A Model-Based Approach to Synthesizing Insulin Infusion Pump Usage Parameters for Diabetic Patients

Sriram Sankaranarayanan, Christopher Miller, Rangarajan Raghunathan, Hadi Ravanbakhsh and Georgios Fainekos

Abstract— **We present a model-based approach to synthesizing insulin infusion pump usage parameters against varying meal scenarios and physiological conditions. Insulin infusion pumps are commonly used by type-1 diabetic patients to control their blood glucose levels. The amounts of insulin to be infused are calculated based on parameters such as insulin-to-carbohydrate ratios and correction factors that need to be calibrated carefully for each patient. Frequent and careful calibration of these parameters is essential for avoiding complications such as hypoglycemia and hyperglycemia.**

In this paper, we propose to synthesize optimal parameters for meal bolus calculation starting from models of the patient's insulin-glucose regulatory system and the infusion pump. Various off-the-shelf global optimization techniques are used to search for parameter values that minimize a penalty function defined over the predicted glucose sensor readings. The penalty function "rewards" glucose levels that lie within the prescribed ranges and "penalizes" the occurrence of hypoglycemia and hyperglycemia. We evaluate our approach using a model of the insulin-glucose regulatory system proposed by Dalla Man et al. Using this model, we compare various strategies for optimizing pump usage parameters for a virtual population of in-silico patients.

I. Introduction

Insulin infusion pumps are commonly used by type-1 diabetic patients to control their blood glucose levels. These pumps supply insulin at programmable rates over time. Typically, the use of insulin infusion pumps has two components: (a) continuous background infusion provided at a fixed *basal rate* to offset the endogenous glucose production, and (b) a fixed amount of insulin bolus provided to cover elevated glucose levels, especially after a meal. The basal rate is set by trial and error until the level of glucose remains steady during fasting conditions (eg., overnight). Likewise, the bolus dosage is decided by a fixed insulin to carbohydrate ratio (icRatio) and a correction factor (Cor). The parameter icRatio is used to calculate the amount of insulin bolus required to address the increase in blood glucose levels following a meal, based on the amount of carbohydrates in the meal. Likewise, the parameter Cor can be used to help reduce higher than desired blood glucose levels.

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

S. Sankaranarayanan C. Miller, R. Raghunathan and H. Ravanbakhsh are with the Department of Computer Science, University of Colorado, Boulder. E-mail: last.firstname@colorado.edu.

G. Fainekos is with the School of Computing, Informatics and Decision Systems Engineering at Arizona State University. E-mail: fainekos@asu.edu.

In this paper, we propose the use of mathematical models of the insulin glucose regulatory system to find ideal values of the basal rate, the insulin-carbs and the correction factor. Mathematical models of the insulin glucose regulatory system are quite sophisticated and can capture many key physiological processes that govern the insulin glucose regulation [1], [2], [3]. Furthermore, these models include a large number of parameters that can potentially be adjusted to fit the data available for an individual patient including their carbohydrate intake, insulin infusion and blood glucose readings over an extended period of time. The models have the potential to account for short-term and medium-term changes in insulin sensitivity and changes to the physical fitness levels that can require changes in the pump usage parameters.

Assuming the availability of a mathematical model with parameters fitted to a particular patient, we consider the problem of finding optimal parameters for insulin infusion pump usage. The criterion for finding the parameters include the absence of hypoglycemia (the glucose concentration remains above a minimum value), the absence of hyperglycemia (the blood glucose concentration remains below a maximum value) and the settling of the glucose concentration within a narrow range roughly 3 hours post-meal. Our approach involves the formulation of a penalty function that measures the undesirability of the blood glucose concentrations over time resulting from a scenario with fixed values of the pump usage parameters, the amount of meal carbohydrates and the starting value of the blood glucose concentration. Starting with a natural penalty function that measures the "robustness" of the glucose concentrations w.r.t. correctness properties specified in Metric Temporal Logic (MTL) [4], we modify the robustness metric to provide proper weightage to hypoglycemia, which is much more undesirable than hyperglycemia. Furthermore, we penalize hypo-/hyper-glycemia that persist for a long amount of time as opposed to transient violations.

Finally, we define optimization problems to discover pump usage parameters that minimize the penalty objective. However, the resulting objective function cannot be expressed in a convenient closed form. In this paper, we use various heuristic global optimization techniques such as simulated annealing, cross-entropy and genetic algorithms [5].

We present an implementation of the setup based on the model of insulin-glucose regulation proposed by Dalla Man et al. [3]. This model is part of the commercially available UVa-Padova simulator that was originally designed

for testing control algorithms for the artificial pancreas concept [1]. Our simulation environment includes models of commonly recommended pump usage strategies governed by the parameters to be optimized. The simulation environment allows us to compute the objective function for the optimization of the pump usage parameters. This optimization is performed under various configurations for a virtual set of in-silico patients available as part of the simulator. We conclude that our approach is viable for synthesizing pump usage parameters automatically in a relatively short amount of time. Our approach can be incorporated into tools that can analyze a patient's pump usage logs and automatically recommend pump usage parameters for patient.

