



LUND UNIVERSITY

A practical guide to PID controller implementation

Sundström, Emil; Bauer, Margret; Guzmán, José Luis; Hägglund, Tore; Soltesz, Kristian

Published in:
International Journal of Control

2026

Document Version:
Peer reviewed version (aka post-print)

[Link to publication](#)

Citation for published version (APA):
Sundström, E., Bauer, M., Guzmán, J. L., Hägglund, T., & Soltesz, K. (in press). A practical guide to PID controller implementation. *International Journal of Control*. <https://arxiv.org/abs/2604.15918>

Total number of authors:
5

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

A Practical Guide to PID Controller Implementation

E. Sundström^a, M. Bauer^b, J. L. Guzmán^c, T. Hägglund^a and K. Soltesz^a

^aDept. of Automatic Control, Lund University, Box 118, SE-22100 Lund, Sweden
(email: {emil.sundstrom, tore.hagglund, kristian.soltesz}@control.lth.se)

^bDept. of Process Engineering, Hamburg University of Applied Sciences, Hamburg, Germany
(email: margret.bauer@haw-hamburg.de)

^cDept. of Informatics, Universidad de Almería, ceiA3, CIESOL, Ctra. Sacramento s/n, 04120 Almería, Spain (email: joseluis.guzman@ual.es)

ARTICLE HISTORY

Compiled 13th July 2026

Abstract

How difficult can it be to implement a PID controller? The answer is twofold. Implementing the PID control law is simple and computationally inexpensive. However, this basic form will not work in practical applications. The primary reason for this is the various physical limitations of the actuator. Measurement noise, different implementations depending on the various structures (P, PI, PD or PID), bump-less transfer, and varying sampling interval also result in problems rendering the basic form inoperable. PID implementation is therefore more difficult than meets the eye. This paper introduces a reference implementation of the PID controller which considers these practical issues. It includes pseudo-code, discussion of the implementation choices and simulation of carefully selected, important test cases.

KEYWORDS

PID control, Implementation, Programming code, Incremental form.

1. Introduction

The future of PID control was correctly predicted in an article 25 years ago (K. J. Åström & Hägglund, 2001): it is still the most prevalent control strategy in real world applications. PID is a superior strategy when assessed in terms of performance, tuning, ease of use and maintenance, resulting in a good solution for many process dynamics. Then and now, the PID implementation has been buried in proprietary archives of control technology suppliers. From MathWorks' MATLAB implementations to automotive, flight control, process control systems, etc.: the PID controller is generally implemented as a black box, and the inside workings are not disclosed to the user. In some instances, the archives are not well maintained and the PID controller is poorly implemented with serious problems. The knowledge of how to implement the PID control law correctly is neither communicated nor developed.

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. All authors from Lund University are members of the ELLIIT Strategic Research Area. Moreover, the paper was also partially funded by the Spanish Ministry of Science within the project PID2023-150739OB-I00.

The classical PID control law in the form

$$u(t) = K \left(e(t) + \frac{1}{T_i} \int_0^t e(\tau) d\tau + T_d \frac{de(t)}{dt} \right) \quad (1)$$

is easy to implement in hardware and software (Hägglund & Guzmán, 2024). While the nomenclature of Equation (1) is well-established, specifications of involved signals and parameters are listed in Appendix C, alongside those of other nomenclature introduced throughout the paper. While it is possible to implement the PID algorithm in analogue electronic circuits, mechanics, pneumatics and even biological systems, the most prevalent implementation is in programming code. The ease of implementation may explain why there appears to be only one publication solely dedicated to PID implementation (Clarke, 1984), covering some basic steps focusing on the z -domain. Several textbooks cover some basic aspects, too, e.g., K. Åström and Hägglund (2006); Shinskey (1994, 1996); Visioli (2006). In recent years, the void of PID implementation publications has been filled with a large number of internet educational resources providing implementations of the PID controller in various languages (C, C++, Python or MATLAB). Unfortunately, these implementation resources realise Equation (1) but rarely include the practical considerations required when implementing a real-world PID controller. Some publications describe student projects or experimental setups that implement the PID controller, see (Bhandari & Csurscia, 2022; Krejcar, Špička, & Frischer, 2011; Uzunović, Zunić, Badnjević, Mioković, & Konjicija, 2010). A clear description of how to implement the real-world PID controller is not available. Some efforts have been made to standardise the PID controller. The IEC has published a standard for evaluating the performance of PID control (IEC, 2010). The International Society of Automation (ISA) has issued a technical report on PID algorithms and performance that standardises PID control, focusing on nomenclature (ISA, 2023). There is currently no standard covering PID controller implementation. These circumstances make it difficult for control engineering professionals to find a concise and thoroughly explained reference implementation. As a result, PID implementations in commercial solutions are often ad hoc and may not include important features required to deal with practical problems. Anecdotally, the authors have heard of several industrial controllers without anti-windup that could not be used during process startup. There are numerous extensions to the standard controller, including fractional-order controllers (Chen, Petráš, & Xue, 2009), PID controllers that additionally consider the second derivative of the error (Huba & Vrančić, 2018), etc. These fall outside the scope of this work. Neither do we consider PID tuning. Instead, the sole focus of this paper lies in arriving at an algorithmic structure for a digital reference implementation of the PID controller. There are two alternative implementation forms: positional and incremental, the latter also referred to as velocity form (Clarke, 1984). Both are explained in detail in Section 2. Section 3 first describes the most common problems found in practice, extending the basic discretisation of the PID control law to deal with all of them effectively. The result, and our main contribution, is the reference implementation given in Section 4, which combines both positional and incremental form. The positional form is needed when there is no integral action, while the incremental form inherently provides bumpless behaviour when there is. Unlike the incremental form, the positional form could in principle reproduce the full functionality of the combined form, but only at the cost of special-case handling and storing numerous past signal values. The combined form provides this functionality without these drawbacks.

The reference implementation is provided in pseudo-code, discussed line by line, aided by a graphical representation of the code. While this means that the code cannot be taken directly into one development environment such as MATLAB, C, or Python, the readability is improved because the code provides a clearer description of the algorithm, without any distractions. Section 5 introduces additional implementation considerations such as the state initialisation, the execution interval and the runtime environment. The most common problems that were described in Section 3 are then simulated in Section 6 to show the benefit of extending the basic PID algorithm to practical working code. The simulation codes are available in a GitHub repository by Sundström, Bauer, Guzmán, Hägglund, and Soltesz (2026a) located at <https://github.com/copybit/PID.code>. A tutorial-oriented webpage by Sundström, Bauer, Guzmán, Hägglund, and Soltesz (2026b) that documents and exemplifies our reference implementation is found at <https://w3.ual.es/personal/joguzman/pid/>.

2. Digital PID implementation

To implement the controller in a computer, Equation (1) must be converted to a discrete time version. The principle of discrete time sampling is explained in this section. Both the integral and the derivative parts of the PID controller must be approximated for the conversion. There are two alternative implementations of the PID controller: the positional form and the incremental form (Clarke, 1984). Both will be described in the following. The positional form is more intuitive, but the incremental form has important practical advantages. The implementation presented in this paper will be based on a combination of the two to exploit the advantages of each form. The naming of variables and parameters in this implementation is based on the standard form described in Equation (1) where $u(t)$ is the control signal and $e(t) = r(t) - y(t)$ is the control error, that is, the difference between the setpoint $r(t)$ and the measured process variable $y(t)$. The basic control loop is depicted in Figure 1. Note that setpoint $r(t)$ and process variable $y(t)$ enter the PID block instead of control error $e(t)$. The reason for this depiction will become apparent in the following sections. The parameters of the PID controller




Figure 1.: The basic control loop.

in Equation (1) are controller gain K , integral time T_i , and derivative time T_d . Our implementation uses the representation

$$u(t) = k_p e(t) + k_i \int_0^t e(\tau) d\tau + k_d \frac{de(t)}{dt} \quad (2)$$

with parameters $k_p = K$, $k_i = K/T_i$, and $k_d = KT_d$, since it is more general in the sense that it allows for $T_i = 0$ in Equation (1) without causing a division-by-zero error. If other forms are preferred, such as the series implementation (K. Åström & Hägglund, 2006), the implementation codes presented in this paper can be easily adapted.

2.1. Discrete time sampling

Time series sampling is the prerequisite for digital implementation of Equation (1). The continuous-time process variable $y(t)$ and the continuous-time setpoint $r(t)$ have to be sampled at discrete time intervals Δt . The discrete-time control error is then

$$e(t) = r(t) - y(t), \quad (3)$$

with $t = n\Delta t$, where n is the discrete time index. The control signal $u(t)$ is then calculated at discrete times $n\Delta t$. Equation (1) contains the integral and the derivative of the control error $e(t)$. A common way of discretising integral and derivative is to approximate the integral with a sum and the derivative with a difference such that

$$\int_0^t e(\tau) d\tau \approx \Delta t \sum_{j=0}^n e(j\Delta t), \quad (4a)$$

$$\frac{de(t)}{dt} \approx \frac{e(t) - e(t - \Delta t)}{\Delta t} = \frac{\Delta e(t)}{\Delta t}. \quad (4b)$$

These approximations are illustrated in Figure 2. The integral is approximated using the forward rectangular rule, and the derivative is approximated using the forward difference approximation. Note that also other approximations such as backward difference or tangent line approximation can be considered without affecting the results presented in this paper.




Figure 2.: Approximations of the integral and the derivative.

2.2. Positional form

Introducing the approximations from Equation (4) in Equation (1), the discrete version of the PID controller results in the positional form

$$u(t) = k_p e(t) + k_i \Delta t \sum_{j=0}^n e(j\Delta t) + k_d \frac{e(t) - e(t - \Delta t)}{\Delta t}, \quad (5)$$

where proportional gain $k_p = K$, integral gain $k_i = K/T_i$, and derivative gain $k_d = KT_d$ have been introduced. Splitting Equation (5) into corresponding terms yields

$$u(t) = u_p(t) + u_i(t) + u_d(t), \quad (6)$$

where $u_p(t)$, $u_i(t)$, and $u_d(t)$ represent the proportional, integral, and derivative terms, respectively. A more efficient way to implement the PID controller is to update the integral $u_i(t)$ recursively, and rewrite the individual terms as follows

$$u_p(t) = k_p e(t), \quad (7a)$$

$$u_i(t) = u_i(t - \Delta t) + k_i \Delta t e(t), \quad (7b)$$

$$u_d(t) = k_d \frac{e(t) - e(t - \Delta t)}{\Delta t} = k_d \frac{\Delta e(t)}{\Delta t}. \quad (7c)$$

PID controllers are often implemented based on the positional discrete representation of Equation (6) and Equation (7). This representation is also sometimes called the direct form. In each iteration, the control signal is computed from the two most recent setpoint and process variable values, with the previous value of the error $e(t - \Delta t)$ and the previous value of the error integral $u_i(t - \Delta t)$ as controller states.

