2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

A New Online Continuing Education Course on Control Engineering

Jan Maciejowski
Life Fellow, IEEE

Abstract— **We describe an online 'continuing education' course on Control Engineering for graduates of any branch of Engineering or other scientific discipline. The course runs mostly asynchronously over 8 weeks, with the expectation that each student will devote 7–9 hours per week to it. Initially, motivational material is presented in the form of examples of control systems and their benefits. The feedback structure is emphasised, but feedforward, cascade and multivariable structures are also discussed. Sensors and actuators are introduced, several examples of each being given. Mathematical analysis is introduced after a qualitative understanding of feedback has been established. Emphasis is given throughout the course to PID controllers, including their implementation and limitations, as well as approaches to tuning them. 'Classical' frequency-domain analysis and design methods for SISO systems are presented and emphasised. Later parts of the course cover more advanced material such as state feedback, observers, and LQG controllers. More advanced material, including MPC, adaptive and robust control, is introduced very briefly. The course is assessed by graded assignments based on a nonlinear model of an industrial process. Students develop a working knowledge of *Matlab* and *Simulink* software during the course.**

Keywords: Continuing education, online course, mature students, control engineering.

I. INTRODUCTION

The World Bank's 2019 World Development Report on the future of work [1] states that 'continuing education' opportunities (also known as 'lifelong education' or 'lifelong learning') that allow workers to retrain and retool are vital in order for labour markets to adjust to the future of work.

This paper describes an online 'continuing education' course on Control Engineering that has been developed for graduates from any Engineering – or other scientific – discipline who need or wish to learn about Control. The course is delivered over 8 weeks, and assumes that each student devotes 7–9 hours per week to the course. This is a demanding requirement, since it is assumed that students taking this course are in full-time regular employment.

The course is 'asynchronous' in the sense that students study it in their own time, although they are expected to follow the weekly modular structure. Also there is a weekly one-hour 'live session', which is hosted by the course leader and/or a tutor, and which all students can join in a Zoom session. The students can also discuss online the course and their work with each other, the course tutor(s) and the course leader.

The course is assessed by means of a 5-part exercise which involves the use of *Matlab* and *Simulink* software. A certificate is awarded to students who perform sufficiently well in this exercise. The awarding authority is *University of Cambridge Online* [2]. The course is marketed commercially, and costs approximately 2000 GBP.

The author developed this course in 2021, with about 6 man-months of effort. The course is delivered four times per year, and to date (September 2023) has run 7 times. It is implemented in the *Canvas* platform [3].

II. TARGET AUDIENCE

The primary target audience for the course is graduate Engineers who are in employment, and need to learn about Control Engineering for their job. No assumption is made about which engineering discipline they graduated from. It is envisaged that the course may also be of use to engineering project managers, who may be supervising Control Engineers. And of course some may wish to take the course speculatively, in order to enhance their skill-set and marketability.

Being an online course, it is available to students located anywhere in the World, and who may have graduated from a university anywhere in the World, an unknown number of years ago. It is therefore necessary to make some assumptions about material with which the students already have (or had!) some familiarity.

A. Prerequisites

The published mathematical prerequisites for the course have been kept as minimal as possible, and are:

1) Basic familiarity with differential equations.
2) Basic knowledge of complex numbers.
3) Basic knowledge of linear algebra.

How basic is 'basic'? This is not defined precisely, but some examples are given to prospective students of the kind of material that will appear in the course.

Students do not really need to be proficient in any of these areas, because almost all calculations can be done by software. The knowledge is needed so that they can follow mathematical arguments, and relate them to Control Engineering problems.

Note that this target audience is very different from traditional university students. We do not need to concern ourselves with developing their intellects or problem-solving skills, or assessing how clever they are. We assume that they do not need to see formal proofs of mathematical results, but that it is enough to indicate informally how results are arrived at, without producing mysterious results 'by magic'.

Jan Maciejowski is with the Department of Engineering, University of Cambridge, UK. `jmm1@cam.ac.uk`

B. Learning objectives

The learning objectives of the course were that, by the end of the course students should be able to:

1) Recognise needs/opportunities for Control in their Engineering projects/products.
2) Assess which Control method/technology is appropriate to their project/product.
3) Design and analyse a Control solution.
4) Select and use appropriate software tools, including simulation software.
5) Read and understand graduate-level Control textbooks and some research papers.

Of course it was not expected that students should be proficient at any of these activities by the end of the course, but that they could undertake them without supervision and gradually become proficient. This skill-set seemed to be the most useful for the target audience, and was achievable given the constraints of course duration and workload.