A. Related Work

Our ongoing project on *robustness-guided model checking* studies the use of optimization algorithms for finding input signals and parameter values to falsify correctness properties of system designs [6], [7]. The tool S-Taliro provides an implementation of these ideas to analyze MTL properties of Simulink/Stateflow (tm) diagrams [8]. S-Taliro was used previously to study insulin infusion pump usage models. Therein, we examined the effect of various types of system failures and user mistakes on the occurrence of hyper- and hypoglycemia [9]. The key difference in this paper is the focus on synthesizing parameters that minimize the overall robustness. Furthermore, we modify the robustness metric to obtain a penalty function that addresses some of the characteristics of this problem.

Recently, Jha et al. presented the use of statistical model checking to tune the parameters for a Proportional-Integral-Derivative (PID) controller that regulates insulin infusion based on glucose sensor readings [10]. Their approach searches for the proportional, integral and differential gain parameters that satisfies a given set of temporal properties with a given probability at a high level of confidence. The search is guided by the number of simulations required before the statistical model checker rejects the correctness criterion. A higher number indicates likely property satisfaction. In contrast, our approach uses a robustness metric in lieu of Boolean property satisfaction to search for pump usage parameters that are optimal with respect to the chosen robustness metric. The approach of Jha et al. is applicable to the problem proposed here if additional information in the form of probability distributions over the initial physiological state and meal intake profiles are available. A detailed comparison of the two approaches will be carried out in the future.

Our work is similar in spirit to the idea of program sketching proposed originally by Solar-Lezema et al. [11]. In particular, the patient's usage strategy shown in Figure 4 can be seen as a simple program with "holes" specified by the parameters $\mathsf{basal}, \mathsf{icRatio}$ and Cor. A recent extension to sketching uses program smoothing, wherein a discrete program is modeled as a continuous function by adding noise to the program variables and computing the expected output [12]. In contrast to the work on sketching, the technique proposed here does not provide any guarantees of correctness by construction. This is primarily due to the complexity of the non-linear hybrid system model being treated here, whereas work on correct-by-construction sketching is mostly restricted to programs with linear guards and updates.

Model predictive control algorithms proposed for the artificial pancreas [13] use optimization techniques to control infusions in real time. However, the presence of delays in sensing glucose values and the action of insulin hinders the ability to control blood glucose levels in the presence of unannounced meal disturbances. The problem of retrofitting artificial pancreas with meal disturbance prediction and estimation has been studied recently by Lee et al. [14]. They report substantial improvements in the ability of their retrofitted technique to handle meal disturbances.

A number of other works study the parameter synthesis problem for biological systems [15], [16], [17]. In general, the problem is posed as follows. Given a hybrid or nonlinear dynamical system and a temporal logic specification, find the parameter ranges for which the resulting system trajectories satisfy the specification. In particular, in [16], the authors use sensitivity analysis in order to quantify neighborhoods of trajectories with the same qualitative behavior under uncertain system parameters and initial conditions. The authors in [17] study the parameter synthesis problem for discrete-time piecewise affine systems with parametric uncertainties. Rizk et al. [15] provide an alternative definition of robustness for temporal logic specifications. In addition, they use evolutionary optimization methods in order to find biochemical kinetic parameter values satisfying properties in temporal logic.

II. Background

In this section, we briefly describe models of insulin-glucose regulatory system and insulin infusion pumps. Further details are available from numerous surveys and monographs on this topic [18], [19].

A. Diabetes

The healthy human body has a sophisticated closed-loop control mechanism to maintain the level of glucose in the blood within upper and lower limits (roughly 60 mg/dl to 100 mg/dl under the fasting state). This is achieved mainly by the action of the pancreas, using the hormones insulin and glucagon. Insulin regulates blood glucose levels in many ways including the promotion of glucose uptake by the liver and skeletal muscles, the inhibition of glucagon and conversion of glucose by the fat cells.

Diabetes Mellitus is a condition wherein this control system is disrupted either by damage to the β-cells in the pancreas that secrete insulin (type-1 diabetes) or by reduced sensitivity of the cells in the body to insulin (type-2 diabetes). As a result, the blood glucose levels are chronically elevated, damaging many organs including the kidneys, eyes and nerves.

A common treatment for chronic type-1 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. This can be done through "finger stick" blood glucose monitors, or continuous glucose monitors (CGMs) that provide frequent estimates of the blood glucose concentration by measuring the subcutneous glucose concentration.

B. Insulin Infusion Pumps

Insulin infusion pumps are commonly used by type-1 diabetic patients to infuse artificial insulin subcutaneously. Commercially available pumps include numerous features, including the ability to deliver insulin at a preset rates through the day (basal insulin) and the ability to deliver a programmable amount of insulin (bolus insulin) upon request. The amount, timing and shape of the bolus can also be adjusted by the user.

The use of insulin infusion pump poses numerous challenges to the patient. Too much insulin poses the risk of *hypoglycemia*, wherein the patient's blood glucose levels are dangerously low. This condition seriously impairs the patient, causing coma in extreme situations. On the other hand, too little insulin poses the risk of *hyperglycemia*, wherein the patient's blood glucose levels are too high causing dangerous conditions such as ketacidosis. Likewise, blood glucose levels that are chronically elevated can also cause damage to important organs such as the kidneys, eyes and heart. Thus, the patient using the insulin pump has to choose appropriate basal rates and bolus amounts at the appropriate times to maintain their blood glucose level inside the ideal range.

