High-fidelity operator training simulators

Integrating a high-fidelity simulation model with a real plant control system requires some finessing to produce a true high-fidelity operator training

MARTIN SNEESBY
APESS

The current trend in operator training simulators (OTS) is to connect a high-fidelity dynamic simulation model to a copy of the actual plant distributed control system (DCS) that is running the same software as the plant and uses the same human-machine interface (HMI). This approach is attractive and usually offers the best opportunity to produce a high-fidelity OTS that will maximise training value. However, success is not guaranteed by the project model, and there are often some complex and difficult issues to solve as part of the project delivery, particularly for existing plants where the process and control dynamics are well known.

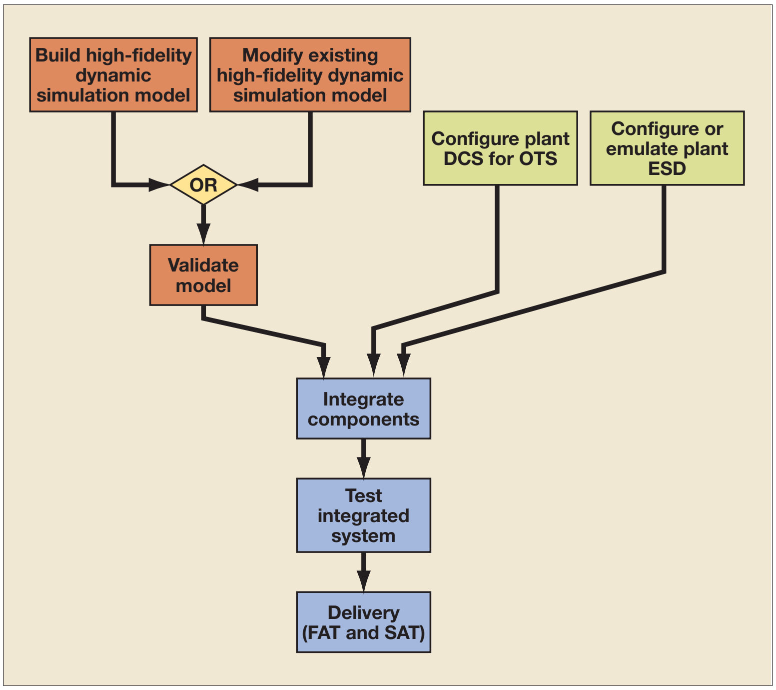

Figure 1 Standard project model for an operator training simulator

Project phases

Project award only arrives after a lot of preliminary work by the sales team and the buyer. Issues are discussed and decisions taken. Project award typically marks the end of this phase and is usually a happy occasion. Those involved (vendor and buyer) will congratulate themselves on their decisions and look forward to an OTS that meets all their expectations. Thereafter, the project implementation team will be put to work following a fairly standard project model (see Figure 1).

The first major implementation task is model development. A new dynamic process model will require an enormous amount of data to be processed and configured within a model framework. This needs the attention of simulation experts in order to ensure robustness, accuracy and adequate simulation speed. Several months (or more) are likely to be needed. Short-cuts are rarely available or effective because they almost inevitably reduce robustness and/or accuracy and, therefore, compromise the project's objectives.

A pre-existing simulation model can be used, but it is still likely to need significant work to adapt it to the specific OTS requirements, current plant configuration and operating conditions, and/or to make it consistent with the current simulation technology. Reuse of an existing model also risks depriving the project of the deep process understanding that comes from building a high-fidelity model. Process understanding is almost always crucial to any problem solving on an OTS, including issues with the DCS and ESD functionality, and it is critical this is not ignored. Common mistakes include:
- Separating the model builders from those with real plant experience
- Trying to apply a one size fits all execution model
- Ring-fencing parts of an existing process model against any change.

These approaches should be avoided, even if they appear superficially cost effective.

