Advanced PID architectures for tracking changing active constraints

Sigurd Skogestad *

* *Department of Chemical Engineering, Norwegian University of Science and Technology (NTNU), Trondheim (e-mail: sigurd.skogestad@ntnu.no)*

Presented at IFAC World Congress 2026, Busan

Abstract: Advanced regulatory control (ARC), also known as advanced PID architectures, is a simple and robust way of controlling processes with changing and possibly conflicting constraints, where it previously was believed - at least in academia - that model-based solutions, such as MPC, were the only effective solution. To illustrate this, ARC is applied in two case studies. The first is a gas-liquid separation process, in which selectors and split-parallel control are combined to achieve bidirectional inventory control in which the throughput manipulator moves automatically to the most optimal position. The second case study is on keeping acceptable air quality (CO2-level) and temperature in a room (in this case, a barn for cows). The CO2 and temperature constraints can be conflicting, leading to a hierarchical switching network of PID controllers.

Keywords: process control, PID control, decentralized control, selectors, hybrid control

1. INTRODUCTION

An active constraint is a constraint which, in the current situation, should be kept at its limiting (maximum or minimum) value to achieve optimal (economic) operation at steady state. The "current situation" is defined by the current value of disturbances, parameters, and prices in the cost function. To achieve optimal economic operation, the most important is usually to track and control the active constraints (Maarleveld and Rijnsdorp, 1970).

For about 40 years, since the 1980s, there has been a myth, at least in the academic community, that this requires model-based schemes, such as model predictive control (MPC) and real-time optimization (RTO). Of course, academic researchers had some knowledge about industrial schemes for dealing with constraints, such as selectors and split range control, but such "override" schemes (e.g., as described in Maarleveld and Rijnsdorp (1970) for an industrial distillation column) were thought to be *ad hoc*, old-fashioned and not optimal, and not worth the effort of more detailed studies. Until about 10 years ago this was also my view.

Selectors. Consider the case where a single manipulated variable (MV) (input) should control more than one controlled variable (CV) (output). MAX- and MIN-selectors are simple logic elements (sometimes referred to as overrides) for switching the CV (output) in response to a change in the active output constraints. Only one CV should be controlled at a time, which is known as *CV-CV switching*.

Split range and split parallel control. In this paper, we also consider the case where two or more MVs are used to control a single CV. Only one MV should be used at a time, which is known as *MV-MV switching*. The traditional solution is to use split range control, but this requires the user to specify the MV values (say, 0% and 100% in the simple case of valve saturation) where switching should occur. This is avoided with split parallel control where each MV has its own controller with a different CV setpoint. Here, the switching occurs by feedback when an MV saturates (or is taken over by another controller) and control of the CV is lost.

These and other ARC schemes were discussed in a recent paper (Skogestad, 2023). The objective of the present paper is to extend these findings (with some partly new rules) and show with two case studies, including case study I on bidirectional inventory control, the power of selectors and split parallel control.

2. CASE STUDY I: GAS-LIQUID SEPARATOR

2.1 General about inventory control and TPM

Maintaining throughput and managing and controlling inventories (e.g., levels) are of vital importance for process operation. An important concept is the throughput manipulator (TPM), which is the variable that sets the throughput (production rate) for the entire process. Most processes have only one TPM, which also sets the feed rate(s) and product rate(s). Additional feeds are usually set by ratio control to the main feed. The location of the TPM has a very strong impact on the resulting control system. The TPM is traditionally located at the (main) feed or at the (main) product, but in general it could be located anywhere in the process. To maximize the throughput, which often has a large economic benefit, the TPM should be located at the production bottleneck so that the bottleneck constraint can be tightly controlled (with minimum "back-off").

