Seamless integration of target-controlled infusion and closed-loop anesthesia

Wahlquist, Ylva; Soltesz, Kristian

Published in:
The 2025 American Control Conference (ACC)

2025

Document Version:
Publisher's PDF, also known as Version of record

Seamless integration of target-controlled infusion and closed-loop anesthesia

1st Ylva Wahlquist
Department of Automatic Control
Lund University
Lund, Sweden
ylva.wahlquist@control.lth.se

2nd Kristian Soltesz
Department of Automatic Control
Lund University
Lund, Sweden
kristian@control.lth.se

***Abstract*—The anesthetic drug propofol is commonly used to control hypnotic depth (suppression of awareness) in patients undergoing surgery or intensive care. In addition to manual titration, a model-based open-loop feed-forward strategy called target-controlled infusion (TCI) has attained some clinical popularity. Research on closed-loop control, with awareness estimates derived from an electroencephalogram (EEG), has proven feasible through several extensive clinical studies over the past decades. While TCI is vulnerable to model imperfections, closed-loop control is susceptible to corrupt measurements. By combining Kalman-filter-based state estimation with model predictive control (MPC), we introduce a novel anesthetic dosing regimen that can transition seamlessly between TCI and closed-loop control, thus constituting an adequate trade-off between model and measurement reliance. We introduce this regimen and provide a realistic simulation example that highlights its strengths compared to pure TCI or closed-loop control of propofol infusion.**

***Index Terms*—Biomedical systems, Model Predictive Control, Kalman filtering**

I. Introduction

Hypnosis is the pharmacologically induced anesthetic component utilized to temporarily repeal awareness, as mandated by certain surgical procedures. Intravenous drugs, such as propofol, are becoming increasingly popular and are commonly administered manually and titrated by an anesthesiologist based on monitor readings, patient signs, and experience.

Target-controlled infusion (TCI) is a model-based augmentation to manual titration. When TCI is used, the anesthesiologist enters a desired set point (reference) drug concentration in the blood plasma [1]. Based on this reference and an underlying patient model of the assumed pharmacokinetics, the TCI system computes an optimized infusion trajectory that is then applied [1]. The patient model is adapted to covariates such as age, weight, and sex. Several such pharmacokinetic covariate models have been used as the basis for TCI systems, e.g., [2, 3]. TCI is a feed-forward open-loop control strategy, which means that it cannot account for model errors or disturbances that divert the true plasma concentration from that assumed by the TCI system model. However, the anesthesiologist can update the plasma concentration setpoint if observations make it plausible that the patient is subjected to inadequate anesthetic depth, possibly caused by external disturbances.

This work was partially funded by the Hans-Gabriel and Alice Trolle Wachtmeister foundation for medical research. The authors are members of the ELLIIT Strategic Research Area at Lund University.

When TCI relies solely on a dynamic patient model, closed-loop controlled anesthesia constitutes a dosing regimen where the infusion rate is updated based on an online measurement [4, 5]. This measurement is typically an estimate of the depth of hypnosis (DoH), calculated from a non-invasive electroencephalogram (EEG) signal. The anesthesiologist sets a desired reference anesthetic depth, usually given on the BIS scale, where a value close to the maximum of 100 BIS indicates that the patient is fully aware, and the range 40–60 BIS is adequate for many clinical use cases [6]. Too deep anesthesia may result in adverse effects, such as post-operative nausea, while insufficient anesthesia can lead to awareness during surgery.

While the lack of clinical feedback in TCI makes it vulnerable to model error, the presence of clinical feedback makes closed-loop anesthesia vulnerable to measurement disturbances and sensor noise. This motivates a unified framework that combines TCI and closed-loop control to obtain a tunable trade-off between the two. This paper presents such a framework, based on a combination of Kalman filtering state estimation and model predictive control (MPC). We provide illustrative and motivating examples and conclude with a realistic simulated use case.

II. Modeling and control

Ordinary TCI is schematically illustrated in Figure 1, and our novel hybrid approach is illustrated in Figure 2. The patient is represented by a PKPD model, defined in section II-A. The measurement y and disturbances d_1, d_2 are explained in section II-B and the Kalman filter state estimator in section II-C. Finally, the MPC that governs dosing in both the TCI and the closed-loop case is presented in section II-D.

