Trends in process systems engineering

JOHN PERKINS†

1. Introduction

Process Systems Engineering (PSE) is a discipline which addresses the problems of design, operation, and control of manufacturing facilities for the process industries, that is those industries whose products arise from changes of physical and chemical composition not only from changes in geometry. The process industries represent a large, and historically very successful, sector of manufacturing, including, among others, chemicals, oil and gas, pharmaceuticals, food and minerals.

The discipline of PSE emerged from chemical engineering in the 1960s. The pioneers of the discipline were keen to emphasize a role for engineers investigating generic problems in the design and operation of plant. In this way, it was hoped that the quality of engineering education would be improved, as well as the discovery of new, more effective methods to address these important problems.

The fifteenth anniversary of *Modeling, Identification and Control* provides a good opportunity to look back over the first thirty years of PSE, to highlight key developments, and to try to identify some current and future trends. Inevitably, the following represents my personal and biased view. I am flattered to have been asked by Jens G. Balchen to prepare it, and to be included with such distinguished company in this very special issue of *MIC*.

2. Process design

The design of processing facilities has always been recognized as a challenging and exciting problem for process engineers. A process design problem is an essential component of the formation of process engineers. In fact, it is one of the few elements required by the professional institutions before they recognize university degrees as providing suitable initial training for process engineers.

As in all design problems, process design involves the conception and analysis of alternatives in order to devise a suitable plant for the task at hand. Some features of the process design problem that make it especially challenging are the large number of alternatives that are available (estimated as between 10^4 and 10^9 for a typical large-scale continuous plant) and the complexity of the components that make up the process. The very earliest work in process systems engineering on generic process design tools focused on the latter aspect. Recognizing that plants are designed by

combining a few general types of operations (the *unit operations* of reaction, distillation, heat exchange etc.), and the emergence of the digital computer as a potentially powerful device for handling routine calculations effectively, development began in the early 1960s on large software systems called *flowsheeting packages*. Development has continued ever since so that today these packages are sophisticated tools routinely used by process engineers throughout industry. The packages provide support to the analysis of design alternatives by allowing mathematical models suitable for simulation to be built up from component models stored in a model library. Once assembled, a variety of calculations can readily be performed on the model and the results presented to the user.

The earliest packages had only a limited library of models and supported prediction of the steady-state behaviour of continuously-operated plants. Developments to enhance calculational flexibility within a steady-state framework came quite early, as did evolution of the model libraries. More recently, flowsheeting packages which also support dynamic modeling and simulations have been developed and are finding use in a variety of industrial organizations.

Of course, mathematical modeling and simulation underpin much of the work of process engineers, not merely the design activity. Thus, as these packages have matured, they have found use in the support of operations as well as design. We shall return to the subject of modeling, after all one of the themes of this journal, later.

The other aspect of process design which was highlighted above is the large number of possible alternatives available to the design engineer tackling a typical process design problem. The construction of overall processes from component building blocks was christened *process synthesis* in the 1960s. Since that time, work on this aspect of process design has proceeded in parallel with the support of process analysis described above. The combinatorics associated with these problems means that they are very challenging, and a large variety of different approaches has been tried in an effort to provide the process engineer with some assistance. The techniques employed have ranged from attempts to formulate simple guidelines (heuristics) which capture good design practice, to formulation of process design as a combinatorial optimization problem.

Each of these approaches has their difficulties. With heuristics, the problem is to formulate suitably general and yet useful rules, and to reconcile the inevitable contradictions between different heuristics for the same problem. To address the latter difficulty, a variety of methods has been proposed, the most elaborate being the building of a number of increasingly sophisticated expert systems for process design problems. But it is not obvious how the basic difficulty that for every heuristic there is a counter-example can be overcome.

With the optimization approach, there are two kinds of problem. First, it is hard to convince experienced design engineers that what they do can be captured purely in algebra. Second, it is clear that even idealized representations of process design lead to very demanding computational problems. Of course, the computer industry is working very hard to overcome the second limitation, and advances in computing technology have fuelled large improvements in numerical techniques to tackle large combinatorial optimization problems.

