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

Safety, constraints and anti-windup in closed-loop anesthesia

K. van Heusden * **N. West** ** **A. Umedaly** **
J.M. Ansermino ** **R.N. Merchant** ** **G.A. Dumont** *

* *Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC , Canada*
** *Department of Anesthesiology, Pharmacology & Therapeutics, University of British Columbia, Vancouver, BC, Canada*

Abstract: Feasibility of closed-loop anesthesia has been shown in a number of clinical studies. Demonstration of patient safety will be essential to convince regulatory authorities of the benefits of such systems. This paper considers safety constraints for closed-loop propofol anesthesia based on its therapeutic range. Simulation scenarios are proposed for evaluation of control strategies in the presence of these constraints. The scenarios reproduce realistic situations encountered in clinical practice. Using the proposed scenarios, the performance of $\mathcal{L}_2$ anti-windup is compared to sliding mode reference conditioning and to back-calculation anti-windup. It is concluded that $\mathcal{L}_2$ anti-windup might not be appropriate for this problem. The sliding mode solution results in behaviour comparable to the Hanus conditioned controller and there seems to be no need for fast switching. The back-calculation anti-windup performs well in a variety of situations.

Keywords: Control in anesthesia, anti-windup, PID control

1. INTRODUCTION

In closed-loop control of anesthesia, drug infusion is adjusted based on feedback of the measured clinical effect. Clinical evaluations with experimental systems have shown the feasibility of closed-loop anesthesia (e.g. Liu et al. [2006], Struys et al. [2004], Gentilini et al. [2001], Sawaguchi et al. [2008], West et al. [2013]), but no such system is currently commercially available. To convince regulatory authorities of the benefit of automated systems, both patient safety and improved outcome will have to be demonstrated. The control system needs to be robust to realistic patient variability, and be safe in all situations including failure modes (Dumont [2013]). Typical failure modes in control of depth of hypnosis (DOH) are related to the infusion pump or the DOH monitor. The pumps have physical limitations and drug cannot be removed from the body, limiting the control input to positive values. Device failure results in a loss of feedback or loss of the actuator. Tao et al. [2013] recently proposed an observer based solution to loss of the feedback signal. Viability theory has been proposed for safe control of anesthesia (Maidens et al. [2013]).

This paper considers safety constraints based on the therapeutic window of propofol, an anesthetic drug typically used to maintain anesthesia. Examples from clinical evaluation show that situations where such bounds are essential are readily encountered in clinical practice. On the other hand, in simulation studies constraints are rarely active and safety issues are not encountered when controllers are evaluated based on population-based or nominal models. In this paper, realistic simulation scenarios are proposed that are motivated by examples from clinical practice. Note that these scenarios are not sufficient to validate designs, they are not exhaustive, nor do they include the most challenging scenarios.

The proposed scenarios are used to evaluate the performance of three anti-windup (AW) strategies for PID control of propofol infusion. The $\mathcal{L}_2$ anti-windup framework (Zaccarian and Teel [2002]) proposes AW strategies that guarantee asymptotic stability, a desirable property in medical applications. An $\mathcal{L}_2$ AW compensator is added to a clinically evaluated PID controller. Secondly, sliding mode reference conditioning, a safety addition that was recently proposed for the artificial pancreas (Revert et al. [2013]), is applied to control of anesthesia. The performance of these two AW strategies is compared to classical AW based on back-calculation (Åström and Hägglund [2005]). It is concluded that the $\mathcal{L}_2$ AW objective might not be appropriate for this problem. The sliding-mode AW is shown to approximate a Hanus conditioned controller, and approximation or fast switching is not required to achieve the same result. Back-calculation AW performs well and is predictable in a variety of situations.

The paper is organized as follows. Section 2 describes the safety constraints, presents examples from clinical practice and proposed realistic simulation scenarios. The three AW strategies are described in Section 3, simulation results comparing the performance of the AW strategies are presented in Section 4. Section 5 concludes the paper.