A. Patient model

At the core of our approach lies a pharmacokinetic-pharmacodynamic (PKPD) model. The structure of this model is well established in the context of closed-loop controlled anesthesia [7, 8]. The PK part of the model is a mammillary three-compartment system, relating drug infusion rate

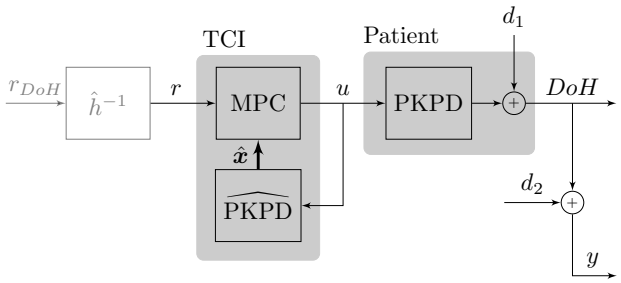

Fig. 1: Schematic illustration of open-loop control with TCI. The DoH reference r_{DoH} is transformed by $\hat{h}^{-1}$ into a plasma concentration which is inverted through (5) to create the reference r for the MPC, which computes a control signal u based on the state estimate $\hat{\boldsymbol{x}}$ from the underlying PKPD model. This transform is in gray, as it is typically conducted implicitly by the monitoring anesthesiologist, who sets r directly, in the case of TCI control. The disturbances d_1 and d_2 act on the DoH and its measurement y, respectively.

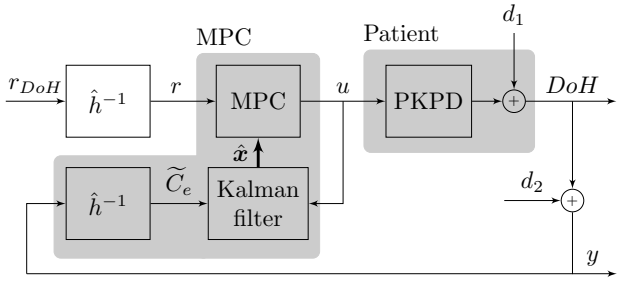

Fig. 2: Schematic illustration of closed-loop control with MPC. The reference for DoH, r_{DoH}, is inverted through $\hat{h}^{-1}$ in (5) to create the reference for the MPC, which computes a control signal u based on the state estimates from the Kalman filter. The disturbances d_1 and d_2 act on the DoH and the measurement y, respectively.

u [$\mathrm{mg\,s^{-1}}$] to the blood plasma drug concentration $x_3 = C_p$ [$\mathrm{mg\,L^{-1}}$] [7]. The PD model, connected in series to the output of the PK model, consists of a linear and a non-linear part. The linear part is a first-order lag filter relating the blood plasma concentration $x_3 = C_p$ to the effect-site (brain cortex) concentration $x_4 = C_e$. A zero-order-hold discretization, in our examples with sampling period $T_s = 10$ s (sufficient to resolve the dynamics) can thus be expressed

$$\boldsymbol{x}_{k+1} = A\boldsymbol{x}_k + Bu_k, \tag{1a}$$

$$(C_e)_k = C\boldsymbol{x}_k = (x_4)_k, \tag{1b}$$

where

$$\boldsymbol{x}_k = \left[(x_1)_k, (x_2)_k, (x_3)_k, (x_4)_k\right]^\top, \tag{2}$$

as detailed in for example [9].

The non-linear part of the PD model relates effect-site concentration $C_e = x_4$ to DoH via the Hill sigmoid

$$DoH = h(C_e; E_0, C_{e50}, \gamma) = E_0\left(1 - \frac{C_e^\gamma}{C_e^\gamma + C_{e50}^\gamma}\right), \tag{3}$$

C_{e50} corresponds to a $DoH = 50$ BIS, and γ determines the steepness of the sigmoid.

The variability in the DoH response between individuals is accounted for using a pharmacometric covariate model. In this work, we use the Schnider population model [3], which expresses the parameters of the continuous-time counterpart of (1) as a function of the covariates age, weight, and sex. The remaining variability, not explained by these covariates, is modeled using a log-normal distribution, so that

$$\theta = \theta_0 \exp(\eta_\theta), \tag{4}$$

where θ_0 is the covariate-adjusted parameter value, and the random effect η_θ is a normal stochastic variable with zero mean and variance σ_θ^2. These random effects thus capture intra-individual variability and the part of the inter-individual variability not explained by the covariates.

B. Measurements and disturbances

Disturbances pose a central challenge for both TCI and closed-loop control. Surgical stimulation is commonly [8] modeled by an additive disturbance d_1, affecting the actual DoH, as shown in Figure 1.

It is not possible to measure the full patient state or the true DoH directly. Instead, the only available measurement is $y = DoH + d_2$, where d_2 corrupts measurement noise, as shown in Figure 1. That is, d_1 affects the actual DoH and consequently also the measurement y, while d_2 affects the measurement y, without influencing the true DoH in the case of TCI shown in Figure 1. However, in the case of closed-loop feedback shown in Figure 2, d_2 will influence the true DoH through the controller.