Pre-existing models often come with baggage, too. Rose-tinted glasses can make old models unrealistically attractive, even though the project team accepts that there have been changes to the plant, changes to simulation software and the need for a new project at all. Too many projects have floundered on the unrealistic assumption that a pre-existing model will be good enough to meet elevated standards.

The dynamic process model must then be validated. Tests are usually set up on the standalone model to prove its accuracy to one or more steady state conditions before testing the model's dynamic behaviour. The model responses should be 'reasonable and realistic' for a wide range of transients. Judging reasonableness and realism is difficult. Typically, experienced operating staff are asked to make this assessment and, while this is usually the best option, it is not quantifiable and is subject to preconceptions and other human factors. For example, it is very common to have experienced operators reporting that a model reacts more slowly than the real plant. This misconception is well known by the idiom that a watched pot never boils. Real plant data that are free of excessive noise and unmeasured disturbances would be ideal, but is difficult to collect and is not usually included in the project planning.

Model validation is, occasionally (and wrongly), considered an unnecessary luxury. Maybe those involved are sufficiently confident of the model's fidelity, or they feel the schedule is more pressing than quality control. However, skipping model validation (or failing to take sufficient notes during the process) is usually the prelude to a project's disaster.

In parallel with model development (or starting slightly later) is the preparation of an OTS version of the real plant DCS. This step is necessary to incorporate basic OTS functionality such as start, stop, save, load and possibly faster than real-time operation. It may also be a necessary (or cost effective) mechanism to ensure the project hardware costs are minimised so that redundancy (unnecessary in an OTS) is not built in accidentally. Not all vendors have the same approach to preparing the OTS version of the real plant DCS, and it is not guaranteed that all aspects of the DCS functionality will be fully supported. In particular, logic sequences and alarm monitoring systems should be scrutinised to ensure that they function correctly under all circumstances, including save and load. This phase is not trivial, and the effort involved should be measured against the likely statement of requirements that it should be possible to quickly and easily import that current plant DCS database into the OTS.

So far, so good. The project has successfully developed a high-fidelity dynamic simulation model of the process and an OTS version of the real plant DCS. The integration of these two systems will typically take a few weeks to a few months. This task is predominantly mechanistic, although there will always be a set of more taxing problems to solve, and the difficulty should not be dismissed lightly. Care needs to be taken to ensure that signals have the right sense (on/off or healthy/unhealthy), control actions and valve actions match, and signals arrive at the correct landing site with the DCS. The complexity of this task is often multiplied by the need to include an ESD system that is tightly coupled to the DCS, or to manipulate process signals to meet plant formats. Aspects of this project phase can be automated, but only after establishing the rules that apply for this particular project.

Surprises

The finish line would seem to be in sight. However, it is at this stage of the project that someone will suddenly discover that a particular control loop on the integrated system (model plus simulated DCS) is unstable or slow to converge. They will be disappointed and an investigation will begin. The first port of call is likely to be checking the controller tuning constants, perhaps by ringing the control room to make sure the project team has the latest data! Next, someone will conclude that the tuning is right so that the model is wrong. The simulation engineer will then be asked to fix the high-fidelity dynamic simulation model to better match actual plant performance. And this is where the problems really start.

Process control engineers tend to have the best understanding of the offset between a model and reality. They are often reluctant to use the results of a dynamic simulation model to build or tune their controllers. Instead, they prefer to use the plant as a test bed and tune their controllers to the actual plant responses. In most cases, this approach is pragmatic, sensible and efficient (although it is often possible to get good, noise-free data from a high-fidelity simulation, too). What should be clear is that combining plant controller tuning constants with a simulation model built from theory is not guaranteed to work perfectly. However, despite this, it often comes as a surprise to the project team and is unlikely to have ever been discussed back at the post-award dinner.