In practice, physicians and diabetic educators use a widely accepted system based on three parameters that are calibrated for each patient individually [20]. These include: (a) the basal rate (basal), (b) an insulin-to-carbohydrates ratio ($\mathsf{icRatio}$) and (c) a correction factor (Cor). These parameters dictate the user's pump usage strategy for covering two major sources of glucose: (1) endogenous glucose production by the liver covered by the basal infusion and (2) meal glucose absorbed by the digestive system covered by bolus infusions. Finally, a bolus infusion is called for periodically to correct high blood glucose levels, as measured by a glucose monitor.

Consequently, an infusion pump is typically used as follows:

1) The patient is expected to infuse a background insulin at the basal rate basal. Basal rates can be varied depending on the time of the day and physical activity level.
2) The patient frequently measures the blood glucose level G and infuses a correction bolus using the correction factor Cor as follows:

$$\text{correction} = \mathsf{Cor} * (G - G_{\text{desired}}) \text{ if } G > G_{\text{desired}} + \Delta\,.$$

Here G_{desired} is the normal glucose level (eg., $80mg/dl$) and Δ represents a tolerance factor (eg., $\Delta = 20mg/dl$).

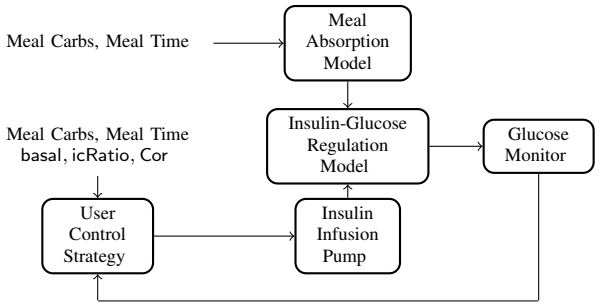

Fig. 1. Key components in the meal insulin infusion pump usage scenario.

3) Roughly ΔT minutes prior to each meal, the patient infuses a meal bolus to compensate for the amount carbohydrates due to the meal (mealCarbs). The value of ΔT depends on the meal composition. Typical values for ΔT are in the range of 15 to 20 min.

$$\text{mealBolus} = \mathsf{icRatio} * \text{mealCarbs}\,.$$

Typically, the usage parameters $(\mathsf{basal}, \mathsf{icRatio}, \mathsf{Cor})$ are calibrated individually through trial and error, starting from an initial guess given by the patient's weight and daily carbohydrate consumption. These parameters are periodically readjusted to account for longer term changes in the patient's insulin sensitivity due to factors such as illness, medications and physical fitness levels.

III. Modeling

We will now discuss the process of modeling the various parts of the insulin infusion scenario in order to optimize the pump parameters for the user. Figure 1 shows the main parts of the overall model for the insulin infusion pump usage scenario. Commonly used models of insulin glucose regulatory system use non-linear ODEs to predict the blood glucose concentration. On the other hand, the model of insulin infusion pump and the patient's usage strategies involve discrete actions. The overall composed model is a non-linear hybrid system.

A. Modeling Insulin-Glucose Regulation

In this paper, we seek to synthesize pump usage parameters using models of the patient's insulin glucose regulatory system that are periodically updated based on data that includes the patient's insulin infusion log, food intake and blood glucose data. We first briefly review the state-of-the-art for modeling the physiological processes involved in the regulation of glucose. Details on the models and the model identification process are available elsewhere [18], [19], [3], [2].

There are numerous modeling approaches for the insulin glucose regulatory system [18]. These models attempt to predict the key physiological state variables including the blood glucose levels and the insulin levels in various tissues, along with key physiological process including the action of insulin on glucose consumption, endogenous glucose

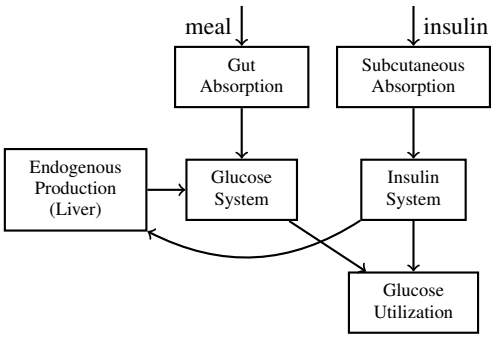

Fig. 2. Block diagram showing the major physiological processes involved in the insulin-glucose regulatory system for a type-1 diabetic patient using a subcutaneous infusion pump. Adapted from Dall Man et al. [21].

production, renal clearance of glucose, glucose secretion, gut absorption of glucose from meal and the transfer of subcutaneously infused insulin into the plasma. Figure 2 shows a block diagram that depicts these processes and the interactions between them.

Modeling approaches can be broadly classified into *minimal models* that attempt to capture the basic trends without modeling the underlying physiological processes that gives rise to these trends and *comprehensive models* that predict the blood glucose levels using models of the physiological processes mentioned above. A popular example of a minimal model is the Bergman minimal model [22]. Two recent examples of comprehensive models used in developing artificial pancreas controllers include the Hovorka model [23], [2] and the Dalla Man model [21].

We will use the Dalla Man model in our work following descriptions available elsewhere [21], [3]. This model, along with the patient parameters from the UVa-Padova simulator will be used in this paper as the reference model for insulin glucose regulation. The model is a non-linear ODE with ~ 13 continuous variables. The dynamics are hybrid due to the discrete action of renal clearance which is activated whenever the blood glucose level exceeds a threshold.