The inventory control system needs to be consistent, which means that the mass balances are satisfied at all locations at steady state. We have the following consistency rules regarding inventory control and TPM location:

- *Rule C1. One cannot control (set the flowrate) the same flow twice.*
- *Rule C2. Controlling the inlet or outlet pressure indirectly sets the in- or outflow (and indirectly makes the MV (valve) used for pressure control a TPM).*
- *Rule C3. Follow the **radiation rule** whenever possible.*
- *Rule C4. No inventory loop can cross the TPM.*

The radiation rule says that, to get a consistent inventory control system using local control loops, the inventory control loops must be radiating around the TPM (Price and Georgakis, 1993). This means that inventory control must be in the direction of flow downstream of the TPM, and opposite the direction of flow upstream the TPM, see Figure 1.

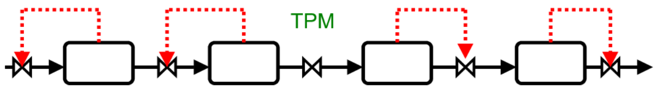

Fig. 1. Radiation rule for inventory control

Traditionally, the TPM location and thus the inventory control system is fixed. However, a fixed inventory control system means that changes in the operation, including reaching new production bottlenecks, temporary stops for maintenance, batch operation, and changes in the possible feed and product rates, will require manual intervention. This puts a heavy burden on the operators and leads to non-optimal operation. Shinskey (1981) has proposed a very clever bidirectional control architecture for automatic reconfiguration of inventory control loops and which also automatically moves the TPM to the current bottleneck. This approach is next applied the important industrial case of gas-liquid separation of a two-phase stream from a well.

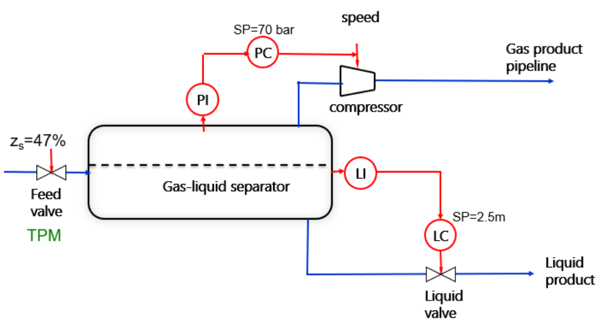

Fig. 2. Nominal control system for case study I with TPM at feed (no constraints encountered).

2.2 Application of bidirectional control to case study I

The nominal control system where the operator manually sets the feed rate (choke valve position z) is shown in Figure 2. In this case, the TPM is at the feed, so inventory control of the downstream separator tank is in the direction of flow: The gas inventory is controlled using the pressure controller (PC) with the compressor speed as the manipulated variable (MV), and the liquid inventory is controlled using the level controller (LC) with the liquid outflow valve as the MV.

However, if the operator opens the feed valve z too much, it may happen that the upstream well pressure becomes too low (p_{min} = 170 bar) which may damage the reservoir. This may be solved using a minimum pressure control override as shown in Figure 3. From Rule C2, since we are setting the inlet pressure, the TPM remains at the feed valve, and the downstream inventory control system for level and pressure is unchanged. Note that there is a MIN-selector that allows the operator to reduce the feed rate when desired by setting a smaller value for z_s. On the other hand, if the objective is to maximize production, the operator may set z_s= 100% (or, equivalently, the operator may set $z_s = \infty$, since the valve anyway has a build-in maximum constraint at $z_{max} = 100\%$).

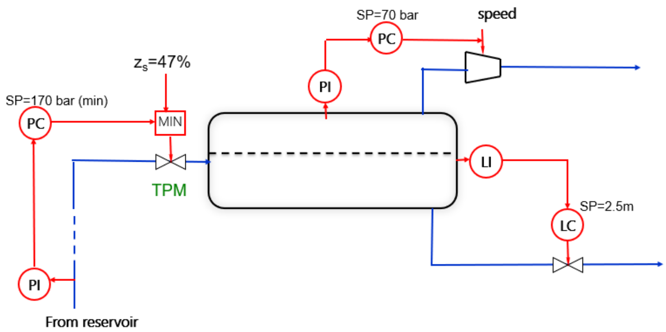

Fig. 3. TPM at feed with override selector to satisfy minimum well pressure of 170 bar at reservoir.