Operators (and engineers) tend to look at the overall system response when comparing a simulator to their reality. Both the model and the control system contribute substantially to the overall response. They are complementary rather than supplementary. When tuning a controller, the process control engineer will look for a set of tuning constants that work best with the plant responses, with a goal of achieving a quarter-decay response or some other (arbitrary) target that makes sense at the time and for that particular system. There is not a right answer to controller tuning. It is a problem to be optimised against a wide range of constraints. If the simulation model differs from plant behaviour, a different set of tuning constants will be optimal.

At this stage of the project, it might be time for some introspection. The use of the actual plant controllers (or simulated version thereof) and actual plant tuning constants is a good thing, but the real goal should be the overall system response. Exactly how good is the high-fidelity dynamic simulation model?

Sources of model error

High-fidelity dynamic simulation models are now widely accessible. Computing power has increased and the tools have become easier to use. The unit operation models typically contain the best available solutions. The old adage of garbage in equals garbage out still applies, but simulation software, when used by the right people and in the right way, is more powerful than ever. Engineers should be able to trust the results from well built models. Despite this, there remains a (small) gap between a model and reality. The size of the gap (the model error) might be decreasing, but it is foolish to ignore it.

The process modelling engineers that build and configure these models are required to accept and make many assumptions and estimates. The better modelling engineers will be inherently aware of the assumptions they are making (either explicitly or implicitly) and will be well equipped to make appropriate engineering estimates where necessary. However, this is not always sufficient to match plant behaviour for many reasons (see Figure 2). These include:

- Imperfect unit operation models that do not capture all of the real process behaviour and/or the real process dynamics
- Invalid modelling assumptions, either explicitly made by the modelling engineer, or implicit to the modelling approach
- Inadequate physical properties and/or thermodynamic methods
- Inaccurate design performance data (ie, the unit does not operate exactly according to design), or degraded system operation (fouling or other limitations)
- Inaccurate data estimates
- Unmodelled malfunctions and phenomena (flooding, side reactions, and so on)
- Stability and robustness constraints imposed by the

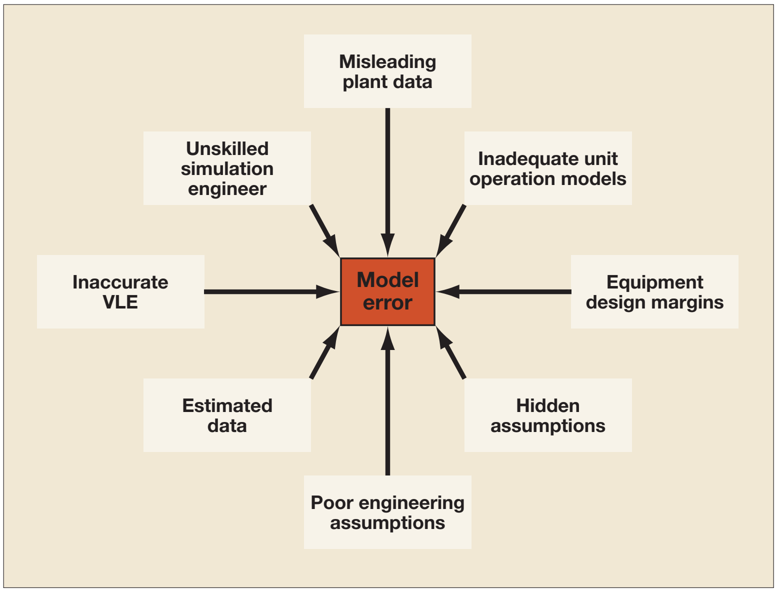

Figure 2 Sources of plant model mismatch with high-fidelity OTS models

simulation mathematics

• Misleading plant data that suggest correlations between actions and effects where no direct connection exists or responses influenced by unmeasured (or unobserved) disturbances.

The basic unit operation models have changed very little over the last 50 years. This is not to say that they are already perfect, but it suggests that we have found ways to live with their weaknesses, particularly for equipment design, which remains the dominant application of process simulation technology.

