and efficacy of medical devices. However, as closed loop devices become more complicated with increasing reliance on software-based control, it is clear that clinical trials can be inadequate for testing the safety and reliability of the closed loop system. This is evidenced by the growing number of instances of medical device failures due to software errors. Continuous post-market testing and monitoring over hundreds of thousands of users will be necessary to obtain information about rare, and potentially dangerous defects [28]. Unfortunately, defects found at this late stage are quite expensive to fix. As a result, a lot of emphasis has been placed on *in-silico* simulation of closed-loop devices with increasingly sophisticated models of human physiology and the device's operating environment.

The *in silico* study in this paper focuses on a *predictive hypo/hyperglycemia minimizer device* for treating patients with type-1 diabetes by regulating insulin delivery. The device employs a Kalman filter to predict the future values of blood glucose from past noisy samples. A series of rules are applied on the predicted future glucose values to decide whether to turn off the pump, continue normal delivery, or increase the normal basal rate. Additionally, the device includes rules to detect conditions such as sensor dropouts, and pressure induced sensor attenuation. Phase I and II clinical trials have been successfully carried out on an earlier prototype of the system being studied [10, 39]. These trials have demonstrated the effectiveness of the devices, and in particular, an increased time in euglycemic range for the participants on the closed loop system.

Our approach to verification first builds a closed loop simulator in Matlab(tm) to simulate meal and insulin bolus patterns against closed loop and open loop control. The Dalla Man et al. model is used to capture the effect of meal and insulin inputs on the patient's blood glucose levels [20, 41].

We formulate interesting properties about the closed loop system behavior use simulation-based falsification approach in the tool S-Taliro to drive the simulations [3]. S-Taliro uses a metric called the robustness of a simulation to predict a distance between a given simulation output and a property of interest [26, 27]. In general, as the robustness value becomes smaller, the simulation output approaches a property falsification. This principle is used inside a stochastic optimization solver to select a sequence of inputs that result in decreasing values of the robustness metric, possibly leading to a violation of the property.

For the closed loop system under study, we formulate ten properties of interest that concern the behavior of the closed loop system when the patient's blood glucose levels are low, the behavior under high blood glucose levels and comparisons between closed and open loop performance. We use S-Taliro to search for violations that can represent *corner case* behaviors of interest to the designers and clinical experts. S-Taliro provides violations for 8 out of the 10 proper-

ties formulated. It also provides the output that approaches "closest" to violation for the remaining properties.

We conclude by identifying two potentially important directions for future work: (a) assign likelihood scores to violations to enable designers to decide if a design should be modified in response to such a violation and (b) improve physiological models by incorporating features that are commonly seen in actual patient data.

1.1 Related Work

A growing body of work focuses on modeling and analysis of closed loop medical devices, including pacemaker and implantable cardiac defibrillators (ICDs). This has included a range of ideas from using specification formalisms for physiological models [49], formal verification techniques to verify closed loop models and the use of physiological models to test control software [47, 35, 34]. Our work here focuses on the analysis of an artificial pancreas control system using models of human insulin-glucose regulatory systems.

The development of detailed mathematical models for the human insulin-glucose regulation system has led to a number of widely used models. These include the Bergman minimal model [8, 7], Dalla Man et al [41, 20, 42] and Hovorka et al models [33, 54]. The success of these modeling efforts has led to the concept of *in silico* clinical trials that use these models to test control algorithms [40, 48]. These approaches use a *virtual clinical protocol* to specify the external inputs (meal timing, amount and bolus) to the simulation. The simulation is performed for varying patient parameter sets and predictions on the performance measures of interest (eg., time in euglycemic range) are obtained. Our approach deploys more exhaustive search techniques that search over a large space of possible inputs to the simulation. Also, we search for worst case scenarios with respect to given property of interest, formulated by the user and expressed in a specification language such as Metric Temporal Logic (MTL) [38].

This paper also builds on our previous work [11] that performs an exhaustive analysis of a PID control algorithm [52]. However, the study described in this paper involves the use of the software implementation in the loop rather than a model constructed from descriptions in the published literature. Working with the software implementation raises some challenges for closed loop simulation that includes the need to construct interfaces between the plant model and the control software. Furthermore, we focus on a wider set of properties that are more specific to the control system under analysis.

Our work is also related to that of Chen et al, wherein symbolic decision procedures are applied to find patient parameter ranges for which a PID controller can be shown to be safe [16]. Beyond the choice of a different verification approach, Chen et al focus on capturing a range of variations of patient parameters whereas our approach captures variations in the inputs (meals, bolus, CGM noise). Furthermore, our approach works with the actual software-in-the-loop setup rather than using a model of the controller.

2. BACKGROUND & MOTIVATION

We provide a brief background on artificial pancreas controllers that motivates the need for verification [28]. We refer the reader to monographs on this topic for further details [19, 30, 53]. Our previous work on this topic also provides background on artificial pancreas control systems [11].

Table 1: Pathway to the artificial pancreas project with representative papers showing technological feasibility. Source: Juvenile Diabetes Research Foundation (JDRF). See [37] for a recently proposed revised pathway.