2. CONSTRAINTS IN CLOSED-LOOP ANESTHESIA

2.1 Safety constraints in anesthesia

Pharmacokinetics (PK) describe the transport and metabolism of a drug, pharmacodynamics (PD) relate plasma

drug concentrations to a hypothesized effect site concentration C_e and the clinical effect. The lower limit of the therapeutic window of propofol is determined by minimal concentrations that assure adequate anesthesia, while the upper bound is probably limited by its hemodynamic effects (Vuyk et al. [1997]): In combination with the opioid remifentanil, effect site concentrations $\approx$ 2.7 mg/l were associated with a 95% probability of no response. $C_e \approx 1.6$ has been associated with awakening after termination of the drug infusion. In the absence of opioids, $C_e = 3.4$ mg/l has been associated with loss of consciousness in 50% of patients, and 10-15 mg/l with suppression of response to stimulation. This required C_e strongly reduces in combination with opioids.

The therapeutic window, as described above, can be used to define safety bounds for a closed-loop system by limiting the predicted propofol C_e. In this paper, safety constraints are defined for closed-loop propofol anesthesia in combination with a range of remifentanil analgesia levels:

$$1.5\text{mg/l} < C_e(t) < 8\text{mg/l}, \qquad (1)$$

where the effect site concentration C_e is predicted using the Schnider model (Schnider et al. [1998]).

The range defined in equation (1) is not a clinical hard bound. While the constraints reduce the risk of under- or overdosing for most patients, the bounds are expected to be reached for outliers. Reaching these bounds indicates that the anesthetic requirements or the response to drug infusion of that particular patient is far from the population average. If the safety system is activated, the anesthesiologist can make a clinical decision, and the closed-loop system will not compromise patient safety.

The PKPD model is known, and constraints on the predicted C_e can be implemented directly or indirectly through constraints on $u(t)$. Definition of constraints that correspond to bounds on C_e require open-loop prediction.

2.2 Examples from clinical practice

The following examples are taken from 45 cases of closed-loop controlled propofol infusion guided by the WAV_{CNS} measure provided by the NeuroSENSE monitor (NeuroWave Systems Inc., Cleveland Heights, USA). The system is described in Dumont et al. [2011]. The software was approved by Health Canada for clinical evaluation. Following approval from the Research Ethics Board, and informed consent, 45 adults requiring anesthesia for elective surgery were enrolled (17 F, age (median (range)) 63y (32-82), weight 81kg (47-116), height 175cm (151-190)). Anesthesia was closed-loop controlled in 42 of these cases, target controlled infusion was used during the whole procedure or part of it in 3 cases.

Fig. 1 shows an example of a typical case of closed-loop propofol anesthesia. Induction of anesthesia was completed in 2min35s. Electrocautery was used when the procedure started (after 14 minutes), resulting in a loss of feedback from the DOH monitor. After 27 minutes, the propofol syringe was empty and had to be replaced, resulting in a temporary loss of the actuator. Propofol infusion rates and resulting predicted C_e remained within the predefined therapeutic window.

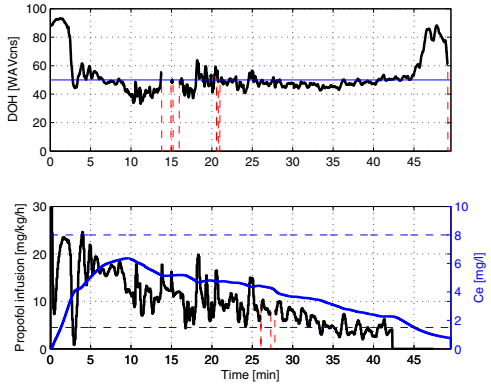

Fig. 1. Example of a typical case. Top figure: Measured WAV_{CNS}. Bottom figure: Black, closed-loop controlled propofol infusion. Blue, corresponding predicted C_e based on the Schnider model.

