Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Sensitivity Analysis of a Predictive Pump Suspension System to Treat People with Type 1 Diabetes

Fraser Cameron*, Bruce A. Buckingham, Darrell M. Wilson**, B. Wayne Bequette***

*Rensselaer Polytechnic Institute, Troy, NY 12180 USA (e-mail: fmccamer@gmail.com)
**Stanford University, Stanford, CA 94305 USA

Abstract: The primary goal of a low glucose suspend system is to reduce the risk of overnight hypoglycemia (low blood glucose) in individuals with type 1 diabetes by reducing/suspending insulin infusion. We have developed a Kalman filter-based algorithm, combined with a number of safety rules, to implement a predictive low glucose suspend system that shuts off an insulin pump based on a prediction of hypoglycemia 30-70 minutes in the future. This system has been studied in over 2,000 nights in an outpatient-home environment. In this paper, based on an analysis of this data, we isolate the effects of the individual rules in part by simulating their removal from the existing data. Specifically, we decompose the basal insulin into small boluses and, using a model of insulin pharmacodynamic action (the time effect of insulin on blood glucose), alter the real data corresponding to the addition or removal of basal insulin via simulation. Our results show that limiting the total suspension to 180 minutes per night prevents excessive suspension in cases where the average calibration is an excessive 58 mg/dl, above the mean of 18 mg/dl. Further, we also show that a simple threshold algorithm that suspends below 100 mg/dl if the glucose level is flat or falling, is comparable in performance. Lastly, we show that the Kalman filter at the heart of this algorithm reduces the time spent below 70 mg/dl by 50% at the expense of a mean rise of 12 mg/dl in morning glucose levels.

Keywords: **Biomedical systems, Biomedical control, Pump Suspension, Diabetes, Artificial Pancreas**

1. INTRODUCTION

Type 1 diabetes (T1D) is an autoimmune disease that directly destroys the body's ability to produce insulin and indirectly the body's ability to regulate blood glucose concentrations. Individuals with T1D must use either multiple daily injections of insulin (one bolus of long-acting insulin each day, and boluses of rapid-acting insulin at meal/snack time and when blood glucose needs to be reduced), or continuous infusion of rapid-acting insulin using an insulin pump. While insulin therapy can lead to lower blood glucose levels and reduce the risk of complications due to high blood glucose, there is an underlying risk of hypoglycemia (low blood glucose) if insulin is over-administrated. Indeed, one of the greatest fears of a parent of a child with T1D is extended overnight hypoglycemia, which could lead to a coma or, in rare cases, death (known as "dead in bed" syndrome). The development of continuous glucose monitoring (CGM) technology, allowing a near continuous measurement of glucose levels, enabled the use of alarms to warn individuals of low (or high) glucose levels. Unfortunately, these alarms have been found to be insufficient, since individuals and their caregivers often sleep through the alarms (Buckingham et al., 2005).

Low glucose suspend (LGS) or pump shut-off (PSO) systems have been developed specifically to shut-off pumps to reduce the risk of hypoglycemia, based on real-time CGM signals; they are also a natural first step towards the development of a fully closed-loop artificial pancreas (Kowalski, 2009; Harvey et al., 2010; Cobelli et al., 2011; Bequette, 2012). Initial LGS systems were threshold-based (the pump is turned off when the threshold is violated), while much current effort has been on predictive low glucose suspend (PLGS) systems that turn-off a pump when a hypoglycemic event is predicted to occur (usually 30-70 minutes in the future; Bequette, 2014).

The PLGS algorithm that we have developed involves the use of a Kalman filter predict future glucose values, combined with a set of rules to reduce the risk of prolonged periods of pump shut-off. This algorithm was first tested in in-clinic studies (Cameron et al., 2012), followed by extensive out-patient (in-home) studies (Maahs et al., 2014). The objective of this paper is to analyze the results from over 2,000 nights of out-patient studies to understand the effect of various algorithm parameters and rules on the blood glucose control. We first review the results of the outpatient study, then describe our hybrid experiment/simulation approach, and finally discuss the results.

