AN INTEGRATED APPROACH FOR DISTILLATION COLUMN CONTROL DESIGN USING STEADY STATE AND DYNAMIC SIMULATION

Donald P. Mahoney
Hyprotech, Inc.
501 Silverside Road
Wilmington, DE 19809

Paul S. Fruehauf
DuPont
P.O. Box 6090
Newark, DE 19714-6090

ABSTRACT

Steady state techniques have been used for decades to develop control strategies for distillation columns. While theses techniques are effective for screening out clearly undesirable control structures and suggesting viable candidates, they provide incomplete and sometimes misleading information. To accurately assess the performance and suitability of alternative control schemes, rigorous dynamic simulation is required. This paper presents an integrated distillation column control design methodology that involves both steady state and dynamic simulation. The limitations of steady state techniques are discussed, and the need for rigorous dynamic simulation for final selection of a workable and robust strategy is illustrated. An integrated simulation environment that encourages the proposed design methodology is also described.

KEYWORDS

Distillation Control, Steady State Modeling, Dynamic Modeling, Computer Simulation, Multiple Steady States, Distillate-Bottoms Control

INTRODUCTION

Steady state techniques are used extensively in the development of distillation column control strategies. For complex, multivariable control systems, the Relative Gain Array (Bristol, 1966) has become quite a popular steady state control analysis technique. Other methodologies that involve steady state sensitivity analysis for control purposes have been proposed by Tolliver and McCune [8], and extended by Fruehauf and Mahoney [3]. While all of these techniques are useful for screening out unattractive control structures, the results they provide about the remaining alternatives are often incomplete and sometimes misleading. This, along with the fact that a large number of industrial columns still operate in manual or with ineffectual controls, illustrates that there is a need for improved distillation control design techniques.

Rigorous dynamic simulation is clearly a more accurate way of evaluating process response and the dynamic performance of various control structures. However, it is not the most efficient way of sifting through the often numerous possible control candidates. By making use of its solving efficiency, we can employ a steady state technique to screen out unattractive control structures and suggest viable candidates. Then, using rigorous dynamic simulation, we can discriminate among the smaller number of remaining alternatives.

In this paper we present an integrated distillation control design procedure that involves both steady state and dynamic simulation. The steady state design methodology presented is useful for suggesting viable distillation control candidates and screening out clearly unworkable schemes. While this methodology is an effective control design technique and has been applied to many industrial operations, when used alone it has a number of limitations — as do all steady state control techniques. We will examine some of these limitations and illustrate how rigorous dynamic simulation can be used to complete the analysis by revealing important operability and control performance information. Having presented the integrated design procedure, we propose criteria for a suitable simulation environment.

STEADY STATE PROCEDURE

The initial steady state screening procedure is similar to a design technique proposed by Tolliver and McCune [8]. However, there are a number of significant differences in the methodologies that warrant some discussion. First, we propose the use of mass flows for model specifications as opposed to the “typical” molar flows. As we will illustrate in a later example, there can be significant differences in the results when molar flows are used. Furthermore, most industrial columns measure and control mass or mass equivalent flows, not molar flows.

Second, when examining temperature sensitivity for composition control, we impose the actual control structure via appropriate selection of steady state specifications for the model. For example, if we are proposing the use of mass reflux to control composition, we would explicitly set mass reflux in the model specifications. Most techniques simply vary molar distillate in order to gauge temperature sensitivity regardless of the proposed composition control variable.

Third, this technique may be used for multicomponent systems to quantify the incremental benefit of using on-line analyzers over temperature control.

This technique deals exclusively with the design of single-point composition controls. By “single-point” we mean that the composition of only one end of the column is controlled directly. Dual-point control involves controlling top *and* bottom compositions. The main benefit of dual-point control is energy savings, however, since these savings are often too small to justify the added complexity of the design, single-point control seems to predominate the industry.

Having introduced the basis for the steady state design and screening procedure, we present the detailed methodology below.

Step 1 - Developing the design basis.

As with any design effort, it is important to begin by establishing all of the important criteria that the final design must satisfy. This includes, but may not be limited to identifying and understanding:

- what the product draw composition specs are

- whether the specifications are one- or two-sided[1]
- which stream is the demand stream[2]
- what the expected disturbances to the column are
- what the operating constraints are
- what the base case or "normal" operating condition for the column is.