As to the first difficulty, it is clear that process design problems are rather tricky to formulate mathematically in a convincing way. The traditional mathematical formulation considers choosing a structure for the plant from a 'superstructure' containing all the options of interest based on an economic objective evaluated using steady-state

analysis. Now, process design represents the first step in a much longer project culminating in the building and commissioning of the plant in preparation for full-scale operation. At some stage in this project, all significant aspects of design and operation will need to be considered. The appropriate aspects to consider at any particular stage, and in the process design stage in particular, are not straightforward to identify. A number of studies have shown that consideration of operational aspects (safety, environmental impact, flexibility, controllability, availability etc.) during the early stages of process design leads to different conclusions as to the best design to carry forward to later stages of the project. It seems that, to do an effective job, the operational requirements for the plant need to be incorporated into the early stages of design where the plant structure is being chosen. The implication of this conclusion is that ideally the choice of structure should be based on the dynamic characteristics of the plant, which in turn are affected by the particular control structure to be implemented. This last observation is particularly relevant to the design of multipurpose plants, of services plants, and, more generally, of plants which are required to be responsive to changes in demand. It is likely that this last attribute will be increasingly expected of plant in the process industries in the future (see below).

A further complication arises from the need to cope with uncertainties in the design data and operational requirements in the mathematical formulation. Again, case studies have highlighted the issues. The best deterministic design is often unable to deal with variability in an effective way.

Summing up, it seems that in order to be an effective approach to process design, it will be necessary to enhance the optimization models to include uncertainties explicitly, to include the choice of control system as well as process structure, and in many cases to utilize dynamic as well as the traditional steady-state analysis in evaluating alternatives. Whilst initial attacks have been made on many of these aspects, there is still a lot more work to be done to make the optimization approach a convincing tool for general process design problems. Increased computing power has greatly enhanced our ability to handle realistic process design issues, and to explore the various aspects of the problem. While computing power is a necessary condition for progress using an optimization approach, we are also in need of more refined formulations of the criteria to be used and of the options available. Recent advances in industrial process design hint at the need to move away from traditional unit operations as basic building blocks. As well as increasing the scope for innovative process designs, this may also lead to efficiencies in the representation of the options available.

Recently, interest has grown in the design of manufacturing facilities in the context of complete supply chains. As a result, some of the assumptions underlying the traditional approach to process design have come into question. If this analysis turns out to be correct, and there is significant scope to design and operate more responsive process plant, there will be a need to develop a new generation of process design tools.

There have been great strides in process design in the last three decades. Computer-based methods for the analysis of design options are now in routine use, and have in themselves enabled the consideration of a wider range of design options. The emergence of tools for dynamic process modeling has permitted the analysis of aspects of dynamics and control during design, enabling major structural decisions about the control system to be made alongside equipment decisions where desired. In process synthesis, tools have been developed which are now being used to good effect by industry. We should be optimistic about the likelihood of similar progress in the decades to come. Certainly, there remains a lot of work to do!

3. Process operation and control

Manufacturing in the process industries covers a broad spectrum—from the large-scale, continuously operated plants of the petroleum refining, petrochemicals and minerals industries, through large dedicated batch plants in sectors of the chemical and food industries to multipurpose batch facilities in many sectors. Each sector and each plant type brings its own special problems, but there are a number of general issues which I wish to discuss.

The problem of *measurement* of key variables, especially those associated with product quality, is a common one. Whether it is the chemical composition of a complex mixture of hydrocarbons, and the relationship of this composition to performance as a fuel oil, or the taste of a food product, the analysis of product quality remains problematic. A determination will usually involve some sort of off-line test, be it a laboratory analysis, or the opinions of a taste panel. To get around the difficulty of on-line measurement the usual approach that has been adopted is to control more readily measurable process conditions, and to seek by maintaining those at desired values to ensure consistent product quality. In some cases, soft sensors can be, and have been, used to provide more direct control of product quality.