III. SELECTION OF MATERIAL

A. Motivational material

It is assumed that students know nothing about Control Engineering. The course therefore starts with motivational material, before getting down to technical details. Examples are given of feedback systems used in domestic heating, cruise control, ship steering, and electrical power generation. Also some historical examples are presented, to emphasise that Control was possible before modern technology existed:

- 18th-century windmill fan-tail (see Fig.1),
- Watt's centrifugal speed governor,
- Black's feedback amplifier.

Fig. 1. Windmill with fantail-driven feedback

Block diagrams are introduced informally as aids to visualising how feedback operates in all the examples. Students are also urged to read the less technical parts of the survey article [4], which presents a number of recent applications of Control in various fields.

Some material is presented about sensors and actuators. The more common sensors such as LVDTs, accelerometers and orifice plates are introduced with very brief mentions of their principles of operation. A few more complex sensors, such as inertial measurement platforms and Coriolis mass-flow meters (see Fig.2) are also presented, but without any attempt at explaining how they work. Flow valves and control surface servos are given as examples of actuators. The important point is made that actuators and sensors are often feedback systems in themselves. Students are encouraged to report on what sensors and actuators are used in their industrial sectors, especially if they are unusual.

Fig. 2. A Coriolis mass-flow rate meter

The domination of Control by feedback systems is explained, but the existence and role of feedforward is also mentioned, and explained in an intuitive manner. Cascade feedback structures are also introduced early, because they are very commonly encountered in practice.

B. Prominence of PID control

The course recognises that the predominant paradigm is the PID controller. This is introduced one term at a time (ie P, then PI, then PID), again informally at first, and students are encouraged to tune such controllers 'by hand' with some simple plants. The objective is to motivate the need for analytical tools, especially as the number of tunable parameters increases, and to illustrate the effects of poor tuning. Performance criteria such as IAE, ITAE and integrated quadratic error are introduced, as well as more qualitative criteria such as degree of damping and speed of response.

The PID controller is referred to throughout the course, even when mathematical analysis and design techniques are introduced. Students become familiar with Ziegler-Nichols tuning, and with relay-based autotuning [5] of PID controllers. They also encounter integrator wind-up and how to protect against it.

Some attention is given to the detailed implementation of PID controllers. For example, positioning the derivative term in the feedback path but not in the set-point path is motivated and discussed, as is the need to combine the derivative term with a high-frequency cut-off filter.

C. *Mathematical material*

Undergraduate-level mathematics is not introduced until weeks 3 and 4, which are challenging for students who have not seen any of it before – but the hope is that most of them will already be familiar with some of it, or something very similar. At this point we introduce Laplace transforms, transfer functions, block-diagram algebra, notions of stability and their relations to pole locations, frequency response, the Nyquist thorem (without formal proof), Bode plots, and stability margins.

There is also a little material on modelling of physical systems by ODEs, including simple nonlinear models of easily-accessible examples such as the damped pendulum. State-space models are introduced, and the notion of linearising such models by using Taylor's theorem.

The analysis considers Linear Time-Invariant (LTI) systems only. The consequences of the inherent trade-off

$$S(s) + T(s) = 1 \tag{1}$$

(where S is the sensitivity and T the complementary sensitivity) are discussed, particularly in relation to limitations on the return-ratio $L(s)$. This is briefly generalised to multivariable systems, with the use of transfer-function matrices and singular values (as functions of frequency). Limitations imposed by right half-plane zeros are also discussed.

All of this mathematical equipment is used to examine the limitations of PID controllers and to examine more general controller structures. It is also used to give the students some proficiency in (SISO) frequency-domain design using phase-lead and phase-lag compensators.

Internal stability is discussed, and the Youla parametrisation is introduced for the simple case of a stable plant (when the Youla parametrisation is equivalent to Internal Model Control).

Discrete-time systems are introduced (z-transforms, pulse responses, etc), the main point made being that there is an almost-complete analogy to the continuous-time case. The need to avoid 'aliasing' in sampled-data systems is emphasised.

Some advanced topics are covered, necessarily briefly but with enough detail for the students to be able to follow the arguments. In particular, multivariable control, state feedback and linear quadratic control, observers and Kalman filters, and LQG control, are introduced in enough detail for students to be able to perform some exercises with them.

At a more superficial level, the penultimate week of the course is devoted to a 'broad but shallow' horizon scan of other topics in Control (in each case with reference to one or more books which would allow the students to follow up in more detail if they wished). The topics covered here are nonlinear systems (including an introduction to Lyapunov functions), optimal control and MPC, gain scheduling, adaptive control, and robust control.