ID	Description	Refs.
1	Low Glucose Pump Shutoff Pump shutoff during hypo.	[44]
2	Hypoglycemia Minimizer Pump shutoff for predicted hypo.	[13]
3	Hypo./Hyper. Minimizer #2 + additional insulin when glucose above threshold	[4, 46, 29]
4	Hybrid Closed Loop Closed loop insulin delivery with manual bolus	[33, 32, 31]
5	Fully Automated Closed Loop #4 with no manual boluses	[9, 10, 12, 18, 40, 36, 21]
6	Multi-hormone Closed Loop Use glucagon and insulin	[24, 23]

2.1 Background: Artificial Pancreas

Patients with type-1 diabetes (T1D) rely on external administration of insulin to manage their blood glucose levels. The ideal range of blood glucose levels (euglycemic range) is taken to be $[70, 180]$mg/dl. *Hypoglycemia*, caused by glucose values that fall below 70mg/dl, can lead to coma or even death. Similarly, glucose values that persist above 180mg/dl is considered *hyperglycemia*. The short term risks of hyperglycemia above 300mg/dl include *diabetic ketacidosis*. The longer term risks of hyperglycemia above 180mg/dl include damage to eye, kidneys, heart and blood vessels. The overall goal of treating T1D is to maintain the blood glucose levels in the euglycemic range, avoiding hypoglycemia and minimizing the time under hyperglycemia.

Artificial Pancreas (AP) refers to a series of increasingly sophisticated devices used to treat patients with type-1 diabetes through the external administration of insulin (and other hormones). An AP system's core function is to continuously adjust the insulin infusion into the patient through an insulin infusion pump. Typical AP systems use continuous glucose monitors (CGMs) to measure the blood glucose levels in the patients. Additionally, some AP systems may use inputs from the patients such as impending meals and physical activity. AP systems may also output warnings/alarms to patients to announce hypoglycemia and suggest using glucagon or rescue carbohydrates. Table 1 summarizes the existing approaches to realizing the AP concept. The low glucose pump shutoff product has been incorporated in some commercially available pumps [44]. Other stages are undergoing various phases of clinical trials. It must be noted that all stages have been shown to be technologically feasible. Note that the more recent road map by Kowalski [37] emphasizes "automated insulin delivery" and "multihormonal" approaches as parallel pathways rather than successive stages.

Figure 1 presents the structure of a closed loop artificial pancreas at a high level. As mentioned earlier, the closed loop involves the action of a software-based controller that decides on the insulin delivery rate. The high level goal of the system is to maximize the time for which the patient's blood glucose level $G(t)$ remains in the euglycemic range

Figure 1: Closed loop schematic diagram of the overall artificial pancreas system.

Figure 2: Schematic diagram of the hypo/hyper minimizing controller.

of $[70, 180]$mg/dl. Furthermore, the system seeks to avoid hypoglycemia $G(t) < 70$mg/dl and minimize time under hyperglycemia $G(t) > 180$mg/dl.

As shown in Fig. 1, the patient's glucose regulatory system is subject to external disturbances such as meals and exercise. The value of the blood glucose level is estimated by the *continuous glucose monitor* (CGM) to yield a sensed glucose level $G_s(t)$. Note that the CGM is subject to external disturbances $n(t)$. The controller is run in a time-triggered fashion with time period Δ. Typically Δ is in the order of minutes $(1 - 5$ minutes). The controller periodically senses the value of $G_s(t)$ from the CGM at $t = j\Delta$ and computes an insulin level $u_c(t)$, which is held constant in the time interval $t \in [j\Delta, (j + 1)\Delta)$. Furthermore, in many systems, the user can provide an external bolus $b(t)$. Typically this bolus is provided before meals. The overall insulin infusion $u(t) = u_c(t) + b(t)$ is the sum of the controller and externally administered insulin.

The key challenges that make artificial pancreas control hard include:

(a) Excessive insulin can cause dangerously low blood glucose levels leading to coma or even death. On the other hand, too little insulin can cause prolonged high glucose levels leading to short term consequences such as ketacidosis and longer term damage to eye, kidneys, heart and blood vessels.

(b) Insulin is (typically) the only available control. However, many AP systems cannot directly counteract insulin under normal circumstances. As a result, the insulin already administered persists in the system with an onset of action 20 minutes after administration, a peak effect around 90 minutes after administration and persists until 4-6 hours after administration. Furthermore, the insulin action profile can differ by person.

(c) The system must counteract significant disturbances in the form of meals and physical activities, without advance knowledge of these disturbances. Other disturbances include short term changes in patient's physiology due to illnesses, prescription drugs, and alcohol consumption.

(d) CGMs provide an estimate of the blood glucose levels. However, they are subject to noise, calibration errors, dropouts and pressure induced sensor attenuation [25, 14]. As a result, the available sensor measurements can be incomplete or erroneous leading the controller to make a wrong decision.

(e) The system is subject to various delays including sensor delays and actuation delays caused by the delayed action profile of insulin.

Need for Verification: As mentioned earlier, external control of blood glucose levels is a challenging problem. Controller malfunctions that lead to excessive insulin infusion can risk severe consequences to the patient. Typically, software systems often carry the risk of malfunctions due to typical programming errors such as buffer overflow, numeric overflow, and divide by zero. Software verification techniques and better programming language design have focused on eliminating these errors. However, a larger set of functional correctness properties of the overall closed loop remain quite important. The easiest functional correctness property is that the blood glucose levels remain in the euglycemic range $[70, 180]$mg/dl. However, due to the significant disturbances present, it is always possible for the unanticipated user actions to cause the blood glucose levels to go out of range.

Nevertheless, verification approaches are needed to answer many functional correctness questions: (a) Will the insulin infusion always be turned off when the sensor glucose value is below 70mg/dl, the limit for hypoglycemia? (b) Can the controller infuse extra insulin when the patient's blood glucose level is low? If yes, what is the lowest blood glucose level for which the controller may decide to infuse extra insulin? (c) Can the patient remain in hypoglycemia longer than 3 hours? (d) Can the patient remain in hyperglycemia longer than 5 hours? A larger list of such properties will be examined in Section 6.