2.3. Incremental form

The incremental form, also called the velocity form (Clarke, 1984), computes the control signal increment since the last iteration, and uses the control signal itself to store the controller state. The control law of the incremental form is

$$u(t) = u(t - \Delta t) + \Delta u(t), \quad (8)$$

where $\Delta u(t)$ represents the increment of the signal u between time $t - \Delta t$ and time t :

$$\Delta u(t) = u(t) - u(t - \Delta t). \quad (9)$$

(This use of Δ applies also to other signals than u .) The control signal increment $\Delta u(t)$ can be split into the individual terms

$$\Delta u(t) = \Delta u_p(t) + \Delta u_i(t) + \Delta u_d(t), \quad (10)$$

which can be obtained from Equation (7) as

$$\Delta u_p(t) = k_p \Delta e(t), \quad (11a)$$

$$\Delta u_i(t) = k_i \Delta t e(t), \quad (11b)$$

$$\Delta u_d(t) = k_d \left(\frac{\Delta e(t)}{\Delta t} - \frac{\Delta e(t - \Delta t)}{\Delta t} \right), \quad (11c)$$

where $\Delta e(t - \Delta t) = e(t - \Delta t) - e(t - 2\Delta t)$. In each iteration, the term increments are calculated according to Equation (11) from the three most recent error values, with the actual control signal itself, $u(t)$, saved as a controller state between iterations.

2.4. Incremental versus positional form

The incremental form has two important practical advantages over the positional form: bumpless transfer and clamping anti-windup are easily implemented.

Bumpless transfer—a requirement in most applications—means that switching between manual and automatic or between tracking and automatic mode, or changing controller parameters, does not cause large jumps in the controller output u . Positional form implementations achieve this by detecting such changes and adjusting the integrator state accordingly. In the incremental form, the output only changes by the increment Δu , so no such checks and adjustments are needed.

The integral part can give rise to integral windup: when the actuator and process cannot achieve the desired setpoint, the persistent control error may keep increasing the control signal beyond the physical limitations of the actuators. Three common remedies are control signal clamping, back calculation and integrator clamping. All can work in practice, but back calculation requires tuning an additional parameter, a tracking-time constant. In the incremental form, control signal clamping is obtained simply by keeping the control signal inside the control limits.

On the other hand, the incremental form requires the integral term and therefore cannot be used for P- or PD-controllers (Section 3.1); in these cases the positional form is required, handled in the implementation of Section 4. This is a minor drawback, since neither anti-windup nor bumpless transfer are relevant without integral action.

Pseudo-code for the two forms, based on Equation (7) and Equation (11), is presented in Listing 1 and Listing 2, respectively. We drop the time index, replacing $u(t)$ with u , etc. Previous values, needed by the increments, carry the prefix x ; for example, $e(t - \Delta t)$ is written as xe . The letter D denotes Δ , so $\Delta e(t) = e(t) - e(t - \Delta t)$ becomes $De = e - xe$. The pseudo-code language is detailed in Appendix A. The purpose of Listing 1 and Listing 2 is a simple comparison between the positional and incremental forms; several practical features, including initialisation of state variables, have therefore been omitted. These are covered in-depth for the combined form PID controller of Section 4.

Listing 1: Basic positional form.

```

1 function u=control(r,y)
  % Compute control signal
2   e=r-y
3   up=kp*e
4   ui=xui+ki*Dt*e
5   ud=kd*(e-xe)/Dt
6   u=up+ui+ud

  % Update state
7   xe=e
8   xui=ui
9 end

```

Listing 2: Basic incremental form.

```

1 function u=control(r,y)
  % Increment control signal
2   e=r-y
3   De=e-xe
4   Dup=kp*De
5   Dui=ki*Dt*e
6   ud=kd*(e-xe)/Dt
7   Dud=ud-xud
8   u=xu+Dup+Dui+Dud

  % Update state
9   xe=e
10  xud=ud
11  xu=u
12 end

```

3. Implementation features

Implementing the basic form of the PID controller is easy, but practical problems complicate the basic code. These problems often only become apparent during development and are usually not taught in a course or included in textbooks. Some but not all problems may occur simultaneously. Conversely, not all adjustments may be necessary. However, there is no major downside in any of the discussed modifications. Table 1 gives an overview of the problems that are addressed in the reference implementation presented in Section 4 and are discussed in the remainder of this section.

3.1. Control without integral action

The PID implementation must be able to handle the situation of a PID control law without integral action, such as pure P- or PD-control. As discussed earlier, the incremental form cannot be used for controllers without integral action because the incremental form results in a non-changing control signal when the error is constant. To see this phenomenon, observe Equation (8), Equation (10) and Equation (11) when $k_i = 0$ and the error is constant. Then $\Delta u(t)$ is zero, no matter the value of the control signal $u(t)$ and the constant error $e(t)$. We therefore risk getting “stuck” far from $e(t) = 0$. Also, consider a parameter update of k_p for a P-controller when the error is constant, but non-zero. In incremental form, the control action stays the same after the update, and the stationary error of the P-controller is unaffected, which should not be the case. Therefore, the desired strategy for our PID implementation to handle P- and PD-control is to operate the controller in positional form whenever $k_i = 0$. This is illustrated in the example given in Section 6.2. A bias term u_0 must be added to reduce stationary errors, see e.g. (K. Åström & Hägglund, 2006). The bias term is nothing more than a constant control signal, that is normally adjusted to the control signal corresponding to zero error. This is illustrated in Figure 3 for a P-controller, which shows how the value of u_0 encodes the offset from $u(t) = 0$ when the error is zero. The slope in the figure corresponds to the proportional gain k_p , and the limits $u_{\min}$ and $u_{\max}$ represent the actuator limitations.




Figure 3.: Relationship between control error $e(t)$ and control signal $u(t)$ for a P-controller.

Thus, when there is no integral action in the controller, the integral term $u_i(t)$ is replaced by u_0 , such that the control signal in this situation becomes

$$u(t) = u_p(t) + u_0 + u_d(t). \quad (12)$$

Table 1.: Overview of practical implementation problems and their solutions addressed in Section 3.

Problem	Solution	Section
If the controller has no integral action, $k_i = 0$, then the stationary control error is zero only when control signal $u = 0$.	Add a control signal bias term u_0 when $k_i = 0$.	3.1
Setpoint changes lead to large proportional and derivative control signal terms, potentially damaging the actuator.	Setpoint weights	3.2
Adding a feed-forward signal to the control signal of the PID controller results in integral windup.	Add the feed-forward signal to the control signal before applying the anti-windup strategy.	3.3
Controller signal u may exhibit larger than permissible jumps from one time step to the next.	Rate limitation	3.4
Integral windup : The integral part increases, winding up the control signal beyond the actuator's saturation limit. Once wound up, the control signal will take a long time to wind down again. The control signal will stay saturated during wind-down.	Anti-windup	3.5
A controller in tracking mode , whose control signal is not applied, may experience integral windup.	Controller output tracking	3.6
Control signal may jump when switching mode or when changing parameter values.	Bumpless transfer	3.7
Execution interval Δt may vary and the control signal may be computed incorrectly.	Jitter compensation	3.8
High frequency noise of the process variable and in rarer cases the setpoint can lead to undesired movements of the derivative and proportional parts of the control signal.	Process variable and setpoint filtering	3.9

In other words, when not in manual mode, our proposed controller has a bias term that equals the user-specified parameter u_0 when there is no integral action ($k_i = 0$), and that equals the (implicitly defined) integral control signal term $u_i(t)$ when integral action is enabled ($k_i \neq 0$). The user can change u_0 also when integral action is enabled, but the effect of such a change will only become visible once integral action is disabled.

3.2. Setpoint handling

With the positional form control laws of Equation (1) and Equation (5), a step in the setpoint $r(t)$ results in a corresponding step in the control error $e(t) = r(t) - y(t)$ and therefore a large impulse in the derivative of the control error $de(t)/dt = dr(t)/dt - dy(t)/dt$. This is illustrated in Figure 4, from which it easily can be seen that we directly after the setpoint change get a large spike in the control signal, sometimes referred to as *derivative kick*. This behaviour is normally not desired, so to obtain a smoother control




Figure 4.: Responses resulting from a step change in the setpoint.

signal, it is common to introduce setpoint weights in the proportional and derivative parts, so that the control error in the proportional part is replaced by $b r(t) - y(t)$ and the error in the derivative part by $c dr(t)/dt - dy(t)/dt$. The setpoint weights b and c are restricted to $0 \leq b \leq 1$ and $0 \leq c \leq 1$. With these modifications, Equation (1) becomes

$$u(t) = K \left(b r(t) - y(t) + \frac{1}{T_i} \int_0^t e(\tau) d\tau + T_d \left(c \frac{dr(t)}{dt} - \frac{dy(t)}{dt} \right) \right), \quad (13)$$

and implementations Equation (7) and Equation (11) are

$$u_p(t) = k_p (br(t) - y(t)), \quad (14a)$$

$$u_i(t) = u_i(t - \Delta t) + k_i \Delta t e(t), \quad (14b)$$

$$u_d(t) = k_d \left(c \frac{\Delta r(t)}{\Delta t} - \frac{\Delta y(t)}{\Delta t} \right), \quad (14c)$$

for the positional form, and

$$\Delta u_p(t) = k_p (b\Delta r(t) - \Delta y(t)), \quad (15a)$$

$$\Delta u_i(t) = k_i \Delta t e(t), \quad (15b)$$

$$\Delta u_d(t) = k_d \left(c \frac{\Delta r(t) - \Delta r(t - \Delta t)}{\Delta t} - \frac{\Delta y(t) - \Delta y(t - \Delta t)}{\Delta t} \right), \quad (15c)$$

for the incremental form. Note also that for controllers without integral action there will normally be stationary control errors. To minimise these, the choice $b = 1$ is recommended to make the P part act on the true control error. When integral action is considered, the parameter b can be used to reduce the impact of sudden changes in the setpoint signal on the control action, mitigating overshoot without compromising the performance of the disturbance rejection. Use of setpoint weighting is illustrated in the example given in Section 6.3. For the PID reference implementation proposed in this article, setpoint weights are included, as well as the choice $b = 1$ when no integral action is present.

3.3. Feed-forward control

Feed-forward control is a complementary control strategy to the feedback strategy used in PID-control to compensate measurable load disturbances, as illustrated in Figure 5. For several reasons including anti-windup treatment, the feed-forward control signal should be added to the feedback control signal inside the PID controller, since it should affect anti-windup in the same way as other control signal terms; see (Guzmán & Hägglund, 2024). The feed-forward input signal can be used for other purposes than load




Figure 5.: PID controller with feed-forward from load disturbance $v(t)$.

disturbance compensation. It can for instance be used for improved setpoint handling and for decoupling in MIMO systems (see Liu et al., 2019). For a PID implementation in incremental form, the desired behaviour is that feed-forward enters with the increments

$$\Delta u_{ff}(t) = u_{ff}(t) - u_{ff}(t - \Delta t), \quad (16)$$

such that the total update equation for the control signal becomes

$$\Delta u(t) = \Delta u_p(t) + \Delta u_i(t) + \Delta u_d(t) + \Delta u_{ff}(t). \quad (17)$$

In scenarios in which the positional form is used, the feed-forward signal then enters as a direct signal, and not through increments. Feed-forward control is illustrated in the example given in Section 6.4.