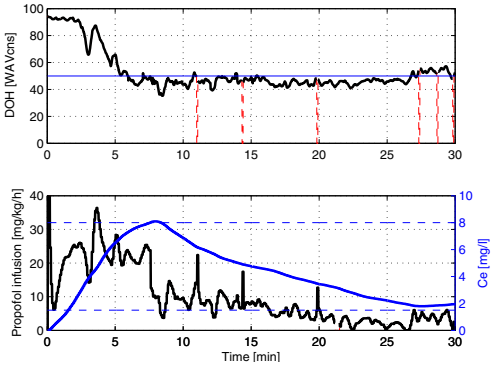

Fig. 2. Example of closed-loop controlled propofol infusion. Top figure: Measured WAV_{CNS}. Bottom figure: Black, closed-loop controlled propofol infusion. Blue, corresponding predicted C_e (Schnider model).

Fig. 2 shows a case where an initial decrease of the measured WAV_{CNS} is followed by some arousal. Airway insertion resulted in another small increase in WAV_{CNS} (around 5 min). The resulting prolonged induction of anesthesia of 5min5s required a higher than average propofol dose, resulting in a C_e of 8 mg/l. This large induction dose was due to response to stimulation and limitation of propofol infusion through the safety bounds reduces the risk of overdosing.

Fig. 3 shows a different outlier. This patient required a low propofol dose for induction of anesthesia. Some response to intubation was followed by an overshoot. The measured WAV_{CNS} did not increase upon decreasing propofol infusion. The predicted C_e decreased to 1.5 mg/l, the predefined lower bound. The measured WAV_{CNS} remained below 40 also after the start of the procedure. Due to a reduced propofol clearance, this patient response does not match predicted behaviour based on PKPD models. The lower bound was reduced to 1 mg/l to better reflect the needs of this patient. After 70 minutes the WAV_{CNS} shows a response to stimulation. The controller responds to this stimulation and increases the propofol rate as required.

These examples show that temporary hardware failure can be expected in most clinical cases. Electrocautery is used for a variety of interventions and the syringes

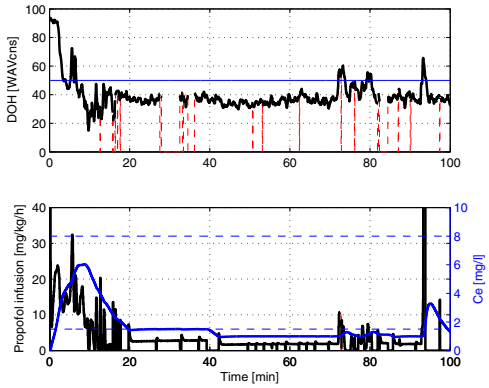

Fig. 3. Example of closed-loop controlled propofol infusion for a patient with below average clearance. Top figure: Measured WAV_{CNS}. Bottom figure: Black, closed-loop controlled propofol infusion. Blue, corresponding predicted C_e based on the Schnider model. Only part of the case is shown, the complete case lasted 2h40.

will run out of drug at some point. The therapeutic window defines sufficient levels of anesthesia for most patients, however, the bounds are expected to be met for outliers and in situations where stimulation or other clinical circumstances affected the drug requirement.

2.3 Realistic scenarios in simulation

The following simulations are based on the nonlinear model set presented in Bibian et al. [2006]. This model set includes 44 individual patient models and covers a large variability of dynamic behaviour. These models were identified from data from induction of anesthesia and do not provide an accurate estimate of the low frequency gain. This does not affect model-based controller design (see van Heusden et al. [2013]), it does limit the scenarios that can be reproduced in simulation without model adjustments.

In these scenarios propofol infusion is controlled by the PID controller used in the clinical study discussed in Section 2.2 and described in Dumont et al. [2011]. The setpoint is fixed to 50 , monitor dynamics are included in the simulation (see Bibian et al. [2006]). Measurement noise is not taken into account.

Scenario 1, linear model Patient model #5 is linearized. Closed-loop induction of anesthesia is simulated including bounds on the predicted C_e (Schnider model): $1.5 < C_e(t) < 8mg/l$. The length of the simulation is 30 min. In this scenario the constraint is active in part of the simulation. The model is linear and known, which allows for evaluation of model-based solutions in the case of an exact model match.

Scenario 2, stimulation This scenario is based on the case shown in Fig. 2. Closed-loop induction of anesthesia is simulated for the nonlinear patient model #35, with bounds on the predicted $C_e(t)$: $1.5 < C_e(t) < 8mg/l$. A disturbance is added 5 minutes after the start of induction, see Fig. 4. Due to this disturbance, the upper bound of the therapeutic window is reached. Simulation length is 30 minutes.