B. Modeling Insulin Infusion Pumps

Insulin infusion pumps are responsible for delivering insulin to the patient at a programmable rate. For the purposes of this paper, we model two key aspects of the insulin infusion pump: (a) the delivery of a continuous infusion at a fixed basal rate (basal) and (b) the delivery of a fixed bolus at the maximum rate, specified by the bolus amount. Our model assumes a fixed delivery profile for the bolus ("sine wave bolus") with a small width. Figure 3 shows the basic model of the pump. This model is a hybrid automaton consisting of two modes for basal and bolus delivery.

C. User Control Strategy

Another important part of our overall model is an *idealized model* of pump usage by the patient. We specifically focus on the process of controlling glucose levels surrounding a meal. Let us assume that the meal is started at time $t = T_m$

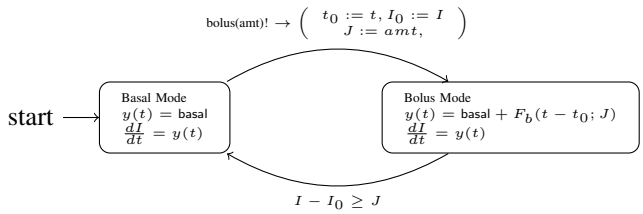

Fig. 3. Hybrid automaton modeling basal and bolus infusion delivery by an infusion pump. Variable y refers to current infusion rate, I: amount of infused so far, t: current time. The input event "bolus(amt)" is parameterized by the amount of bolus requested amt. Function F_b models the shape of the bolus as a function of time from start of bolus.

assert($t \in [T_m - 20, T_m - 15]$);
mealBolus := mealCarbs $*$ icRatio;
G := glucoseSensor(t);
if ($G > G_{\text{desired}} + 10$)
  correctionBolus := $(G - G_{\text{desired}}) *$ Cor;
else
  correctionBolus := 0;
deliverBolus(mealBolus + correctionBolus);

assert($t \in [T_m + 120, T_m + 150]$);
G := glucoseSensor(t);
if ($G > G_{\text{desired}} + 10$)
  correctionBolus := $(G - G_{\text{desired}}) *$ Cor;
else
  correctionBolus := 0;
deliverBolus(correctionBolus);

Fig. 4. An idealized user control strategy for insulin bolus centered around a meal taken at time $t = T_m$ minutes. (Top) Pre-meal and correction bolus $15-20$ minutes prior to the meal and (Bottom) post-meal correction $120-150$ minutes post-meal. Note that the strategy is dependent on the pump usage parameters basal, icRatio and Cor in addition to the pre-meal glucose levels and the estimated amount of carbohydrates in the meal.

minutes. Following the widely prescribed guidelines, we assume that the user infuses a "pre-bolus" at time $t = T_m - [15, 20]$ minutes. This combines a bolus to cover the planned meal, using the insulin-carbs ratio icRatio and a correction factor in case the pre-meal glucose levels are elevated. Similarly, we assume that the user checks their post-meal blood glucose levels roughly $[120, 150]$ minutes post meal and uses a correction bolus if required. Figure 4 shows the calculation of bolus amounts as part of the user control strategy.

D. Model Implementation

The models for the various components of the infusion were implemented inside the Simulink/Stateflow (tm) programming environment. As mentioned earlier, we implement the model proposed by Dalla Man et al. [21], [3]. An implementation is commercially available as part of the University of Virginia (UVa) -Padova simulator along with representative parameter sets for 30 "in-silico" patients. We reimplemented the publicly available Dalla Man model, as described in papers by Dalla Man and co-workers [3], [21],

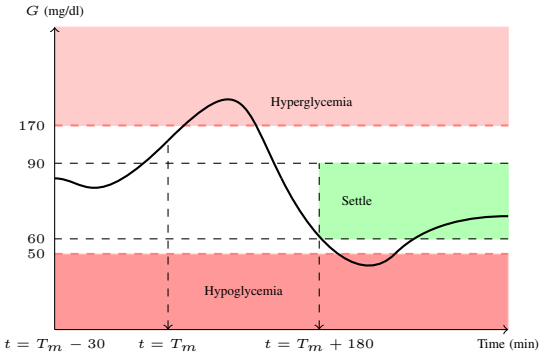

Fig. 5. Correctness requirements for ideal glycemic control at a glance. The requirements include (a) No hypoglycemia, (b) No hyperglycemia, and (c) Settle within a narrow range. Note that axes on the graph are not to scale.

[24] with in silico parameters available from the commercial simulator. This was performed in order to provide a more convenient integration of the resulting model inside the overall optimization scheme for finding pump parameters.

The final model requires a few parameters to be specified by the user at the start of the simulation, including (a) the amount of carbohydrates to be consumed during the meal (mealCarbs), (b) the starting glucose value for the simulation and (c) the pump usage parameters ($\mathsf{basal}, \mathsf{icRatio}, \mathsf{Cor}$).

The simulation involves a single meal consumed at $T_m = 30$ with a fixed duration of 15 minutes. The simulation is carried out for a total time period of $T = 720$ minutes (12 hours), assuming no meals in the intervening period. Other than the first meal at $t = T_m$, we assume no other meals during the simulation. However, the model can be easily modified to support multiple meals over a longer time period.