It is particularly important to understand the nature of the disturbances that are likely to upset the column. Accurate predictions of feed rate and feed composition disturbances are a key element to eventually developing a robust and workable control structure. It is perhaps worth mentioning here as well, that if the design is for an upgrade or expansion to an existing process, it is important to understand the existing control structure and why it is implemented the way it is. Reasons for a particular control structure may be very subtle, yet critically important to the plant-wide operability. On the other hand, many controls are "left-overs" from old designs and have not been changed simply because no one has bothered to improve them.

Step 2 - Selecting candidate control schemes

The second step in the steady state procedure is to select a candidate control structure. Two-product distillation is typically viewed as a 5X5 control problem. There are five degrees of freedom in a typical two-product distillation column[3], represented by:

1. feed valve
2. reflux valve
3. distillate valve
4. heat input valve
5. bottoms valve.

We do not consider pressure controls here since we are normally able to achieve tight pressure control via inerts blanketing and venting, or with low-boiler venting. Under such conditions, pressure may be considered fixed, and thus does not affect the degrees of freedom available for the remaining controls.

In a 5X5 system, there are 5! or 120 possible single-input, single-output control combinations. However, once all of the constraints of the process are considered, normally only a few combinations are left.

First, we must determine which of the streams is the "demand" stream: the one which is set by some upstream or downstream process and thereby sets the production rate for the column. Typically the feed is the demand stream, however, occasionally we see the distillate or the bottoms being set independently.

Next we examine the overhead and bottoms inventory controls in light of the base case, or normal operating conditions. We compare the relative magnitude of the reflux vs. distillate, and bottoms vs. boilup. If there is a difference of 10:1 or more, we typically select the larger of the streams to control level. An example where this ratio applies is a tar still. Here we are typically trying to remove a small quantity of high boiler. It is not uncommon for the boilup to be 100 times the bottoms flow. In this case, the bottoms flow is too small to compensate for many disturbances, thus the boilup must be used to control level.

The next step is to consider the economics and constraints of the system and select a suitable feed-split control scheme for composition control. If the feed is the demand stream, and we do not have a tar still, the structure often reduces to one of the two schemes shown in Figures 1 and 2.

In Figure 1 we have what is called a direct feed-split scheme. In this case, distillate is manipulated directly to control the composition profile. This structure is often used when heat input to a column is limited or must be fixed.

In Figure 2 we have an indirect feed-split control scheme. Here, the distillate is changed indirectly through the heat input-composition controller. If the control temperature is too low, heat input flow is increased. This increase throws more vapor overhead and results in an increase in distillate as the condensate drum level increases. There are two basic advantages to this configuration. First, the temperature-boilup loop has a faster closed-loop response, and provides superior disturbance rejection. The second advantage is that we can make use of the condensate tank inventory to

---

1 One-sided spec's must remain at or below a certain value, two-sided specs must remain within a specified range.

2 The demand stream is set independently by an upstream or downstream process and is therefore not available for control purposes.

3 Because steady state calculations do not take into account the three inventory variables of condensate level, bottoms level, and pressure, steady state models for binary distillation have only two degrees of freedom. For real control purposes, however, we clearly have five.