D. *What is* not *included*

The course attempts to be 'technology neutral' and 'industry neutral' in order to be as widely applicable as possible. So it does not cover the use of Programmable Logic Controllers, for example.

It also does not attempt to teach modelling beyond an elementary level, since that typically requires domain-specific expertise. Of course it shows how mathematical models are used in Control, and it directs students to the topic of System Identification (with mention of the role of Machine Learning).

Perhaps the most questionable omission is any mention of hybrid or cyber-physical systems. The author will reconsider this if a major revision of the course is undertaken - but it would require some other topic to be dropped.

E. *Books*

References to a selection of appropriate books is given for each topic that is covered in the course. It is suggested that if only one book is procured, it should be [6], which is available as a free download. The course does not map one-to-one onto this book, but the coverage is broadly similar, with more details and examples being available in the book than in the course.

IV. CHOICE OF SOFTWARE

The course makes much use of *Matlab* and *Simulink* software [7], which is made available to all course participants. Students who are unfamiliar with these products are encouraged to work through some basic self-teaching material at the start of the course. Advanced use is not expected. It is important to emphasise that the course does not aim to teach programming in *Matlab*, but rather uses it as an advanced control-oriented 'calculator'. Thus students never need to draw a Bode plot or step response or compute a stability margin by hand, for instance, because *Matlab*'s Control System Toolbox does it for them.

One element of *Matlab* programming that *is* taught to students is the use of the LTI (= Linear Time-Invariant) object class. This allows, for example, a Bode plot of a return-ratio to be obtained by a statement such as

bode(Plant*ActuatorLag*Controller)

even if the Plant, ActuatorLag, and Controller have different individual representations (such as transfer-function, state-space, or pole-zero). It also allows statements such as

S = 1/(1+L)

to form the sensitivity function

$$S(s) = \frac{1}{1 + L(s)} \tag{2}$$

where **S** and **L** are both LTI objects.

With *Simulink* the case is a bit different. Students are encouraged to model systems using *Simulink* block diagrams, so in that sense they *are* being taught programming. These diagrams are also used in the course to show how algorithms such as anti wind-up protection, or relay-based autotuning, are implemented (see Fig.3).

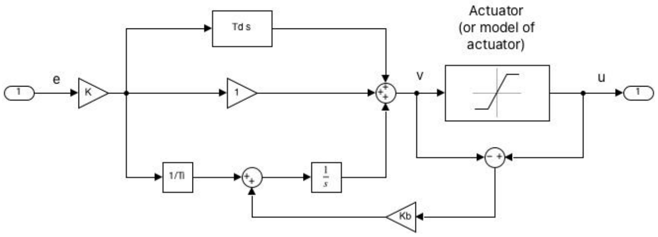

Fig. 3. A *Simulink* block diagram showing an implementation of anti wind-up protection

The 'close coupling' between *Matlab* and *Simulink* allows students to obtain linearisations of nonlinear *Simulink* models as LTI objects, which can then be subjected to further analysis using the Control System Toolbox of *Matlab*.

Familiarisation with up-to-date modelling, simulation and analysis tools, such as those provided by *Matlab* and *Simulink*, is considered to be a very tangible and demonstrable addition to the skill-set of the students that take this course.

This particular set of software tools was chosen because it is commonly used for control system analysis and design in various industrial and commercial sectors, and because it has, to the author's knowledge, more extensive support for control engineering than any other available product.

V. DELIVERY MODES

The great majority of the course is delivered as text with graphics, broken down into small modules, on the *Canvas* platform. There is also a small number of videos with audio, usually one each week on a specific topic (see Fig.4).

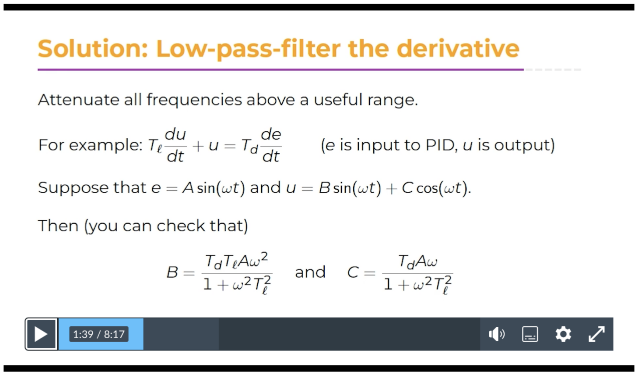

Solution: Low-pass-filter the derivative

Attenuate all frequencies above a useful range.