This work considers disturbances in double-steps of magnitude 20 BIS. Step disturbances are commonly used [10], as they can be considered to be the worst case in the anesthesia control context.

C. Patient state estimation

Since the full state of (1) is not directly measurable, we use the Kalman filter to obtain state estimates. It utilizes the system input u, together with the effect-site concentration estimate

$$\hat{C}_e = \hat{h}^{-1}(DoH; E_0, C_{e50}, \gamma) = C_{e50}\left(\frac{E_0 - DoH}{DoH}\right)^{1/\gamma}. \tag{5}$$

to produce a state estimate $\hat{\boldsymbol{x}}$.

To begin with, we will, somewhat optimistically, assume perfect knowledge of the PD dynamics, $\hat{h} = h$, resulting in $\hat{h}^{-1}$ constituting a perfectly linearizing transform. However, in our concluding realistic simulation scenario of section IV, we use a previously published population average for the parameters in (5) to obtain a corresponding, but non-perfectly linearizing $\hat{h}^{-1}$. This was also done in [11], where the same population average resulted in satisfactory results.

The Kalman filter equations are presented below, with the following notation: $\hat{\boldsymbol{x}}_{k,k-1}$ denotes the estimate of the state vector $\boldsymbol{x}_k$, based on data up to and including sample $k-1$, and $\hat{\boldsymbol{x}}_{k,k}$ is the updated estimate at sample k. The same notation

is used for the covariance estimate P. The state update is governed by

$$L_k = P_{k,k-1}C^\top \left(CP_{k,k-1}C^\top + R_k\right)^{-1}, \tag{6a}$$

$$P_{k,k} = (I - L_kC)P_{k,k-1}(I - L_kC)^\top + L_kR_kL_k^\top, \tag{6b}$$

$$\hat{\boldsymbol{x}}_{k,k} = \hat{\boldsymbol{x}}_{k,k-1} + L_k\left((\tilde{C}_e)_k - C\hat{\boldsymbol{x}}_{k,k-1}\right), \tag{6c}$$

and the prediction is given by

$$\boldsymbol{x}_{k,k+1} = A\boldsymbol{x}_{k,k} + Bu_k, \tag{7a}$$

$$P_{k+1,k} = AP_{k,k}A^\top + Q_k, \tag{7b}$$

where A, B, C are the system matrices of (1). Kalman filtering is a well-understood and documented technique, and we refer to [12] for further details and insights.

The scalar R_k and the matrix Q_k quantify measurement and state uncertainty in the sample k. While they can be viewed to represent covariances of Gaussian disturbances within the Kalman filtering framework, we instead consider them as free tuning parameters that enable a trade-off between measurement and model reliance.

The Kalman filter state estimate is a weighted sum of its most recent measurement and the previous state estimate, where the Kalman gain L is the proportionality constant. When the measurement uncertainty R is large, the Kalman gain L will decrease, shifting trust from measurement to model, and vice versa.

D. MPC formulation

The objective of our MPC– which is to be used both for TCI and closed-loop controlled infusion– is to produce an optimal drug infusion trajectory to follow an effect-site concentration reference. Its roles in each of these use cases are shown in Figure 1 and Figure 2, respectively.

In TCI, the MPC relies solely on a PKPD patient model, without utilizing the DoH measurement y. In closed-loop mode, the MPC instead utilizes Kalman-filter state estimates based on y, as explained in section II-C.

We denote the effect-site concentration profile across a horizon of N samples by $\boldsymbol{x}_4$ so that $\boldsymbol{x}_4 = \left[(x_4)_1 \dots (x_4)_N\right]^\top$, and the reference trajectory $\boldsymbol{r} = \left[r_1 \dots r_N\right]^\top$. The objective of our MPC is to find the infusion trajectory $\boldsymbol{u} = \left[u_1 \dots u_N\right]^\top$ that minimizes the quadratic cost function

$$J'_e(\boldsymbol{x}_4) = \sum_{k=1}^{N} \left((x_4)_k - r_k\right)^2. \tag{8}$$

To enable minimization over $\boldsymbol{u}$, we rewrite (8) in terms of $\boldsymbol{u}$. With the initial state $\boldsymbol{x}_0$ and the use of (1), we get

$$J_e(\boldsymbol{u}) = \frac{1}{2}\boldsymbol{u}^\top F^\top F\boldsymbol{u} + (\boldsymbol{x}_0^\top E^\top F - \boldsymbol{r}^\top F)\boldsymbol{u}, \tag{9}$$

where

$$E = \begin{bmatrix} A_4^1 \\ \vdots \\ A_4^N \end{bmatrix}, \; F = \begin{bmatrix} A_4^0B_4 & & & \\ A_4^1B_4 & A_4^0B_4 & & \\ \vdots & \vdots & \ddots & \\ A_4^{N-1}B_4 & A_4^{N-2}B_4 & \dots & A_4^0B_4 \end{bmatrix}, \tag{10}$$

