PID Tuning – science, art or both?

In the last article we have honoured the PID controller and discussed the key factors that made it such a remarkable success. On the other side, we have also taken a critical look at the weaker points of this controller and in this context we pointed also at the still existing problems and difficulties with its tuning – despite the millions of applications and the many, many tuning methods that were developed over all these years. Let this be motivation enough for taking a closer look at the different tuning approaches and the issues involved.

The art

Even today most PID controllers are still tuned manually, by trial and error. One reason why this approach is favoured by many control professionals is that it requires practically no preparation work. Another one is that reasonable results can be achieved without doing any plant testing, which is typically required when we use a calculation method. No matter whether these tests are done with the loop open or closed, this is a clear plus because tests are very much disliked by the operation staff and thus we need to be quite convincing and persistent to get the permission to carry them out.

On the negative side, trial and error tuning can take a tremendous amount of time. There is no systematic path to the final answer and thus it is left up to the experience and the imagination of the user, with other words: to his or her ability to extract the information contained in the behaviour of the PV and the controller output and to turn these findings into meaningful adaptations of the controller settings. This is obviously an art – which inevitably leads us right to the weakest part of that approach: Not everybody is a true artist and therefore on average the results are rather mediocre – a fact that is confirmed by all the studies that dealt with the performance of the controllers in industrial plants. Interestingly and strangely (and with respect to the power and the features of today’s control systems even most embarrassing), these figures have also not changed much over time. One conclusion that can be derived is that most users today do not have much more knowledge or experience than those 20 or 30 years ago, their level of skills is practically the same despite all the research and the advances in the underlying theory over time. This points right at a weakness that has been identified already long ago but still is present today, namely, shortcomings in the education of the users. I can confirm this from my own teaching experience: “Now I finally understand …” is certainly the most frequently made statement from the trainees.

The science

Alternatively to trial and error tuning one could also take an approach that is more scientifically and technically founded, namely the use of one of the dozens of tuning methods that have been developed over time.

With respect to the data to be supplied as the basis for the calculations we can divide these methods into two classes: Those that are based on the process parameters and the (rather few) other ones that are based on the outcome of a closed loop test, a setpoint change. The process parameters needed for the first class are typically derived from an open loop test and therefore these methods are commonly also referred to as open loop methods. Let us now compare the two approaches without going much into details, just by looking at their pros and cons.

Open loop methods

Regarding the advantages of the methods that use the process parameters we can say
that they force us to gather more and especially quantitative information about the process behaviour - which always is a big plus: The better we know the process, the better we will be able to control it. Besides, they give us the possibility to calculate the tuning parameters for different situations before the occur, e.g. in cases where there is a change in the process or the operations coming up and we can estimate the effect of this change. The deadtime e.g. is often inversely related to the throughput of the plant. Thus if we know that the throughput will be reduced by x % we also know immediately that the deadtime will increase by the same percentage. With that we are already prepared and have new controller settings at the ready when the change in the operation occurs – without any further tests.

On the negative side we have to note down the need for an open loop test – which is never liked by operations staff and in some cases even not possible: Some loops they would never allow us to open. Estimating the process parameters alternatively from ‘normal’ operating data can be quite difficult as they may not have enough information content in them or they may be too much distorted by influences outside the loop.

Also, the results of these methods will be never perfect - by design - as almost all of them (I only know only one exception, namely Haalman) use a process description that is reduced to first order plus deadtime or even integrator plus deadtime. Of course processes of high order, let alone those with open loop overshoot or inverse response, cannot be dealt with very accurately at all.

Another however not so obvious fact is that we need to examine quite carefully if the method we intend to use is indeed suitable. Some of them (example: Cohen-Coon) have been developed at a time when digital systems did not exist and every single controller was working in a range of 0-100%. Of course, at that time there was no need to take the measurement range of the controller (and also of the secondary in case of a cascade) into account. With today’s DCS systems, however, this is a must, as they all allow us to specify the range in engineering units. Funny enough, however, almost all of them still calculate the PID action ***not*** in engineering units but normalized with respect to the range - just as the ‘good old’ analogue controllers, an anachronisms.

On top of that, to make things even more puzzling, some tuning methods are simply not suited for more difficult processes as they use only the time constant for the calculation of the integral time and completely miss the very severe effect of the deadtime – a weakness that is hard to understand. Sometimes this is camouflaged by intensive use of the controllability ratio, the ratio between deadtime and the time constant, in connection with either the time constant or the deadtime, but when carrying out all the calculations one can see that it still boils down to the time constant only.

And finally we must state that, since the test was done open loop, these methods do not reflect the specifics of the controller used, the differences in the PID formulation from system to system, from vendor to vendor.

Closed loop methods

Closed loop methods, i.e. methods that allow us to calculate the PID tuning based on the outcome of a setpoint step with the controller in automatic mode, do not have this shortcoming: In this case the results include of course all the elements and specifics that are present in the loop, may they concern the process or the controller. Therefore closed loop tests give us a more complete picture of the actual situation.

Unfortunately, they also have some severe weaknesses: The principle of the original Ziegler – Nichols method is stunningly simple but is based on one requirement that makes it practically prohibitive for use in the process industry: The loop has to be brought into stable oscillation. Some enhancements were made later and one approach allows us now to make a setpoint test such that PV shows a strongly dampened oscillatory response with two peaks and of course a minimum in between. From just 5 data points that we take from this response it is possible to calculate the PID tuning as well as the process parameters in one go (more details on the method can be found on our Web site). This underlying calculation is by no means simple but there are tools available that can do this for us in a second.

This sounds great but still has one severe limitation: This methods delivers only accurate results when the test is done with a pure P-controller. Since a few month it can be done now approximative with a PI-controller as well, but of course it would be ideal to do this also in an accurate way for the PI-controller: We could retain full control during the test. Unfortunately there are no signs that this will be possible in the near future: It is hard to imagine that extending the P-only method is an insurmountably difficult task, but so far all my efforts to interest researchers to work on this extension have not found open ears. Thus we still have to wait patiently for science to deliver.

Back to the art?

Because of the mentioned shortcomings of the closed loop methods, in most cases the open loop calculations are used. But no matter what approach we use there is one fact which we have to be clearly aware of: All these various methods give us only starting points, typically pretty good ones, but never the real true optimal settings for the situation. The reasons are quickly found.

Concerning the open loop methods, all are based on various assumptions and simplifications: As said before, practically all of them just use a first-order-plus-deadtime process representation.

Also, all methods assume more or less idealistic conditions: They do not take any ‘imperfections’ into account like noise on the measurement, sticking valves or the strength of the interaction between two loops. The effect of the latter one on the PID tuning can be also calculated but it comes as a side-result from the Relative Gain Analysis (RGA) and is not inherently done in the tuning calculations itself.

Using such methods saves a lot of time that is otherwise required to get to such a good starting point. The consequence, however, because of the assumptions and simplifications made, is that the fine-tuning of the controller is still up to the users.

And this demands from them the same special skills and the same good portion of experience as for trial and error tuning – which brings us back to the beginning, to the art of the tuning of the PID. So even when we decide to choose a suitable method and thus to proceed in a more systematic and more efficient way, we still have to use our knowledge and also our imagination. Therefore, at the end, PID tuning is both: A science and an art.

Hans H. Eder

ACT
hans.eder@act-control.com
www.act-control.com