3.4. Control signal limitation

In almost all PID control systems, the controller output is limited due to actuator constraints. In many cases, the lower bound is $u_{\min} = 0 \%$, which reflects a motor being turned off or a valve being fully closed, and the upper bound is $u_{\max} = 100 \%$, which reflects the motor running at maximum speed or the valve being fully open. The values of these limits depend on the actuator for each particular control problem, and the limitation is expressed as

$$u_{\min} \leq u(t) \leq u_{\max}. \quad (18)$$

In some PID control systems, the rate at which the control signal is allowed to change may also be limited due to safety or due to a slow actuator. This puts the restriction

$$\dot{u}_{\min} \leq \dot{u}(t) \leq \dot{u}_{\max}, \quad (19)$$

which is normally referred to as rate limitation. The amplitude and rate constraints can be combined at every sampling instant to form an admissible interval for the control signal. The rate limitation implies that the control signal at the current instant must satisfy

$$u(t - \Delta t) + \Delta t \dot{u}_{\min} \leq u(t) \leq u(t - \Delta t) + \Delta t \dot{u}_{\max}. \quad (20)$$

This interval is intersected with the actuator saturation limits. The resulting admissible bounds used by the controller are therefore

$$u_{s,\min} = \max(u_{\min}, u(t - \Delta t) + \Delta t \dot{u}_{\min}), \quad (21)$$

$$u_{s,\max} = \min(u_{\max}, u(t - \Delta t) + \Delta t \dot{u}_{\max}). \quad (22)$$

The control signal is finally constrained to satisfy

$$u_s(t) = \begin{cases} u_{s,\max} & \text{if } u(t) > u_{s,\max}, \\ u_{s,\min} & \text{if } u(t) < u_{s,\min}, \\ u(t) & \text{else.} \end{cases} \quad (23)$$

This ensures that the control signal simultaneously respects the actuator saturation limits and the maximum allowed rate of change between consecutive sampling instants. Saturation handling is illustrated in the examples given in Section 6.2 and Section 6.4.

3.5. Integrator anti-windup

Enforcing limits on the control signal means that the saturated control signal $u_s(t)$ may differ from the calculated nominal control signal $u(t)$. If a control error remains, the magnitude of the integral state will keep increasing—winding up—without resulting in any change of the already saturated control signal. This gives rise to the well-understood phenomenon of integrator wind-up, as explained in, e.g., K. Åström and Hägglund (2006). There are several established strategies to address integrator wind-up (Bohn & Atherton, 1995; da Silva, Flesch, & Normey-Rico, 2018; Hoyo, Hägglund, Guzmán, & Moreno, 2023; Skogestad, 2023; Visioli, 2003), but since naming of anti-windup strategies is somewhat ambiguous in the literature, we list here the terminology used throughout this article of the most commonly used anti-windup strategies, and what those strategies mean for an incremental form implementation:

- *Control signal clamping*: If the calculated control signal falls outside the interval $[u_{s,\min}, u_{s,\max}]$, it is adjusted to the saturated value, i.e., $u(t) = u_s(t)$.
- *Back calculation*: The calculated control signal u is dynamically adjusted such that the saturation error $u(t) - u_s(t)$ decays as a first-order process with a time constant T_t , whenever the control signal is saturated. The time constant T_t is typically set to the integral time T_i .
- *Integrator clamping*: In its simplest form, integrator clamping entails not updating the integrating part whenever there is saturation. In incremental form, this corresponds to setting $\Delta u_i(t) = 0$ whenever the control signal saturates.

As stated in Section 2.4, control signal clamping results naturally with the incremental form by saturation of the control signal according to Equation (23). However, we propose back calculation for this article due to its flexibility. In this context, flexibility means that with $T_t = \Delta t$ we have control signal clamping, but by adjusting T_t , different anti-windup behaviours can be obtained. As discussed in Section 3.4, enforcing the saturation limits achieves both absolute control signal limits and rate limitations. A slightly more advanced form of integrator clamping allows for integration exactly up to a saturation limit instead of setting the increment $\Delta u_i = 0$ if it would nominally pass the limit. This is covered by the two first points in the slightly more advanced combined anti-windup strategy specified below:

- *Integrate to saturation*: Consider the situation where the control signal $u(t)$ is saturated, but not the control signal without $\Delta u_i(t)$. Then the desired behaviour is to integrate to the limit. Thus, $\Delta u_i(t)$ is reduced.
- *Allow integration in the right direction*: When $u(t) > u_{\max}$ then integration further into saturation ($\Delta u_i(t) > 0$) is not allowed, but integration away from saturation ($\Delta u_i(t) < 0$) is. In this case, $\Delta u_i(t)$ is not set to zero, but remains as calculated in integration. Of course, the corresponding logic applies to the lower limit $u_{\min}$.
- *Back calculation*: The saturation error is updated according to the description of *Back calculation* above.

The two first points can be used alone to achieve integrator clamping that occurs exactly at the saturation limit. They can also be part of a combined anti-windup strategy, that includes also the third point, constituted by back-calculating anti-windup.

The rationale behind applying integrator clamping prior to back-calculation is to prohibit an initial increment integrating further into a saturated region, before the back-calculation correction is applied. By doing this, the saturation error will decay exactly following the back-calculation time constant T_t , whereas allowing for adding the integral increment prior to applying back-calculation may slightly shift the decay rate. Anti-windup is illustrated in the example given in Section 6.4.

3.6. External signal tracking

There are situations in which the controller output should track an external signal instead of trying to make the process variable follow a setpoint. Examples are override control, where another controller may take over the control, often handled by selectors, and gain scheduling using several controllers. In these instances, the controller, which currently does not affect the process, must enter *tracking mode* (*track*) to avoid windup and bumps; see (Hägglund, 2023). A tracking example including a simulation when tracking is desired is provided in Section 6.7. In tracking mode, the controller assumes that the control signal currently applied to the actuator was produced by itself, even if it was not. In other words, the controller updates its internal states so that its computed output is consistent with the actual actuator signal. This means that the control signal is calculated by

$$u(t) = u_{\text{track}}(t) + \Delta u_p(t) + \Delta u_i(t) + \Delta u_d(t) + \Delta u_{ff}(t), \quad (24)$$

where u_{track} is the control signal from the controller that was actually in control of the actuator in the last iteration. Controller output tracking is illustrated in the examples given in Section 6.6 and Section 6.7. Note that control signal increments are added to the tracking signal u_{track} in Equation (24). If this was not done, the (unsaturated) output of the tracking controller would be identical to the tracking signal u_{track} , preventing controller switching in selector scenarios as the one of Section 6.6.

3.7. Bumpless mode and parameter changes

There are three control modes: automatic or switched on (auto), tracking (track) as described in the previous section, and manual (man), where the control signal is specified by the user. Switching between these modes should not result in a large change or *bump* in the control signal $u(t)$. Bumpless transfer is a fundamental concept in PID control and is therefore a must in a good implementation. Bumpless transfer is desired for all controller mode changes (when possible) and for controller parameter updates. A common implementation of bumpless transfer in positional form is to adjust the integral term to obtain the desired control signal. In incremental form, since the integral state is encoded into the control signal, bumpless transfer results naturally in many situations. However, there are some scenarios where bumpless transfer is not possible without large workarounds. For example, bumpless parameter updates when $k_i = 0$ are usually neither possible nor desired. Table 2 describes when bumpless transfer is desired for the PID reference implementation and when not.

Since bumpless transfer from manual mode to auto is desired, the manual control signal u_{man} must enter the controller for this to be implementable; see Figure 6.

It is often also desired to have bumpless transfer from automatic mode to manual mode. This can be accomplished outside the controller by setting u_{man} to u_{auto} at these mode transitions. Mode switching between man and auto is illustrated in Section 6.1.

Table 2.: Bumpless transfer requirements during mode changes and parameter updates for our PID reference implementation. Tracking mode is described in Section 3.6.

Change	k_i	Bumpless
man $\rightarrow$ auto	$k_i \neq 0$	Yes
auto $\rightarrow$ man	$k_i \neq 0$	No
man $\rightleftharpoons$ auto	$k_i = 0$	No
track $\rightleftharpoons$ man	—	No
track $\rightarrow$ auto	$k_i \neq 0$	Yes
auto $\rightarrow$ track	$k_i = 0$	No
k_p or k_d update	$k_i \neq 0$	Yes
k_p or k_d update	$k_i = 0$	No
k_i update to $k_i \neq 0$	—	Yes
k_i update to $k_i = 0$	—	No




Figure 6.: Switching between manual and automatic mode.

3.8. Sampling rate and jitter

The sampling interval Δt is normally set to a constant value so that periodic sampling is obtained. Jitter refers to the deviation from this nominal execution period, whether it was introduced intentionally or accidentally. In most applications, the jitter is small and neglected in the calculations.

However, in cases where the jitter is a problem or intentional, for instance in event-based control (Årzén, 1999), Δt can be calculated as the actual time elapsed since the last execution. Therefore, the possibility to run a PID controller with jitter compensation is a required functionality for the reference PID controller implementation. Note that a large change in the sampling interval Δt also affects the saturation limits $u_{\min}$ and $u_{\max}$ through the rate limits defined by the parameters $\dot{u}_{\min}$ and $\dot{u}_{\max}$.

3.9. Process variable and setpoint filtering

The process variable $y(t)$ is often corrupted by high-frequency noise or quantisation effects, which may cause undesired variations in the control signal $u(t)$. These variations are to some extent caused by the proportional term, $u_p(t)$, but more importantly by the derivative term, $u_d(t)$. The integral term $u_i(t)$ averages the noise across samples, and thus serves as a low-pass filter, blocking noise propagation. Our reference implementation uses the backward difference approximation

$$\frac{dy}{dt} \approx \frac{y(t) - y(t - \Delta t)}{\Delta t}, \quad (25)$$

with dr/dt approximated in the same way, and used in controllers with setpoint derivative weight $c \neq 0$, as described in Section 3.2. To illustrate noise amplification of the derivative, let $n(t)$ be the additive noise term in $y(t)$. This means that while we consider $y(t)$ to be the process variable, the true, and unknown, process variable is in fact $y(t) - n(t)$. Letting $\Delta n = n(t) - n(t - \Delta t)$ we thus have a noise amplification of $\Delta n / \Delta t$, which is inversely proportional to the sampling interval. For this reason, it is often necessary in practice to low-pass filter the noisy process variable $y(t)$ before computing the finite derivative approximation. This becomes increasingly important as the sampling interval Δt decreases. Similarly, while a true step function exhibits an unbounded derivative, a setpoint step of size Δr results in a spike $\Delta r / \Delta t$ in our derivative approximation Equation (25). To limit noise amplification, which may otherwise result in unnecessary actuator wear or excite resonant modes in the dynamics of the process, it is customary to low-pass filter the process variable. In situations where a setpoint weight $c > 0$ is used in combination with possible abrupt (step) changes in the setpoint $r(t)$, it might also be adequate to low-pass filter $r(t)$. While a first-order low-pass filter is most often used in existing implementations and textbooks alike, a second-order filter is often advisable. This is because the latter results in derivative term noise amplification gain asymptotically approaching zero for high frequencies. This desired property is commonly referred to as high-frequency roll-off (K. Åström & Hägglund, 2006). This strongly motivates the use of second-order filters, although they introduce slightly more phase-loss than their first-order counterparts. For our reference implementation, we introduce a simple filter that approximates either the first-order filter with transfer function

$$F_1(s) = \frac{1}{sT_f + 1}, \quad (26)$$

or the second-order filter with transfer function

$$F_2(s) = \frac{1}{(sT_f + 1)^2}, \quad (27)$$

in discrete time, with a sampling interval Δt . The parameter T_f is the filter time constant, meaning that the magnitude of the filter transfer function has its cutoff at $\omega_c = T_f^{-1}$, and thus effectively attenuates (noise) components in $y(t)$ with angular frequencies larger than ω_c . We use a discretisation that approximates zero-order-hold (ZOH) well, while being computationally beneficial, and allowing for a filter with bumpless output at changes in either of the filter time constant T_f or the sampling interval Δt . It is based on a cascaded Tustin approximation, further explained in Appendix B. The first-order filter Equation (26) then results in an update equation

$$y_f(t) = y_f(t - \Delta t) + a(y(t) - y_f(t - \Delta t)), \quad (28a)$$


with

$$a = \frac{\Delta t}{T_f + \Delta t/2}, \quad (28b)$$


where y is the filter input and y_f is the filter output. The second-order filter Equation (27) is approximated by cascading two first-order filters, as shown in Listing 4. The incremental form of Equation (28a) allows for bumpless changes in T_f and Δt during operation. Note that there is no formal requirement that the filter run with the same sampling interval as the controller. Noise filtering is illustrated in the example given in Section 6.5.