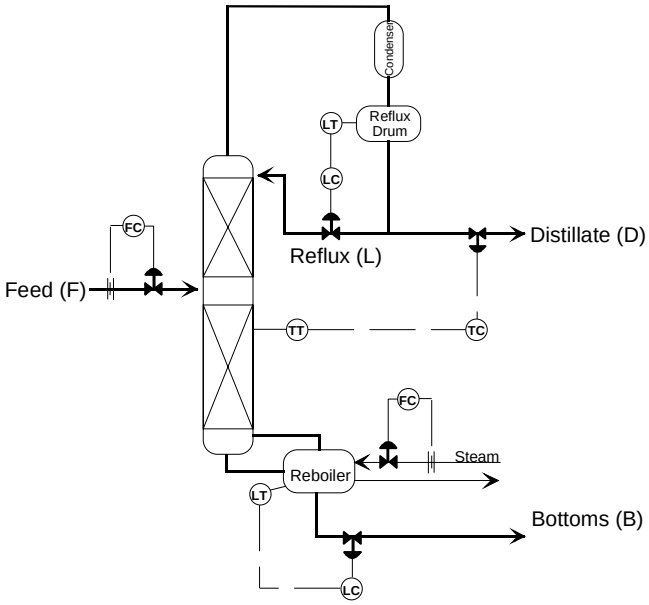

Figure 1 - Direct Feed-Split Control

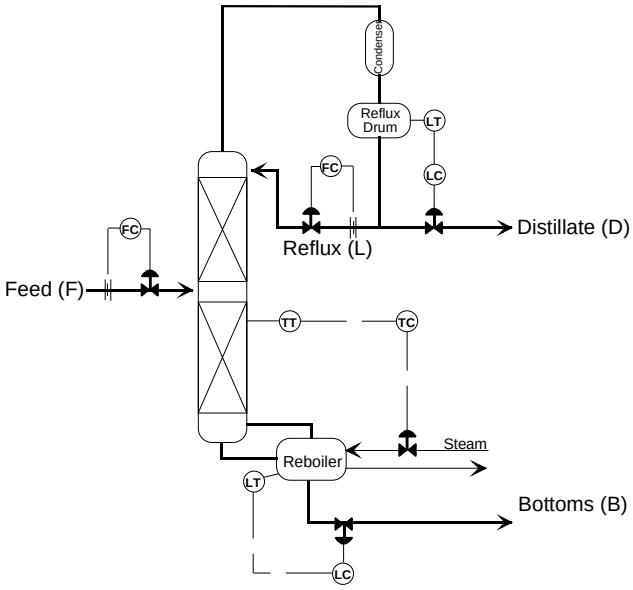

Figure 2 - Indirect Feed-Split Control

achieve substantial flow smoothing in the distillate to the benefit of downstream processes[4].

Finally, we consider a ratio alternative that might reduce utility costs or improve the dynamic response. One alternative for the scheme shown in Figure 2 would be a controller that keeps the ratio between the feed and the reflux streams constant. This scheme is likely to use less energy since for smaller feed rates, we have less reflux, and thus less heat input is required.

While we present a procedure here for two-product distillation, this technique may easily be extended to multiple-draw, multiple-feed columns. If we use a partial condenser with a vent stream, we have and additional degree of freedom, represented by the vent valve. If this valve is used for pressure control, then the analysis is the same. If we have an additional draw not used for pressure control, then we simply have one more degree of freedom. In that case, we may chose to try to manipulate this valve in order to control another variable, or we may chose simply to ratio this stream to another key stream. Recognize, however, that as we increase the number of manipulated variables, we tend to increase the degree of interaction between controls. This makes the column more difficult to control and requires a longer time to recover from disturbances. Ratioing additional draws to other key streams avoids such interaction and is often adequate.

Step 3 - Open loop testing

The third step in the process involves evaluating the open loop sensitivity of temperature (composition) to the candidate composition control variable. The goal here is to identify a sensitive region in the column for temperature measurement. By holding all other inputs constant and running a number of case studies with different values of the manipulated variable, we can generate a family of temperature profiles around the base case. Typically, changes in the manipulated variable of ±1%, ±2%, ±5%, and ± 10% are sufficient.

Examining the curves, we look for sensor locations where temperature changes are *significant*, and *linear*. We can often control temperatures accurately to within ±0.5°C. Thus a temperature change of 1°C, while not ideal, is often significant enough. By "linear", we mean that temperature changes are roughly equal in magnitude when the manipulated variable is changed by the same amount in both directions around the base case.

As mentioned earlier, we advocate using mass flows when examining the temperature sensitivity. In practice, we control mass flows, volumetric flows, or flows measured by a pressure differential. The latter two are essentially the same as mass flow, and are different from molar flows if molecular weight varies.

[4] To achieve the best flow smoothing, the condensate tank level control should not be tightly tuned. Proportional-only or "averaging" level control tuning is preferred.

Figure 3 shows a group of temperature profiles for an industrial column currently in operation. Using the criteria described above, we might select tray 38 for temperature control. Tray 10 would be a poor location as there is no sensitivity to negative changes in the manipulated variable.