Scenario 3, low clearance This scenario is based on the case shown in Fig. 3. Nonlinear model #15 is multiplied by $G_2 = 1 + \frac{0.8}{(700s+1)(800s+1)^2}$. This increases the gain at low frequencies, resulting in a time response similar to reduced clearance. Closed-loop induction of anesthesia is simulated, the disturbance profile is shown in Fig. 4. Bounds on the predicted $C_e(t)$ are set to $1.5 < C_e(t) < 8mg/l$. Due to the reduced clearance the lower bound is reached. The disturbance leads to an increased measured WAV_{CNS} that the control system should be responsive to. The simulation length is 115 minutes.

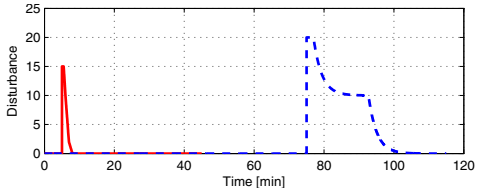

Fig. 4. Red: Disturbance profile used in scenario 2, representing a response to stimulation upon intubation. Blue dashed: Disturbance profile used in scenario 3, representing a response to surgical stimulation.

3. ANTI-WINDUP STRATEGIES

3.1 Classical anti-windup

A simple AW solution is given by the classical strategy based on back-calculation (Åström and Hägglund [2005]). The following term is added to the integrator input:

$$\frac{1}{T_t}(u_{sat} - u_{PID}), \tag{2}$$

where u_{sat} is the saturated control input, u_{PID} is the controller output and T_t is the time constant for the dynamical integrator reset. The AW is tuned manually, $T_t = 60s$. This method can be implemented both in continuous and discrete time. Section 4 is based on a discrete time implementation.

3.2 Sliding mode safety element

Revert et al. [2013] proposed a safety element for closed-loop control in diabetes. In control in type 1 diabetes, so called 'Insulin-On-Board' (IOB) constraints have been used to avoid overdosing (Ellingsen et al. [2009]). This constraint is based on an insulin absorption model. The PKPD based constraint for anesthesia used in this paper is similar to this IOB constraint and the proposed safety element can be implemented directly as an addition to the PID controller for propofol anesthesia.

The safety element uses sliding mode reference conditioning. The main idea is that the sliding mode feedback shapes the reference to the closed-loop system such that the constraints are not violated (see Revert et al. [2013] for details). Fig. 5 shows the corresponding block diagram. The method requires measurement of the constrained variable as well as the derivatives up to order $l-1$ where l is the relative degree between the reference $r(t)$ and the constrained variable $C_e(t)$. A low pass filter F is required to smoothen the switching signal $w(t)$.

The PKPD model is known, the relative degree between $r(t)$ and $C_e(t)$ is 2 and the derivative of $C_e(t)$ is available from the state-space PKPD model. The switching signal

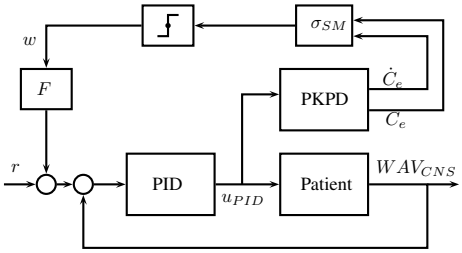

Fig. 5. Block diagram for sliding mode reference conditioning for closed-loop anesthesia. The constrained variable is the predicted effect site concentration $C_e(t)$.

is defined as:

$$w(t) = \begin{cases} w^+ & \text{if } \sigma_{SM}(t) > 0 \\ 0 & \text{if } \sigma_{SM}(t) \leqslant 0 \end{cases} \quad (3)$$

where $w^+ > w^\sigma$, a value that can be calculated as described in Revert et al. [2013]. $\sigma_{SM}(t)$ is defined as:

$$\sigma_{SM}(t) = C_e(t) - C_{e_{MAX}}(t) + \tau_\sigma(\dot{C}_e(t) - \dot{C}_{e_{MAX}}(t)). \quad (4)$$

τ_σ is a constant gain. A similar switching signal can be defined for a lower bound on the predicted $C_e(t)$. In the simulation results reported in Section 4, $\tau_\sigma = 0.3$, $w^+ = 50$ and $F(s) = \frac{1}{s+1}$. This sliding mode scheme requires continuous time implementation, and is added to a continuous time simulation of the PID controlled patient model.