Given the need to develop improved understanding of the often complex relationships between operating conditions and product quality, there is a long tradition in the use of statistical methods in process operations as an aid to process improvement and optimization. The development of empirical models has been given extra impetus through the exploitation of neural network techniques (sometimes called 'curve fitting for adults') which appear to provide a convenient way of fitting non-linear models to plant data.

For *regulatory control* of key variables the workhorse of process control has been the three-term controller. In general, regulatory control in the process industries is a multivariable problem, and it has taken a fair amount of ingenuity to deliver acceptable performance (although in the past, performance requirements have probably been quite modest) using sets of single-loop controllers driven by simple algorithms. The structuring of the control system has been a key issue, seeking to ensure that any interaction between loops is not detrimental to overall performance.

Contrary to popular belief, multivariable control techniques have been in use for decades in the process industries, especially in the large-scale plants in oil and petrochemicals where the economic benefits of good control are largest. However, the need in many cases to remain close to process constraints in order to chase optimal economic performance meant that a lot of engineering was necessary to make traditional multivariable techniques, which took little account of process constraints, work. In the late 1970s, techniques were developed in the USA (Dynamic Matrix Control) and in Europe (IDCOM) which take explicit account of process constraints, and are thus better suited to the needs of the process industries. As a result of these techniques, the use of multivariable control by industry has expanded greatly, and the quality of regulatory control has improved quite significantly. Indeed, it has been argued recently that performance is now as good as it need be, and that therefore further research on regulatory control for the process industries can only be justified in terms of reducing engineering effort to deliver a working controller, not in terms of improvements in control quality!

The new methods discussed above have been given the generic title of **model-predictive control,** which is somewhat confusing given that all control is based on models, and that the particular models used in commercial systems fall a long way short

of the physically-based models used in other aspects of PSE. It is a truism that all processes are non-linear; what is surprising at a superficial level is the ability of control techniques based on linear models to deliver adequate performance. However, it should be recognised that the main interest is in regulation about a fixed operating point, or sometimes over a small range. For wider variations, it appears still to be necessary to provide a fair amount of 'jacketing' for the methods. Whilst techniques to allow the incorporation of non-linear, physically-based models into regulatory control algorithms have been proposed and tested on pilot-scale plant, currently their use in industry is not widespread.

Above the regulatory control level, there is a need to determine set points for the controllers. Traditionally, these were manually set by operations personnel based on process experience or off-line modelling studies. The use of process models on-line to determine optimal performance is increasing. At present, industrial applications are most common in large-scale scontinuous plants based on steady-state process models. This approach has its drawbacks, notably in the need to monitor the plant data used for model updating to try to ensure that the data set corresponds to a steady-state operating condition. As a result, the performance of the plant under transient conditions will not be improved by the optimisation, and large amounts of information about the current state of the plant are discarded. To overcome these shortcomings, it seems desirable to employ a dynamic framework for on-line optimisation. Two issues arise. We should be able to rely on continuing rapid development in the performance of computer hardware to address the need for increased computing power to handle the dynamic optimisation. However, more reliable dynamic models of process plant will need to be developed. This development would be greatly aided by provision of more advanced tools for model building (see below).

The impact of process systems techniques on the operation of plant during upset conditions has been less great than it might have been. During start-ups, shutdowns and mode changes, it is still very common to have the plant revert to manual operation. Control algorithms often require protection ('jacketing') from these 'abnormal' conditions. Even the vocabulary used to describe these events suggests that we would be better off without them! While this may be so, these upsets are a fact of life, and it is arguable that even more of plant operation will be spent in significant transients in the future. Again, this issue would be partially addressed by the use of dynamic models in on-line optimization. However, for more involved operations (e.g. startup) it is necessary to take proper account of discrete and discontinuous actions. Optimization of such cases requires a 'multiprocess' formulation in general. To support investigation of these cases, there is a need for more general modeling tools. On the optimization side, there has been recent progress both in theory and algorithms to handle optimal multiprocess control. It is to be expected that these techniques will become more widely used in future, as we become more ambitious about the class of operations we wish to study.