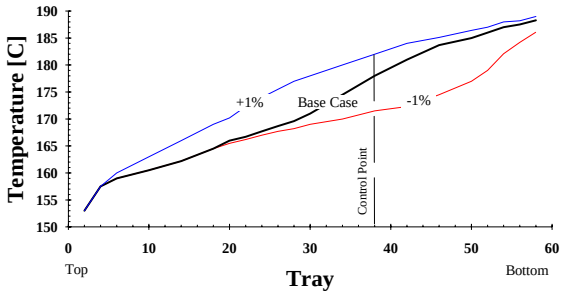

Figure 3 - Temperature Sensitivity

If we had additional draws, we could evaluate the sensitivity of temperature to changes in the draw as well. We may find that controlling temperature with the side draw is better than using any of the other available streams. As mentioned earlier, we may choose to try to control an additional variable with the extra draw, or simply ratio the flow to another key stream. If we want to control an additional variable (another temperature for example) we would need to examine the sensitivity of that variable to changes in the draw. This kind of analysis should reveal whether the extra draw represents an attractive control candidate, or whether it is best to simply ratio the stream to another. Recall again, that the more variables we try to manipulate independently, the more interactions we have and the more difficult the column will be to control in practice. Dynamic simulations in this case are particularly useful.

An important issue to consider when selecting a temperature sensor location is its proximity to the end of the column with the more important purity specification. Normally we prefer to measure temperature as close to the composition of interest as possible. While there are a number of techniques in the literature for selecting an appropriate temperature sensor location [1], we believe that further research in this area is needed. As we will illustrate later, dynamic simulation is particularly helpful for identifying a suitable temperature sensor location.

A practical implementation detail is in order here. Once a control scheme has been designed, tested, and selected, we recommend specifying two additional temperature sensors; one theoretical stage above, and one theoretical stage below the primary nozzle. This accounts for any inaccuracies that may exist in the model, and is far cheaper to specify at this point than after start-up.

Step 4 - Closed loop testing

Once the candidate control structure has been defined and tested for open loop sensitivity, we are ready to perform closed loop testing by subjecting the model to the expected feed rate and composition disturbances. Before we do so, however, we must determine the operating conditions that maximize the demand for the fixed flow or flow-ratio variable. To use the structure of Figure 2 as an example, we must determine the value of fixed reflux that is required to maintain the purity specifications for both ends of the column when the most severe expected feed conditions are encountered. In the case of a ratio scheme, we must determine the ratio required to satisfy the specifications at the most extreme feed conditions. This may be an iterative process that involves making changes to the feed rate and feed composition and observing which conditions place the highest demands on the column when the composition specs are met or exceeded. In many instances, we may use our knowledge of distillation to determine these conditions.

Once this is determined (the maximum demand reflux rate in the case of Figure 2), we must note the temperature at the selected control location as this will be used as our setpoint. At this point, the scheme is fully defined and we are ready to test the structure on the full range of operating conditions. It is important here that the actual control structure be enforced on the model by careful selection of the steady state specifications. For example, the control structure shown in Figure 2 would require a temperature specification for the selected tray, and mass reflux specification equal to the mass reflux rate equal to the maximum load rate. Here, the mass reflux will be fixed, and the value of heat input will be manipulated by the steady state solver in order to achieve the temperature setpoint. In the case of a ratio scheme, the desired ratio may be maintained with the use of an “adjust” or “set” operation. The same specifications would apply.

To fully test the candidate structures, we recommend testing the model with all combinations of low, mid, and high values for both feed and composition. Figure 4 shows the 9 cases that would

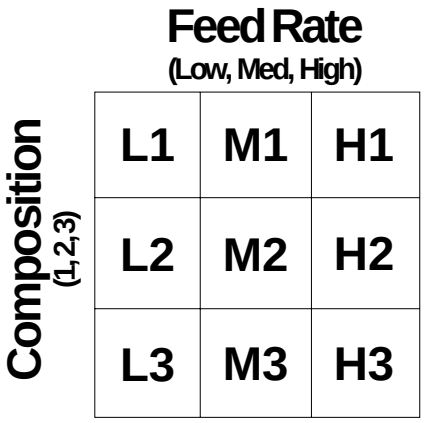

Figure 4 - Nine Cases for Binary Distillation

be required for a simple binary distillation column. For multiple components, more cases are recommended.