The equilibrium stage model is a good example of both the strengths and weaknesses of the established approaches to unit operation. It includes assumptions about perfect mixing and instantaneous vapour-liquid equilibrium, which we know to be dubious, and can use efficiency factors that are hard to justify or specify a priori. Despite the limitations, the equilibrium stage model has a solid track record in process design and it remains a cornerstone of process simulation. However, there is less evidence that it meets all the demands of a high- fidelity dynamic model. For example, variation in throughput with stage efficiency (or weir height) is often seen on process plants, but is not predicted. Similarly, flooding (like surge on a compressor) involves pseudo-random phenomena that are beyond most models. Even the hold-up on a tray can be difficult to predict accurately because of the need to calculate frothing factors and so on. Nevertheless, distillation models based on the equilibrium stage model usually predict time constants that are good enough for most purposes, except possibly advanced controller design, where step testing is still the de facto approach.

It is easy to overlook the thermodynamic models used in a simulation as a potential source of uncertainty. Hopefully, the modelling engineering will be sufficiently skilled to choose an appropriate method, but the results still contain a margin of error. Thermo will generally affect process gains rather than time constants, but may affect both in distillation and some other systems.

Many models are based on assumptions of lumped (homogeneous) volumes even though we know that perfect mixing is unlikely to be realised. Pure dead-time is often poorly captured because of this assumption. However, lumped volumes are necessary to meet speed and robustness demands.

Equipment is often designed and delivered with a margin of error. We can never be absolutely sure how an item of equipment will perform. Even a simple valve is less than definitive – different valve manufacturers will use different equations to calculate flow, and the results can vary by more than 10%. Heat exchangers and distillation columns have similar degrees of uncertainty associated with the design and subsequent performance.

Plant data are usually held up to be fact. However, transmitters and measurements are not infallible. Uncertainty is always present. Sometimes plant measurements can mislead and suggest phenomena that are not actually there. Data reconciliation is one technique for guarding against these errors, but it is not always viable.

Many excellent chemical engineering books have been written on troubleshooting plant problems. Books by Norman Lieberman and Henry Kister are especially noteworthy. The prevalent message in these books is that things can go wrong on an operating plant. Sometimes, process equipment does not operate the way it was intended or designed. So why would we expect all models to be perfect over all operating conditions? Of course, we do not. We just occasionally get seduced by the accuracy of high-fidelity models and forget that all model predictions are subject to uncertainty.

Controller tuning parameters

Process control engineers understand that all control loops have to be tuned for their specific characteristics. Some loops are simple and need little attention. For example, a simple flow controller is almost guaranteed to work well with a low gain and low integral time (high reset rate). Other loops are more complex and require the right combination of gain and integral time to match the process (see Figure 3).

Figure 4 shows the effect of the process gain and time constant on the overall loop performance for a specific set of controller tuning parameters (responses are indicative only). The situation is the inverse of the controller tuning. If the model time constant is much less than the actual process time constant, the loop may be unstable, slow to settle or very sluggish. Mismatches in gain also affect the response, but less critically if the process gain is predicted accurately.

It should also be noted that loops with little apparent dead time are more robust against model controller mismatch than those with more dead time. Unfortunately, model predictions of apparent dead time tend to be much less accurate than those for gain and the process time constants.

Solutions

A high-fidelity process simulation is built from unit operation models that are based on accepted theory and configured with design performance parameters. These models are said to be predictive because they do not rely on operating data for their development. Results are not regressed into the model. There are no tuning parameters to quickly change the gain or time constant of a particular loop. This, in itself, might come as a surprise to some members of the project team.

The simulation engineer may be tempted to modify the tested model with a gain or bias on a particular instrument, or to add a lag (or lead) to the system to meet the immediate need. These actions, even if superficially effective, are unlikely to address the real problem and often create further issues in themselves because a high-fidelity OTS needs to work over a very wide range of operating conditions, and this type of tuning at one operating point could be detrimental at another.