3.3 Model-based anti-windup

The $\mathcal{L}_2$ anti-windup framework (Zaccarian and Teel [2002]) provides an AW solution that guarantees asymptotic stability. In medical applications this guarantee provides a much desired characteristic. The scheme is model-based and it is assumed that a reliable model of the controlled system is available.

The model-based $\mathcal{L}_2$ AW scheme is shown in Fig. 6. The AW compensator consists of the patient model, with input $u_{sat} - u_{PID}$ and two outputs, v_1 and v_2:

$$\begin{aligned} \dot{\xi} &= A\xi + B(u_{sat} - u_{PID}) \\ \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} &= \begin{bmatrix} K \\ -C \end{bmatrix} \xi + \begin{bmatrix} L \\ 0 \end{bmatrix} (u_{sat} - u_{PID}). \end{aligned} \quad (5)$$

A, B, C, D is the state-space representation of the patient model without time delay. The feedback terms L and K that determine v_1 are designed using the solution described in equations (10) and (11) in Zaccarian and Teel [2002]. The solution is combined with the additional constraint suggested in Zaccarian et al. [2005], equation (16), to enforce a well-posed solution. Following the approach in Zaccarian et al. [2005] for delayed systems, feedback of v_2 is delayed by τ, the time delay of the patient model, as indicated in Fig. 6.

This model-based AW solution assumes an accurate plant model. In control in anesthesia, an accurate patient model is not available. For practical implementation, the AW solution therefore uses the nominal model for the set of 44 models as defined in Dumont et al. [2009]. To verify the applicability of the approach, the exact linearized model is also implemented and tested in simulation, providing an exact model match. Although this solution does not explicitly take robustness into account in the design,

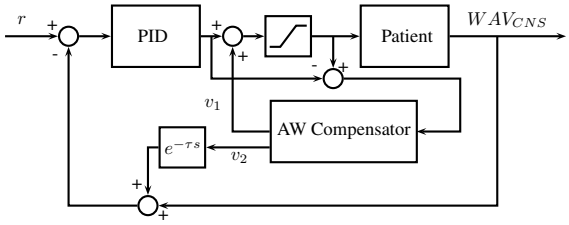

Fig. 6. Model-based anti-windup scheme as presented by Zaccarian and Teel [2002].

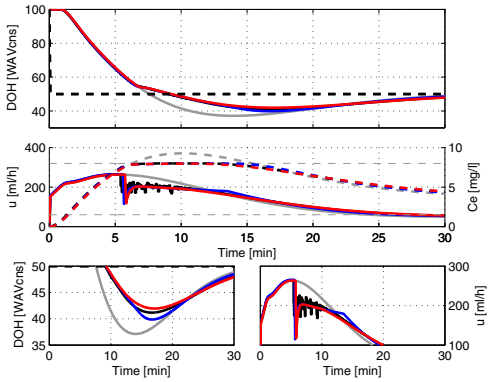

Fig. 7. Responses for simulation **scenario 1**: unconstrained response (grey), constrained response with classical AW (red), sliding mode AW (black), model-based solution where the linear patient model is matched exactly in the AW compensator (blue).

it allows for an evaluation of the applicability of this methodology in handling safety constraints in closed-loop anesthesia. This model-based AW is a continuous time approach and is therefore implemented in a continuous time simulation of the PID controlled patient model.

4. RESULTS

4.1 Scenario 1, linear model

Fig. 7 shows the simulation results for scenario 1. The responses of the three AW strategies are similar and remain within bounds. The sliding mode results in fast switching when the bound is reached, the low pass filter reduces the chattering. The model-based solution with matching model in the AW compensator shows a larger overshoot than the other two solutions, but better performance as defined in the $\mathcal{L}_2$ framework; the $\mathcal{L}_2$ performance objective is defined with respect to the unconstrained solution and the model-based approach leads to a smaller error relative to this unconstrained response.