In this paper, we present a combination of mathematical modeling of the various components of the closed loop, including disturbances. We perform in-silico simulations of these models driven by simulation-based property falsification tools to find potential violations that can inform the designers of these systems about possible worst case behaviors that can result.

3. CONTROLLER

Figure 2 shows the overall schematic for the control algorithm, which is an advanced version of a predictive pump shutoff algorithm, originally described by Cameron et al. [13]. The original system uses Kalman filter-based prediction algorithm to shutoff the pump when a hypoglycemia is predicted. The extended system studied here also commands extra insulin by temporarily increasing the basal insulin delivery rate to mitigate hyperglycemia, as well. Clinical trials of the predictive pump shutoff system include the inpatient clinical trials described by Cameron et al [13], and a more

Table 2: Rules for pump shutoff, additional insulin and resumption of basal insulin delivery. Note that $G(t)$ refers to current CGM value, $G_p(t)$ refers to the Kalman filter prediction at time t, $T_{shutoff}(t)$ is the total amount of time the pump has been turned off until time t.

Condition	Mode
$G_p(t+30) \leq 80\text{mg/dl}$ and $T_{shutoff}(t) < 180$ and $T_{shut}(t-150) < 120$	SHUTOFF
not SHUTOFF and $G_p(t+30) \geq 150$ and $T_{shut}(t) < 180$	ADD INSULIN
SHUTOFF and $\left(\begin{array}{c} G_p(t+30) \geq 70 \text{ or} \\ T_{shut} \geq 180 \text{ or } T_{shut}(t-150) \geq 120 \end{array}\right)$	NORMAL

Figure 3: Structure of a physiological insulin-glucose regulatory model.

recent trial that studied 45 patients over a total of 42 nights, described by Maahs et al [39]. These trials reported promising results, including a longer time in the euglycemic range for the participants.

As mentioned earlier, the system is based on a Kalman filter that analyzes the CGM glucose readings and estimates the first and second derivatives of the blood glucose level $G(t)$. Based on the estimated derivatives, it predicts the value $G_p(t+30)$ of the blood glucose levels 30 minutes into the future. This prediction is used by a rule-based decision support system to command possible pump actions that include (a) SHUTOFF: shut the pump down for a given time interval, (b) ADD INSULIN: infuse extra insulin, (c) NORMAL: continue the current basal rate, and (d) PISA: Alert the user of faulty CGM values caused by *pressure induced sensor attenuation*.

Table 2 briefly describes the major rules that are used by the system to control insulin delivery. However, many of the finer details such as the handling of CGM dropouts and sensor attenuation have been omitted from this discussion. These details will be provided upon request. A shutoff is commanded when the predicted glucose level at $t+30\text{mins}$ is below 80mg/dl, the total shutoff time is less than 180mins, and total shutoff time in the previous 150mins is less than 120mins. Likewise, additional insulin maybe commanded if the pump is not currently shutoff, and the predicted glucose value is above 150mg/dl. Furthermore, additional insulin may only be commanded if there is time remaining for pump shutoff.

The full set of rules that govern the control algorithm considers numerous other factors such as *pressure induced sensor attenuation*, the possibility of loss in sensor signal (dropout), large changes in the sensor value and the time to next sensor calibration. These rules constitute about 900 lines of code written in MATLAB(tm).

4. MODELING HUMAN PHYSIOLOGY

In this section, we briefly describe the process of modeling the insulin glucose regulatory system in patients with type-1 diabetes.

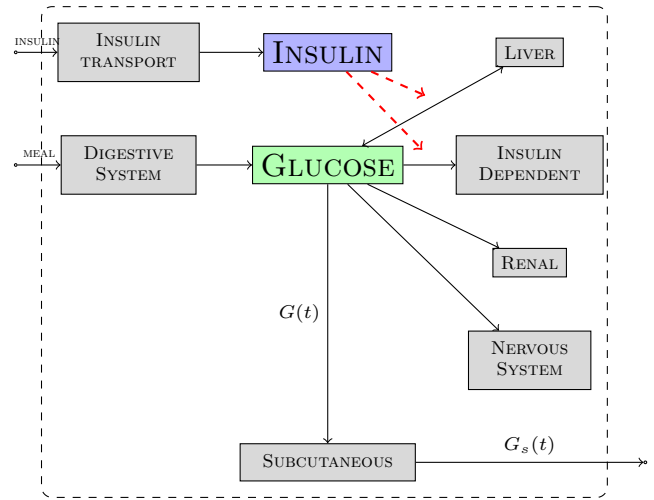**Physiological Models:** The area of physiological modeling of insulin-glucose regulation has received considerable attention, following seminal work by Bergman, Cobelli and Others [7, 15]. Recently, physiological models have been proposed, mainly by Dalla Man et al. [20, 41, 43], and Hovorka et al. [54, 33]. Figure 3 shows the schematic of these physiological models. The main idea is to (a) write balance equations that account for the entry, storage, uptake and excretion of glucose and insulin and (b) the effect that plasma insulin levels have on the uptake of glucose and endogenous production by liver. The resulting model is an ordinary differential equation (ODE). This ODE is often nonlinear due to the nonlinear action profile of plasma insulin levels on endogenous glucose production, and insulin dependent glucose uptake. Also, the gut absorption model used is nonlinear [20]. Finally, the model can exhibit hybrid mode switches due to the action of renal clearance that is typically turned on only when $G(t) \geq G_r$, a renal clearance threshold parameter (typically 180mg/dl).