If the gas flow from the reservoir becomes too large, then the compressor speed may go it its maximum value (100%) and the separator pressure can no longer be kept at its nominal setpoint of 70 bar. The solution is to use an ”override” where we reduce the feed, that is, to let the feed valve take over as the MV for the pressure control. This is a case of MV-MV switching which may be achieved with either split range control (one controller and a figure or table with MV switching values) or with split parallel control (two controllers with different setpoints) (Skogestad, 2023). The first option becomes a bit complicated because the switching value for the feed valve is not 100%, but given through the MIN-selector by the minimum of the operator setpoint (z_s) and the varying “override” value from the reservoir pressure controller (with SP=170 bar).

We therefore use split parallel control as shown in Figure 4, which also has the advantage of allowing different controller tunings for the two pressure controllers (PC$_A$ and PC$_B$). Since the separator pressure will start rising when the compressor saturates at 100%, the setpoint SPH for PC$_B$ should be higher than SPL for PC$_A$, that is, SPH = SPL + Δ. With Δ = 1 bar, this gives SPH = 71 bar.

When split-parallel control is combined with MIN-selectors as shown in Figure 4, this results in ”bidirectional inventory control” (Shinskey, 1981; Zotică et al., 2022) of the separator pressure, where the throughput is maximized

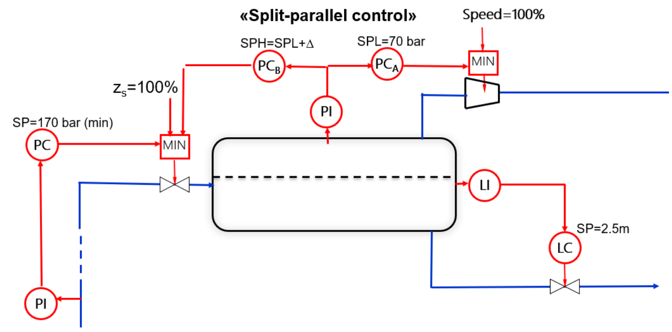

Fig. 4. Bidirectional control scheme where TPM moves automatically to the bottleneck (feed valve or compressor) to maximize throughput.

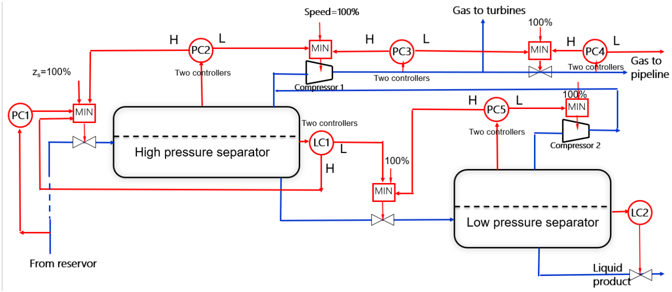

Fig. 5. Final bidirectional control scheme for case study I for more complex process with additional constraints. The TPM can be at all 5 MIN-selectors.

and the TPM is automatically moved to the present bottleneck (feed valve or compressor).

Note the following about the bidirectional scheme in Figure 4:

- We used the term "override" for controller PC_B, but this is a bit misleading because this is the optimal thing to do. This is illustrated by the fact that, if we start from a situation where the TPM is at the compressor (e.g., with the compressor speed set at 100%), then controller PC_A becomes the "override" for pressure control if the feed rate becomes the bottleneck.
- For the split-parallel control, there should normally be a period during the MV-MV switching (between the SPL and SPH setpoints) where none of the two controllers (PC_A and PC_B) are active, because both MVs are either saturated or set by another controller. However, if one uses a too small setpoint difference Δ, then both PC_A and PC_B (which control the same pressure) may be active at the same time during the MV-MV switching and start fighting. This may slow down the time for switching back when we go out of a constraint (e.g., the compressor is no longer a bottleneck) and we want to increase the throughput (this is related to internal instability; see a preliminary analysis in Forsman et al. (2025)). Thus, it is likely that a setpoint difference $\Delta = 1$ bar is too small. A larger setpoint difference also has the advantage of delaying the switch between using the compressor and feed valve as MVs for separator pressure control, which may be an advantage if the constraints are only temporary.