Fig. 8 compares the response of the AW compensator with the exact matching model to the response when the nominal model is used. For this patient model, the overshoot is the same. The response from the AW compensator with the exact model results in a smaller error compared to the unconstrained model.

4.2 Scenario 2, stimulation

Fig. 9 shows the simulation results for scenario 2. The bounds of the therapeutic window are reached due to

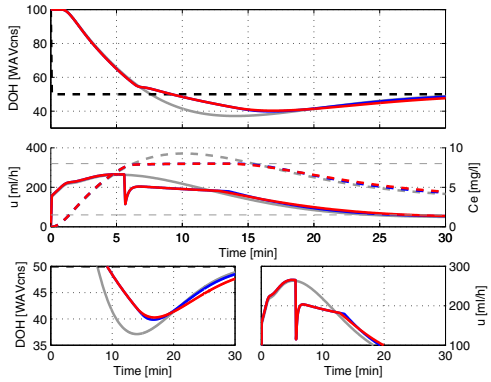

Fig. 8. Responses for simulation **scenario 1**: unconstrained response (grey), model-based solution with a matched model in the AW compensator (blue), model-based solution using the nominal model (red).

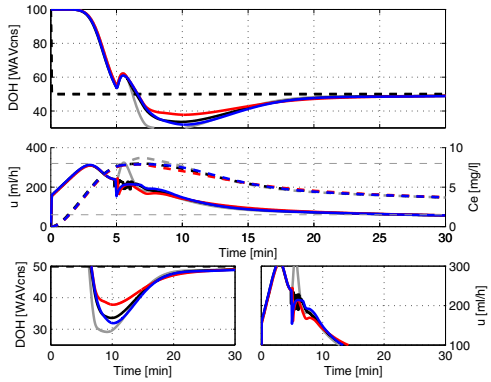

Fig. 9. Responses for simulation **scenario 2**: unconstrained response (grey), classical AW (red), sliding mode AW (black), model-based solution based on the nominal model (blue).

response to stimulation. Note that in both this case and scenario 1, the overshoot upon induction of anesthesia is not desired, the bounds are activated to avoid overdosing. The model-based and sliding mode AW result in higher overshoot than the back-calculation solution.

4.3 Scenario 3, low clearance

Fig. 10 shows the simulation results for scenario 3, representing a patient with below average clearance. In this scenario the lower bound is active for a prolonged period. After 90 minutes, surgical stimulation results in a response observed in the simulated WAV_{CNS}. The controller needs to be responsive to this increased WAV_{CNS} to assure sufficient anesthesia in the presence of this disturbance. Both the back-calculation AW and the sliding mode AW are responsive. The back-calculation AW achieves better disturbance rejection. The response of the model-based AW compensator is insufficient.

4.4 Sliding mode anti-windup and the Hanus conditioned controller

The concept of reference conditioning used in the sliding mode AW is based on Hanus et al. [1987]. This Hanus conditioned controller provides a reference for which the

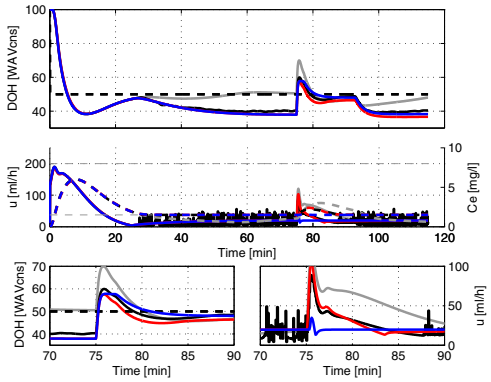

Fig. 10. Responses for simulation **scenario 3**: unconstrained response (grey), classical AW (red), sliding mode AW (black), model-based solution based on the nominal model (blue).