IV. Trace Robustness

We discuss the desired properties of the insulin infusion process and the derivation of a penalty function. These properties can be expressed as formulae in Metric Temporal Logic (MTL) [4] involving the blood glucose concentration $G(t)$ as a function of time. We then define a notion of robustness of a trace that assigns a numerical score that penalizes situations such as hypoglycemia and hyperglycemia (implicitly "rewarding" their absence). The design of the penalty function requires careful consideration of the weights given to hypoglycemia and hyperglycemia. Finally, we pose the problem of calibrating the pump usage in terms of an optimization of the trace robustness, or equivalently a minimization of the trace penalty. We show different ways of formulating this optimization to account for varying meal sizes and initial physiological states.

Figure 5 depicts the three main correctness requirements for meal insulin bolus selection. The correctness specifications refer to the observed values of blood glucose concentration over time $G(t)$. These include

(a) No hypoglycemia: $(\forall t)\ G(t) \geq 50mg/dl$,
(b) No hyperglycemia: $(\forall t)\ G(t) \leq 170mg/dl$, and
(c) The settling of the glucose concentrations to a narrow range three hours after the meal:

$$(\forall t),\ t \geq T_m + 180min \ \Rightarrow\ G(t) \in [60, 90]mg/dl\,.$$

A. Robustness

Given the correctness specifications provided in the previous section, it is an easy problem to check if a given signal $G(t)$ (specified by means of samples at discrete time points) satisfies the specification. As a result, given a trace of glucose levels $G(t)$, we may obtain a true/false answer to whether it satisfies the specifications or not. In this paper, we seek further information about "nearby" traces $G'(t)$ that are within some ϵ distance from G. We first present the notion of robustness, that augments the true/false answer obtained from the formula satisfaction with a value $\epsilon \in \mathbb{R}$ that specifies the *robustness* of $G(t)$ w.r.t the specification.

Definition 4.1 (Robustness Metric): Given a trace $G(t)$ and a property φ involving G, the robustness metric $\mathsf{rob}_\varphi(G)$ satisfies the following main properties:

1) If $G(t)$ satisfies the property φ, then $\mathsf{rob}_\varphi(G(t)) > 0$.
2) If $G(t)$ violates φ, then $\mathsf{rob}_\varphi(G(t)) < 0$.
3) Let $\mathsf{rob}_\varphi(G) = \epsilon$. Consider any trace $G'(t)$ that is contained in an ϵ cylinder around $G(t)$, i.e, for all t, $|G'(t) - G(t)| \leq |\epsilon|$. It follows that G' has the same outcome for the property φ as G. In other words, both G, G' satisfy the property or both violate the property φ.

The general theory of robustness metrics for continuous signals and MTL properties has been described by Fainekos et al. [25], [26]. Their work provides a systematic definition of the robustness metric given a trace and a property described in MTL. Since the conditions for hyperglycemia, hypoglycemia and failure to settle can be described in MTL, the work of Fainekos et al. is directly applicable to obtain a robustness metric. Based, on their work, we obtain the following functions for the robustness of a signal $G(t)$ w.r.t each of the three properties for the ideal control of post-meal blood glucose levels:

1) The robustness R_{hyper} for hyperglycemia is given by
$$R_{hyper}(G) = \max_t(170 - G(t))\,.$$
2) The robustness R_{hypo} for hypoglycemia is given by
$$R_{hypo}(G) = \max_t(G(t) - 50)\,.$$
3) The robustness R_{settle} for failure to settle is given by
$$R_{settle}(G) = \max_{t \geq T_m + 180}(\min(90 - G(t), G(t) - 60))\,.$$

The reader can verify that the robustness metrics for each of the three properties satisfy the requirements for being a robustnes metric.

Robustness metrics directly lead to a penalty function defined as

$$F(G) = -\min(R_{hyper}(G), R_{hypo}(G), R_{settle}(G))\,.$$

A large positive value of this penalty function implies a low, negative value of robustness. In turn, this implies a

"blatant" violation wherein nearby traces are also violations. Likewise, a low negative penalty implies a large positive robustness and thus a robust satisfying trace with nearby traces also satisfying the properties.

B. Modified Robustness Metrics

While robustness metrics provide a natural means for associating a numerical measure of satisfaction with a trace, they are not entirely suitable for the purposes of providing a penalty score for the trace $G(t)$ resulting from a fixed set of parameters basal, icRatio and Cor. This is due to two reasons:

1) The metric penalizes a violation based on the maximum/minimum value of $G(t)$ over a time interval. In practice, a hypoglycemia that persists for a significant period of time may be more harmful than a transient hypoglycemia for a short time period.
2) Secondly, hypoglycemia is generally deemed much more harmful than hyperglycemia. For instance $G(t) = 40mg/dl$ is a significant problem whereas $G(t) = 180mg/dl$ is only a problem if it persists for a long time. As a result, the metric needs to provide appropriate weights for hyperglycemia vs. hypoglycemia.

Weighting Hyperglycemia vs. Hypoglycemia One approach to appropriately penalizing hypoglycemia is to provide an exponential penalty for hypoglycemia. As a result, the penalty function for hypoglycemia may be defined as

$$F_{hypo}(G) = \max_t (e^{\lambda(50-G(t))} - 1)\,, \text{for some fixed } \lambda > 0.$$

The penalty functions for hyperglycemia and failure to settle are given by the robustness metric. The overall penalty is obtained by adding the various penalty functions from the three properties together.