2. OUTPATIENT STUDY DATA

The data used in this paper comes from an outpatient clinical trial of the described pump suspension algorithm. The idea is that by predicting impending hypoglycemia and suspending insulin delivery that the body's natural release of glucose into

the blood stream would mitigate or prevent the hypoglycemia. For each of the 45 patients there was a run-in phase where the algorithm was enabled each night and then a 42-night phase where the algorithm was randomly and blindly enabled or disabled. This design isolated the effectiveness of the algorithm from any behavioural adjustments the patients might make for a final system. This resulted in 925 control nights and 1125 intervention nights.

The system consisted of a continuous glucose monitor (CGM) and an insulin pump communicating with a bedside laptop computer that contained the pump suspension algorithm. Each of the nights has at least 4 hours of CGM data post activation, a morning blood glucose measurement, records of any snacks from bedtime until morning, records of exercise in the previous day, basic demographics, and a full insulin history.

A typical intervention night is shown in Fig. 1. The x's indicate the glucose concentration as measured by the continuous glucose monitor and provided to the algorithm. When these trend downwards at the start of the dataset the algorithm triggers suspensions, as indicated by the orange triangles and the zeroing out of the basal rate. Later in the evening two periods of sensor noise also trigger pump suspensions. Eliminating those is the subject of further study.

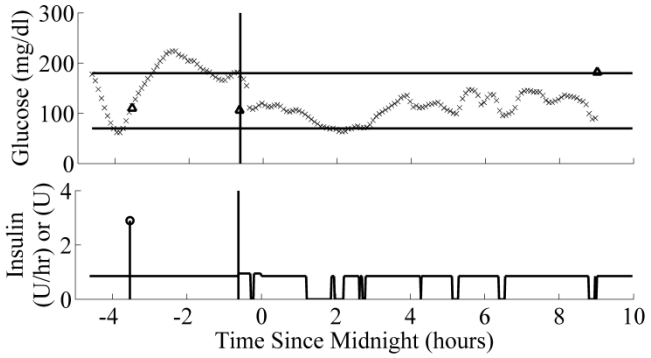

Fig 1. Sample Intervention Night. The top plot shows the CGM values (x), reference/calibration values (triangle), and bounds of desirable glucose values (horizontal line). The bottom plot shows the basal insulin (line) and boluses (bubble). The vertical line is the time when the system was activated.

3. SIMULATOR

Running a separate clinical trial to test the effect of a change to or removal of a rule would be prohibitive in terms of cost, and noisy due to inter- and intra-patient variability. Instead we simulate from the existing data. Specifically, we assume that the patients' insulin sensitivity can be calculated according to the 1800-rule:

$$IS = \frac{1800}{TDD}$$

which is a common heuristic used in clinical practice. TDD is the patients' total daily insulin dose in units/day. Then we use an average of published insulin time action profiles (Frohnhauer et al., 2001; Heinemann and Steiner, 1997; Swan et al., 2009) shown in Fig. 2, to approximate the effect of any insulin subtracted or added. The multiplication of $IS$with the curve in Fig. 2 represents the convolution model of insulin action, $\vec{I}$. Given a glucose and insulin profile represented by glucose values $\vec{g}$ and $\vec{u}$ at regular time intervals the simulator morphs the original profile into a simulated ones by looping over the following steps. 1) $u^* = f(u_{1\ldots k}, g_{1\ldots k})$ where u_i is the i$^{\text{th}}$ element of the vector $\vec{u}$ and $f(\vec{i}, \vec{y})$ is a controller that determines what new input to command given the past history of inputs $\vec{i}$ and outputs $\vec{y}$. 2) $\Delta u = u^* - u_k$ 3) $u_k = u^*$ 4) $g_{k+1\ldots \inf} = g_{k+1\ldots \inf} + \vec{I}\Delta u$ which simulates the effect of the changed insulin administration and 5) $k = k + 1$. After each repetition of the above steps the $\vec{g}$ and $\vec{u}$ vectors represent the simulated value of glucose $\vec{g}$ corresponding to the provided insulin $\vec{u}$. This simple explanation ignores issues corresponding to vector lengths and missing glucose readings that can be easily fixed in practice.

This simulator only makes assumptions about insulin action. It does not make any assumptions about meals, exercise, sleep, or anything else. Consequently, the inaccuracy of the simulator stems only from estimating the effects of large changes to the administered insulin. For the vast majority of simulated cases the glucose levels are changing only within $\pm$ 10 mg/dl, a range for which the assumption of locally linear insulin action likely holds.