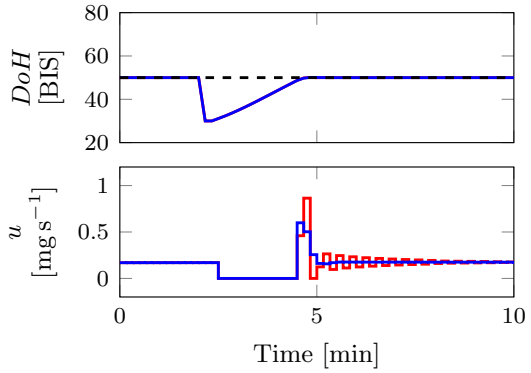

Fig. 3: DoH and infusion rate u with (red, $\alpha = 0$) and without (blue, $\alpha = 0.1$) ringing following a negative step disturbance $d_1 = -20$ BIS in Figure 5. The red and blue DoH curves are not visually distinguishable.

and where A_4 is the fourth row of A, and B_4 the fourth element of B. A complete derivation can be found in [13].

1) Preventing ringing in u: Figure 3 shows the solution that minimizes (9), following a negative step disturbance $d_1 = -20$ BIS. It results in undesired ringing in the infusion rate u (red). To avoid this, we add a cost term that penalizes sample-to-sample differences in u:

$$J'_{\Delta u}(\boldsymbol{u}) = \alpha \sum_{k=2}^{N} (u_k - u_{k-1})^2 \tag{11}$$

where α is a tuning parameter for this penalty.

Rewriting (11) as a quadratic form, and using that (11) is scalar, results in

$$J_{\Delta u}(\boldsymbol{u}) = \frac{\alpha}{2}\boldsymbol{u}^\top G^\top G\boldsymbol{u}, \tag{12}$$

where

$$G = \begin{bmatrix} -1 & 1 & & & \\ & -1 & 1 & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \end{bmatrix}. \tag{13}$$

The blue lines in Figure 3 show the effect of introducing this penalty with $\alpha = 0.1$. As can be seen, the red ($\alpha = 0$) and blue ($\alpha = 0.1$) DoH are not distinguishable, while there is no ringing in the blue control signal.

2) Constraints: We introduced constraints to keep the infusion rate nonnegative, $u_k \geq 0$ for $k = 1, \dots, N$. That is, element-wise larger than zero, $\boldsymbol{u} \succcurlyeq \boldsymbol{0}_{N\times 1}$, where $\boldsymbol{0}_{N\times 1}$ is a zero vector of size $N \times 1$. Similarly, the infusion rate is bounded by the maximum possible infusion rate of the pump, $u_{\max} = 1200\ \mathrm{mL\,h^{-1}}$, representative of several clinical infusion pumps. This corresponds to $u_{\max} = 6.67\ \mathrm{mg\,s^{-1}}$ with a propofol concentration of $20\ \mathrm{mg\,mL^{-1}}$. The corresponding constraint can be written as $\boldsymbol{u} \preccurlyeq u_{\max}\boldsymbol{1}_{N\times 1}$. The combined infusion rate constraints are thus

$$\begin{bmatrix} -I_{N\times N} \\ I_{N\times N} \end{bmatrix} \boldsymbol{u} \preccurlyeq \begin{bmatrix} \boldsymbol{0}_{N\times 1} \\ u_{\max}\boldsymbol{1}_{N\times 1} \end{bmatrix}. \tag{14}$$

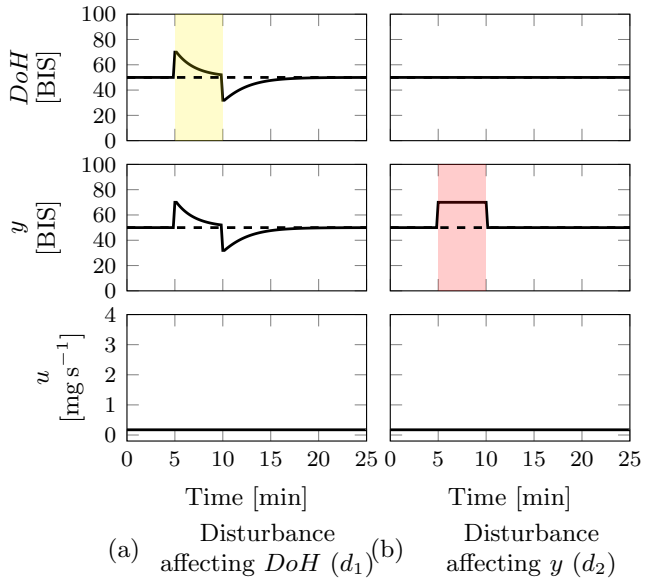

Fig. 4: Disturbance rejection by TCI for two types of disturbances; d_1 affecting the DoH directly (a, in yellow) and d_2 affecting the measurement (b, in red), as introduced in section II-B. The reference is shown in dashed.