The bidirectional inventory control idea may be extended to the more complex case in Figure 5, where we have added a low-pressure separator with another compressor, a gas turbine and additional constraints. In Figure 5, we also use a simplified notation where the two split-parallel controllers (e.g.., PC_A and PC_B in Figure 4) are combined into one controller block (e.g., PC2) with two setpoints L=SPL and H=SPH, where H=L+Δ.

In summary, the following seven constraints are handled in Figure 5:

(1) Min. reservoir pressure: "Override" from PC1 using a MIN-selector.
(2) Saturation (max. speed) for Compressor 1: PC2-H reduces the feed rate (after some transition time from setpoint L to H).
(3) Saturation (max. valve opening) for liquid leaving separator: LC1-H reduces the feed rate (also here after some transition time from setpoint L to H).
(4) Saturation (max. speed) for Compressor 2: PC5-H reduces the liquid feed. This results in loss of level control (LC1-L), so LC1-H will reduce the feed rate.
(5) Min. pressure for gas to turbines: PC3-L closes the extra gas valve (after the split).
(6) Max. pressure at inlet gas pipeline: PC4-H closes the same extra gas valve.
(7) Min. pressure at inlet gas pipeline. PC4-L closes a downstream valve (not shown in Figure 5).

Note the following in Figure 5:

- The bidirectional scheme automatically moves the TPM (after some transition time) to the MV-location (valve/compressor) determined by which of the above seven constraints is active.
- All MIN-selectors have an external setpoint, which allows the operator to reduce the flow and move the TPM to this location, if desired. In Figure 5, the setpoints are set at 100% because the objective is to maximize throughput. More generally, the external setpoint could be the optimal unconstrained value (without constraints).
- All controllers reduce the MV value (valve/speed) when activated (and do the opposite when deactivated).
- Note that all the constraints are consistent in that they all are satisfied (in the end) by reducing the feed rate. Thus, we only need MIN-selectors.
- We have H=L+Δ for the setpoints of the five split-parallel controllers. Note that with a larger setpoint separation Δ, the buffer of gas or liquid between setpoints L and H (while no control loops are active) will delay the transition. Thus, it will take more time before the feedrate needs to be reduced, which is an advantage economically.
- The inlet pressure to the gas pipeline is often "floating", that is, uncontrolled, which means that the valve between PC3 and PC4 is fully open (as in Figures 2-4) which is an economic advantage as it reduces compression power. In this case, the values for H and L for PC4 may be set far apart, for example, H=200 bar and L=100 bar.

The solution in Figure 5 may at first seem complex, but note that the structure is repeating: all valves/compressors have MIN-selectors with an L-setpoint from the upstream unit, an H-setpoint from the downstream unit and a desired setpoint (which can be given up; in Figure 5 it is 100% because we want to maximize production). Correspondingly, all inventory controllers (LC and PC) have an H-setpoint to the upstream valve/compressor and an L-setpoint to the downstream valve/compressor.