4. The combined form PID controller

In this section, we propose our reference implementation that implements all functionalities and features described in Section 3, comprising the combined form controller of Listing 3, and associated signal filter of Listing 4. The code is both generic and general. For example, if a PID controller will never be used in tracking mode, all tracking-specific code can be removed. Similarly, the code can be simplified if, for example, no derivative action or in other words a PI controller, is required. The term ‘combined’ comes from the controller making use of both the positional and incremental forms, described previously in Sections 2.2 to 2.4. The essential idea is to use the incremental form whenever the controller has integral action. This ensures bumpless behaviour at parameter and mode changes, and facilitates tracking of an exogenous control signal. However, for controllers without integral action, the positional form is needed to ensure the desired response in the control signal. In the following sub-sections, we cover features of the proposed implementation to deal with the practical problems described in Section 3. The pseudo-code language we use is meant to be generic and focus on algorithmic aspects. For an in-depth explanation of it we refer to Appendix A. In the following, we go through our proposed implementation of the combined form PID controller (Listing 3), and the signal filter (Listing 4) that was proposed in Section 3.9. We also propose an implementation of the alternative anti-windup strategy that was introduced in Section 3.5 (Listing 5). With the exception of the filter and alternative anti-windup strategy, we will use the schematic drawing of Figure 7 as a basis for introducing the code. The numbered blocks indicate the order in which the corresponding functionality appears in the implementation. Each block is described in a subsection (4.x) correspondingly numbered with x referring to the block number in Figure 7a. The core part of the implementation, including PID with the integral term and anti-windup techniques, is summarised in Section 4.7 and depicted in Figure 7b.



(a) Schematic representation of Listing 3.



(b) Schematic representation of the "PID with I" block in Figure 7a.

Figure 7.: Schematic overview of the combined-form PID controller. Block numbers reflect the order of appearance in our implementation and correspond to the numbered subsections referencing the relevant code lines in Listing 3.

Listing 3: Combined form PID controller.

```

1 function u=control(r,y,uff,uman,utrack,mode)
    % 1 Rate limits
2   usmin=max(umin,xus+Dt*dumin)
3   usmax=min(umax,xus+Dt*dumax)

    % 2 Derivative approximation
4   dr=(r-xr)/Dt
5   dy=(y-xy)/Dt

    % 3 Tracking mode
6   if mode==TRACK
7       xu=utrack
8   end

    % 4 Manual mode
9   if mode==MAN
10      u=uman
    % 5 Integral switch
11  elseif ki==0
    % 6 PID without I
12      u=u0+kp*(r-y)+kd*(c*dr-dy)+uff
13  else
    % 7 PID with I
    % Signal increment
14      Dr=r-xr
15      Dy=y-xy
16      Ddr=dr-xdr
17      Ddy=dy-xdy

    % Control increment
18      Dup=kp*(b*Dr-Dy)
19      Dui=ki*Dt*(r-y)
20      Dud=kd*(c*Ddr-Ddy)
21      Duff=uff-xuff

    % Nominal control signal
22      u=xu+Dup+Dui+Dud+Duff

    % Integrator anti-windup
23      us=max(min(u,usmax),usmin)
24      u=u-Dt/Tt*(u-us)
25  end

    % 8 Saturate
26  xu=u
27  u=max(min(u,usmax),usmin)

    % 9 State update
28  xus=u
29  xr=r
30  xdr=dr
31  xy=y
32  xdy=dy
33  xuff=uff
34 end

```

Listing 4: Second-order signal filter.

```

1 function xf2=filter(y)
    % Filter parameter
2   a=Dt/(Tf+0.5*Dt)

    % State update
3   xf1=xf1+a*(y-xf1)
4   xf2=xf2+a*(xf1-xf2)
5 end

```

Listing 5: Combined integration anti-windup.
Can replace Line 24 in Listing 3.

```

1 if Dui*(u-us)>0
2     u=u-sign(u-us)*min(abs(Dui),abs(u-us))
3 end
4 u=u-min(Dt/Tt,1)*(u-us)

```

In Listing 3, the `control` function itself is invoked on Line 1, and takes as its arguments:

- setpoint `r`,
- process variable `y`,
- feed-forward control signal `uff`,
- manual control signal `uman`,
- tracking control signal `utrack`,
- operating mode `mode`.

The process variable `y` and setpoint `r` are assumed to be adequately filtered, as discussed in Section 3.9, using for example our proposed filter implementation of Section 4.10. The mode variable `mode` has three defined values: automatic mode `AUTO`, manual mode `MAN`, and external signal tracking mode `TRACK`. The controller is in automatic mode whenever `mode` is neither `MAN` nor `TRACK`. Encoding of the mode is further discussed in Section 5.1.3. If no feed-forward action is desired, `uff` is simply set to zero. On a similar note, the values of `uman` and `utrack` are relevant when the controller is in manual and tracking mode, respectively. The implementation also consists of both states and parameters. The parameters, e.g. controller gains and limits, are set externally by some language-dependent unspecified mechanism, further discussed in Section 5. States are used to store controller values between executions. Complete lists of all arguments, parameters, and states are found in Appendix C.

4.1. Rate limits

The parameters `umin` and `umax` define the interval into which the control signal is saturated, cf. Section 4.8. The parameters `dumin` and `dumax` define a lower and upper rate limit for the change of the control signal. They are used together with the previous value of the saturated control signal, `xus` and the past sampling interval `Dt` to compute corresponding saturation bounds `xus+Dt*dumin` and `xus+Dt*dumax` (see Section 3.4). Note that it is important that the previous saturated control signal `xus`, and not the previous nominal control signal `xu` is used to produce a correct rate limitation. The actual saturation limits `usmin` and `usmax` are computed as the more conservative of the absolute limits defined by the parameters `umin` and `umax`, and those imposed by the rate limit parameters `dumin` and `dumax`, as coded on Lines 2–3.

4.2. Derivative approximation

The derivatives used by controllers with derivative action are approximated using finite differences, as explained in Section 2.1. The difference between the current and previous process variable value `y-xy` is divided by the execution interval duration `Dt`, to produce the finite difference approximation `dy` on Line 5. The setpoint derivative approximation `dr` is computed in the same way on Line 4. To avoid undesired jumps and noise amplification in the control signal, it is important that the process variable is adequately filtered, as described in Section 3.9. In the case of nonzero setpoint weight `c`, the setpoint must also be adequately filtered.

4.3. Tracking mode

Tracking of external signals, as explained in Section 3.6 is implemented to overwrite the previous output `xu` with the signal to be tracked, `utrack` on Line 7. Then the control signal increment is computed as normal, cf. Section 4.7. Tracking mode is active when `mode==TRACK`. Note that tracking is only meaningful when there is integral action. In the case of no integral action, tracking mode and manual mode behave identically. In our implementation, feed-forward and integrator anti-windup behave the same way in tracking mode as in automatic mode.

4.4. Manual mode

The controller is in manual mode when `mode==MAN`. In manual mode, the nominal control signal `u` takes on the externally provided value `uman`, as expressed on Line 10. The manual mode control signal is limited as explained in Section 4.1. Feed-forward action is disabled in manual mode. Due to the incremental form when we have integral action, the bumpless mode switches required in Section 3.7 are automatically obtained.

4.5. Integral switch

The combined form switches between positional and incremental form based on whether the controller utilises integral action, as specified in Section 3.1. This is achieved by checking for `ki==0` on Line 11. If the condition is true, positional form is employed, as explained in Section 4.6. Otherwise, the `else` block starting on Line 13 is executed, implementing incremental form as explained in Section 4.7.

4.6. PID without integral action

In the absence of integral action (`ki==0`), the positional form control signal is computed according to Line 12. The constant bias term `u0` can be used to zero-offset the control signal to match the stationary operating point of consideration. Setpoint weighting with weight parameter `c` is used for the derivative term, but for the proportional term the weight is fixed to `b=1`, as explained in Section 3.2. Feed-forward is enabled, and the control signal is saturated according to Section 4.1, but since there is no integral action, no integrator anti-windup is applied. Parameter and mode changes are not bumpless, as there is no integrator state that can be shifted to compensate for bumps.

4.7. PID with integral action

With integral action (`ki` not zero), an incremental form PID control law is applied, which is summarised in Figure 7b according to the code Lines 14–24. The setpoint and process variable increments `Dr` and `Dy` are computed on Lines 14–15, followed by increments of their derivative approximations (see Section 4.2) on Lines 16–17. Taking the process variable as example, the increment is simply computed as the current value `y` minus the previous value `xy` stored in the previous execution of the `control` function. The computed signal increments are then used to compute increments of the individual control signal terms on Lines 18–21. Since the derivatives are calculated from the actual time between executions, `Dt`, and the integral increment also uses `Dt`, we compensate for possible jitter according to Section 3.8.

Weight parameters `b` and `c` are used for the proportional and derivative terms, respectively, to account for setpoint weighting as described in Section 3.2. The rectangular integral approximation and finite difference derivative approximation utilised in computing `Dui` and `Dud` are described in Section 2.1. The nominal (i.e., not yet saturated) control signal `u` is computed on Line 22 by adding the increments from Lines 18–21 to the previous nominal control signal `xu`. Note that the feed-forward signal must enter the signal increment here, as motivated in Section 3.3. Note that it is important that it is the nominal previous control signal `xu` that is used, and not its saturated counterpart. Otherwise, it is not possible to implement integrator anti-windup strategies other than control signal clamping, as both conditional integration and back calculation rely on the possibility for the control signal to go outside of its unsaturated range. As explained in Section 3.5, integrator anti-windup relies on knowing the difference between the nominal control signal `u`, and its saturated counterpart that we call `us`. To this end, the saturation limits computed on Lines 2–3, as explained in Section 4.1, are applied on Line 23 to obtain `us`. The reference implementation of Listing 3 implements the back calculation anti-windup strategy on Line 24 (with control signal clamping as a special case when `Tt` equals `Dt`). A more general anti-windup implementation, that can serve as an in-place replacement for Line 24 of Listing 3 is provided in Listing 5. It fulfils the advanced anti-windup behaviours specified at the end of Section 3.5. For integrator clamping, it can be used as provided with `Tt` set to `Inf`, effectively translating Line 4 into `u=u`. Alternatively Line 4 can be removed entirely when integrator clamping is desired. Lines 1–3 in Listing 5 have no purpose in control signal clamping, and if no other anti-windup mode needs to be supported, Listing 5 can be replaced in its entirety by `u=us`. An even simpler alternative for this case is to move saving of the control signal state (`xu=u`) on Line 26 of Listing 3 to after saturating the control signal on Line 23, and removing the entire integrator anti-windup block on Lines 23–24; replacing `xus` by `xu` on Lines 2–3, and removing Line 28.

4.8. Saturate

Upon saving the nominal, unsaturated, control signal as `xu` on Line 26, the control signal is saturated on Line 27, as described in Section 4.1. It is this saturated control signal that is eventually returned by the `control` function.

4.9. State update

As a final stage of the `control` function, the state variables are updated on Lines 28–33 (in addition to `xu` that was updated on Line 26 before saturating the control signal).