For example: $T_\ell \frac{du}{dt} + u = T_d \frac{de}{dt}$ (e is input to PID, u is output)

Suppose that $e = A\sin(\omega t)$ and $u = B\sin(\omega t) + C\cos(\omega t)$.

Then (you can check that)

$$B = \frac{T_d T_\ell A\omega^2}{1+\omega^2 T_\ell^2} \quad \text{and} \quad C = \frac{T_d A\omega}{1+\omega^2 T_\ell^2}$$

Fig. 4. One frame of a video about filtering the derivative term in a PID controller

There are also quizzes and 'tasks' for students to perform, such as block-diagram algebra, Ziegler-Nichols tuning, or finding an equilibrium condition for a nonlinear model and linearising about that condition (generally using *Matlab* and/or *Simulink*, sometimes 'by hand'). Students are also asked to report their results (publicly, namely visible to other students) and are encouraged to discuss their results and experiences with each other. This work is not graded, although the course tutor and/or the course leader frequently comment on it, especially when students report difficulties.

Although the course generally runs 'asynchronously', there is a one-hour 'live session' scheduled each week, hosted on Zoom by the course leader and tutor, which everyone is encouraged to join.

Certificates of success are given to the students on the basis of 5 graded assignments, which are all related to each other, and are executed using *Matlab* and *Simulink*. The results are submitted for assessment via '*Matlab* Drive', and can consist of screenshots and *Matlab*/*Simulink* files, assembled into a single folder. The assignments are all based on the nonlinear model of an industrial evaporator process which is presented in [8] (see Fig.5). Students are provided with a *Simulink* model of this process, in which 3 variables are to be controlled, using 3 manipulated variables. Three of the assignments require the successive design of 3 single-loop PI/PID controllers for these variables (two by Bode plot analysis, one by Ziegler-Nichols tuning and autotuning), taking into account the addition of actuator lags. The fourth assignment is the design of a multivariable LQR controller for all 3 variables simultaneously (without actuator lags, for simplicity). The fifth (in fact, initial) assignment is a simple check that the open-loop model runs and settles to the expected equilibrium conditions. Those students who submit an assignment before its deadline (typically 2 weeks after seeing the assignment) can get feedback from the course tutor, and are allowed to make a revised submission. Students are assured that the grading is not a 'zero-sum game', it being possible for all students to get top grades. A certificate is awarded if a student gets a grade of 70% or better. Results are expected to be self-consistent to obtain such a grade – namely, no unique 'correct' design is expected, so long as the analysis and design process is consistent, minor slips and errors being discounted.

VI. EXPERIENCE SO FAR

A. Student backgrounds

The students who have taken the course so far come from a variety of industrial sectors. The majority can be loosely classed as 'electronics', and the largest group among these work on semiconductor-based power converters. Other energy industries are also well represented, particularly in the 'renewable' sector, such as wind, solar and wave power generation. Automotive and aerospace also have significant representation, including at least one from Formula 1 car racing. A few participants have come from the healthcare sector.

The great majority of participants have Engineering degrees, as anticipated, most of them in the 'electrical' area, broadly interpreted. But there have also been some with

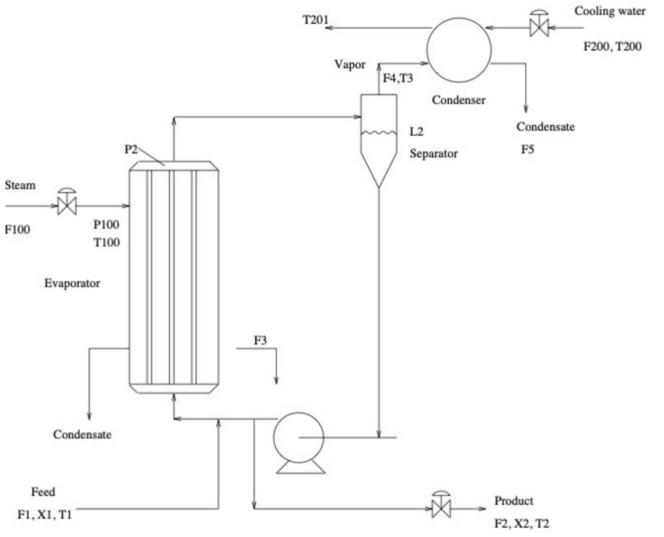

Fig. 5. The evaporator process used as the basis of the graded assignments

physics and mathematics degrees. A few appear to have been through non-university education, such as industrial-based technician training, sometimes followed by a part-time degree. As could be expected, these students have more trouble following the mathematical developments (sometimes for trivial reasons, such as unfamiliar notation), but often make interesting contributions to discussions on topics such as sensors and actuators, or controller implementations and other aspects of industrial practice.