An example simulation for the base algorithm and one where the prediction horizon is extended from 30 to 70 min is shown in Fig. 3. Here, the trial night begins at the vertical black line. The blue is the simulated closed-loop insulin delivery and resultant glucose values. Increasing the prediction horizon leads to earlier suspension and so higher glucose values and less hypoglycemia.

Because we can get negative values in simulation when removing particularly important rules, we use a modified risk measure that is fitted to the Kovatchev risk profile, but that allows for negative values (Cameron, 2010; Cameron et al., 2011).

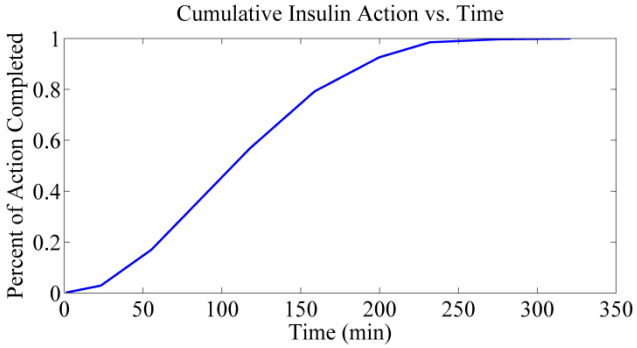

Fig. 2. Cumulative Insulin Action vs. Time

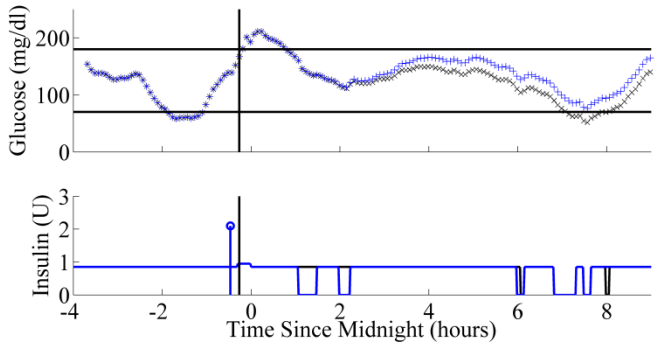

Fig. 3. Sample Simulation with Different Prediction Horizons.

4. RESULTS

4.1 Maximum Minutes of Suspension Per Night

Our first safety rule requires that the pump not be suspended more than 180 minutes in a night. This is intended to guard against the effects of substantial negative sensor bias. This assumes that patients set their insulin regimes well enough that no more than 180 minutes of suspension per night is required to prevent hypoglycemia.

Studying only the mean CGM in the 6777 total minutes after hitting the suspension limit, we see a low value of 71 mg/dl. This suggests that, as measured by the CGM setting a limit is always bad. However, looking at the comparison between the morning CGM and morning BG, our reference measure, for the nights affected by a limit of 180 minutes, we find a bias of 50 mg/dl; see Fig. 4. By comparison, across all nights the bias is 19 mg/dl. Now, the unbiased mean CGM post limit is an acceptable, if slightly high, 121 mg/dl. This in turn suggests that our current limit is preferable to no limit, which would result in an even higher mean glucose. Table 1 shows some summary results with and without this safety rule. The low and high risk measures aggregate the undesireability of the low and high CGM values, respectively. CGM values just above 112 mg/dl will add a negligible amount while those at 400 mg/dl will add substantially more to the high risk value. The low risk value does the same for values deviating below 112 mg/dl.

Table 1. Effects of Allowing no more than 180 minutes of pump suspension per night

Rule is:	Morning BG (mg/dl)	Mean Low Risk	Mean High Risk	Mean CGM (mg/dl)
Active	127.2	19.5	0.15	71.2
Inactive	132.0	18.1	0.16	84.0

To allow comparison, the morning BG numbers are all calculated using just the data affected by the nominal limit of 180 minutes per night. Fig. 4 shows, as expected, a tendency for the morning BG to drop with decreasing limits on total suspension time. It also shows that the higher the limit the more we are selecting for cases with a substantial negative bias in the CGM readings. Lastly, the larger the limit, the fewer nights are affected.