3. CASE STUDY IIA: HAPPY COWS

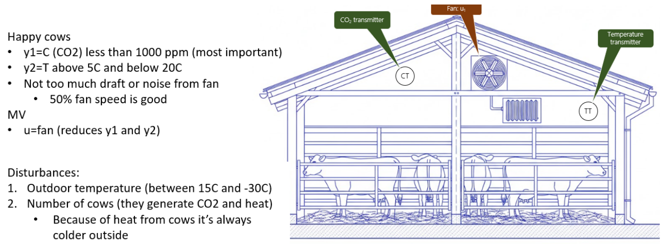

Fig. 6. Case study IIA: Happy cows with fan speed as MV

Consider the happy cow case study in Figure 6 with the fan speed as the manipulated variable (MV). To keep the cows happy, the air quality should be good with CO2-concentration (y_1) below 1000 ppm, and the temperature $y_2 = T$ in the barn should be between 5C and 20C. Nominally (and hopefully most of the time), the fan speed u will be fixed at the desired value (nominal setpoint) of 50%. However, if there are too many cows, the air quality may get bad (with CO2 rising above 1000 ppm) or the temperature may get too high (above 20C) and we need to increase the fan speed above 50%. On the other hand, on a cold winter day, the temperature may get too low (below 5C) and we need to reduce the fan speed below 50($y_2 > 5$C).

In summary, there are three control objectives (constraints) (one on y_1 and two on y_2) and only one manipulated variable (u=fan speed). We will use one controller for each of the three constraints and selectors on the three controller outputs so that only one constraint is actively controlled at a time.

- *Rule S1: In general, we need a MAX-selector for constraints that are satisfied by a large value of the input u, and a MIN-selector for constraints that are satisfied by a small value of the input u (Skogestad, 2023).*

From selector rule S1, we have for the three constraints:

(1) Constraint $y_1 \leq 1000$ ppm: Satisfied by large fan speed $\Rightarrow$ MAX-selector
(2) Constraint $y_2 \geq 5$C: Satisfied by small fan speed $\Rightarrow$ MIN selector
(3) Constraint $y_2 \leq 20$C: Satisfied by large fan speed $\Rightarrow$ MAX selector

In general, when we require both MAX- and MIN-selectors for an input (here, for u=fan speed), there may be cases where the constraints are conflicting and there is no feasible solution.

- *Rule S2: In cases where both a MAX- and MIN-selector is needed, and there is a potential conflict, we must decide which constraint is the most important and put the corresponding selector "at the end" of the sequence (closest to the MV (input) u).*

In our case, the MIN- and MAX-selectors for constraints 2 and 3 will never conflict because they are on the same variable, $y_2 = T$, and temperature cannot be both 5C and 20 C at the same time. A feasible control structure for band temperature control is shown in Figure 7. Here, the MIN-selector is last (at the end), but the order could be interchanged.

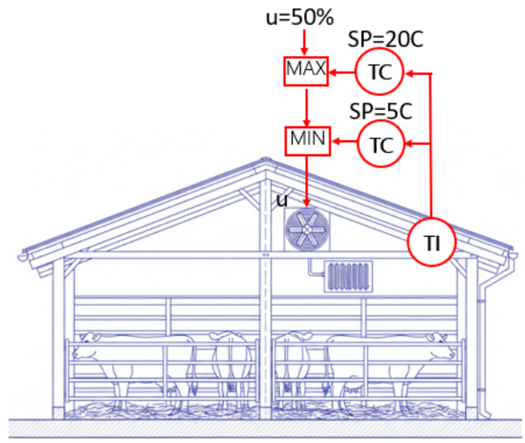

Fig. 7. Case study IIA: Control of temperature y_2 within a band. Since the constraints never conflict, the order of MAX- and MIN-selectors may be interchanged or they may be replaced by a MID-selector.

However, the order of the selectors matters when we include the CO2-constraint. When it is cold outside and many cows, constraint 1 on CO2-level (max. 1000 ppm) (requiring a large fan speed and a MAX-selector) may conflict with constraint 2 on temperature (min. 5C) (requiring a small fan speed and a MIN-selector). Since constraint 1 on CO2 has higher priority (see Figure 6), the MAX-selector should be at the end. This means that we may need to accept that $y_2 = T$ drops below 5C when it is very cold outside. The final proposed control structure, with three controllers and two selectors, is shown in Figure 8. It is suggested to use PI-controllers, and because at most one controller is active at a time, all three controllers need to have anti windup on the integral action.