For our study, we use the Dalla-Man et al. model, *ibid*. This model is a nonlinear ordinary differential equation (ODE) with 10 state variables. The model and corresponding parameters are available as part of the FDA approved T1DM simulator that can now be used as an alternative to animal testing [42]. The model has been increasingly popular inside a simulation environment for "in-silico" or "virtual" clinical trials [48, 40].

Closed Loop Simulation: Closed loop simulation is performed by simulating the patient physiological model in composition with the controller. The overall closed loop model follows the schematic in Figure 1. The inputs to the closed loop simulation include: (a) the initial state of the patient physiological model, (b) the timing of meals and their carbohydrate content, (c) the noise in CGM readings and (d) a set of parameters for the physiological model that are representative of a particular patient's insulin-glucose regulation. Factors such as exercise are not supported by the Dalla Man et al. model used in our simulation. However, a more recently proposed model incorporates exercise and the action of glucagon (counter-regulation) [42].
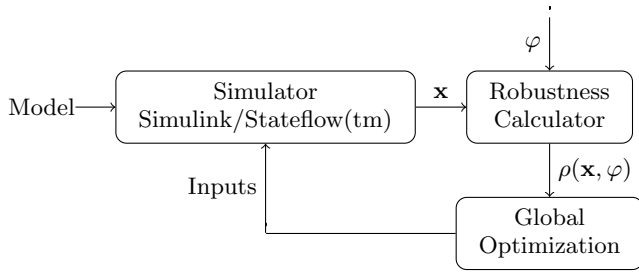
Figure 4: Illustration of the overall robustness-guided falsification setup.

Figure 5: Timeline of simulated events for the closed loop simulation.

5. ROBUSTNESS-GUIDED TESTING

In this section, we briefly describe robustness guided falsification approach to testing properties of closed-loop control systems. Numerous details that are skipped here are available elsewhere [1, 11]. The presentation of robustness-guided falsification below has been excerpted from our previous survey on this topic [11].

Given a mathematical model $\mathcal{M}$ and a property P over its outputs, are there inputs to the model whose outputs can violate P? To answer this problem, model checking approaches search over the space of all possible inputs, stopping when a violation is found [5, 17]. However, in many cases, the models are *infinite state* making the process of exhaustively simulating all inputs quite expensive, if not impossible. As a result, many approaches have been proposed to examine inputs that are "promising" while avoiding inputs that are "unlikely" to yield violations.

Robustness-guided falsification approaches are based on two main ideas: (a) A distance metric from an output trace to a property violation [26, 50, 22]. Such a metric is referred to as the "trace robustness". Intuitively, a trace with a smaller robustness is therefore "closer" to a violation when compared to a trace that has a larger robustness. (b) The robustness metric is used as an objective function to guide the system towards property violations in a systematic manner by seeking trajectories of ever decreasing robustness [45, 1, 3]. This is typically solved using heuristic global optimization algorithms such as simulated annealing [45, 1], ant-colony optimization [2], genetic algorithms and the cross-entropy method [51]. If these techniques discover a negative robustness trace, then a property violation is concluded. Otherwise, the least robust trace often provides valuable information to the designer, as to how close we get towards violating the property.

5.1 S-Taliro Tool

Figure 4 shows a schematic diagram for S-Taliro [1], a robustness guided falsification tool that supports MTL properties [3]. S-Taliro has been implemented inside the Matlab (tm) environment, and can support models described inside Simulink/Stateflow (tm). The tool uses the inbuilt simulator and computes the robustness for a trace. The resulting robustness is used as an objective function by a global optimization engine that seeks to minimize this value. The global optimizer, in turn, decides on future test inputs to the simulator based on the past inputs and the robustness values

of the resulting traces. Currently, the tool supports many optimization engines including uniform random exploration, simulated annealing search, ant-colony optimization, cross-entropy method and genetic algorithms. Since no single optimization engine can guarantee finding a global minimum, the typical practice of using the tool consists of using multiple optimization engines, repeatedly and in parallel. If the tool fails to discover a violation, one of the key advantages of robustness metrics is that the least robust trace can provide a relaxed property that can be violated by S-Taliro. S-Taliro is available as an open source tool [2], and is built to be extensible through the addition of new solvers and alternative robustness computation techniques. The latest version uses multiple cores to perform numerous simulations in parallel. It also supports features such as property-directed parameter tuning for models and requirements. These features will be enhanced in future releases of the tool.

6. PROPERTIES

We now describe the overall setup for our closed-loop *in silico* study. We describe the use of S-Taliro to search for violations of key properties. We report the results of S-Taliro alongside each property.

Study Setup: The overall timeline of events for the simulation is shown in Figure 5. The simulation setup models a common usage scenario, wherein the user is assumed to eat a meal (dinner) and a snack. The meal timings and amount of carbohydrates (CHO) vary over a range, as specified in Table 3. The simulation also models the user's open loop bolusing behavior by selecting a bolus time which can be anywhere between 20 min before the meal or up to 20 min after the meal. The insulin to CHO ratio used to calculate the insulin bolus amount is also varied in a range. Finally, we assume a fixed controller starting time when the closed loop is switched on and include a range of values for the sensor noise, also shown in Table 3. In particular, we also perform an open loop simulation wherein the controller is never turned on, in contrast with a closed loop simulation, wherein the controller is turned on at $t = 50$min.