In this model, assuming $\lambda = 0.3$, a hypoglycemic trace with $G_{\min} = 40mg/dl$ yields a penalty $e^3 - 1 \sim 20$. This is equivalent to the penalty induced by a hyperglycemic trace with $G_{max} = 190mg/dl$. However, the exponential nature of the scale ensures that a significant hypoglycemia $G_{\min} = 30mg/dl$ yields the same penalty as a significant hyperglycemia with $G_{max} = 630mg/dl$.

Integrating Violations As an alternative to the use of extremal glucose values (obtained by the use of $\max$ and $\min$ in the robustness function), we propose integrating the penalty in order to differentiate between a property violation that is corrected soon as opposed to a violation that persists. As a result, we define the penalty function for hypoglycemia as

$$F_{hypo}(G) = \int_{T_{hypo}} (e^{0.3(50-G(t))} - 1)dt$$
$$\text{where } T_{hypo} = \{t \mid G(t) \leq 50\}\,.$$

If $G(t)$ is assumed to be a continuous function of time, then the integral above is well-defined. Note that the function $F_{hypo}(G) > 0$ if a trace $G(t)$ violates the hypoglycemia property, while $F_{hypo}(G) = 0$ if $G(t)$ satisfies the property. Likewise, we may define $F_{hyper}(G)$ and $F_{settle}(G)$ by integrating the penalties over all time intervals that pertain to property violations.

$$F_{hyper}(G) = \int_{T_{hyper}} (G(t) - 170)dt,$$
$$\text{where } T_{hyper} = \{t \mid G(t) \geq 170\}$$
$$F_{settle}(G) = \int_{T_{up}} (G(t) - 90)dt + \int_{T_{down}} (60 - G(t))dt,$$
$$\text{where } T_{up} = \{t \mid t \geq 180 + T_m \ \wedge \ G(t) \geq 90\}$$
$$\text{and } T_{down} = \{t \mid t \geq 180 + T_m \ \wedge \ G(t) \leq 60\}\,.$$

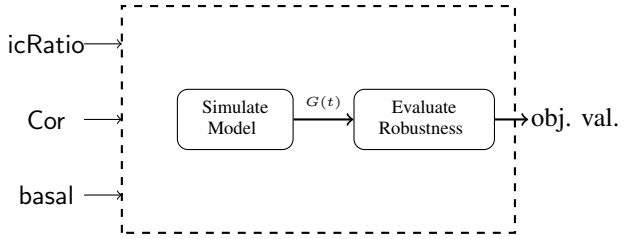

Fig. 6. The objective function $\mathcal{F}$ for the optimization problem takes in three parameters, performs a simulation and reports the penalty function value on the resulting simulation trace.

V. Optimization

We will now present the setup for the optimization problem for setting up the various pump usage parameters. We first present a joint optimization setup wherein all parameters are jointly optimized for a fixed meal scenario.

Figure 6 shows the setup for the optimization problem. The penalty function $\mathcal{F}$ has parameters (basal, icRatio, Cor). The optimization seeks to *minimize* the penalty function over a permissible range of values for the usage parameters (basal, icRatio, Cor).

Formally, let P be the set of search parameters. Namely, given a point (basal, icRatio, Cor) $\in$ P and a function Σ that returns the system trace $G(t)$ for parameters (basal, icRatio, Cor), then the optimization problem that we try to solve is

$$\min_{\vec{p} \in P} \mathcal{F}(\vec{p}) = \min_{\vec{p} \in P} F(\Sigma(\vec{p}))$$

Non-linear Objective Function: We note that the overall model of the insulin infusion process is a hybrid system with discrete mode switches induced by the insulin pump and the usage models. Furthermore, each mode has non-linear dynamics due to the physiological model. As a result, the function F cannot be written down in a closed form suitable for optimization. We can evaluate F numerically to some given degree of precision using a simulation environment such as Simulink/Stateflow (tm).

As a result, the optimization cannot be solved exactly. However, we may use heuristic techniques to obtain pump usage parameters that yield an acceptable value of the penalty function, possibly satisfying the desired properties for ideal control of glucose levels. Examples of heuristic techniques include stochastic optimization techniques such as simulated annealing and the cross-entropy method [5], heuristic global optimization techniques such as genetic algorithms, and gradient descent techniques that estimate an approximate gradient by evaluating the objective function.

Implementation The optimization routines provided in the global optimization toolbox (`fmincon`, `simulannealbnd` and `ga`) were used to carry out the optimization. The ranges of the pump usage parameters were restricted to $\mathsf{basal} \in [0.1, 5], \mathsf{icRatio} \in [0.05, .5]$ and $\mathsf{Cor} \in [0.05, 2]$ for our experiments.

VI. Experimental Results

We will now present an experimental evaluation of our approach to optimizing pump usage. We report on the results of our optimization procedure and the performance over a virtual set of 10 adult patients, namely d1 to d10, available as part of the UVa-Padova simulator [1] [1].