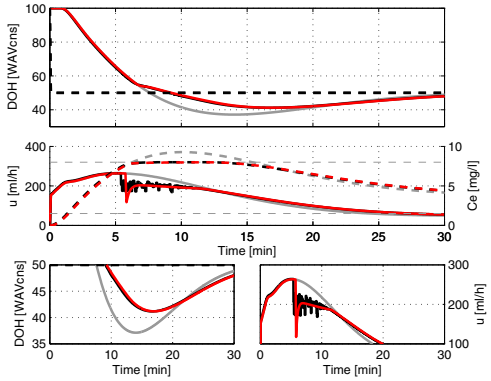

Fig. 11. Responses for simulation **scenario 1**: unconstrained response (grey), sliding mode AW (black), corresponding Hanus conditioned controller (red).

controller output matches the constrained value. The sliding mode AW approximates this reference through high frequency switching. However, since both the PKPD model and the controller are known exactly, and bounds on the input u that correspond to the bounds on C_e can be calculated at all times, the Hanus conditioned controller can be implemented without the need for approximation or high frequency switching. Furthermore, for the PID formulation considered in this paper, the conditioned controller corresponds to a back-calculation AW with T_t defined by the PID parameters, resulting in $T_t > 60$ s.

Fig. 11 compares the response of the sliding mode AW to the response using the Hanus reference conditioning formulation. The input is adjusted earlier in the sliding mode approach due to the approximation. The resulting WAV_{CNS} is equivalent. Similar results can be shown for the other scenarios (not shown). The Hanus conditioned controller can be implemented in discrete time and combination of multiple constraints is straightforward.

5. CONCLUSION

This paper describes scenarios encountered in clinical practice that emphasize the need for safety constraints in closed-loop anesthesia. Simulation scenarios are presented that provide a realistic test environment for constrained control. The performance of three anti-windup strategies is

compared in these situations; classical AW based on back-calculation, sliding mode reference conditioning and $\mathcal{L}_2$ model-based AW.

The stability guarantees in the $\mathcal{L}_2$ AW framework present a desirable characteristic for medical applications. However, the $\mathcal{L}_2$ AW objective might not be appropriate for these applications where the constraints present a safety bound and not actuator limitations. The constraints are active when the unconstrained response is not desired, i.e. to avoid over- or underdosing. Furthermore, bounds based on the therapeutic window are expected to be reached for outliers, i.e. patients with a dynamic response to drug infusion that deviates from the population average. In these situations, the model in the AW compensator will not be an accurate description of the patient response and the performance of the AW scheme will be limited for the cases when it is actually needed.

The sliding mode reference conditioning AW behaviour can also be obtained using a Hanus conditioned controller. Since both the controller and the PKPD model are known, there seems to be no need for approximation and fast switching that may introduce chattering.

The results in this paper are limited to single faults and do not include measurement noise. The proposed scenarios are examples, do not cover all relevant situations encountered in practice, and do not include the most challenging situations encountered in clinical practice. It is concluded that although classical AW based on back-calculation generally does not provide stability guarantees, it performs well in a variety of situations encountered in closed-loop anesthesia. The suggested simulation scenarios can be used to evaluate and compare alternative solutions based on model-predictive control or viability theory.

REFERENCES

K. J. Åström and T. Hägglund. *Advanced PID Control*. ISA, 2005.

S. Bibian, G. A. Dumont, M. Huzmezan, and C. R. Ries. Patient variability and uncertainty quantification in clinical anesthesia: Part I – PKPD modeling and identification. In *IFAC Symposium on Modelling and Control in Biomedical Systems*, Reims, France, 2006.

G.A. Dumont. Feedback control for clinicians. *Journal of Clinical Monitoring and Computing*, ahead of print, April, 2013.

G.A. Dumont, A. Martinez, and J.M. Ansermino. Robust control of depth of anesthesia. *International Journal of Adaptive Control and Signal Processing*, 23(5):435–454, 2009.

G.A. Dumont, N. Liu, C. Petersen, T. Chazot, M. Fischler, and J.M. Ansermino. Closed-loop administration of propofol guided by the neurosense: Clinical evaluation using robust proportional-integral-derivative design. In *American Society of Anesthesiologists Annual Meeting*, Chicago, IL, 2011.