S-Taliro Setup: S-Taliro takes the overall closed loop model and searches over the space of inputs from Table 3 for property violations. The tool formulates a total of 127 inputs that includes 120 sensor noise values. Rather than find a single violation, we repeatedly ran the tool up to 7 times for each property, stopping each run when a violation is discovered. This allows us to discover multiple violations. We simply used the uniform random search heuristic that blindly samples from the set of violations. We observed that simulated annealing was less effective for this benchmark.

---

[1]S-Taliro stands for System TemporAl LogIc RObustness

[2]Cf. https://sites.google.com/a/asu.edu/s-taliro/s-taliro

Table 3: Inputs to the closed-loop simulator.

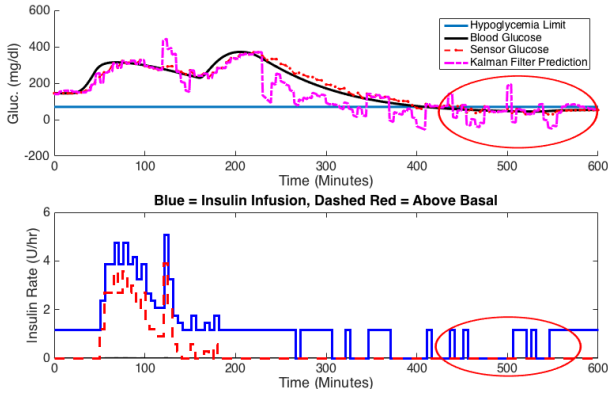| INPUTS	RANGE
Meal #1 Time	$[20, 40]$min
Meal # CHO	$[60, 150]$gms
Meal #2 Time	$[180, 300]$min
Meal #2 CHO	$[0, 60]$gms
Insulin Bolus Delta	$[-20, 20]$min
Insulin-CHO Ratio	$[0.05, 0.2]$U/gm
Open Loop Basal	$[0.01, 0.1]$U/hr
Controller Start Time	$\{50$min$\}$
Sensor noise (~ 120 inputs)	$[-20, 20]$mg/dl

Figure 6: Least robust trace showing violation of property **P1.1**. The basal insulin is resumed when $G(t) \leq 70$mg/dl. The solid (blue) curve on bottom plot the insulin delivered to the patient as sum of the original basal insulin plus the controller commands that can infuse the basal insulin, add extra insulin over the basal rate, or shutoff delivery. The red dotted line shows the extra insulin commanded by the controller. Pump shutoff occurs whenever the total insulin infusion is zero.

The relative large search space is one possible reason for this.

6.1 Properties and S-Taliro Results

We will now describe the classes of properties that we wish to test for the overall closed loop. Note that the properties are meant to expose and understand corner case behaviors of the closed loop. In other words, we currently lack information as to the likelihood of the violations in realistic usage scenarios. However, once understood, it will be essential to find fixes/mitigations for the likely violations described in this section.

Control during Low Glucose Levels: An important objective of the control algorithm is to reliably turn off insulin delivery in advance of an impending hypoglycemia. As a result, important questions include whether the controller will resume basal insulin delivery or even infuse additional insulin when the patient is already under hypoglycemia?

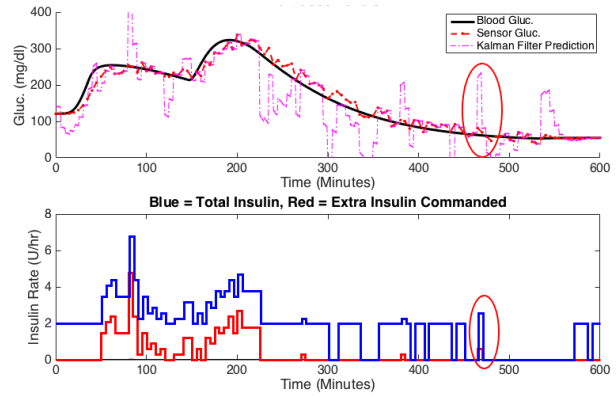P1.1: Is it possible for basal insulin to be resumed when $G(t) \leq 70$mg/dl while the total shutoff time and the

Figure 7: Least robust trace showing a "near violation" of property **P1.2**. Extra insulin is commanded, for 5 minutes at $G(t) \sim 90$mg/dl.

shutoff time within the current time window are still below their upper limits?

It is important to specify that the shutoff times so far are under the maximum permitted limit since the pump will resume automatically when these limits have been exceeded.

S-Taliro ran for nearly 2 hours and 5 minutes and found 5 violations. Figure 6 shows the glucose and insulin for the violation. The circled region shows the violation wherein insulin delivery is resumed even under hypoglycemia. The noise pattern in the CGM affects the future glucose prediction (shown in magenta in Fig. 6) and causes the delivery to resume. Such a scenario can be potentially addressed by either (a) adjusting the gains for the Kalman filter to be more robust to noise and/or (b) requiring glucose levels to cross a minimal threshold before resuming insulin delivery.

P1.2: Is it possible for additional insulin to be commanded when $G(t) \leq 80$mg/dl.

S-Taliro ran for nearly 7.5 hours, performing 750 simulations. It could not violate this property. Nevertheless, we find an interesting near violation, shown in Fig. 7. It demonstrates the infusion of extra above basal insulin commanded by the controller. However, this happens around $G(t) \sim 90$mg/dl. Also, the command is for a very short time period. Also, note that the brief rise in CGM values due to the added disturbances coincides with the additional insulin commanded.