4.10. Process variable and setpoint filtering

As discussed in Section 3.9, the process variable, and sometimes the setpoint, need to be adequately low-pass filtered before being passed to the PID control algorithm of Listing 3. To this end, Listing 4 implements the low-pass filter introduced in Section 3.9 and further explained in Appendix B. Line 2 computes the filter parameter `a` from the filter time constant parameter `Tf` and execution interval `Dt`, as explained in Appendix B. If the filter runs at a constant sampling interval, there is no need to re-compute `a` at each invocation of the `filter` function. Instead, `a` can then be considered a parameter, computed outside the `filter` function and provided in the same fashion as other parameters.

The cascaded filter is implemented on Lines 3–4. If a first-order filter is desired, Line 4 can be removed, and `xf2` changed to `xf1` on Line 1. Note that the filter function must be instantiated separately for every signal being filtered, since each instance requires its own state variables `xf1` and `xf2`.

5. Closed loop implementation

The PID control algorithm may be executed in a real-world, run-time environment or in a simulated environment, but always in closed loop. Both scenarios require more code than the function presented in Section 4. This section first provides PID function execution details including initialisation, default parameters and discussions relating to the control mode. It then provides the code necessary for closed loop implementation as will be used in Section 6.

5.1. PID function execution

The PID function call cannot be stand-alone but is embedded in an environment. The environment can take various forms and is usually application specific, but a number of common questions arise relating to the execution: how are states and signals initialised? How are parameters passed to the function? Answers to these questions are given in the following sections.

5.1.1. Initialisation

The state variables of the controller, as well as any signal filters, must be initialised to appropriate values before the `control` or `filter` functions of Listing 3 and Listing 4 are invoked for the first time. The controller states are `xr`, `xu`, `xus`, `xy`, `xdy`, `xdr` and `xuff`. The latest available signal values can be used to initialise the controller states, so that, for example, `xr=r`, etc. If the controller is started in manual mode, the initialisation of `xu` has no practical effect. If, instead, the controller is started in automatic or tracking mode, it is reasonable to initialise `xu` to the value that the actuator can be expected to have at startup. For this reason, setting `xu=0` and `xus=0` is appropriate in many situations. However, other initialisations may be preferable depending on the application context. As with controller parameter tuning, such choices depend more on the specific use case than on the controller implementation itself, and are therefore not discussed further here. A reasonable initialisation of the low-pass filter is to assign `xf1` and `xf2` the latest available value of the corresponding unfiltered measurement. To avoid transient effects, it is advisable to let filters converge to their steady-state before placing the corresponding controller in automatic or tracking mode. In the closed-loop simulation examples of this paper, the initial variables and subsequently the controller states are all set to zero by default.

5.1.2. PID parameters

The parameters, which are `orange` in colour in the code, are listed in Appendix C. If parameters are passed by reference, it is important to prevent them from changing while critical parts of the functions are executing. For example, switching to a new pair `kp`, `ki` should be done in a way that avoids a single invocation of Listing 3 using the old `kp` value together with the new `ki` value. This can be avoided in several ways.

One approach is to pass all parameters by value as arguments to the function that uses them. Another approach is to call an update function that reads parameters from memory at the beginning of the function. A third option is to use thread-safety mechanisms such as locks or mutexes to prevent parameter updates while critical sections of the function are executing. The reference implementation assumes a runtime capable of performing floating-point arithmetic, either natively through a floating-point unit or via emulation. Established methods exist for porting floating-point implementations to fixed-point architectures. We therefore do not delve further into this topic, but refer the interested reader to the introduction provided in Crisp (1991). One important aspect concerns how signals and parameters are passed to the controller, and how the computed control signal is actuated. In our reference implementation, signals are passed as arguments to the `control` and `filter` functions (see Listing 3 and Listing 4), while parameters (orange) and state variables (pink) are assumed to be provided through some unspecified mechanism. It should be noted that the distinction between signals and parameters is not always clear-cut. For example, `mode` could just as well be regarded as a parameter. We deliberately refrain from prescribing how values are passed, since this depends on programming language and runtime environment. In an object-oriented setting, it is natural to define a `PID` object with `control` and `filter` methods, together with appropriate constructors for initialisation, as the MATLAB and Python implementations in our GitHub repository (Sundström et al., 2026a). However, the proposed algorithms can equally well be implemented without object-oriented constructs. For this reason, we focus on the algorithmic structure rather than on software architectural details.

5.1.3. Control mode

The key function of the control mode is to allow the PID user or operator to switch the controller on or off, which is labelled automatic (`AUTO`) and manual (`MAN`), respectively. As described in Section 3.7, the controller can be in tracking mode when the calculated control signal is not applied but the controller tracks another signal. There is a risk that the controller defaults to automatic mode on Line 13 if `mode` holds a value other than `MAN` or `TRACK`. This risk can be eliminated by introducing the fourth mode ‘Disabled’. The ‘Disabled’ mode means that the control signal u is neither calculated nor applied.

5.2. Basic closed-loop simulation

To illustrate simulation use, Listing 6 shows the code for a basic control loop. The simulation codes for this and for the examples in Section 6 are available on GitHub.

The PID controller parameters, saturation limits, rate limits, sampling interval, and filter time constant are set on Lines 1–2, and the PID controller is created from those parameters on Line 3. The initial values for the different signals are given on Line 4, also setting the controller mode to `AUTO`. These values are used on Line 5 to initialise the PID controller states by calling the `initialisation` function according to the ideas described in Section 5.1.1. The control loop for a given simulation time T_{sim} is defined on Lines 6–15. Inside the loop, Line 10 simulates the process dynamics, a setpoint change is configured on Lines 7–9, the process output is filtered on Line 11 using the `filter` function, and the controller action is calculated on Line 12 by calling the `control` function. The calculated control signal is sent to the plant on Line 13, and the loop waits for the next sampling instant on Line 14. Unless otherwise stated, the initial conditions defined on Lines 1–5 of Listing 6 are used as default values for all examples presented in Section 6, and only the specific modifications are described in the corresponding sections.

Listing 6: Basic control loop.

```

1 kp=1.0 ki=0.5 kd=0.1 b=1 c=0 umin=0 umax=100
2 dumin=-10 dumax=10 u0=0 Tt=0.1 Tf=0.1 Dt=0.1
3 % Creating PID instance with parameter values
4 pid1=PID(kp,ki,kd,b,c,u0,umin,umax,dumin,dumax,Tt,Tf,Dt)

5 % PID states and signals
6 r=0 y=0 dr=0 dy=0 u=0 uff=0 us=0 uman=0 utrack=0 mode="AUTO"
7 % Initialise PID states
8 pid1.initialisation(r,y,dr,dy,uff,u,us)

9 % Execution Loop
10 while t < Tsim

11     % Set-point change
12     if t>=2
13         r=1
14     end

15     % Read from sensor
16     y=read_process_variable()

17     % Filtering process output
18     yf=pid1.filter(y,Tf,Dt)

19     % PID control signal calculation
20     u=pid1.control(r,yf,uff,uman,utrack,mode)

21     % Set control signal
22     set_control(u)

23     % Wait for next iteration
24     wait(Dt)
25 end

```

6. Simulation examples

This section presents seven simulation examples to show how the proposed code handles the different implementation features discussed in Section 3. The first three examples deal with switching between automatic and manual mode, switching the integral action on and off, and showing the effect of rate limitation, respectively. The fourth example is focused on analysing the combination of the PID controller with a feed-forward compensator and different solutions to treat the control signal saturation problem. The measurement noise filtering capabilities are presented in the fifth example. Finally, the last two examples demonstrate the use of the basic PID functionality for the cases that involve tracking a control signal, in particular gain scheduling and selector control. For most of the examples, a first-order plus dead-time process model with transfer function

$$P(s) = \frac{k}{1 + sT} e^{-sL} = \frac{1}{1 + s} e^{-0.5s} \quad (29)$$

is used and controlled by a PI-controller tuned using the Lambda method (Dahlin, 1968), with $\lambda = T = 1$ and sampling interval 0.1.

This results in the controller

$$C(s) = K \left(1 + \frac{1}{sT_i} \right) = 0.667 \left(1 + \frac{1}{s} \right) = k_p + k_i \frac{1}{s} = 0.667 + 0.667 \frac{1}{s}. \quad (30)$$

When different process dynamics or controller parameters are required for some particular example, new information will be provided.

6.1. Example 1: Man/Auto switching

In the proposed simulation, the system starts in manual mode at the beginning of the simulation and then switches to automatic mode in the middle of the simulation time. Figure 8 shows a simulation for this example in which the system starts in manual mode from $t = 0$ to $t = 10$, with a step control signal change from $u_{\text{man}} = 0$ to $u_{\text{man}} = 1$ at $t = 1$, and switches from manual to automatic mode at time $t = 10$. The setpoint is fixed to $r = 3$ from the beginning of the simulation. As observed, the switch between manual and automatic mode is performed properly, and there is no bump at the switching time. Listing 7 shows the control loop code for this example. Before simulating the control




Figure 8.: Switching from manual to automatic mode in Example 1.

loop, the following changes must be made in the initialisation section for parameters and states: $k_p=0.667$, $k_i=0.667$, and $k_d=0.0$, $\text{mode}=\text{MAN}$, and $r=3$. As observed, the step control signal change for manual mode is done on Lines 3–5 setting $u_{\text{man}}=1$, while the switch between manual and automatic modes is performed on Lines 6–8 setting $\text{mode}=\text{AUTO}$.

Listing 7: Example 1 code.

```

1 while t < Tsim
    % Simulation of process dynamics
2     y=process_simulation(u,Dt)

    % Step change on manual input signal
3     if t>=1
4         uman=1;
5     end

    % Switch to AUTO
6     if t>=10
7         mode="AUTO"
8     end

    % PID control signal calculation
9     u=pid1.control(r,y,uff,uman,utrack,mode);

    % Simulation time update
10    t=t+Dt
11 end

```

6.2. Example 2: P- or PD-control and rate limitation

This example shows how the proposed code can be used to implement P- or PD-controllers, that is, controllers without integral action. The rate limitation function is also illustrated. We use the process and controller given by Equation (29) and Equation (30). For the P-controller the integral gain is set to $k_i = 0$ and the bias term u_0 is introduced with $u_0 = 2$. Figure 9 shows the simulation results for this example.




Figure 9.: Switching between PI- and P-control in Example 2.

The simulation starts using the PI-controller for an operating point given by $y = 1$ and $u = 1$, and with a setpoint change from $r = 1$ to $r = 3$ at $t = 1$. At time $t = 10$, the controller is switched to a P-controller by setting $k_i = 0$. Since the control error is zero at time $t = 10$, the control signal will jump down to the bias term $u_0 = 2$, and stay there until the process delay has elapsed. The P-controller will then control the process variable to a new position, but since there is no integral action there will be a control error. At time $t = 15$, another setpoint change is made, from $r = 3$ to $r = 1$. The P-controller behaves as expected, and the two setpoint responses are similar, showing that the transitions from PI to P work well. Finally, at time $t = 21$ we switch back to a PI-controller and the control error is eliminated, obtaining a bumpless transfer response. For this example, the simulation code is not shown since it is similar to the code used in Listing 7. The manual part in Listing 7 is removed, the variable `ki` is set to `ki=0` at time $t = 10$ to introduce the P-controller, and it is set to `ki=0.667` at $t = 21$ to go back to the PI-controller again. At $t = 1$ and $t = 15$ setpoint changes are introduced by setting `r=3` and `r=1`, respectively, in a similar way as shown in Listing 6 on Lines 7–9. Moreover, the following changes are required at the initialisation section: `y=1.0`, `u=1.0`, `u0=2.0`, `r=1.0`, and `mode=AUTO`. This example is also used to evaluate the rate limitation option. Figure 10 shows the same simulation results presented in Figure 9 but for a rate limit of 1. As observed at time instants $t = 1$, $t = 10$, and $t = 15$, the changes of the control signal are now limited and thus a slower response is obtained according to the rate limitation imposed. For this simulation, rate limitation variables must be modified in the code by setting `dumin=-1` and `dumax=1` in the initialisation section.