C. Ellingsen, E. Dassau, H. Zisser, B. Grosman, M.W. Percival, L. Jovanovic, and F.J. Doyle III. Safety constraints in an artificial pancreatic β cell: an implementation of model predictive control with insulin on board. *Journal of Diabetes Science and Technology*, 3(3):536–44, 2009.

A. Gentilini, M. Rossoni-Gerosa, C. W. Frei, R. Wymann, M. Morari, A. M. Zbinden, and T. W. Schnider. Modeling and closed-loop control of hypnosis by means of bispectral index (BIS) with isoflurane. *Biomedical Engineering, IEEE Transactions on*, 48(8):874–889, 2001.

R. Hanus, M. Kinnaert, and J.-L. Henrotte. Conditioning technique, a general anti-windup and bumpless transfer method. *Automatica*, 23(6):729 – 739, 1987.

N. Liu, T. Chazot, A. Genty, A. Landais, A. Restoux, K. McGee, P-A. Laloë, Be. Trillat, L. Barvais, and M. Fischler. Titration of propofol for anesthetic induction and maintenance guided by the bispectral index: Closed-loop versus manual control: A prospective, randomized, multicenter study. *Anesthesiology*, 104(4):686–695, 2006.

J.N. Maidens, S. Kaynama, I.M. Mitchell, M.M.K. Oishi, and G.A. Dumont. Lagrangian methods for approximating the viability kernel in high-dimensional systems. *Automatica*, 49(7):2017 – 2029, 2013.

A. Revert, F. Garelli, J. Pico, H. De Battista, P. Rossetti, J. Vehi, and J. Bondia. Safety auxiliary feedback element for the artificial pancreas in type 1 diabetes. *Biomedical Engineering, IEEE Transactions on*, 60(8): 2113–2122, 2013.

Y. Sawaguchi, E. Furutani, G. Shirakami, M. Araki, and K. Fukuda. A model-predictive hypnosis control system under total intravenous anesthesia. *Biomedical Engineering, IEEE Transactions on*, 55(3):874 –887, 2008.

T. W. Schnider, C. F. Minto, P. L. Gambus, C Andresen, D. B. Goodale, S. L Shafer, and E. J. Youngs. The influence of method of administration and covariates on the pharmacokinetics of propofol in adult volunteers. *Anesthesiology*, 88(5):1170 – 1182, 1998.

M. M. R. F. Struys, T. De Smet, S. Greenwald, A. R. Absalom, S. Bing, and E.P. Mortier. Performance evaluation of two published closed-loop control systems using bispectral index monitoring: A simulation study. *Anesthesiology*, 100(3):640–647, 2004.

Y. Tao, M. Fang, and Y. Wang. A fault tolerant closed-loop anesthesia system based on internal model control and extended state observer. In *Chinese Control and Decision Conference*, pages 4910–4914, 2013.

K. van Heusden, J.M. Ansermino, K. Soltesz, S. Khosravi, N. West, and G.A. Dumont. Quantification of the variability in response to propofol administration in children. *Biomedical Engineering, IEEE Transactions on*, 60(9):2521–2529, 2013.

J. Vuyk, M.J. Mertens, E. Olofsen, A.G. Burm, and J.G. Bovill. Propofol anesthesia and rational opioid selection: determination of optimal EC50-EC95 propofol-opioid concentrations that assure adequate anesthesia and a rapid return of consciousness. *Anesthesiology*, 87(6): 1549–62, 1997.

N. West, G.A. Dumont, K. van Heusden, C.L. Petersen, S. Khosravi, K. Soltesz, A. Umedaly, E. Reimer, and J.M. Ansermino. Robust closed-loop control of induction and maintenance of propofol anesthesia in children. *Pediatric Anesthesia*, 23(8):712–719, 2013.

L. Zaccarian and A.R. Teel. A common framework for anti-windup, bumpless transfer and reliable designs. *Automatica*, 38(10):1735 – 1744, 2002.

L. Zaccarian, D. Nesic, and A.R. Teel. Anti-windup for linear dead-time systems. *Systems & Control Letters*, 54(12):1205 – 1217, 2005.