For batch operations, all operating decisions above the regulatory control level require consideration of discrete decisions and dynamics. There have been significant advances recently in the application of integer programming techniques to the solution of complex multipurpose plant scheduling problems. However, the on-line use of these techniques is still some way off, requiring significant increases in computing power combined with efficiency improvements in the algorithms themselves. In the mean time, the problem of reactive scheduling to cope with unforeseen changes is being tackled by more heuristic methods.

As in the case of process design, there is a need for more systematic handling of uncertainties for many of these operational problems. At the regulatory control level, the issue of robustness to uncertainties has been one of the main concerns of the control community in the past decade or so. New design methods based on H^∞ techniques and their generalizations have been developed, and used to good effect in control of mechanical systems where the dynamics are well understood and uncertainties can be precisely characterized from manufacturing tolerances on components. Unfortunately, quantifying uncertainties in process dynamics is not quite so straightforward. Apart from the uncertainties associated with actuators and sensors, where some idea of the likely errors is available, very often all that is known about uncertainties is that they exist and are probably quite large! Nevertheless, there is a need for modeling tools to allow uncertainties to be represented, and for techniques to take account of these uncertainties more systematically in addressing operational problems.

In summary, there have been significant developments in techniques to handle process operations problems in the past three decades. Most of these are concerned with automation of 'normal' operation. There is a significant scope to expand application of process systems methods to handle large-scale dynamic effects more systematically. Also, as with process design, the issue of uncertainties is a key area for the future.

4. Process modeling

Underpinning the techniques discussed above is a requirement for some mathematical representation of the problem. The effort required to build such mathematical representations has been characterized as a bottleneck restricting the more routine application of process systems engineering techniques.

To support the generation of process models, increasingly sophisticated software tools have been developed. The calculational facilities provided in today's packages encompass steady-state and dynamic simulation, as well as optimization and parameter estimation from plant data. Comprehensive libraries of models are available, backed up by facilities to provide the physical properties necessary to predict plant behaviour. As a result, building models for many processes may be achieved by assembling standard elements into an overall process model.

In spite of the undoubted power of the packages available today, there are still areas where further work is justified in order to improve the efficiency with which process models can be built. First, existing packages provide little support for the development of 'bespoke' models for processes not covered by the model library. While the latest packages reduce the task of model building to assembling the appropriate set of describing equations, all aspects of calculation being handled by the package, it is not clear that equations are really the most natural language for engineers to use to describe their processes. One possibility might be to permit description of processes to be modeled in terms of physical phenomena occurring and the assumptions to be made about these phenomena. A package supporting this approach to model building would handle the construction of the appropriate equation set, as well as the calculations. Another approach might be to consider more elementary building blocks than unit operations in the library of models associated with the package. Work is under way exploring each of these approaches. If successful, such work would go a long way towards alleviating the bottleneck mentioned earlier.

Another area where enhancements to existing techniques would be helpful arises from the observation that in many cases modeling the operating procedure for a plant is

as complex as modeling the process behaviour itself. The latest packages provide quite comprehensive facilities for the representation of process behaviour, but their facilities for modeling operating procedures are more rudimentary. Again, work is under way on a new generation of process modeling package which provides comparable levels of support for process and procedure modeling.

Finally, a recurring theme in this paper has been a more systematic treatment of the uncertainties that inevitably arise in engineering problems. A first step in a more systematic treatment would be to provide facilities for the modeling of uncertainties in a given process model, together with methods to perform calculations taking account of these uncertainties. Whilst there is work under way on different ways of taking account of uncertainties in many aspects of PSE, I am not aware of any work specifically to provide a comprehensive modeling environment incorporating uncertainty representation. I believe this to be an important challenge for the PSE community in the future.

5. Conclusion

The discipline of process systems engineering is about three decades old. In its comparatively short existence, important problems have been identified and addressed, and useful problem solving ideas and tools generated. Continually increasing computing power is opening up new horizons in terms of the complexity of problems that can be tackled. Fortunately, the range of unsolved problems is large enough to keep us all busy for the foreseeable future!