While we do not explicitly call it out here, there is an implied Step-5 in this procedure that involves iterating back to any of the earlier steps 2 through 4 in the event that the closed loop testing is not satisfactory. What we should have after successfully completing Step 4, is one or more control structures that appear to be viable candidates. The procedure to this point will likely have screened out those schemes that are clearly unworkable, leaving only those alternatives that satisfy the control requirements to the extent that steady state modeling is able. As was mentioned earlier, the design methodology presented up to this point has been successfully applied to numerous industrial columns. It does, however, provide incomplete information on much of the dynamic operability of the candidate schemes, and has several other limitations which will be discussed in the next section.

LIMITATIONS OF STEADY STATE TECHNIQUES

There are a number of limitations associated with using a steady state modeling approach for the design of an inherently *dynamic* process. As we have shown, aspects such as sensitivity and steady state response to upsets may be revealed through steady state design techniques. Attractive control options can be identified and clearly unworkable structures eliminated. However, these studies reveal very little about the dynamic operability of the control structures being considered; or the process itself for that matter. Further, it is often difficult to discriminate among the viable control alternatives using steady state techniques alone.

Two very significant shortcomings of steady state techniques are that 1) the effect of holdups in the system are not considered, and 2) the high-frequency, or initial response of the system is not considered. These two issues account for many of the problems, and much of the misleading information that steady state control design techniques can produce. To illustrate the concepts, two very interesting cases will be examined.

Ignoring holdup effects - the Distillate-Bottoms control structure

If we were to devise a control strategy aimed at controlling both the top and bottoms compositions using the distillate and bottoms streams respectively, we might end up with a control structure similar to that shown in Figure 5. While this is clearly a configuration that we can set up physically, is it possible to actually control the process using this scheme?

From a purely steady state standpoint, the answer would be no. At steady state, the relationship between the feed rate, and the distillate and bottoms flows are not independent. Here, the relationship F = D + B must hold, and thus L and D may not be varied

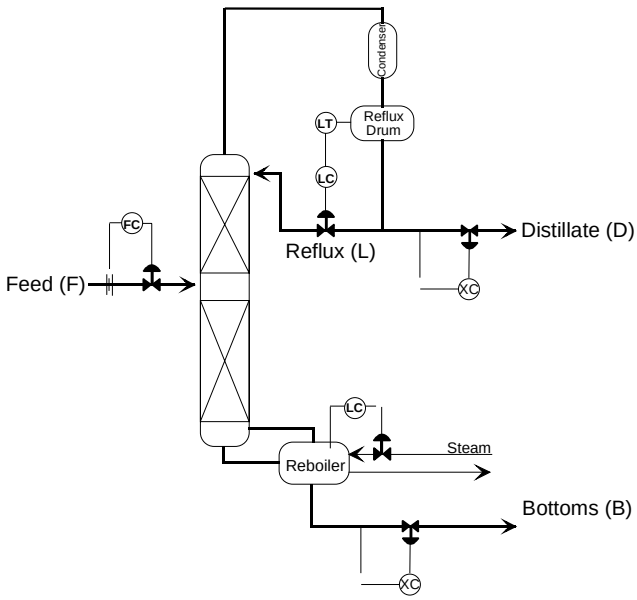

Figure 5 - Distillate-Bottoms Control

independently. This relationship is so fundamental to the steady state description of distillation that, until only recently, control experts considered this strategy "impossible" and "unworkable." Those familiar with the Relative Gain Array, or RGA technique for control design will recognize that the gain matrix using the Distillate-Bottoms strategy is singular; thus "proving" that such a control strategy is not possible. However in 1989, Finco et al. [2] demonstrated for the first time that the Distillate-Bottoms structure is indeed feasible. Papastathopoulou et al. [6] proposed a tuning methodology, and Skogestad et al. [7] have provided an explanation for reason it actually works.

In short, the liquid lag from the top of the column to the bottom due to holdup, effectively decouples the distillate and bottoms responses at high frequency [7]. High frequency, or initial response, is where most of the control action takes place. Thus the Distillate-Bottoms control scheme is not only possible, but it turns out that it is relatively easy to control [7]. The notion that steady state data may provide misleading information for control analysis is somewhat intuitive, yet it is often forgotten when developing multivariable controls.

Steady state techniques are clearly valuable, but must be used with caution and good judgment. The next example illustrates how even steady state sensitivity may be misleading, and that it is important to examine control system initial response using dynamic simulation.