P1.3: Is it possible for additional insulin input to be commanded and subsequently the pump shutoff within 30 minutes?

Whereas, a violation of this property is not a safety violation, it provides us insights into how the process of deciding on extra insulin can interact with the pump shutoff. S-Taliro ran for nearly 2.2 hours, finding nearly 4 violations of this property. Figure 8 shows the least robust violation. Interestingly, we notice that a temporary glitch caused by the

Figure 8: Least robust trace showing a "violation" of property P1.3. Extra insulin is commanded at the point of impending hypoglycemia and the pump shutdown almost immediately.
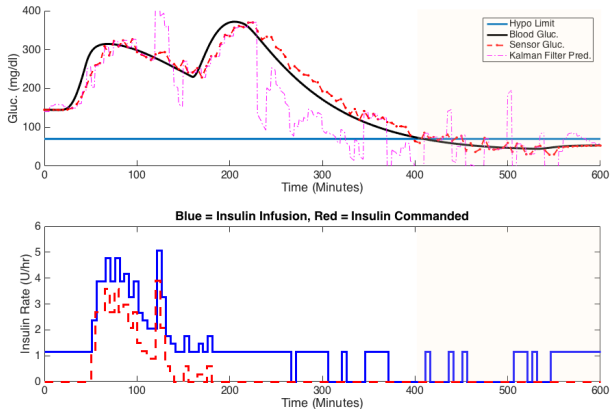

Figure 9: Least robust trace showing a violation of property P1.4. The user enters hypoglycemia around $T = 400$ minutes. The pump is shutoff for less than 50% of this time.

CGM noise causes the controller to briefly command a small amount of extra insulin above basal and shut off the pump *immediately*. The scenario can be potentially addressed by (a) adding an additional rule that would require a minimal threshold for CGM value before additional insulin is commanded and/or (b) adjusting the Kalman filter gain values.

P1.4: Let γ be the ratio of total pump shutoff time divided by the total time under hypoglycemia. Can $\gamma \leq 0.7$? In other words, will the pump be shutoff for less than 70% of the time under hypoglycemia?

S-Taliro ran for nearly 3 hours and 20 minutes, finding 5 violations. The least robust violation is shown in Fig 9.

Control during High Glucose Levels: We now examine a list of questions about the controller behavior under blood glucose levels under hyperglycemia ($G \geq 180$mg/dl) and extreme hyperglycemia ($G \geq 300$mg/dl).

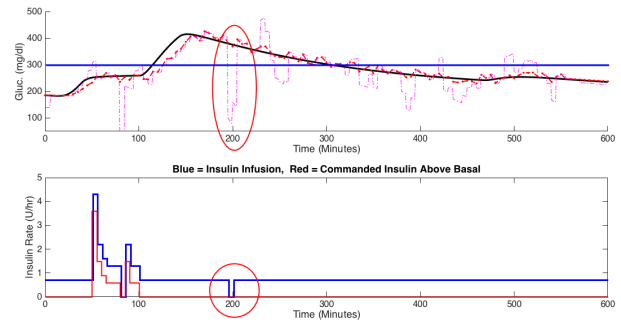**Figure 10: A single trace that violates properties P2.1, P2.2 and P2.3. The pump shuts down for a short interval of 5 minutes when $G \sim 400$mg/dl, the blood glucose levels are in hyperglycemia for almost 95% of the total simulation time and the time above 300mg/dl exceeds 3 hours.**

P2.1: Can the pump be shutoff when $G \geq 300$mg/dl?

P2.2: Can the total time under hyperglycemia $G \geq 180$mg/dl exceed 70% of the total simulation time?

P2.3: Can the total time under hyperglycemia $G \geq 300$mg/dl exceed 3hrs?

S-Taliro ran for nearly 1 hour and 6 minutes to discover 5 violations for property P2.1 The least robust violation is shown in Fig 10. Note that the pump is shutoff for a small duration of 5 minutes while the blood glucose $G(t) \sim 400$mg/dl. Interestingly, we found that all three properties are violated by this single trace! The pump shutdown results due to sensor noise that causes it to shutdown. However, the extended time under hyperglycemia is a result of inadequate meal bolus. The controller does not infuse enough extra insulin to rectify this situation for this simulation.

Comparing Closed and Open Loop Performance: A key class of properties involve questions about open vs. closed loop performance for identical meals, insulin bolus, basal insulin levels and starting patient physiological state. Such a comparison is hard, if not impossible, in a clinical setting, but possible *in silico*, due to mathematical models.

P3.1: Is it possible for the closed loop *hypoglycemia* whereas the open loop blood glucose value remains above 80mg/dl?

S-Taliro ran on this property for nearly 4.5 hours, yielding 2 violations. The least robust violation is shown in Figure 11.

P3.2: Is it possible for the closed loop *hyperglycemia* above 300mg/dl whereas the open loop blood glucose
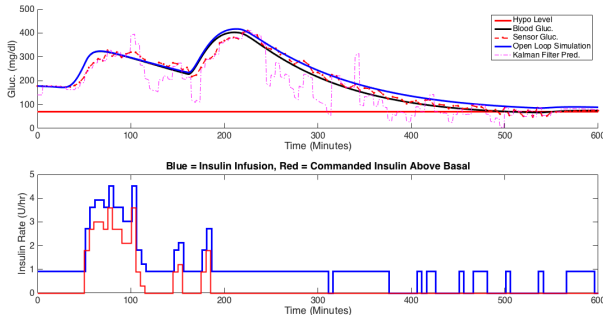
Figure 11: Violation of property P3.1 showing the closed loop simulation under hypoglycemia whereas the open loop stays well above the hypoglycemic limit.