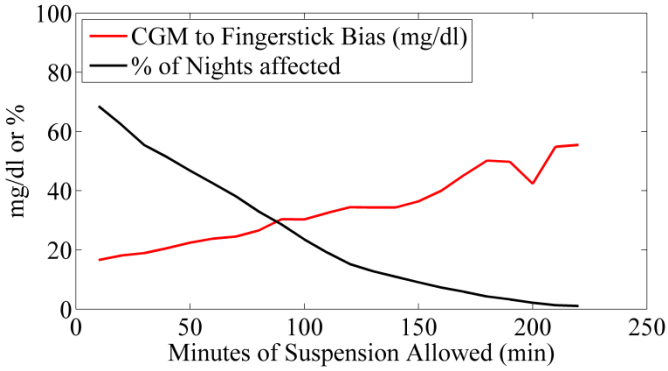

Fig. 4. Calibration Quality reaching a Maximum Suspension Limit vs. the Maximum Suspension Limit

4.2 Maximum Duty Cycle

Our second safety rule is intended to prevent ketone formation from a simple lack of insulin in the blood. When insulin is not present in the blood the body begins to break down fat to get energy, releasing acid as a by-product. This can dangerously acidify the blood. While not directly related to the glucose level, low glucose values heavily correlate with insulin action, and so insulin concentration. So, we only allowed the insulin pump to suspend for a maximum of 120 in every 150 minutes.

Generally, the intervention nights actually showed a non-significant drop of 0.2% to 0.1% of nights with blood ketones greater than 1.0 mmol/L. Studying general performance effects of this rule gives the results in Table 2. By limiting pump suspensions low glucose levels are extended, with a much smaller amount of hyperglycemia mitigated.

Table 2. Performance vs. Existence of a Duty Cycle Limit

Rule is:	Morning BG (mg/dl)	Mean Low Risk	Mean High Risk	Mean CGM (mg/dl)
Active	129.8	20.1	0.82	83.6
Inactive	132.1	19.3	0.95	85.3

Since a limit of 120 in 150 minutes is also an ad-hoc safety rule, we simulated the effect of allowing between 60 and 150 minutes during the same time period of 150 min, as shown in Fig. 5. Again, to prevent a changing dataset from confounding the results, the morning BG and percent of readings less than 70 mg/dl are simulated for just the 8,868 minutes of data affected by the nominal 120 minute limit. Sure enough, higher limits lead to more pump suspension and so less hypoglycemia and higher morning BG levels.

We do not simulate ketones. However, Fig. 5, does show the measured ketones that occurred on nights where the modified rule would have been in effect. This serves to indicate the likelihood of further suspensions when allowed. Sure enough

there is a rise in the blood ketones, with increasing limits on suspension with a knee in the curve at our nominal limit of 120 minutes.

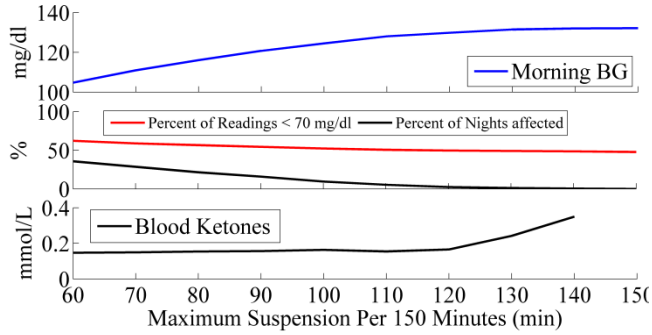

Fig. 5. Performance vs. Maximum Suspension Duty Cycle

4.3 Handling Extended Sensor Dropout

Our third rule deals with not commanding a suspension during periods where we do not have any new sensor readings. The actual implementation of this is highly situation specific, since communication with the insulin pump cannot be assumed. Here, we implement the suspensions as temporary basal rates. So, when communication is lost the pump will resume normal insulin delivery when the temporary basal rate expires. We recommend this rule to prevent trumping the patient's decisions in a situation where they are likely awake.

4.4 Do Not Suspend on Rising Glucose

This rule is intended to mitigate rebound hyperglycemia, though it comes at the cost of dampening pump suspensions before hypoglycemia. This idea assumes that there is some inertia to the second derivative of glucose; once the glucose nadir has occurred that the glucose levels will rise regardless of the control action taken. So, any extra suspension after the nadir only increases hyperglycemia instead of mitigating hypoglycemia. This makes sense given the slow response to pumped insulin, per Fig. 2. An example of this is shown in Fig. 6.