If quick fixes are excluded, the modelling engineer must review the process model with a fine-toothed comb. This involves checking and rechecking all the input data (thousands and thousands of parameters) and assessing all the implicit assumptions. A solution may or may not be available via this route.

In short, tuning the model to match a desired performance is difficult to impossible.

The pragmatic engineering solution is to use a modified set of controller tuning constants so that the overall process response (process dynamics plus control action) is closer to reality. This adds a small amount of complexity to any future upgrades to the simulated DCS, but that process is already sufficiently complicated that it should only be realistically done no more often than once per year.

For the examples in Figure 3 and 4, the central trend (#5) represents a controller that is well tuned to the process. The overall response is good. However, if there is a mismatch between the process time constant predicted by the model and that seen on the actual plant and we insist on using the actual plant tuning constants, the responses will certainly differ from the plant and will possibly show signs of instability if the plant controller is tuned aggressively. However, if the simulated controller can be recalibrated to match the model dynamics, the overall response can be significantly improved.

It is important to realise the limitations of retuning controllers. If the model is significantly slower to respond than the real plant, no amount of controller tuning is going to result in a response as fast as the plant. However, if a loop is stable on the real plant and unstable in the OTS (with plant tuning constants, see Figure 5) then retuning the OTS controllers

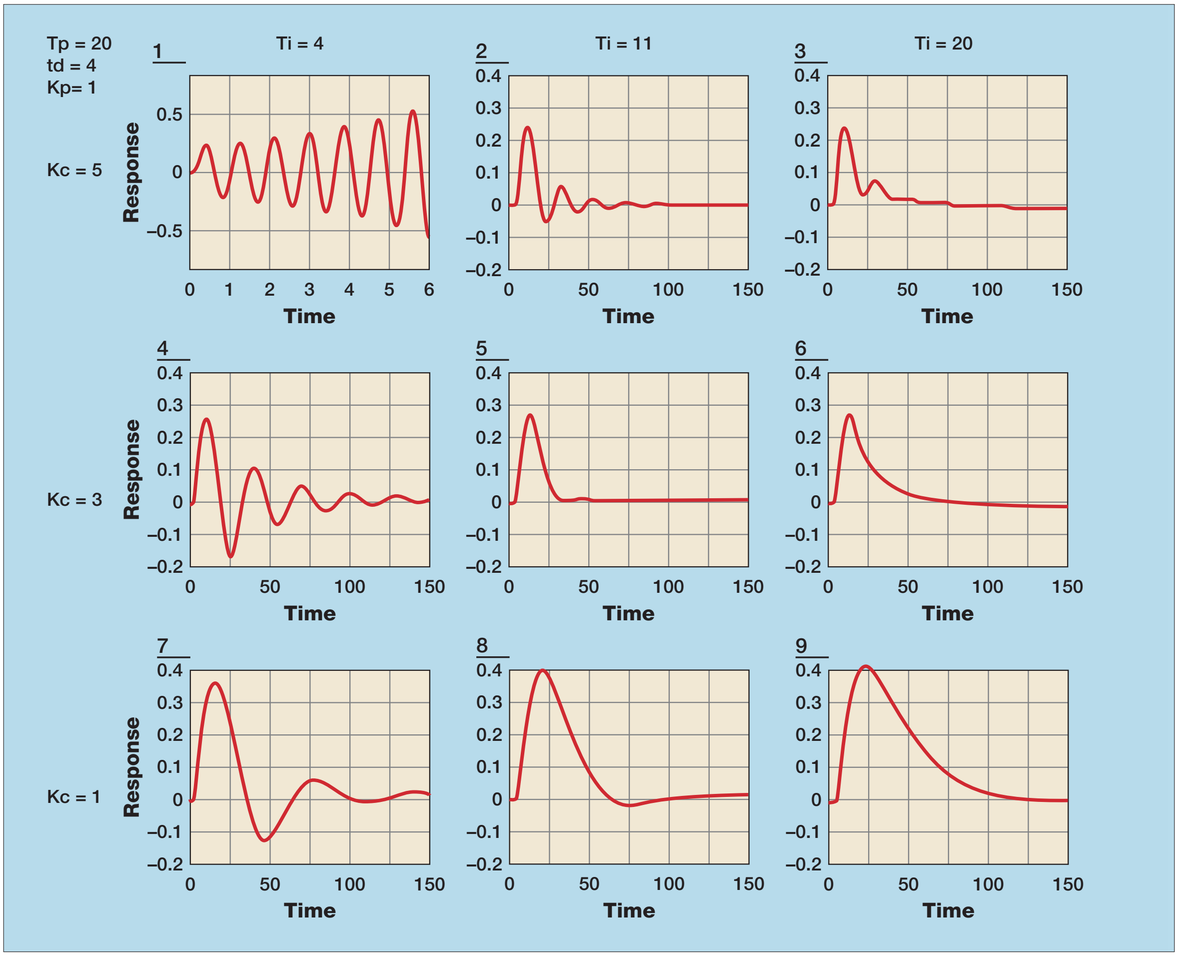