Figure 12: Violation of property P3.3 showing the open loop simulation having a longer time in range. Note that starting from $t \sim 450$ mins, the closed loop simulation enters a prolonged hypoglycemia with $G(t) \sim 70$mg/dl.

---

level remains below 180mg/dl? In practice, it is necessary to check for ketones when $G \geq 300$mg/dl, making it an important limit.

S-Taliro could not obtain a violation for this property even after running S-Taliro for 10 hours during which about 700 simulations were performed.

P3.3: Let ρ represent the ratio of time in range for the closed loop vs. time in range for the open loop. Is it possible for $\rho < 0.7$?

After running for 8 hours, S-Taliro discovers 8 violations, the least robust of which is shown in Figure 12.

6.2 Discussion

At a high level, we find violations to 8 of the 10 properties and a near violation for one more. It must be noted that the CGM noise pattern seems to be the single most important cause of these violations. However, our simulation currently uses a coarse range of $[-20, 20]$mg/dl as bounds on the error for each CGM readings, allowing S-Taliro to choose any error value in this range. Our future work will use existing clinical

data that contrasts CGM data with *gold standard* blood glucose data from instruments such as YSI glucose meters. Incorporating more realistic constraints on the CGM noise patterns introduced by S-Taliro is one way to ensure that the results are valid.

Next, we note that the lack of modeling for counter-regulatory processes makes the physiological model potentially less accurate for hypoglycemia. Incorporating more recently proposed models that incorporate counter-regulation is an important next step.

Finally, we note that simulations carried out by S-Taliro are quite expensive. We are investigating the process of parallelizing the simulations to gain efficiency and perform much larger number of simulations.

7. CONCLUSIONS

In conclusion, we have shown a case-study of a hypo/hyper mitigating controller using S-Taliro to identify corner case property violations. In doing so, we formulated 10 properties, discovering 8 violations and one scenario that comes close to a violation. We also identify possible solutions to mitigating some of these violations by modifying the algorithm. However, before we do so, it is essential to assign likelihood scores to these violations. Such scores will allow us to better triage these violations and prioritize the fixes. Building such scores will require us to design approaches to precisely identify the "root causes" and use data to estimate their likelihood. Another important advance involves the automatic tuning of system parameters to fix these violations. Promising approaches that use sensitivity analysis [22] and robustness of temporal properties have been proposed for this problem [6].

8. REFERENCES

[1] H. Abbas, G. Fainekos, S. Sankaranarayanan, F. Ivancic, and A. Gupta. Probabilistic temporal logic falsification of cyber-physical systems. *Trans. on Embedded Computing Systems (TECS)*, 12:95–, 2013.

[2] 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.

[3] 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.

[4] E. Atlas, R. Nimri, S. Miller, E. A. Grunberg, and M. Phillip. MD-Logic artificial pancreas system: A pilot study in adults with type 1 diabetes. *Diabetes Care*, 33(5):1072–1076, May 2010.

[5] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.

[6] E. Bartocci, L. Bortolussi, L. Nenzi, and G. Sanguinetti. Systematic design of stochastic models

using robustness of temporal properties. *Theoretical Computer Science*, 587:3 – 25, 2015.

[7] R. N. Bergman. Minimal model: Perspective from 2005. *Hormone Research*, pages 8–15, 2005.

[8] 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.

[9] F. Cameron. *Explicitly Minimizing Clinical Risk through Closed-loop Control of Blood Glucose in Patients with Type 1 Diabetes Mellitus*. PhD thesis, Stanford University, 2010.

[10] F. Cameron, B. W. Bequette, D. Wilson, B. Buckingham, H. Lee, and G. Niemeyer. Closed-loop artificial pancreas based on risk management. *J. Diabetes Sci Technol.*, 5(2):368–79, 2011.

[11] F. Cameron, G. Fainekos, D. M. Maahs, and S. Sankaranarayanan. Towards a verified artificial pancreas: Challenges and solutions for runtime verification. In *Proceedings of Runtime Verification (RV'15)*, volume 9333 of *Lecture Notes in Computer Science*, pages 3–17, 2015.

[12] F. Cameron, G. Niemeyer, and B. W. Bequette. Extended multiple model prediction with application to blood glucose regulation. *Journal of Process Control*, 22(8):1422–1432, Sep 2012.

[13] F. Cameron, D. M. Wilson, B. A. Buckingham, H. Arzumanyan, P. Clinton, H. P. Chase, J. Lum, D. M. Maahs, P. M. Calhoun, and B. W. Bequette. Inpatient studies of a kalman-filter-based predictive pump shutoff algorithm. *J. Diabetes Science and Technology*, 6(5):1142–1147, 2012.

[14] 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.

[15] F. Chee and T. Fernando. *Closed-Loop Control of Blood Glucose*. Springer, 2007.

[16] S. Chen, M. O'Kelly, J. Weimer, O. Sokolsky, and I. Lee. An intraoperative glucose control benchmark for formal verification. In 5^{th} *IFAC conference on Analysis and Design of Hybrid Systems (ADHS)*, Oct 2015.

[17] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 1999.

[18] Claudio Cobelli et al. and AP@Home Consortium. First use of model predictive control in outpatient wearable artificial pancreas. *Diabetes Care*, 37(5):1212–1215, May 2014.

[19] 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.

[20] 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.