Figure 10.: Rate limitation effect in Example 2 for `dumin=-1` and `dumax=1`.

6.3. Example 3: Setpoint handling

As discussed in Section 3.2, the code proposed in this paper includes the capabilities for setpoint handling to obtain a smoother control signal when strong changes of setpoints (like step changes) are required. To show this idea, in this example we use the process and controller given by Equation (29) and Equation (30).

As the aim of this example is to show the effect of the b parameter on the responses of the control system, Figure 11 shows the simulation results for the cases with $b = 0$, $b = 0.5$, and $b = 1$. As observed, the smoother control signal is obtained for the case with $b = 0$ and the more aggressive control signal for $b = 1$. The b factor will of course also influence the process variable response, where the slower and faster responses are obtained for $b = 0$ and $b = 1$, respectively. Figure 11 also shows the effect on the disturbance rejection for a step-like disturbance signal of amplitude 1 added to the control signal at time $t = 10$. As expected, identical responses are obtained for any value of b , as the setpoint handling is decoupled from the disturbance rejection problem. For this example, the simulation code is not shown as it is the same code as used in Listing 7 skipping the manual part at the beginning of the simulation and using different values for the b parameter.




Figure 11.: Setpoint handling in Example 3.

6.4. Example 4: Feed-forward and saturation

This example focuses on analysing the anti-windup solutions discussed in Section 3.5 for the control signal saturation problem, and also the combination of the PID controller with a feed-forward compensator to deal with measurable disturbances. The effect of the anti-windup clamping and back calculation solutions implemented in Listing 3 will be analysed and compared. Moreover, the capability of the new PID code that includes the feed-forward control signal inside the PID controller will also be explored. Notice that in classical implementations, the contribution of the feed-forward compensator is added to the output of the PID controller outside the controller code (Guzmán & Hägglund, 2024), and in the proposed code the feed-forward control signal is considered part of the PID controller (see Section 3.3).

For this example, the following process models are used:

$$P_u(s) = \frac{k_u}{1 + sT_u} e^{-sL_u} = \frac{1}{1 + 3s} e^{-0.5s}, \quad (31)$$

$$P_v(s) = \frac{k_v}{1 + sT_v} e^{-sL_v} = \frac{2}{1 + s} e^{-0.5s}, \quad (32)$$

where P_u represents the dynamics that relates the control signal with the process output and P_v represents the dynamics that relates the load disturbance with the process output. In this case, a PID controller was designed using the Padé approximation for the time delay and using the Lambda method with $\lambda = 0.3$, resulting in $K = 2.83$, $T_i = 3.25$ and $T_d = 0.23$. The control signal limits are $u_{\min} = -3.5$ and $u_{\max} = 3.5$, and the feed-forward compensator is designed in the classical manner by dividing P_v over P_u with reversed sign:

$$FF(s) = -\frac{P_v}{P_u} = -\frac{2(1 + 3s)}{1 + s}. \quad (33)$$

Figure 12 shows the simulation results for this example, where the cases without saturation limits (No saturation) and with saturation limits but without using anti-windup techniques (No anti-windup) are included for a better comparison.




Figure 12.: Feed-forward plus saturation in Example 4.

The simulation starts with a large setpoint change from $r = 0$ to $r = 4.5$ at $t = 1$. Then, the setpoint value is set again to $r = 0$ at $t = 20$ and a step-like change is also included for the load disturbance from $v = 0$ to $v = 1.5$ at $t = 45$. Notice that for the no-saturation case, the desired closed-loop response is perfectly achieved for the setpoint changes, and the effect of the disturbance is completely removed.

When actuator saturation is taken into account, differences among the evaluated strategies become evident. Between $t = 1$ and $t = 20$, following the first setpoint change, no noticeable differences are observed, since the control signal in all simulated cases exceeds the saturation limit. However, after the second step change at $t = 20$, the impact for the no anti-windup case becomes apparent. In this case, the system performance is significantly degraded, as the controller requires a longer time to exit the saturation region and reach the new setpoint due to the windup effect. The performance is also significantly degraded in the disturbance rejection scenario, where an undesirable overshoot is observed in the response due to the same underlying cause.

When clamping and back calculation anti-windup techniques are used, the responses are improved in different manners, as observed in the dashed-line and black solid-line curves for clamping and back calculation cases, respectively. As commented in Section 3.5, the PID code is the same for both techniques only changing the value of T_t , with $T_t = \Delta t = 0.01$ for clamping and $T_t = T_i = 3.25$ for back calculation in this example. When clamping is used, the saturation time is dramatically reduced as observed in the control signal plot. It can be seen that the control signal leaves saturation quickly for both the setpoint tracking and disturbance rejection responses. However, this advantage from a saturation point of view results in a slow response at the process output for setpoint tracking and disturbance rejection, as observed in the process output plot. On the other hand, when anti-windup back calculation is used, the saturation time is only slightly reduced, but the performance on the process output for setpoint tracking and disturbance rejection cases is much closer to the ideal response without saturation limits. As discussed in Section 3.5, in this paper a third anti-windup approach is proposed as a combination of clamping and back calculation. Figure 13 shows the same previous example, where clamping, back calculation, and the combined solution (named `Combined solution` in the figure) from Listing 5 are compared. As observed, an intermediate behaviour is obtained in this case, where for setpoint tracking the result is similar to the clamping solution, and for the disturbance rejection problem it is closest to the back calculation solution. The main advantage of this third solution is that the calculated control signal is dynamically adjusted so that the saturation error decays exactly as a first-order process with a time constant T_t , whenever the control signal is saturated. Thus, three different anti-windup options are available to deal with the saturation problem, clamping allowing to leave saturation very fast but slowing down the process output response; back calculation allowing to obtain a trade-off between saturation time and process output performance; and the combination of the two approaches providing an intermediate behaviour according to the trade-off between saturation time and process output performance. Therefore, the user can choose between these three solutions depending on each control problem.

6.5. Example 5: Noise filtering

An interesting feature of the implementation presented in this work is that measurement noise filtering is suggested to be combined with the PID controller, as described in Section 3.9. This example shows how the process variable can be properly filtered by tuning the T_f parameter as input to the filter function. The process transfer function and PI-controller given by Equation (29) and Equation (30), respectively, have been used for this simulation, where a white noise signal was added to the process output. Figure 14 shows the simulation result for three cases: without filter, corresponding to $T_f = \Delta t/2$, and with filter for $T_f = 0.01T_i$ and $T_f = 0.1T_i$, respectively.




Figure 13.: Feed-forward plus saturation in Example 4 comparing anti-windup techniques.




Figure 14.: Noise filtering in Example 5.

Table 3.: Process and controller parameters for the gain scheduling Example 6.

	Zone 1	Zone 2	Zone 3
Process model parameters			
k [cm/(cm ³ /s)]	0.0420	0.0727	0.0938
T [s]	16.36	28.34	36.59
PI-controller parameters			
K [(cm ³ /s)/cm]	14.39	14.39	14.39
T_i [s]	16.36	28.34	36.59

The simulation starts with the controller in manual mode with $u = 0$. The setpoint is changed from $r = 0$ to $r = 3$ at $t = 1$. At time $t = 1.25$ the controller is switched to automatic mode. As observed in the bottom plot, for the case with $T_f = 0.01T_i$, the original control signal is mainly kept and only the very high frequency components of the signal are filtered. For the case with $T_f = 0.1T_i$, the original control signal is highly filtered by cutting off high-frequency components. Notice that for all three cases, almost the same process output responses are obtained, which means that the filter is mainly affecting only the control signal.

6.6. Example 6: Gain scheduling

This example shows how to use the proposed code for the implementation of gain-scheduling approaches. We consider the classical tank level control problem, whose process dynamics are given by the following differential equation:

$$\frac{dy(t)}{dt} = -\frac{a}{A}\sqrt{2gy(t)} + \frac{u(t)}{A}, \quad (34)$$

where y is the tank level, u is the inlet flow, $g = 981$ cm/s² is the gravitational acceleration, $a = 2.15$ cm² is the cross section of the outlet hole, and $A = 390$ cm² is the cross section of the tank. Since the process is nonlinear, it is divided into three different operating zones, where a linear model approximation is derived in each zone. The following three operating points are considered for this study:

Zone 1 for $y \in [4, 12)$ cm

Zone 2 for $y \in [12, 20)$ cm

Zone 3 for $y \in [20, 25]$ cm

In each zone, the process is modelled as

$$P(s) = \frac{k}{1 + sT}. \quad (35)$$

The model parameters are given in Table 3.

The gain-scheduling control scheme can be implemented in two different ways: by running all the controllers in parallel and switching among them based on the current operating point; or by using a single controller and updating its parameters based on the current operating point. Both solutions have been used in this example, and Listing 8 and Listing 9 show the corresponding codes for the multiple controllers and single controller case, respectively. Notice that the initialisation of the variables was omitted and only the control loop is presented. In both cases, the setpoint is increased from $r = 4$ to $r = 22$ cm at $t = 10$ s, and the process is initialised with a value of $y = 4$ cm and $u = 190.66$ cm³/s. When using multiple controllers as presented in Listing 8, the key point is to correctly manage the mode status of the controllers, as observed from Lines 6–21. Once the corresponding controller is selected to be active (`mode=AUTO`), the rest of the controllers must be in tracking mode, and thus the corresponding variables `mode` must be set to `TRACK`. Moreover, the tracking signal for all the controllers must be the actual input to the process in order to ensure bumpless transfer when switching among the controllers (see Line 22). In the case of using a single controller as presented in Listing 9, the key point is the adaptation of the controller parameters according to the corresponding operating point observed on Lines 4–16. Notice that in this case, the controller mode is set to `AUTO` since the beginning of the simulation. Both implementations provide identical results. Figure 15 shows the simulation results for both cases where the setpoint was changed along the operating points to force the switching among the different controllers. As observed, the same results are obtained in both implementations and the curves in the two upper plots overlap. It can also be seen how the same closed-loop behaviour is obtained despite the changes on the operating range. Moreover, note that no bumps are observed in the control signal at the switching moments. The lower plot shows the number of the active controller, representing the switching state.




Figure 15.: Gain scheduling control results for the tank level control problem in Example 6 for codes Listing 8 and Listing 9.