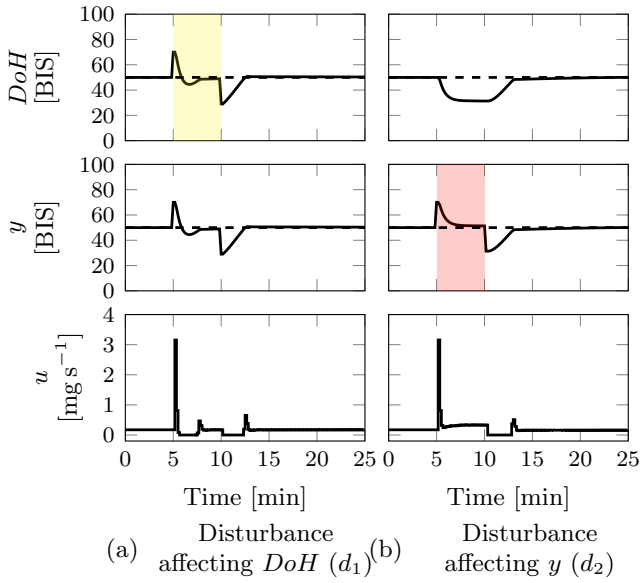

Fig. 5: Disturbance rejection by MPC for two types of disturbances; d_1 affecting the DoH directly (a, in yellow) and d_2 affecting the measurement (b, in red), as introduced in section II-B. The reference is shown in dashed.

3) MPC as a quadratic program: The MPC cost, combining (9) and (12), is

$$J(\boldsymbol{u}) = J_e(\boldsymbol{u}) + J_{\Delta u}(\boldsymbol{u}) \tag{15}$$
$$= \frac{1}{2}\boldsymbol{u}^\top (F^\top F + \alpha G^\top G)\boldsymbol{u} + (\boldsymbol{x}_0^\top E^\top F - \boldsymbol{r}^\top F)\boldsymbol{u}.$$

Minimizing (15) subject to (14) corresponds to solving a quadratic program (QP), for which there exist solvers such as `quadprog` in MATLAB [14] that were used in our examples.

In the closed-loop MPC case, one such QP is solved at each sample as a new measurement and the corresponding Kalman-filter state estimate arrive. In the TCI case, the optimization problem will only be solved once over a horizon, unless the anesthesiologist changes the reference and a new trajectory is computed. In this work, we use a practically sufficient prediction horizon of 10 min, corresponding to $N = 60$.

III. Disturbance rejection comparison

We evaluated the performance of TCI and closed-loop MPC subject to disturbances d_1 and d_2, as introduced in section II-B, and initially under the assumption of a perfect model, $\widehat{\text{PKPD}} = \text{PKPD}$ and $\hat{h}^{-1} = h^{-1}$ in Figure 1 and Figure 2.

The patient state is initialized to $DoH = 50$ BIS, i.e., $\boldsymbol{x}_0 = -A^{-1}Bu_{\text{ref}}$, where u_{ref} is the corresponding stationary control signal. At $t = 5$ min, a double step in d_1 or d_2 is introduced, with a duration of five minutes. The rejection of disturbances by TCI and closed-loop MPC is shown in Figure 4 and Figure 5, respectively.

Since we assume perfect model knowledge and no additional noise, the Kalman filter is tuned to behave like there is (almost) no state or measurement noise, i.e., R and Q are chosen to be (almost) zero (but not exactly zero for numerical reasons). The chosen values were $R = 10^{-6}$ mg^2/s^2 and $Q = 10^{-6} I_4$ mg^2/s^2, where I_4 is the identity matrix of size 4×4.

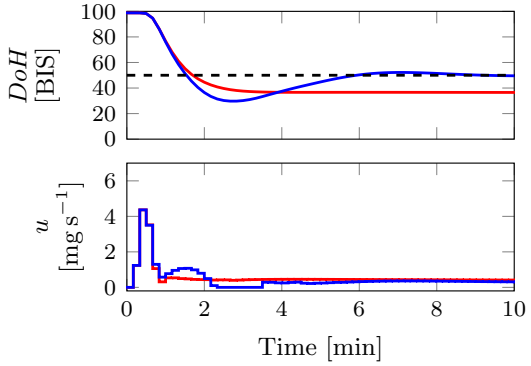

Fig. 6: Induction phase DoH and corresponding infusion rate u with (blue) and without (red) integral action. The reference $r_{DoH} = 50$ BIS is shown in dashed.