[21] E. Dassau, H. Zisser, H. R.A., P. M.W., B. Grosman, W. Bevier, E. Atlas, S. Miller, R. Nimri, L. Jovanovic, and D. F.J. Clinical evaluation of a personalized artificial pancreas. *Diabetes Care*, 36(4):801–9, 2013.

[22] A. Donzé and O. Maler. Robust satisfaction of temporal logic over real-valued signals. In *FORMATS*, volume 6246 of *Lecture Notes in Computer Science*, pages 92–106. Springer, 2010.

[23] F. El-Khatib, J. Jiang, and E. R. Damiano. Adaptive closed-loop control provides blood-glucose regulation using dual subcutaneous insulin and glucagon infusion in diabetic swine. *J Diabetes Sci Technol.*, 1(2):181–92, 2007.

[24] F. H. El-Khatib, S. J. Russell, D. M. Nathan, R. G. Sutherlin, and E. R. Damiano. A bihormonal closed-loop artificial panceras for type 1 diabetes. *Sci. Transl. Med.*, 2, April 2010.

[25] 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.

[26] G. Fainekos and G. J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science*, 410:4262–4291, 2009.

[27] G. E. Fainekos. *Robustness of Temporal Logic Specifications*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, 2008.

[28] G. P. Forlenza, S. Sankaranarayanan, and D. M. Maahs. Refining the closed loop in the data age: Research-to-practice transitions in diabetes technology. *Diabetes Technology & Therapeutics*, 17(5), 2015.

[29] B. Grosman, E. Dassau, H. Zisser, L. Jovanovic, and D. F.J. Zone model predictive control: A strategy to minimize hyper- and hypoglycemic events. *J Diabetes Sci Technol.*, 4(4):961–75, 2010.

[30] R. Hovorka. Continuous glucose monitoring and closed-loop systems. *Diabetic Medicine*, 23(1):1–12, 2005.

[31] 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.

[32] R. Hovorka, V. Canonico, L. Chassin, U. Haueter, M. Massi-Benedetti, M. Frederici, 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.

[33] 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.

[34] Z. Jiang, M. Pajic, R. Alur, and R. Mangharam. Closed-loop verification of medical devices with model abstraction and refinement. *International Journal on Software Tools for Technology Transfer*, pages 1–23, 2013.

[35] Z. Jiang, M. Pajic, S. Moarref, R. Alur, and R. Mangharam. Modeling and verification of a dual chamber implantable pacemaker. In *Tools and*

Algorithms for the Construction and Analysis of Systems (TACAS), volume 7214 of *Lecture Notes in Computer Science*, pages 188–203. 2012.

[36] B. Kovatchev, C. Cobelli, E. Renard, S. Anderson, M. Breton, S. Patek, W. Clarke, D. Bruttomesso, A. Maran, S. Costa, A. Avogaro, C. Dalla Man, A. Facchinetti, L. Magni, G. De Nicolao, J. Place, and A. Farret. Multinational study of subcutaneous model-predictive closed-loop control in type 1 diabetes mellitus: summary of the results. *J Diabetes Sci Technol.*, 4(6):1374–81, 2010.

[37] A. Kowalski. Pathway to artificial pancreas revisited: Moving downstream. *Diabetes Care*, 38:1036–1043, June 2015.

[38] R. Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.

[39] D. M. Maahs, P. Calhoun, B. A. Buckingham, and Others. A randomized trial of a home system to reduce nocturnal hypoglycemia in type 1 diabetes. *Diabetes Care*, 37(7):1885–1891, July 2014.

[40] L. Magni, D. Raimondo, L. Bossi, C. D. Man, G. D. Nicolao, B. Kovatchev, and C. Cobelli. Model predictive control of type 1 diabetes: an *in silico* trial. *J. Diabetes Science and Technology*, 1(6):804–12, 2007.

[41] C. D. 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.

[42] C. D. Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli. The UVA/PADOVA type 1 diabetes simulator: New features. *J. Diabetes Science and Technology*, 8(1), January 2014.

[43] 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.

[44] Medtronic Inc. "paradigm" insulin pump with low glucose suspend system, 2012. Cf. `http://www.medtronicdiabetes.ca/en/paradigm_veo_glucose.html`.

[45] 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.

[46] R. Nimri, I. Muller, E. Atlas, S. Miller, O. Kordonouri, N. Bratina, C. Tsioli, M. Stefanija, T. Danne, T. Battelino, and P. M. Night glucose control with md-logic artificial pancreas in home setting: a single blind, randomized crossover trial-interim analysis. *Pediatric Diabetes*, 15(2):91–100, March 2014.

[47] M. Pajic, R. Mangharam, O. Sokolsky, D. Arney, J. Goldman, and I. Lee. Model-driven safety analysis of closed-loop medical systems. *Industrial Informatics, IEEE Transactions on*, 10(1):3–16, Feb 2014.

[48] 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.

[49] Y. Pei, E. Entcheva, R. Grosu, and S. Smolka. Efficient modeling of excitable cells using hybrid automata. In *Proc. Computational Methods in Systems Biology*, pages 216–227, 2005.

[50] 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.

[51] 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.

[52] G. M. Steil. Algorithms for a closed-loop artificial pancreas: The case for proportional-integral-derivative control. *J. Diabetes Sci. Technol.*, 7:1621–1631, November 2013.

[53] R. E. Teixeira and S. Malin. The next generation of artificial pancreas control algorithms. *J. Diabetes Sci. and Tech.*, 2:105–112, Jan 2008.

[54] 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.