We consider the optimization for a meal scenario with the starting values of the blood glucose level $G_{start} = 140mg/dl$ and the amount of meal carbohydrates $\mathsf{mealCarbs} = 120gms$. Table I shows the tuned parameters for the single meal scenario. We compare two scenarios: (a) all three parameters $\mathsf{basal}, \mathsf{icRatio}, \mathsf{Cor}$ are jointly optimized for the scenario and (b) The parameters $\mathsf{icRatio}, \mathsf{Cor}$ jointly optimized with a fixed value of the basal rate that is calibrated separately. The calibration of the basal rate was performed by running an optimization over basal assuming no meal input to search for a basal rate that held the blood glucose level stable within the range $[75, 85]$mg/dl.

With the exception of three patients (d3, d6 and d7), the joint optimization produces markedly lower penalty values than the separate optimization of the basal parameter. We note that joint optimization produces acceptable maximal/minimal values of the blood glucose levels in almost all cases, with a possible severe hypoglycemia in one case and potential hyperglycemia for patient d9. The performance of separate basal optimization is slightly worse producing both severe hyperglycemia and hypoglycemia for patient d9, and hyperglycemia for patients d1, d10.

The running times for the optimization are mostly within 15 minutes. While these running times are not suitable for practical implementation, we can reduce them considerably given more efficient simulation algorithms (eg., compiling the simulator down to native code), parallel simulations and a better choice of optimization algorithm.

VII. 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 [21], providing evidence for their validity.

This work considers the joint optimization of pump parameters under a *single meal scenario* that consists of a single meal with fixed amount of carbohydrates taken at some time $t = T_m$ and no futher meal disturbances for the next 12 hours. While this scenario is feasible, a realistic scenario involves three meals at times that correspond to breakfast, lunch and dinner times. Discrepancies between announced meal carbohydrates and actual meal consumed are very common. Another limitation of the scenario is the assumption that the meal times are fixed. The choice of the pump parameters should consider some variability in the meal times. However, the framework presented here can be extended naturally to cover a more complex meal scenario.

The model is currently simulated starting from a fixed initial physiological state (blood glucose concentration, blood insulin concentration, insulin infusion history, meal history etc.). This is an unrealistic assumption in practice. Our optimization needs to consider the choice of parameters that perform well under varying physiological states prior to the meal.

Finally, our work does not consider the problem of fitting parameters to the patient's glucose monitor and insulin infusion logs, which is essential to build personalized models for the patients. In practice, model parameters are fitted using an expensive tracer study under physiological controlled conditions that cannot be carried out frequently on a patient [21]. We are currently investigating the use of parameter fitting techniques that start from an assumed prior set of parameter values measured for a patient with similar body weight and daily insulin requirements.

VIII. Conclusions

We have provided a technique to optimize insulin infusion pump usage parameters based on repeated simulations to minimize a penalty function. Our preliminary evaluation demonstrates the advantage of a joint optimization of the three pump usage parameters against a meal scenario, as opposed to the more commonly used separate optimization of the basal infusion rate. Our ultimate goal is to provide a model-based analysis tool that can fit models to patient data and use the resulting models to identify optimal pump usage parameters against meal scenarios. As noted in our discussion on threats to validity (Section VII), a lot needs to be done before such a tool can be made available to patients. Our future research will focus on piecewise affine abstractions of the non-linear model, which will enable us to simplify the dynamics of the model. The optimization of pump usage parameters against uncertain initial physiological states and multiple meal scenarios will also be investigated.

[1] The results on the entire available virtual set of 10 adults, 10 adoloscents and 10 children are available upon request

References

[1] B. P. Kovatchev, M. Breton, C. Dalla-Man, and C. Cobelli, "*In Silico* preclinical trials: A proof of concept in closed-loop control of type 1 diabetes," *J. Diabetes Science and Technology*, vol. 3, no. 1, January 2009.

[2] 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*, vol. 4, January 2010.

[3] C. Dalla-Man, D. M. Raimondo, R. A. Rizza, and C. Cobelli, "GIM, simulation software of meal glucose-insulin model," *J. Diabetes Sci. and Tech.*, vol. 1, no. 3, May 2007.

ID	PARAMS			T	PEN	minG	maxG	After $T_m + 180$ min		
	BS	IC	COR					minG	maxG	avgG
Joint Optimization of Parameters										
d1	1.42	0.13	0.08	747	6886	49	**201**	49	141	81
d2	1.57	0.16	0.2	746	744	54	140	54	117	75
d3	1.97	0.11	0.12	727	573	48	140	49	110	86
d4	1.22	0.07	0.07	946	0	60	168	61	90	80
d5	2.15	0.12	0.27	723	115	54	158	54	98	80
d6	1.8	0.15	0.08	741	15127	49	153	49	136	96
d7	1.75	0.06	0.06	770	560	49	140	49	77	68
d8	1.46	0.06	0.08	1101	554	53	154	53	101	81
d9	0.8	0.08	0.4	721	47338	**43**	184	43	182	107
d10	1.9	0.06	0.4	720	3121	48	193	48	123	81
Basal Separately Calibrated										
d1	1.93	0.09	0.06	731	12021	49	**210**	49	160	72
d2	2.13	0.16	0.08	838	1852	50	140	50	127	66
d3	2.15	0.13	0.06	664	322	49	140	50	108	76
d4	1.30	0.05	0.10	698	5	59	174	59	89	75
d5	2.44	0.10	0.26	712	310	51	161	51	102	73
d6	2.19	0.12	0.06	775	1152	50	169	50	129	72
d7	1.73	0.05	0.05	1072	87	54	140	54	82	72
d8	1.72	0.06	0.05	693	886	50	161	50	112	72
d9	1.91	0.07	0.06	696	244300	**39**	**311**	39	295	94
d10	2.28	0.13	0.19	702	4878	49	**202**	49	134	74