To compare disturbance rejection for open- and closed-loop control (TCI and closed-loop MPC), we calculate the mean square error (MSE) from the actual DoH and its reference $r_{DoH} = 50$ BIS. These are shown in Table I for the simulations in Figure 4 and Figure 5. The TCI is expected to perform well under measurement disturbances (d_2) but worse under disturbances that affect the DoH as it does not feedback on the measurement signal y. In contrast, the MPC was able to reduce the impact of disturbances acting on DoH, but overdoses when there is a measurement error, unable to differentiate between the two disturbances.

TABLE I: Mean squared errors (MSE) between actual DoH and the reference value of $r_{DoH} = 50$ BIS for each of the disturbance scenarios in Figure 4 and Figure 5, respectively. The disturbances d_1 and d_2 are detailed in section II-B.

Method	d_1	d_2
TCI	39.7	0
MPC	28.2	75.9

IV. Simulation scenario

While the examples of section III illustrate two extremes, we also include a more realistic scenario, where we study induction and maintenance for a representative example patient. The demographic data for this patient are the first patient in the data set presented in [15].

In this scenario, we assume non-perfect model knowledge of the PKPD model and corresponding Hill function. The MPC model $\widehat{\text{PKPD}}$ and the Kalman filter assume nominal parameter values (the same as in section III), while the PKPD model that constitutes the true dynamics of the patient (see the patient block in Figure 2) is now different. The true patient model is created by drawing from the log-normal distribution for volumes and clearances that make up the A matrix, as explained in section II-A and shown in (4). For the Hill function, $\hat{h}$, we use population averages $\hat{E}_0 = 95.9$ BIS, $\hat{C}_{e50} = 4.92\ \text{mg}\,\text{L}^{-1}$, and $\hat{\gamma} = 2.69$, as provided in [16].

The imperfect model would lead to a stationary error that can be eliminated in the closed-loop case by introducing integral action. This is done by adding a correction term to the control signal computed by the MPC, u_{MPC}, so that

$$u_k = (u_{\text{MPC}})_k + \beta(u_i)_k, \tag{16}$$

where the correction term is the sum of the integral error

$$(u_i)_k = \sum_{i=1}^{k-1}((C_e)_i - r_{Ce})_i, \tag{17}$$

and β is a tunable parameter determining the amount of integral action. For this simulation scenario, we use $\beta = 0.2$, which was found to be a suitable value to eliminate the stationary error. Figure 6 illustrates the need for integral action with model errors during the induction phase, where a stationary error is obtained and is not corrected by the MPC without integral action.

To further increase realism, we superimpose a noise sequence presented in [17] onto the measurement y used to drive the Kalman filter in the closed-loop case. As in [18] second-order filter with time constant 8 s, being the zero-order-hold discretization of

$$F(s) = \frac{1}{(8s+1)^2} \tag{18}$$

is employed to attenuate this noise.

During the induction phase, the patient goes from fully awake ($DoH \approx 100$ BIS) to a reference $r_{DoH} = 50$ BIS. We study a scenario where two disturbances affect the patient during maintenance. These are introduced at $t = 20$ min and $t = 40$ min and affect DoH and y, respectively, as specified in section III.

In this scenario, we assume that we have information about signal quality in terms of a signal quality index of 0-100 SQI. This index is provided in conjunction with the measurement by the most commonly used clinical monitors DoH, including the BIS monitor.

When the second disturbance d_2 affects the measurement, it is reflected as a drop in SQI from 100 (perfect measurement) to 50 (poor measurement). Then, the signal quality is poor until the disturbance disappears after five minutes. This simulates a scenario of using electrocautery devices, where electrical inference affects the BIS monitor and introduces measurement errors [19]. Details of how SQI is assumed to affect the measurement signal can be found in [11].

In [11], we developed a method to adjust R depending on the signal quality of the BIS signal, SQI, to seamlessly move between trusting the measurement or the model. Then, R was varied through an affine relationship between a minimum value $R_{\min}$ and a maximum value $R_{\max}$, so that

$$R(SQI) = R_{\min} + (1 - SQI/100)\,(R_{\max} - R_{\min}). \tag{19}$$

When $SQI = 100$, the signal quality is perfect, resulting in a small covariance $R_{\min}$. In contrast, a poor signal quality of $SQI = 0$ results in a large covariance $R_{\max}$.