Figure 3 Predicted control responses with various controller parameters for a FOPDT model with process gain (K) = 1, process time constant (Tp = 20) and process dead time (td = 4)

can certainly help. Of course, this should always be considered in conjunction with the possibilities of retuning the model, as discussed below.

Model tuning opportunities

Despite all of the above, which hopefully comes across as logical and reasonable, some project teams still insist that the plant tuning constants should be treated as sacrosanct and that it is up to the modelling engineer to 'fix' the model. What, realistically, can be done?

Some data are more solid than other data. Recognising the more questionable data is a good starting point for model tuning. Avoid adjusting anything that is held to be true, including the laws of physics and the first law of chemical engineering (the mass balance).

The operating plant often provides clues: an exchanger seemingly operating at better than design performance, a compressor not delivering the expected head, possible flooding in a column, and so on. These issues can often be addressed by tuning the equipment parameters, for instance, (increasing the specified UA, adjusting compressor factors) so that a closer match to the steady state condition is achieved. The process gains can be manipulated via this route, but the impact on the process dynamics is likely to be small.

If the problem is an unstable controller, it is necessary to look at the model parameters that directly affect the process dynamics. First and foremost are the system volumes. Volumes are often estimated inaccurately because they do not affect the steady state solution and larger volumes help model stability. It is tempting to over-estimate volumes in order to make the model