TABLE I

OPTIMIZED PUMP PARAMETERS AND THE PERFORMANCE. THE TABLE REPORTS THE PERFORMANCE WHEN (A) THE BASAL RATE IS OPTIMIZED JOINTLY WITH THE OTHER TWO PARAMETERS AND (B) BASAL IS CALIBRATED SEPARATELY. LEGEND: **BS**: BASAL RATE (U/HR), **IC**: INSULIN-CARBS RATIO (U/GM), **COR**: SENSITIVITY FACTOR (U/(MG/DL)), **T**: TIME TAKEN FOR OPTIMIZATION (SEC), **PEN**: PENALTY VALUE (NEGATION OF FITNESS), **MING**: MINIMAL VALUE OF BLOOD GLUCOSE (MG/DL), **MAXG**: MAXIMAL VALUE OF GLUCOSE (MG/DL) AND **AVGG**: AVERAGE VALUE OF BLOOD GLUCOSE (MG/DL).

[4] R. Koymans, "Specifying real-time properties with metric temporal logic." *Real-Time Systems*, vol. 2, no. 4, pp. 255–299, 1990.

[5] R. Y. Rubinstein and D. P. Kroese, *Simulation and the Monte Carlo Method*. Wiley Series in Probability and Mathematical Statistics, 2008.

[6] 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*. ACM Press, 2010, pp. 211–220.

[7] S. Sankaranarayanan and G. E. Fainekos, "Falsification of temporal properties of hybrid systems using the cross-entropy method," in *HSCC*. ACM, 2012, pp. 125–134.

[8] 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*, ser. LNCS, vol. 6605. Springer, 2011, pp. 254–257.

[9] S. Sankaranarayanan and G. Fainekos, "Simulating insulin infusion pump risks by *In-Silico* modeling of the insulin-glucose regulatory system," in *Computational Methods in Systems Biology (CMSB)*, ser. Lecture Notes in Computer Science, vol. 7605, 2012, pp. 322–339.

[10] 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.

[11] A. Solar-Lezama, R. M. Rabbah, R. Bodík, and K. Ebcioglu, "Programming by sketching for bit-streaming programs," in *PLDI*. ACM, 2005, pp. 281–294.

[12] S. Chaudhuri and A. Solar-Lezama, "Smoothing a program soundly and robustly," in *CAV*, ser. Lecture Notes in Computer Science, vol. 6806. Springer, 2011, pp. 277–292.

[13] 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*, vol. 25, pp. 905–920, 2004.

[14] H. Lee, B. Buckingham, D. Wilson, and W. Bequette, "A closed-loop artificial pancreas using model predictive control and sliding meal size estimator," *J. Diabetes Science and Technology*, vol. 3, no. 5, Sept 2009.

[15] A. Rizk, G. Batt, F. Fages, and S. Soliman, "Continuous valuations of temporal logic specifications with applications to parameter optimization and robustness measures," *Theor. Comput. Sci.*, vol. 412, no. 26, pp. 2827–2839, 2011.

[16] A. Donze, E. Fanchon, L. M. Gattepaille, O. Maler, and P. Tracqui, "Robustness analysis and behavior discrimination in enzymatic reaction networks," *PLoS ONE*, vol. 6, no. 9, p. e24246, 09 2011.

[17] B. Yordanov and C. Belta, "Parameter synthesis for piecewise affine systems from temporal logic specifications," in *Proceedings of the 11th international workshop on Hybrid Systems: Computation and Control*, ser. LNCS, vol. 4981. Springer, 2008, pp. 542–555.

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

[19] C. Cobelli, C. Dalla-Man, G. Sparacino, L. Magni, G. D. Nicolao, and B. P. Kovatchev, "Diabetes: Models, signals and control (methodological review)," *IEEE reviews in biomedical engineering*, vol. 2, pp. 54–95, 2009.

[20] G. Scheiner, *Think like a pancreas: A Practical guide to managing diabetes with insulin.* Da Capo Press, 2011.

[21] C. Dalla-Man, R. A. Rizza, and C. Cobelli, "Meal simulation model of the glucose-insulin system." *IEEE Transactions on Biomedical Engineering*, vol. 1, no. 10, pp. 1740–1749, 2006.

[22] R. N. Bergman, "Minimal model: Perspective from 2005," *Hormone Research*, pp. 8–15, 2005.

[23] 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.*, vol. 282, pp. 992–1007, 2002.

[24] C. Dalla-Man, M. Camilleri, and C. Cobelli, "A system model of oral glucose absorption: Validation on gold standard data," *Biomedical Engineering, IEEE Transactions on*, vol. 53, no. 12, pp. 2472 –2478, dec. 2006.

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

[26] G. E. Fainekos, "Robustness of temporal logic specifications," Ph.D. dissertation, Department of Computer and Information Science, University of Pennsylvania, 2008.