Listing 8: Parallel controllers in Example 6.

```

1 while t < Tsim
2     % Simulation of process dynamics
3     y=process_simulation(u,Dt)

4     % Controller 1
5     u1=pid1.control(r,y,uff,uman,utrack,mode1)
6     % Controller 2
7     u2=pid2.control(r,y,uff,uman,utrack,mode2)
8     % Controller 3
9     u3=pid3.control(r,y,uff,uman,utrack,mode3)

10    % Zone 1
11    if y < 12
12        mode1="AUTO"
13        mode2="TRACK"
14        mode3="TRACK"
15        u=u1
16    % Zone 2
17    elseif y >=12 && y < 20
18        mode1="TRACK"
19        mode2="AUTO"
20        mode3="TRACK"
21        u=u2
22    % Zone 3
23    elseif y >=20
24        mode1="TRACK"
25        mode2="TRACK"
26        mode3="AUTO"
27        u=u3;
28    end

29    % Tracking signal
30    utrack=u;

31    % Set-point change
32    if t>=200 && t<400
33        r=3
34    elseif t>=400 && t<600
35        r=15
36    end

37    t=t+Dt
38 end

```

Listing 9: Single controller in Example 6.

```

1 while t < Tsim
2     % Simulation of process dynamics
3     y=process_simulation(u,Dt)

4     % Controller
5     u=pid1.control(r,y,uff,uman,utrack,mode)

6     % Zone 1
7     if y < 12
8         pid1.kp=kp1
9         pid1.ki=ki1
10        pid1.kd=kd1
11    % Zone 2
12    elseif y >=12 && y < 20
13        pid1.kp=kp2
14        pid1.ki=ki2
15        pid1.kd=kd2
16    % Zone 3
17    elseif y >=20
18        pid1.kp=kp3
19        pid1.ki=ki3
20        pid1.kd=kd3
21    end

22    % Set-point change
23    if t>=200 && t<400
24        r=3
25    elseif t>=400 && t<600
26        r=15
27    end

28    t=t+Dt
29 end

```

6.7. Example 7: Selector or override control

This example shows how to handle the tracking state when using a selector control approach. The important thing here is to keep the P-part of the controllers as we need their current P-contribution to properly select (for instance, with MIN or MAX selectors) the current control signal. To analyse this case, we use the example proposed in Figure 16. The process has one input signal u and two output variables, y_1 and y_2 . The two output variables are input signals to the two controllers C_1 and C_2 . Only one of the controllers is active at a time, in this case the controller with the smallest control signal. The control signal corresponding to the non-active controller is tracking the active control signal.




Figure 16.: Selector control scheme.

This structure is common for handling safety constraints. One of the controllers is working in normal operation, but when a certain process variable becomes too small or too large, the other controller takes over, overrides, to ensure that the constraint is not violated. In this example, $P_1 = P_2$ and $C_1 = C_2$, where we use the process and controller given by Equation (29) and Equation (30). The setpoint values are $r_1 = 0.3$ and $r_2 = 0.5$. Figure 17 shows the simulation results for this example. A disturbance signal v is subtracted from the input of P_1 to generate variations in the operating conditions. The plot at the bottom represents the disturbance signal v . As observed from the third plot from the top (representing $u_1 - u_2$), the MIN selector works properly. The active control signal is always the smallest one. Moreover, the controller which is in tracking state calculates the control signal as the current control signal plus its P-part.




Figure 17.: Simulation results for the selector control scheme in Example 7 and for code Listing 10.

Listing 10 shows the code in which the handling of the tracking mode signals is observed based on the selector choice. As for the parallel controllers in Example 6, the tracking signal for the two controllers must be the actual input to the process to ensure bumpless transfer when switching among the controllers (see Line 15).

Listing 10: Simulation code for selector control approach.

```

1 while t < Tsim
    % Simulation of process 1 dynamics
2    y1=process1_simulation(u-v,Dt)

    % Simulation of process 2 dynamics
3    y2=process2_simulation(u,Dt)

    % Controller 1
4    u1=pid1.control(r1,y1,uff,uman,utrack,mode1)

    % Controller 2
5    u2=pid2.control(r2,y2,uff,uman,utrack,mode2)

    % Selector
6    if u1 < u2
7        mode1="AUTO"
8        mode2="TRACK"
9        u=u1;
10    else
11        mode1="TRACK"
12        mode2="AUTO"
13        u=u2
14    end

15    utrack=u

16    t=t+Dt
17 end

```

7. Conclusions

This paper can be used as a practical guide to understand and implement a PID controller in a software environment. The proposed reference implementation is intended to be used as a standard for many applications from the process industry to medical, aeronautical and automotive applications. The contribution lies in the consideration of practical problems that will be encountered when dealing with real-world problems, all of which have been carefully considered and addressed. The contribution further lies in the discussion of several simulated examples where the different features are tested in closed loop. The provided insights can be used to implement individual features of our proposed reference implementation, in use cases where it is undesirable to replace the entire PID implementation. While we expect that small adjustments will have to be made for most applications, we believe that this reference implementation closes the book on PID implementation of the basic form and opens the next when comparing advanced control strategies. In particular, feed-forward control as well as strategies requiring tracking such as selectors or override, gain scheduling and cascade control can be studied in different ways. We hope that this implementation will also make the use of advanced strategies more popular as these rely on the correct basic implementation, which is not always given in industrial settings.

Conflict of Interest: The authors declare no conflict of interest.

References

- Årzén, K.-E. (1999). A simple event-based PID controller. In *IFAC proceedings volumes* (Vol. 32, pp. 8687–8692). (Presented at the 14th IFAC World Congress, Beijing, China)
- Åström, K., & Hägglund, T. (2006). *Advanced PID control*. Research Triangle Park, North Carolina, USA: ISA—The Instrumentation, Systems and Automation Society.
- Åström, K., & Wittenmark, B. (1997). *Computer-controlled systems: Theory and design* (3rd ed.). Upper Saddle River, New Jersey, USA: Prentice Hall.
- Åström, K. J., & Hägglund, T. (2001). The future of PID control. *Control Engineering Practice*, 9(11), 1163–1175.
- Bhandari, P., & Csurscia, P. Z. (2022). Digital implementation of the PID controller. *Software Impacts*, 13, 100306.
- Bohn, C., & Atherton, D. (1995). An analysis package comparing PID anti-windup strategies. *IEEE Control Systems Magazine*, 15(2), 34–40.
- Chen, Y., Petráš, I., & Xue, D. (2009). Fractional order control—A tutorial. In *Proceedings of the 2009 american control conference* (pp. 1397–1411).
- Clarke, D. W. (1984). PID algorithms and their computer implementation. *Transactions of the Institute of Measurement and Control*, 6(6), 305–316.
- Crisp, J. (1991). *Fixed-point signal processing*. Prentice Hall.
- da Silva, L., Flesch, R., & Normey-Rico, J. (2018). Analysis of anti-windup techniques in PID control of processes with measurement noise. *IFAC-PapersOnLine*, 51(4), 948–953.
- Dahlin, E. B. (1968). Designing and tuning digital controllers. *Instruments and Control Systems*, 42, 77–83.
- Guzmán, J. L., & Hägglund, T. (2024). *Feedforward control: Analysis, design, tuning rules, and implementation*. Berlin, Germany: De Gruyter.
- Hägglund, T. (2023). *Process control in practice*. Berlin, Germany: De Gruyter.
- Hägglund, T., & Guzmán, J. L. (2024). Give us PID controllers and we can control the world. *IFAC-PapersOnLine*, 58(7), 103–108.
- Hoyo, Á., Hägglund, T., Guzmán, J. L., & Moreno, J. C. (2023). A practical solution to the saturation problem in feedforward control for measurable disturbances. *Control Engineering Practice*, 139, 105636.
- Huba, M., & Vrančić, D. (2018). Comparing filtered PI, PID and PIDD2 control for the FOTD plants. *IFAC-PapersOnLine*, 51(4), 954–959. (3rd IFAC Conference on Advances in Proportional-Integral-Derivative Control PID 2018)
- IEC. (2010). *IEC 60546-1:2010, Controllers with analogue signals for use in industrial-process control systems* (Tech. Rep.). Geneva, Switzerland: International Electrotechnical Commission.
- ISA. (2023). *ISA-TR5.9-2023, Proportional-Integral-Derivative (PID) Algorithms and Performance* (Tech. Rep.). Research Triangle Park, NC, USA: International Society of Automation.
- Krejcar, O., Špička, I., & Frischer, R. (2011). Implementation of full-featured PID regulator in microcontrollers. *Elektronika ir Elektrotechnika*, 113(7), 77–82.
- Liu, L., Tian, S., Xue, D., Zhang, T., Chen, Y., & Zhang, S. (2019). A review of industrial MIMO decoupling control. *International Journal of Control, Automation and Systems*, 17(5), 1246–1254.
- Shinskey, F. (1994). *Feedback controllers for the process industries*. New York, USA: McGraw-Hill.
- Shinskey, F. (1996). *Process-control systems. Application, Design, and Tuning* (4th ed.). New York, USA: McGraw-Hill.
- Skogestad, S. (2023). Advanced control using decomposition and simple elements. *Annual Reviews in Control*, 56, 100903.
- Sundström, E., Bauer, M., Guzmán, J. L., Hägglund, T., & Soltesz, K. (2026a). *PID.code: Reference implementation of PID controllers*. <https://github.com/copybit/PID.code>. (GitHub repository, accessed February 27, 2026)

- Sundström, E., Bauer, M., Guzmán, J. L., Hägglund, T., & Soltesz, K. (2026b). *PID reference implementation website*. <https://w3.ual.es/personal/joguzman/pid/>. (Website, accessed February 27, 2026)
- Uzunović, T., Zunić, E., Badnjević, A., Mioković, I., & Konjicija, S. (2010). Implementation of digital PID controller. In *Proceedings of the 33rd international convention MIPRO* (pp. 1357–1361).
- Visioli, A. (2003). Modified anti-windup scheme for PID controllers. *IEE Proceedings—Control Theory and Applications*, 150(1), 49–54.
- Visioli, A. (2006). *Practical PID control*. Berlin, Germany: Springer.

Appendix A. Pseudo-code language

Listings 1–5 are written in a pseudo-code language similar to MATLAB, a programming language that many control engineers are familiar with, and that the examples, Listings 6–10, are written in. Comments are **green**. Keywords are **blue**. The keywords are:

- **function** to signal that the line contains the signature of a function definition;
- **if**, **else**, **elseif**, and **end** used for logical branching. (Should the target language dictate, **else** and **elseif** can be replaced by using consecutive **if-end** clauses.) For clarity, we also use the **end** keyword to delimit functions.

Variables are colour coded, with state variables in **pink**. These are variables that must retain their values between function calls. They are managed outside the function to allow multiple controller or filter instances, and may be stored using object-oriented or other mechanisms (e.g., as in the MATLAB and Python implementations in our GitHub repository (Sundström et al., 2026a)). Parameters of the control (Listing 3) and filter (Listing 4) functions are **orange**. They can be set and updated in various ways, such as being passed as arguments (by value or reference) or retrieved from memory within the function. The specific mechanism is left to the implementer. To facilitate readability, block indentation is used, but not strictly needed to avoid ambiguity. Our pseudo-code assumes availability of the standard binary operators `==` to check equality and `>` to check ordering. Furthermore, the language is assumed to be equipped with floating point arithmetic operators `+`, `-`, `*`, `/`, and the absolute value function `abs`. To facilitate portability, we have intentionally avoided the use of compound data structures such as matrices or lists. Exceptions to this are lists (or tuples) used for arguments and returns of functions. To define values of the `mode` variable we have used macro variables `MAN` and `TRACK` for readability. A safer option for encoding the mode is proposed in Section 5.1.3. Were it not for readability, we could have used the more memory-efficient option of integers (e.g., 0 for manual and 1 for tracking mode), possibly defined through macros `MAN` and `TRACK`, language allowing. We also adhere to a list of semantic rules, to facilitate readability:

- A leading lowercase **d** in a variable name, represents a time derivative. For example `dy` is the time derivative of the measurement `y`.
- State variables have variable names with a leading **x**. This is to avoid ambiguity between for example the control signal `u` and its previous value, stored in `xu`.
- A leading uppercase **D** in a variable name represents a difference with respect to the value the variable held in the previous invocation of the function. For example `Dy=y-xy` assigns to `Dy` the difference between the current measurement `y` and its previous value `xy`.
- With the exception of uppercase **D** (explained above), variable names are all lowercase.