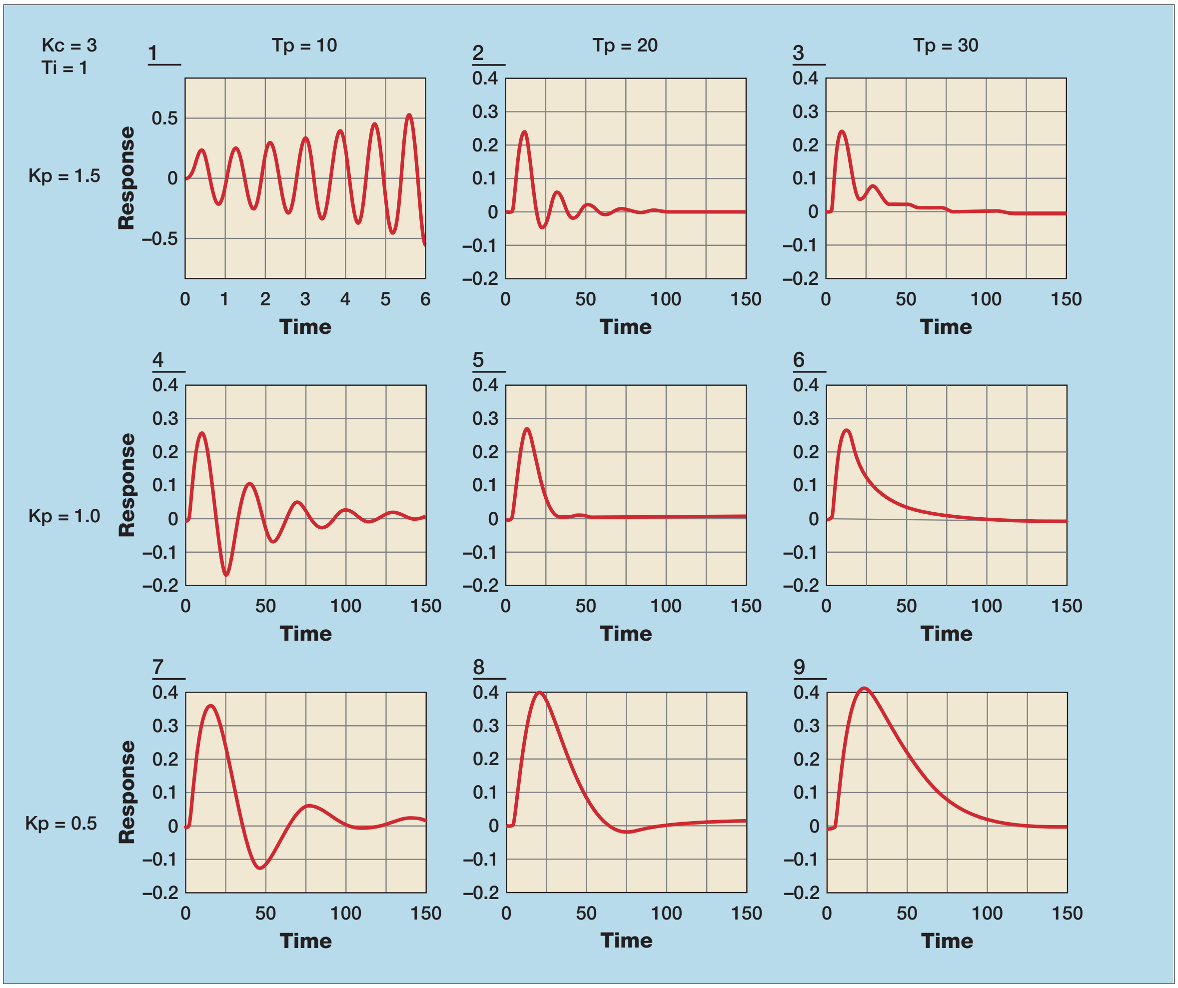

Figure 4 Predicted control responses with fixed controller parameters and variable process responses

robust. A corollary of this is that reducing model volumes might make the model unstable under some conditions!

Transmitter time constants are typically estimated. In fast loops, these time constants may be significant in the overall loop response.

The approach to equilibrium (for example, efficiency factors) can also affect responses. Again, these values are typically estimated and some tuning can be justified.

Metal masses are important for start-up and cool-down, but do not generally have a substantial impact on process dynamics around the normal operating point. Heat losses to ambient should be treated similarly — there is uncertainty in these calculations, but excessive tuning is not warranted because they only affect long term dynamics and has negligible impact on the model mismatch at normal operating conditions.

Tuning should always be sensible. Changing a value by 10% is almost always okay. Changing a good estimate by 50% might be acceptable sometimes. Changing estimates by orders of magnitudes suggests either very poor engineering in the model construction phase or that the engineer is tuning the wrong input altogether.