Ignoring high-frequency response - the control of columns exhibiting multiple steady states

As part of the steady state screening procedure, we described the need to use mass flows as inputs to the steady state sensitivity analysis. The reasons being 1) we typically control mass or mass equivalent flows in practice, and 2) we can get significantly different results when molar flows are used.

Since molar flows enter directly into the material balance calculations used in the column solver and thus determine the separation, we are often inclined to work on a molar basis. Further, we often consider mass reflux to increase monotonically with molar reflux, thus making them rough equivalents for control purposes. While this is often true, an interesting example that illustrates where it is not, is the case of distillation columns that exhibit multiple steady state solutions when mass inputs are specified. The development explaining the multiple steady state behavior for binary distillation with fixed boilup is shown in Appendix A.

Figure 6 shows the steady state sensitivity of overhead purity for changes in mass reflux with fixed heat input to the column. Notice how between about 0.001% and 0.35% overhead impurity, we see a reversal in sign of the mass reflux effect. If we used this sensitivity plot to decide how to control overhead impurities in the neighborhood of 0.01%, we might be inclined to set up our controls to decrease mass reflux, in order to increase overhead purity. After all, our steady state analysis shows that decreasing reflux here should reduce overhead impurity.

As it turns out, all of this steady state data has very little impact on the way the column actually controls. Dynamic simulation of such a system reveals that the high frequency response is much different from the low frequency, or steady state response. In fact, the high frequency, initial response is almost the same regardless of whether mass or molar reflux is manipulated [5]. To see what actually happens, it may be helpful to consider a case disturbance.

Consider a column that has "lined out" to some steady operating point. We now make a step increase in the overhead impurity set-point. If our

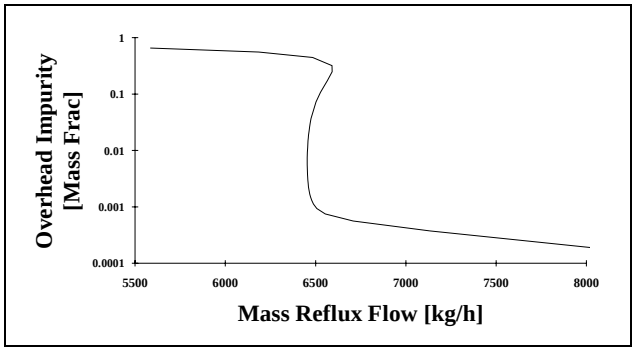

Figure 6 - Overhead Impurity vs. Mass Reflux

control action is to decrease mass reflux, L_W, we will see an initial decrease in molar reflux, L, since the molecular weight is unchanged ($L_W = L\ MW$). This happens quickly and correctly increases overhead impurity by some amount, y_{D1}, regardless of the steady state compositional effects at play. As the increase in overhead impurity increases the molecular weight, the molar reflux actually begins to decrease further. If this second decrease in molar reflux causes a further increase in overhead impurity by another y_{D1} or more, the control action may *increase* mass reflux back above its original value.

Thus as the controller lines out again, we may actually end up with a higher mass reflux with the increased overhead impurity. Whether this occurs or not depends largely on the values of $dy_D / \delta L$ and L (see Appendix A).

Figures 7 and 8 show the simulated[5] response of a mass reflux controller to a step increase in the overhead impurity set-point. Figure 7 shows the "typical" response which we observe in the low impurity region (below 0.001% overhead impurity). Here, as we might intuitively expect, the mass reflux decreases initially, and lines out to a lower value as the higher impurity set-point is achieved. Figure 8 shows the response of the same control structure, however this time, in the region of steady state sign reversal (between 0.001% and 0.35% impurity). The initial response here is the same as that shown in Figure 7. However in this case, we line out to a *higher* value of reflux after the higher impurity is reached. This inverse response clearly illustrates the difference between the initial, high frequency response, and the steady state response. Since control is based extensively on high-frequency behavior, it is important to use steady state information intelligently, and examine the dynamics of the system before selecting the controls.