Typing is not explicit, but all numeric values are assumed to automatically be cast. For example 0 in `ki==0` will be cast into a float 0.0 if `ki` is of type float.

Appendix B. Signal filter properties

This appendix derives the filter update Equation (28a), and discusses how the time constant T_f should be interpreted in the contexts of a first-order filter Equation (26), and its second-order counterpart Equation (27). If it is assumed that the filter input $y(t)$ is (held) constant between samples, known as zero-order-hold (ZOH), approximation-free closed-form discretisations of Equation (26) and Equation (27) can be obtained. This is done using ZOH discretisation, described in for example K. Åström and Wittenmark (1997). Observing that $F_2(s)$ in Equation (27) is the series connection of two filters of the type $F_1(s)$ in Equation (26), it is sufficient to consider discretisation of $F_1(s)$, to then cascade two such discretisations. We can think of discretisation as a means to approximate the derivative operator d/dt as a function of the shift operator q , where $qx(t) = x(t + \Delta t)$. A simple example is the backward difference approximation, which replaces $\dot{x}(t)$ by $(x(t) - x(t - \Delta t))/\Delta t$. While computationally attractive, this approximation may introduce poor stability and accuracy when Δt is not sufficiently small relative to T_f . Therefore, we instead employ the Tustin (bilinear) approximation, which can be interpreted directly in the time domain as applying the trapezoidal rule to the time derivative. For a signal $x(t)$ sampled with interval Δt , the Tustin approximation $\hat{\dot{x}}$ of the derivative $\dot{x}(t)$ can be defined recursively as

$$\hat{\dot{x}}(t) \approx 2 \frac{\Delta x(t)}{\Delta t} - \hat{\dot{x}}(t - \Delta t). \quad (\text{B1})$$

In the transform domain, this corresponds to the bilinear mapping

$$s \approx \frac{2}{\Delta t} \frac{z - 1}{z + 1} \quad (\text{B2})$$

between the continuous-time Laplace transform variable s and the discrete-time transform variable z . The continuous-time first-order filter Equation (26) can be equivalently written in the time domain as the differential equation

$$T_f \dot{y}_f(t) + y_f(t) = y(t), \quad (\text{B3})$$

where $y(t)$ is the measured signal and $y_f(t)$ denotes the filtered output. Substituting the Tustin approximation Equation (B1) for $\dot{y}_f(t)$ into the continuous-time filter of Equation (B3), and rearranging terms to isolate $y_f(t)$, yields the discrete-time recursion Equation (28a), relating the current filtered value to its previous value and the current measurement. While being computationally favourable, and bumpless when implemented in incremental form, the cascade of two Tustin-approximated filters Equation (28) will have slightly higher noise amplification in the derivative term than the ZOH discretisation of a filtered derivative obtained by applying the time-domain counterpart of $s/(sT_f + 1)^2$ to $y(t)$. However, if $n(t)$ is white noise, the difference is negligible for the practically most relevant regime of $\Delta t < 0.1T_f$. (As T_f decreases to approach Δt , it can be shown that the noise amplification, in terms of variance, is at most 50 % better for the ZOH filter.) It is also worth noting that while T_f is the time constant, in terms of $(1 - \exp(-1)) \cdot 100 \% \approx 63 \%$ rise time, of F_1 in Equation (26), the same is not true for the relation between T_f and rise time of F_2 in Equation (27).

If one has tuned a first-order filter Equation (26) to have a (rise) time constant T'_f , and wants to maintain that property for the second order filter Equation (27), the following substitution should be made:

$$T_f \approx 2T'_f, \tag{B4}$$

where the factor of two approximates the solution of

$$\left(1 + \frac{T'_f}{T_f}\right) \exp\left(-\frac{T'_f}{T_f}\right) = \exp(-1). \tag{B5}$$

This is of practical importance if employing a tuning method that assumes a first-order filter, when in fact a second-order filter is used, and vice versa.

Appendix C. Nomenclature

This appendix summarises the nomenclature used throughout the paper. Tables C1–C4 list variables from the reference implementation of Listings 3–5, grouped into parameters, states, function arguments, and internal variables. Table C5 lists the mathematical symbols used in the theoretical development. For each entry, the “Listing” or “Equation” column, together with the “Section” column, indicates where it appears. An em-dash in the “Equation” column means that the symbol is introduced in the running text.

Table C1.: Parameter variables used in our reference implementation. All parameters are typeset in orange.

Parameter	Description	Listing	Section
b	Setpoint weight b for proportional term	3	3.2
c	Setpoint weight c for derivative term	3	3.2
Dt	Execution interval Δt	3	3.8
dumax	Upper control signal rate limit $\dot{u}_{\max}$	3	3.4
dumin	Lower control signal rate limit $\dot{u}_{\min}$	3	3.4
kd	Derivative gain $k_d = k_p T_d$	3	2.2
ki	Integral gain $k_i = k_p/T_i$	3	2.2
kp	Proportional gain $k_p = K$	3	2.2
Tf	Filter time constant	4	3.9
Tt	Back calculation anti-windup time constant	3	3.5
u0	Control signal bias term u_0	3	3.1
umax	Upper control signal saturation limit $u_{\max}$	3	3.4
umin	Lower control signal saturation limit $u_{\min}$	3	3.4

Table C2.: State variables used in our reference implementation.

State	Description	Listing	Section
xdr	Time derivative of setpoint	3	4.2
xdy	Time derivative of process variable	3	4.2
xf1	Filter state 1	4	4.10
xf2	Filter state 2	4	4.10
xr	Setpoint	3	4.9
xu	Unsaturated control signal	3	4.7
xuff	Feed-forward control signal	3	3.3
xus	Saturated control signal	3	4.8
xy	Process variable	3	4.9

Table C3.: Function argument variables passed to the control function.

Argument	Description	Listing	Section
<code>mode</code>	Controller operating mode	3	5.1.3
<code>r</code>	Setpoint	3	4
<code>uff</code>	Feed-forward control signal	3	3.3
<code>uman</code>	Manual control signal	3	4.4
<code>utrack</code>	Tracking control signal	3	3.6
<code>y</code>	Process variable	3	4

Table C4.: Other variables used in Listing 3 and Listing 4.

Variable	Description	Listing	Section
<code>a</code>	Filter coefficient	4	4.10
<code>Ddr</code>	Setpoint derivative increment	3	4.7
<code>Ddy</code>	Process variable derivative increment	3	4.7
<code>Dr</code>	Setpoint increment	3	4.7
<code>Duff</code>	Feed-forward increment	3	4.7
<code>Dud</code>	Derivative increment	3	4.7
<code>Dui</code>	Integral increment	3	4.7
<code>Dup</code>	Proportional increment	3	4.7
<code>Dy</code>	Process variable increment	3	4.7
<code>dr</code>	Setpoint derivative	3	4.2
<code>dy</code>	Process variable derivative	3	4.2
<code>u</code>	Control signal	3	4.7
<code>us</code>	Saturated control signal	3	4.8
<code>usmax</code>	Upper saturation bound	3	4.1
<code>usmin</code>	Lower saturation bound	3	4.1

Table C5.: Mathematical symbols used throughout the paper.

Symbol	Description	Equation	Section
a	Filter coefficient in discrete-time low-pass filter	28b	3.9
b	Setpoint weight for proportional action	13	3.2
c	Setpoint weight for derivative action	13	3.2
$C(s)$	Controller transfer function	30	6
$\Delta e(t)$	Increment of control error	4	2.1
$\Delta e(t - \Delta t)$	Previous error increment	11	2.3
Δn	Increment of measurement noise	—	3.9
$\Delta r(t)$	Increment of setpoint	15	3.2
Δt	Execution interval / sampling interval	4	2.1
$\Delta u(t)$	Control signal increment	8	2.3
$\Delta u_d(t)$	Derivative contribution increment	11	2.3
$\Delta u_i(t)$	Integral contribution increment	11	2.3
$\Delta u_p(t)$	Proportional contribution increment	11	2.3
$\Delta u_{ff}(t)$	Feed-forward control increment	16	3.3
$\Delta y(t)$	Increment of process variable	15	3.2
$\dot{u}(t)$	Time derivative of control signal	19	3.4
$\dot{u}_{\max}$	Upper control signal rate limit	19	3.4
$\dot{u}_{\min}$	Lower control signal rate limit	19	3.4
$\dot{x}(t)$	Time derivative of signal $x(t)$	B1	B
$\dot{y}_f(t)$	Time derivative of filtered process variable	B3	B
$e(t)$	Control error	1	2
$F_1(s)$	First-order filter transfer function	26	3.9
$F_2(s)$	Second-order filter transfer function	27	3.9
$FF(s)$	Feed-forward compensator transfer function	33	6.4
g	Gravitational acceleration	34	6.6
j	Summation index	4	2.1
k	Process gain	29	6
k_d	Derivative gain	5	2.2
k_i	Integral gain	5	2.2
k_p	Proportional gain	5	2.2
k_u	Gain of process from control signal to output	31	6.4
k_v	Gain of process from disturbance to output	32	6.4
K	Controller gain	1	2
L	Process dead time	29	6

Symbol	Description	Equation	Section
L_u	Dead time of process from control signal to output	31	6.4
L_v	Dead time of process from disturbance to output	32	6.4
n	Discrete-time index	—	2.1
$n(t)$	Measurement noise signal	—	3.9
$P(s)$	Process transfer function	29	6
$P_u(s)$	Process transfer function from control signal to output	31	6.4
$P_v(s)$	Process transfer function from disturbance to output	32	6.4
q	Shift operator	—	B
$r(t)$	Setpoint signal	1	2
r_1	Setpoint for controller 1	—	6.7
r_2	Setpoint for controller 2	—	6.7
s	Laplace transform variable	26	3.9
T	Process time constant	29	6
T_d	Derivative time	1	2
T_f	Filter time constant	26	3.9
T'_f	Equivalent first-order filter time constant	B4	B
T_i	Integral time	1	2
T_t	Anti-windup tracking time constant	—	3.5
T_u	Time constant of process from control signal to output	31	6.4
T_v	Time constant of process from disturbance to output	32	6.4
t	Continuous time	1	2
$u(t)$	Control signal	1	2
$u_d(t)$	Derivative control signal term	7	2.2
$u_i(t)$	Integral control signal term	7	2.2
$u_p(t)$	Proportional control signal term	7	2.2
u_0	Control signal bias term	12	3.1
$u_{ff}(t)$	Feed-forward control signal	16	3.3
$u_{\text{man}}(t)$	Manual control signal	—	4.4
$u_{\text{track}}(t)$	Tracking control signal	24	3.6
u_{max}	Upper control signal limit	18	3.4
u_{min}	Lower control signal limit	18	3.4
$u_s(t)$	Saturated control signal	23	3.4
$u_{s,\text{max}}$	Effective upper saturation limit	22	3.4

Symbol	Description	Equation	Section
$u_{s,\min}$	Effective lower saturation limit	21	3.4
$v(t)$	Load disturbance	—	3.3
$x(t)$	Generic signal used in filter derivation	B1	B
$y(t)$	Process variable (measurement)	1	2
$y_f(t)$	Filtered process variable	28a	3.9
y_1	First process output	—	6.7
y_2	Second process output	—	6.7
z	Discrete-time transform variable	B2	B
$\hat{x}(t)$	Tustin approximation of $\dot{x}(t)$	B1	B
λ	Lambda tuning parameter	30	6
ω_c	Filter cutoff frequency	26	3.9