With the given constraints and priorities, the proposed control structure will be optimal at steady state, and also the dynamic responses are expected to be good. MPC is an alternative, but it requires a dynamic model and online optimization and it seems unnecessary complicated.

Note that we have put a nominal desired value $u = 50\%$ into the first selector block in Figures 7 and 8. More generally, we have:

- *Rule S3: It is recommended to always put a "desired input" $u = u_0$ (50% in Figures 7 and 8) into the first selector, and sometimes into selectors later in the sequence. Otherwise, there may not be a unique solution in cases when no constraints are active. Here, u_0 could be a given desired value (e.g., set by the operator), an optimized value (resulting from*

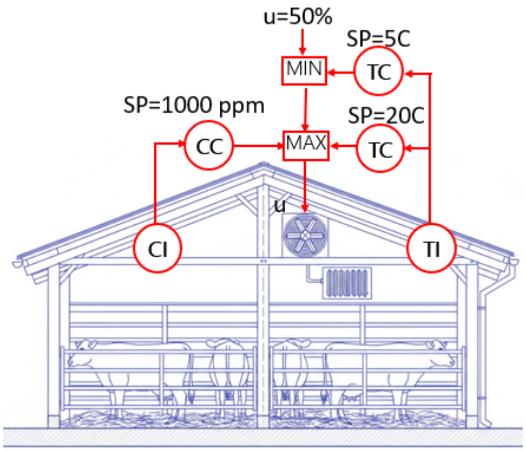

Fig. 8. Case study IIA: Control of CO2 and temperature (MAX-selector must be at end because max. CO2 of 1000 ppm has priority relative to min. T of 5C).

an unconstrained optimization), or the output from a controller which has a setpoint which can be given up.

Comment: In some cases, the desired u_0 is the maximum value for the MV (100%) (e.g., maximize production), which is relevant for a MIN-selector, and in this case we may omit the desired value $u_0 = 100\%$ (and still get a unique solution) because all MVs (valves) have a built-in MIN-selector at $u_{max} = 100\%$ (selector Rule 3 in Skogestad (2023)). Nevertheless, it may be good to keep the desired value $u_0 = 100\%$ for clarity. An example is given in Figure 4 on bidirectional control for case study I.

In other cases, the desired u_0 is to use the minimum input (0%), which is relevant for a MAX-selector, and in this case we may omit the specification $u_0 = 0\%$ (and still get a unique solution) because all MVs (valves) have a built-in MAX-selector at $u_{min} = 0\%$. Nevertheless, it is good to keep the desired input $u_0 = 0\%$ for clarity.

4. CASE STUDY IIB: AVOIDING UNHAPPY COWS

In this extension of case study IIA, we are more worried about cold days, so we add a heater (which is rather expensive to use) as an MV. However, on really cold days, the heater may saturate at maximum, and to avoid unhappy freezing cows (with $y_2 = T$ below 0C), we will allow reducing the fan speed and let the CO2 level go above 1000 ppm. However, to avoid very unhappy (sick) cows, we cannot let the CO2 level go above 3000 ppm. The specifications for case study IIA are summarized in Table 1.

The final proposed control structure is shown in Figure 10 and its design is explained in the following.

Split-parallel control using heater. First, note that both manipulated variables (MVs) u_1 and u_2 affect the temperature, but we only want to manipulate one MV at a time for temperature control. This is a case of MV-MV switching which may be achieved with split range control (one controller plus a figure or table with MV switching values) or split parallel control (two controllers with different setpoints) (Skogestad, 2023). The first option becomes a bit complicated because the switching value for the fan is not 0%, but given by a varying "override" value from the CO2-controller (SP=1000 ppm) through the MAX-selector. We therefore use split parallel control, which also has the advantage of allowing different controller tunings for the two TCs. Since reducing the fan speed (u_1) is the cheap alternative to keep the temperature above 5C, it

Table 1. Cows' comfort conditions and manipulated variables for Case study IIB.