The tuning of the Kalman filter ($R_{\min}$, $R_{\max}$, and Q) was performed by minimizing the mean square error (MSE) between the Kalman prediction of the DoH $\hat{y} = h(\hat{C}_e)$ and the true DoH over the simulation scenario with induction and the two disturbances. We assume that $R_{\min}$ and $R_{\max}$ are scalar and that Q is a diagonal matrix where the diagonal was identified in optimization. The MATLAB function `fmincon` with the interior-point algorithm was used for the optimization. The Kalman filter was initialized in $P_{0,0} = I_4$.

V. Results

Figure 7 shows the full simulation scenario of the induction and maintenance phase affected by disturbances as outlined in section IV The DoH, the measurement y, the infusion rate u, and SQI are presented. The signal quality is perfect ($SQI = 100$) throughout the time except at the time of the second disturbance d_2, which affects the measurement. The obtained values of $R_{\min}$, $R_{\max}$, and Q from optimization were $R_{\min} = 274, R_{\max} = 19706$, and $Q = \text{diag}(0.85, 0.83, 2.5 \cdot 10^4, 1.4 \cdot 10^3)$.

As seen in Figure 7, the MPC can control DoH closely to the reference except when disturbances occur but is immediately able to regulate back. When SQI drops from 100 to 50 at $t = 40$ min, the MPC can instead run on the underlying (incorrect) PKPD model in the meantime to prevent overdosing.

VI. Discussion

TCI performs well when there are measurement errors, but struggles when errors directly impact the patient's DoH. This issue arises particularly in scenarios that involve disturbances

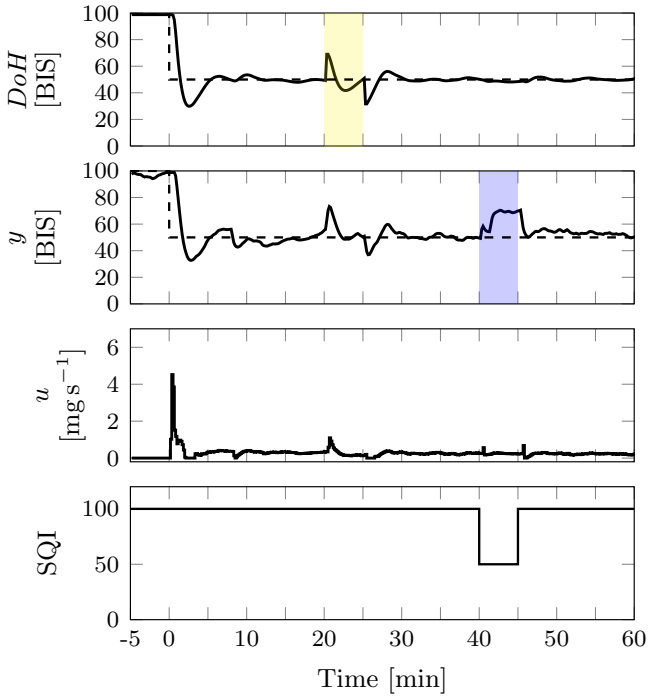

Fig. 7: Simulation example with induction and maintenance phase, subjected to noise and two disturbances (yellow, affecting the DoH, and blue affecting the measurement). These correspond to d_1 and d_2 in Figure 2 and are detailed in section II-B. At the same time as the second disturbance takes place, the SQI drops from 100 to 50 (lower plot). A Kalman filter has been optimized and is detailed in section IV. The reference $r_{DoH} = 50$ BIS is shown in dashed.

and model-error mismatches. However, the opposite is true for MPC, which can manage measured changes in hypnotic depth through feedback, even when these changes are not accounted for in the model. This capability is beneficial in situations where disturbances, noise, or patient-model mismatches are present. In summary, both regimens have strengths and weaknesses and are often complementary. In [11], it was shown that the suggested method was robust to model uncertainties regarding inter- and intra-patient variability.

In this work we have introduced a control structure that seamlessly integrates TCI and closed-loop control, allowing for a continuous re-positioning between the two based on operating circumstances. This framework utilizes a measurement-driven state observer in the form of a Kalman filter, that is re-tuned to rely more or less on model and measurement respectively, thus moving seamlessly between TCI and closed-loop MPC behavior.

REFERENCES

[1] Anthony R. Absalom and Keira P. Mason, eds. *Total Intravenous Anesthesia and Target Controlled Infusions: A Comprehensive Global Anthology*. Cham: Springer, 2017. DOI: 10.1007/978-3-319-47609-4.