Conclusions

Model building is a skilled exercise and should be treated as such. A high-fidelity model requires deep process understanding and experience. It is not an off the shelf item. Modelling engineers must recognise the sources of modelling error and act to limit them. There can be a reluctance to tune predictive models (those built only from data inputs), but some selective tuning, if done with sound engineering

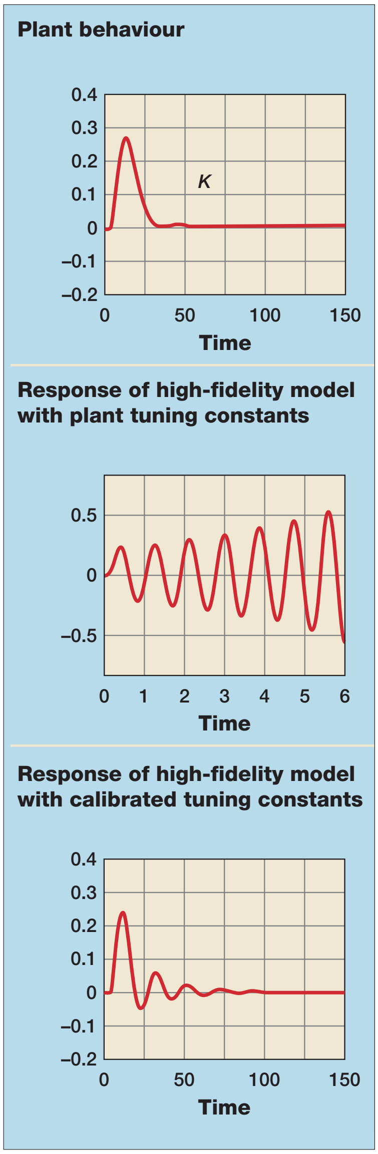

Figure 5 Potential benefits of retuning controllers to work with a high-fidelity simulation instead of the actual plant dynamics

judgement (that is, choosing the right things to tune), can be beneficial. Engineering is reliant on theory, but the best engineers are also pragmatic and results oriented, and should recognise occasions where empirical evidence (data) trumps the accepted theory.

The overall system behaviour should be considered as the highest priority, and is certainly more important than any specific data set. Plant tuning constants will often need adjustment to work well with even the most accurate simulation models. Failure to allow for controller tuning on the simulator system may be detrimental to the overall project's success.

Prognostications

High-fidelity process simulation models are already very accurate, if built well with the right data and good engineering assumptions, and offer a lot of value, particularly for OTSs. However, not everything can be modelled perfectly. Modelling tools and principles can still be improved. There is still room for further development to reduce the granularity of models (such as smaller volumes, smaller time-steps, more detailed unit operation models), although there are diminishing returns.

The increase in fidelity of DCS emulation that is achieved by using 'real' hardware and control configurations is laudable, but it should be remembered that a control system is complementary to a plant or simulation. A low-fidelity process model with a complementary (well tuned) control layer might provide better overall responses that a high-fidelity model with a non-complementary control layer imported from the operating plant. Use the high-fidelity control emulation, but be tolerant of the process model and accept that some controller tuning is likely to be beneficial.

OTSs will continue to become more accurate, but this will not be achieved through software developments alone. The more detailed the simulation, the more data and understanding that it holds, and, consequently, the more engineering effort required. As with any engineering activity, the time and effort should be assessed against the value. The best OTS will not be the best choice for everyone.

Green-field projects share many of the same characteristics as OTS projects built for existing plants, but they tend to be less technically challenging and more schedule driven. This is because there is not the same wealth of detailed knowledge about the operating plant and assumptions and estimates become more acceptable. Data are usually more readily available and 'reasonableness' will be judged by less stringent criteria. The project team generally recognises that value is more closely linked to the delivery date (to maximise training time) than it is to accuracy. Indeed, the incremental benefit of a high-fidelity OTS for training is small — much can be achieved with relatively simple, stable and robust models.

Martin Sneesby is an independent dynamic simulation consultant with more than 20 years of experience in process simulation and modelling, including many operator training simulators and detailed engineering studies. *Email: martin.sneesby@apess.co.uk*

LINKS

More articles from the following categories:
Automation & Process Control
Process Modelling & Simulation