Happy cows	
$y_1 = c_{CO2}$	less than 1000 ppm
$y_2 = T$	between 5 °C and 20 °C
Nominal (desired)	50% fan speed
Unhappy (freezing) cows	
$y_2 = T$	less than 0 °C
Very unhappy (sick) cows	
$y_1 = c_{Co2}$	above 3000 ppm
MVs (inputs)	
u_1 = fan speed	decreases y_1 and y_2 (cheap)
u_2 = heater	increases y_2 (expensive)

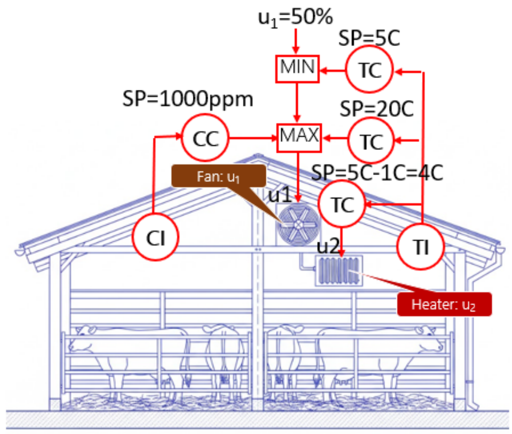

Fig. 9. Intermediate control structure for case study IIB with addition (compared to Figure 8) of split-parallel control using heater u2.

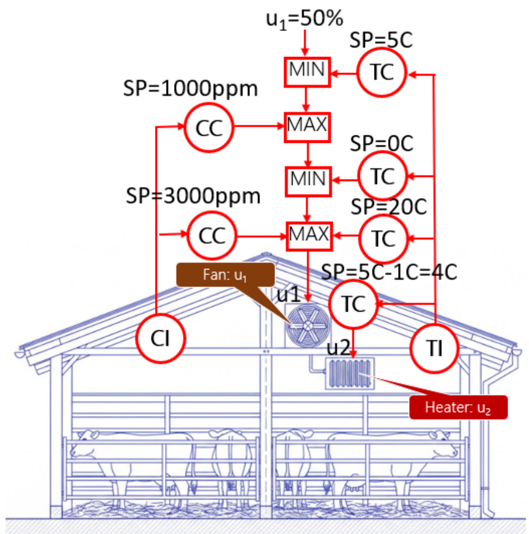

Fig. 10. Final control structure for case study IIB: With additional constraints to avoid unhappy cows.

should be used first, with a setpoint of 5C. The setpoint for the heater (u_2) is then 5C-Δ, where we choose a setpoint difference Δ =1C. The heater will then activate on cold days when the CO2-controller with SP=1000 ppm is using the fan (through the MAX-selector), and the temperature drops to 4C. The proposed intermediate control structure with additional of split-parallel control is shown in Figure 9.

Really cold days with additional constraints. The heater is not very large so if it gets really cold outside it may saturate at maximum (100%) and temperature control is lost and the temperature drops below 4C. However, if it reaches 0C, we override the CO2-controller (and give up keeping SP=1000 ppm) by reducing the fan speed with a MIN-selector. The CO2-concentration is then uncontrolled, but we cannot allow the CO2-level to become too high because the cows get sick. If it reaches 3000 ppm, we override the temperature control (and give up keeping SP=0C) by increasing the fan speed with a MAX-selector.

In this case, the temperature will drop below 0C, so we probably should cover the cows with blankets (as a third manual MV). Similarly, on hot days where the temperature exceeds 25C we may spray water on the cows (as a fourth manual MV).

The final control structure in Figure 10 may seem complicated, but note that we are using one controller for each constraint, with the highest priority (and non-conflicting) constraints located at the end of the sequence, closest to the input u_1 (fan speed).