[2] A. Absalom et al. "Accuracy of the 'Paedfusor' in Children Undergoing Cardiac Surgery or Catheterization". In: *British Journal of Anaesthesia* 91.4 (2003), pp. 507–513. DOI: 10.1093/bja/aeg220.

[3] T. W. Schnider et al. "The influence of method of administration and covariates on the pharmacokinetics of propofol in adult volunteers". In: *Anesthesiology* 88.5 (1998), pp. 1170–1182. DOI: 10.1097/00000542-199805000-00006.

[4] Klaske van Heusden et al. "Optimizing Robust PID Control of Propofol Anesthesia for Children: Design and Clinical Evaluation". In: *IEEE Transactions on Biomedical Engineering* 66.10 (2019), pp. 2918–2923. DOI: 10.1109/TBME.2019.2898194.

[5] Michele Schiavo et al. "Optimized Feedforward Control of Propofol for Induction of Hypnosis in General Anesthesia". In: *Biomedical Signal Processing and Control* 66 (2021), p. 102476. DOI: 10.1016/j.bspc.2021.102476.

[6] J. Vuyk and M. Mertens. "Bispectral index scale (BIS) monitoring and intravenous anaesthesia". In: *Advances in Modelling and Clinical Application of Intravenous Anaesthesia* (2003), pp. 95–104. DOI: 10.1007/978-1-4419-9192-8_9.

[7] Marko M. Sahinovic, Michel M. R. F. Struys, and Anthony R. Absalom. "Clinical Pharmacokinetics and Pharmacodynamics of Propofol". In: *Clinical Pharmacokinetics* 57.12 (2018), pp. 1539–1558. DOI: 10.1007/s40262-018-0672-3.

[8] Kristian Soltesz, Klaske van Heusden, and Guy A. Dumont. "5 - Models for Control of Intravenous Anesthesia". In: *Automated Drug Delivery in Anesthesia*. Ed. by Dana Copot. Academic Press, Elsevier, 2020, pp. 119–166. DOI: 10.1016/B978-0-12-815975-0.00010-2.

[9] Ylva Wahlquist, Amina Gojak, and Kristian Soltesz. "Identifiability of Pharmacological Models for Online Individualization". In: *IFAC-PapersOnLine* 54.15 (2021), pp. 25–30. DOI: 10.1016/j.ifacol.2021.10.226.

[10] K. Soltész. "On automation in anesthesia". PhD thesis. Lund, Sweden: Lund University, 2013. ISBN: 978-91-7473-484-3.

[11] Ylva Wahlquist et al. "Kalman Filter Soft Sensor to Handle Signal Quality Loss in Closed-Loop Controlled Anesthesia". In: *Biomedical Signal Processing and Control* (2025). DOI: 10.1016/j.bspc.2025.107506.

[12] G. Welch and G. Bishop. *An introduction to the Kalman filter*. Tech. rep. UNC-Chapel Hill, 1995. DOI: 10.5555/897831.

[13] Ylva Wahlquist, Amanda Gustafson, and Kristian Soltesz. "Exploring the Influence of Patient Variability on Propofol Target-Controlled Infusion Performance". In: *Proceedings of the 22nd European Control Conference* (2024).

[14] MathWorks. *quadprog*. https://se.mathworks.com/help/optim/ug/quadprog.html. Downloaded: 2024-09-18. 2024.

[15] Ionescu Clara M. et al. "Robust Predictive Control Strategy Applied for Propofol Dosing Using BIS as a Controlled Variable During Anesthesia". In: *IEEE Transactions on Biomedical Engineering* 55.9 (2008), pp. 2161–2170. DOI: 10.1109/TBME.2008.923142.

[16] A. Vanluchene et al. "Spectral entropy as an electroencephalographic measure of anesthetic drug effect: A comparison with Bispectral Index and processed midlatency auditory evoked response". In: *Anesthesiology* 101 (2004), pp. 34–42. DOI: 10.1097/00000542-200407000-00008.

[17] Andrzej Pawlowski et al. "MPC for Propofol Anesthesia: The Noise Issue". In: *2022 IEEE Conference on Control Technology and Applications (CCTA)*. 2022, pp. 1087–1092. DOI: 10.1109/CCTA49430.2022.9966102.

[18] Klaske van Heusden et al. "Quantification of the Variability in Response to Propofol Administration in Children". In: *IEEE Transactions on Biomedical Engineering* 60.9 (2013), pp. 2521–2529. DOI: 10.1109/TBME.2013.2259592.

[19] Matthew T. V. Chan, Sin Shing Ho, and Tony Gin. "Performance of the Bispectral Index during Electrocautery". In: *Journal of Neurosurgical Anesthesiology* 24.1 (2012), pp. 9–13. DOI: 10.1097/ANA.0b013e31823058bf.