Surprisingly few (not more than 10%) of the students come from the process industries, including oil and gas. This may reflect the structure of those industries, in which Control expertise is very much concentrated in the hands of equipment and software vendors rather than process operators.

A few of the students have been enrolled as PhD students working on applications of Control Engineering at various universities, and are using the course to quickly acquire some knowledge of Control, or to refresh/improve their existing knowledge. Sometimes these have non-Engineering backgrounds, so that this is the first Control course that they have taken.

B. Student participation

Typical enrolment for each run of the course is 10–20 students. However, little more than half of these engage with the course tutor, the course leader, and other students. We usually see much the same subset of students at each weekly Live Session, and the same subset of this subset asks most of the questions and contributes to most of the discussion. Each Live Session is recorded, so that all the students can watch it in their own time.

The reasons for this low degree of engagement are not clear. Sometimes there are clashing work commitments, but we are usually told about those, the remainder remaining a mystery. One can speculate that some students have low self-confidence and do not wish to betray in public that they are struggling with the material. Another possibility is that some students are signed up to the course by their employer, but have no real interest in it personally. And of course there are some students who find the course relatively easy, like to work on their own, and who do not feel that they need any support.

Despite the low apparent engagement, the great majority of students submit assignments which are good enough to warrant issuing a certificate of success. Typically only one or two of the participating students in each run fail to obtain a certificate. Since the qualifying standard is quite high (see section V) we infer that the great majority of the students benefit from participating in the course.

C. Difficulty of the course

The course is clearly not an easy one to complete. A lot of material is presented, and pressure of time leads to a rather sparse presentation with little repetition and fewer examples than would be expected in a text-book, say. There is also considerable variation in the amount of time required to work through each week, because weeks are organised mostly around connected topics, rather than by the amount of time each one involves.

Most students report spending more than the expected 7–9 hours per week on average, with a few reporting as much as double that. This seems to be the result of students trying to understand all the details of the course during their first pass through the material. This is probably neither necessary nor advisable. We should consider suggesting to them that learning is not a linear process, and that sometimes material will be clarified by later material or experience. (This kind of advice would normally be supplied face-to-face by a tutor, but is not so readily supplied in an online course.)

Some students have more difficulty with some of the *Matlab* and *Simulink* functionality than the author expected. A particular stumbling-block is the function **linmod**, which is used for linearising *Simulink* models. Its documentation is complete and correct, but needs to be read very carefully in order to be used correctly.

VII. CONCLUSION

Some practitioners will probably regard this course as excessively theoretical. On the other hand, it is clearly not a ‘Control theory’ course, because of the inclusion of practical matters such as anti wind-up protection, emphasis on the use of simulation software, material on sensors and actuators, etc. Hence the decision to call the course ‘Control Engineering’.

The course is a novel attempt to provide ‘continuing education’ specifically for Control Engineering. It is longer than most ‘continuing professional development’ programmes, but much shorter than part-time degree courses. There have been about 85 alumni of the course to date, most of whom have expressed satisfaction with the course. It is believed that the course usefully fills a gap in the range of available courses about control engineering.

REFERENCES

[1] `https://documents1.worldbank.org/curated/en/816281518818814423/pdf/2019-WDR-Report.pdf` (accessed 17 March 2023).

[2] `https://advanceonline.cam.ac.uk` (accessed 16 March 2023).

[3] `https://canvaslms.instructure.com` (accessed 10 March 2023).

[4] Samad, T., Bauer, M., Bortoff, S., Di Cairano, S., Fagiano, L., Odgaard, P.F., Rhinehart, R.R., Sánchez-Peña, R., Serbezov, A., Ankersen, F. and Goupil, P., 2020. 'Industry engagement with control research: Perspective and messages', *Annual Reviews in Control*, **49**.

[5] K.J. Åström and T. Hägglund (2006). *Advanced PID Control*, The Instrumentation, Automation and Systems Society.

[6] K.J. Åström and R.M. Murray (2021). *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd edition, Princeton University Press.

[7] `https://www.mathworks.com/products` (accessed 15 March 2023).

[8] R.B. Newell and P.L. Lee, *Applied Process Control, A Case Study*, Prentice-Hall, 1989.