Again, we simulate the overall effects of this rule by simulating its absence. Since this rule affects CGM values on the majority of nights, we calculate the metrics for the entire dataset. Table 3 shows that while this rule does reduce the mean and morning glucose levels alongside the hyperglycemia, it does so at the cost of extra hypoglycemia. This makes this rule most appropriate for systems, like this one, where there is no capacity to mitigate hyperglycemia.

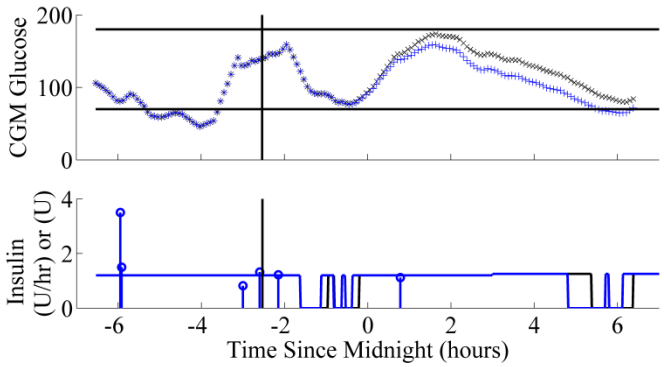

Fig. 6. Sample Night with and without Suspensions on Rising Glucose Levels.

Table 3. Performance with and without Suspension on Rising Glucose Levels

Rule is:	Morning BG (mg/dl)	Mean Low Risk	Mean High Risk	Mean CGM (mg/dl)
Active	151.2	3.1	5.8	137.2
Inactive	155.6	2.6	6.0	139.4

4.5 Automatic Pump suspension Below a Threshold

We implemented a failsafe rule that suspends the pump when the glucose level is falling and is below a threshold of 70 mg/dl. Since we also have a functional Kalman filter, eliminating this rule has little effect on the overall results. However, since other systems do not use Kalman filters, we evaluated a pure threshold system where the pump suspends if the CGM is below a threshold subject to the limits of no more than 120 of suspension in 150 minutes and for a total length of no more than 180 min. Fig. 7, shows these results for different values of the threshold, while Table 4 shows the summary metrics for a threshold of 80 mg/dl, which gives the most similar morning BG values. The numbers show that for similar morning BG values there is much more risk due to low glucose levels. This suggests that this method is much less efficient in its pump suspensions.

We then explored adding the rule where the pump only suspends on falling glucose levels. This rule is particularly applicable here since there is no prediction allowing resumption of insulin delivery for an impending glucose threshold crossing. The revised algorithm showed much lower glucose rises for a given threshold, so we chose a threshold of 100 mg/dl to best match the existing algorithm.

These results are quite similar to those for the basic algorithm and indicate that this is a valid choice for a pump suspension algorithm, with the benefits of simplicity and less susceptibility to sensor noise.

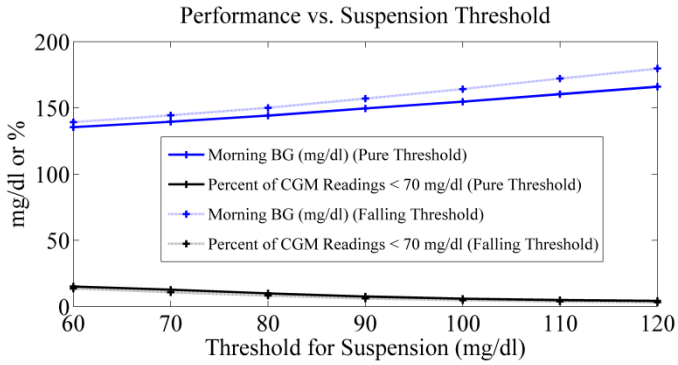

Fig. 7. Performance vs. Threshold for Threshold Based Systems.

Table 4. Comparison of Performance with Threshold Based Algorithms

Rule is:	Morning BG (mg/dl)	Mean Low Risk	Mean High Risk	Mean CGM (mg/dl)
Existing Algorithm	151.2	3.1	5.8	137.2
Falling Threshold (100 mg/dl)	154.6	2.9	5.9	138
Pure Threshold (80 mg/dl)	150.0	3.9	5.7	135

4.6 Sensor Anomaly Avoidance

Our next set of rules seeks to prevent false suspensions due to sensor anomalies. Specifically, it disallows pump suspensions if the glucose value is dropping by more than 8 mg/dl/min. 8 mg/dl/min is a non-physiologic rate of change and indicates that the sensor is, at least temporarily, not giving accurate readings.