5. DISCUSSION

There are many theoretical issues to be resolved about the dynamic behavior and stability of switching schemes, such as the one shown for case study IIB in Figure 10. For example, for certain parameter values and noisy measurements, one may get chattering (e.g., Caponigro et al. (2018)) and even chaotic behavior. This complicates the theoretical analysis, and may be one reason why there has been little work in academia on such schemes. In practice, industrial experience over more than 60 years (Maarleveld and Rijnsdorp, 1970; Smith, 2010), have shown that with proper design, the selector-based swithcing schemes work well with smooth transitions between changing active constraints. In some cases, to avoid chattering, one may put restrictions on the switching frequency (e.g., dead-band or dwell time) but with a good anti-windup scheme (e.g., the tracking scheme of Astrom described in Skogestad (2023)) it is usually not necessary.

Although a hierarchical (and decentralized) switching network of PID controllers may be used in complex cases (case study IIB), there are certainly other cases where it does not apply, and where model-based schemes such as MPC/RTO will be a better, and sometimes even simpler, solution. Specifically. the switching schemes presented in this paper, assume that one makes a pairing between each constraint and a single manipulated variable (input), and with multiple constraints (associated with one input) requiring both a MAX- and MIN-selector, this may result in an apparent infeasibility, which could have been avoided with a multivariable scheme. Bernardino and Skogestad (2024) show that if we have at least as many inputs as constraints (which is restrictive and not satisfied for case study IIB) then it is possible to find a decentralized switching scheme that achieves optimal operation at steady state. Rules and recommendations for more general cases with many constraints are lacking.

6. CONCLUSION

In the academic community, there ha been a myth that you need a detailed nonlinear model and an on-line optimizer (RTO) if you want to optimize the process, and you need a dynamic model and model predictive control (MPC) if you want to handle constraints. The objective of this paper, is to show by some simple case studies that this is not true: Because the most important for optimal operation is usually to implement the active constraints, one can in many cases obtain optimal operation by measuring and controlling the constraints. This is fairly obvious, but what is probably less clear is that the MIN-selectors in the ARC schemes also switch back, that is, they stop controlling the constraint when the disturbance goes away. Today advanced PID solution (ARC) and model-based control and optimization (MPC/RTO) are in parallel universes, but both solutions are needed in the future control engineer's toolbox.

ACKNOWLEDGEMENTS

Discussions with Krister Forsman and Mohammed Adlouni are gratefully acknowledged.

REFERENCES

Bernardino, L. and Skogestad, S. (2024). Decentralized control using selectors for optimal steady-state operation with changing active constraints. *Journal of Process Control*, 136, 103194.

Caponigro, M., Ghezzi, R., Piccoli, B., and Trélat, E. (2018). Regularization of chattering phenomena via bounded variation controls. *IEEE Transactions on Automatic Control*, 63, 2046–2060.

Forsman, K., Adlouni, M., and Skogestad, S. (2025). Split parallel control - a little known control structure. *IFAC-PapersOnLine (DYCOPS conference)*, 59(6), 552–557.

Maarleveld, A. and Rijnsdorp, J. (1970). Constrained control of distillatiomn columns. *Automatica*, 6, 51–58.

Price, R.M. and Georgakis, C. (1993). Plantwide regulatory control design procedure using a tiered framework. *Ind. Eng. Chem. Res.*, 32, 2693–2705.

Shinskey, F.G. (1981). *Controlling Multivariable Processes.* Instrument Society of America.

Skogestad, S. (2023). Advanced control using decomposition and simple elements. *Annual Reviews in Control*, 56, 100903.

Smith, C.L. (2010). *Advanced Process Control - Beyond single-loop control.* Wiley, New York.

Zotică, C., Forsman, K., and Skogestad, S. (2022). Bidirectional inventory control with optimal use of intermediate storage. *Computers and Chemical Engineering*, 159, 107677.