It is difficult to evaluate the success of this type of measure without some measure of when the sensor anomalies occurred. Since we do not have enough reference glucose values to definitively identify the sensor anomalies, we leave this discussion for another paper.

4.6 Kalman Filter for Predictive Suspension

The Kalman filter is the core of this method. We suspend the insulin pump when the 30 minute glucose prediction is below 80 mg/dl and restarting when the same prediction is above 100 mg/dl. Simulating the removal of this rule shifts the overall metrics most of the way to those for the control nights. Specifically, this rule mitigates hypoglycemia with the side effects of a raised morning BG, increased mean glucose, and slightly more hyperglycemia. Even without this rule our failsafe threshold suspension still occurs, explaining some of the differences with the Control Nights. This is shown in terms of dataset wide metrics in Table 5.

Previously, we studied the effect of changing the prediction horizon in a pilot study (Buckingham et al., 2013). Here, we simulate that effect on a much larger dataset without the confounding effects of other algorithmic changes. Fig. 8 shows that changing the prediction horizon has a large effect on the 25th percentile of the time spent below 70 mg/dl, reducing it from 6.3% at a prediction Horizon of 30 min to 1% at 70 min. Also in the figure, there is only an 8 mg/dl effect on the median morning BG values between the two extremes.

Hidden from this plot is the effect of an increasing prediction horizon on our response to sensor anomalies. The longer prediction horizons amplify the sensor noise, leading to more erroneous pump suspensions.

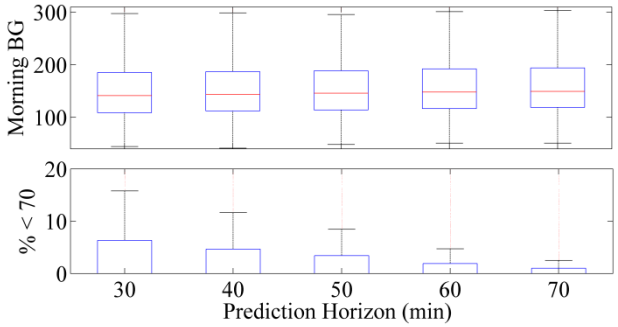

Fig. 8. Effect of Varying the Prediction Horizon.

Table 5. Effect of the Kalman Filter

Rule is:	Morning BG (mg/dl)	Mean Low Risk	Mean High Risk	Mean CGM (mg/dl)
Active	151.2	3.1	5.8	137.2
Inactive	139.4	5.8	5.4	130.7
Control Nights	132.6	5.0	5.8	133.4

5. DISCUSSION

The resounding observation from this analysis is that a much simpler system that suspends only when the glucose is falling and below 100 mg/dl does just as well as our current system. For cases where this level of aggressiveness is desired, this result suggests not using the simpler option. However, it is difficult to consider using a more aggressive version of the threshold. The chosen threshold almost guarantees that the morning glucose level will be greater than the threshold. Since patients would ideally like a morning glucose level of about 100 mg/dl, we are nearly at the limit of how aggressive the simpler system can get.

Fig 1 shows that the CGM data contains a number of sensor anomalies (around 5 AM, 6 AM), where there are extended non-physiologic fluctuations that are likely to cause false suspensions. We have developed algorithms to detect and remove the sensor anomalies characterized by starting with sharp, sudden drops. These algorithms were implemented in the system and are being tested in further clinical trials in younger patients.

Since this system has no ability to reduce high glucose levels, it is particularly cautious to avoid false suspensions. This means that we can increase the aggressiveness of hypoglycaemia mitigation if we also mitigate the high glucose levels. We have developed and are scheduled to test a system with added hyperglycemia mitigation in the spring of 2014. Specifically, we were able to remove the prohibition against suspending on rising glucose level, and to extend the prediction horizon back to 50 min.

6. CONCLUSIONS

Since this paper evaluates the effectiveness of the rules, we close with a judgement on the cost and benefit of each rule. The limit of 180 minutes of total suspension time selects well for poorly calibrated sensors. Limiting the duty cycle to a maximum of 120 minutes in every 150 minutes may not be making a significant difference in the ketone levels since this rule is applied when the patients' glucose levels are low. The use of the rule preventing suspensions on rising glucose levels is a judgement call trading off the morning glucose versus low glucose risk. In this case, where we do not have the capacity to mitigate hyperglycemia, this rule appears beneficial. The threshold rule can be easily omitted without affecting performance. Lastly, the Kalman filter does the bulk of the reduction in the risk from low glucose levels, at the expense of increased mean morning glucose levels.

7. ACKNOWLEDGEMENTS

This work is supported by funding from the Juvenile Diabetes Research Foundation (JDRF) Artificial Pancreas Project (22-2011-647, 22-2009-795 and 22-2007-1801), and the National Institutes of Health (5R01DK085591-03).

REFERENCES

Bequette BW. 2012. Challenges and Progress in the Development of a Closed-loop Artificial Pancreas. *Annual Reviews in Control* 36, 255-266.

Bequette BW. 2014. Hypoglycemia Prevention using Low Glucose Suspend Systems. Chapter 3 in Data-driven Modeling for Diagnosis and Treatment of Diabetes. Marmarelis V, Mitsis G (eds.). Springer (in press).

Buckingham B, Block J, Burdick J, Kalajian A, Kollman C, Choy M, Wilson DM, Chase HP. 2005. Diabetes research in Childrens Network: Response to nocturnal alarms using a real-time glucose sensor. Diabetes Technol Ther 7:440-447.

Buckingham, B, Cameron, F, Calhoun, P, Maahs, D, Wilson, D, Chase, HP, Bequette, BW, Lum J, Beck RW, Kollman C. 2013. Outpatient Safety Assessment of an in-Home Predictive Low-Glucose Suspend System with T1D Subjects at Elevated Risk of Nocturnal Hypoglycemia. Diabetes Technology & Therapeutics 15(8):622–627.

Cameron F. 2010. Explicitly Minimizing Clinical Risk through Closed Loop Control of Blood Glucose in Patients with Type 1 Diabetes Mellitus. PhD Thesis, Stanford University Department of Aeronautics and Astronautics.

Cameron F, Bequette BW, Wilson DM, Buckingham BA, Lee H, Niemeyer G. 2011. A Closed-Loop Artificial Pancreas Based on Risk Management. *J. Diabetes Sci. Technol*., 5(2), 368-379 (2011).

Cameron F, Wilson DM, Buckingham BA, Arzumanyan H, Benzsi K, Chase HP, Lum J, Maahs DM, Calhoun PM, Bequette BW. In-patient studies of a Kalman filter based predictive pump shut-off algorithm. J. Diabetes Sci. Technol. 2012;6(5):1142-1147.

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

Frohnauer MK, Woodworth JR, Anderson Jr JH. 2001. Graphical Human Insulin Time-Activity Profiles Using Standardized Definitions. Diabetes Tech & Ther 3 (3): 419–429.

Harvey RA, Wang Y, Grosman B, Percival MW, Bevier W, Finan DA, Zisser H, Seborg DE, Jovanovic L, Doyle FJ III, Dassau E. 2010. The quest for the artificial pancreas: Combining technology with treatment. *IEEE Eng Med Biol* March/April: 53-62.

Heinemann L, Steiner S. 2011. Subcutaneous Injection Versus Subcutaneous Infusion of Insulin: Are the Rates of Absorption Truly the Same? J Diabetes Sci Tech *2011;5*(5): 1027–9.

Kowalski AJ. 2009. Can we really close the loop and how soon; accelerating the availability of an artificial pancreas: A roadmap to better diabetes outcomes. Diabetes Technology and Therapeutics, 11(S1), S113–S119.

Maahs DM, Calhoun P, Buckingham BA, Chase HP, Hramiak I, Lum J, Cameron F, Bequette BW, Aye T, Paul HT, Slover R, Wadwa P, Wilson DM, Kollman C, Beck RW. 2014. A randomized trial of a home system to reduce hypoglycemia in type 1 diabetes. *Diabetes Care* (in press).

Swan K, Steffen A, Dziura J, Martin, M, Steil G, Tamborlane W, Voskanyan G, Weinzimer S, Sikes K. 2009. Effect of Age of Infusion Site and Type of Rapid-Acting Analog on Pharmacodynamic Parameters of Insulin Boluses in Youth With Type 1 Diabetes Receiving Insulin Pump Therapy. Diabetes Care 32 (2): 240–244.