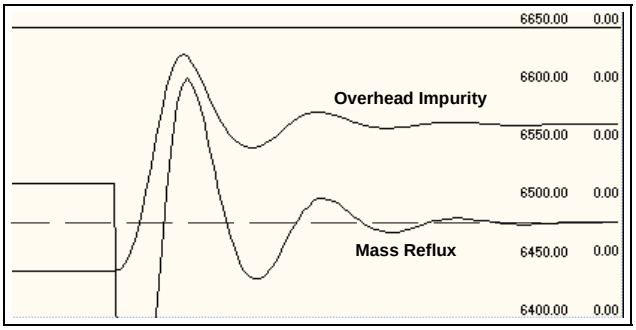

Figure 7 - "Typical" response to step decrease in overhead purity set-point

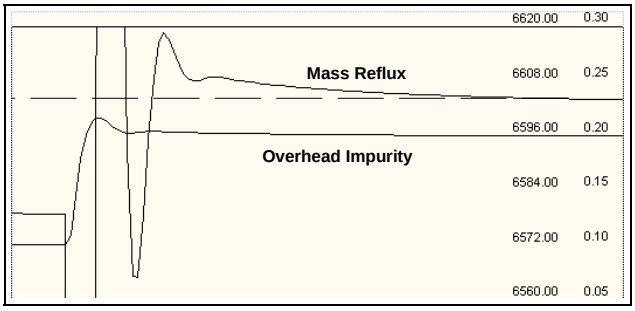

Figure 8 - Inverse response to decrease in overhead purity set-point

[5] These trend charts were taken from a rigorous dynamic simulation of a real process using HYSYS® from Hyprotech, Ltd.

THE DESIGN APPROACH CONTINUED

As these examples have shown, steady state analysis for control, when used alone, can provide incomplete and sometimes misleading information. We have proposed a steady state screening and design procedure which is useful for eliminating undesirable structures, and suggesting viable candidates. Using dynamic simulation to rigorously model the remaining control structures, we have an integrated design approach that is both efficient and accurate.

The dynamic simulation part of the design procedure does not lend itself as much to a detailed step-by-step process as the steady state part does, however, we can trace some general steps.

Step 5 - Supplying holdup information

Steady state simulations do not typically involve holdup information in the solutions. Thus, in the dynamic analysis, we typically begin by assigning the relevant holdup information to condensers, tray sections, reboilers, and any other ancillary equipment that possesses a material holdup. Since holdup volumes play a critical role in defining the system time constants, and thus influence response times, disturbance rejection, and overall controllability, it is important to specify holdups properly.

Step 6 - Identifying the dominant dynamics

Controls and instrumentation are also not normally part of a steady state model, thus we must consider the measurements and actuators that we are likely to have at our disposal when we go to operate the plant. While it is important to impose each candidate control structure on the model in the same way that it is likely to be implemented in the plant, we do not want to introduce unnecessary complexity. For example, we may need to add lags or dead times to process measurements if they are likely to be updated only periodically from lab samples. Dead

times have a very large effect on closed loop dynamics and therefore need to be accurately represented. We may also need to add cascade controls in loops where they will apply in the real plant. These are likely to influence the dynamics to a sufficient extent that they ought to be included in the model. Other pieces of instrumentation and equipment will not significantly affect the dynamics of interest, and thus need not be modeled. The dynamic response of instruments such as valves or on-line sensors, for example, are seldom included due to their negligible effect on the dynamics that dominate the process. The level of detail used in the simulation must be driven by a judicious balance between accuracy and usability.

Step 7 - Applying the control structure

Once we have supplied holdup information, and identified the important additional dynamics, we may apply the candidate control structure to the process model. One of the benefits of having supplied accurate holdup information and significant instrumentation dynamics, is that we may now use one of many tuning techniques[6] to generate preliminary controller tuning constants.

Step 8 - Exercising the model

Once we have the control structure implemented, and the controllers tuned, we are ready to exercise the model. At this point we are able to evaluate a number of different scenarios. Start-up and shut-down performance may be studied. Also, feed rate and composition upsets, as well as other likely load disturbances may be applied to the model to test the dynamic operability and disturbance rejection capabilities of each control scheme under consideration. It is a good practice to design a suite of test disturbances to which all of the candidate control structures will be subjected. This provides a common basis useful when comparing the performance of competing schemes.

Step 9 - Selecting the strategy; and beyond

Steps 7 and 8 may be repeated for each candidate control strategy. While the best strategy may emerge quite clearly in some cases, the differences in control performance are often subtle, and may involve trade-offs between things like ease of start-up and shut-down, disturbance rejection capabilities, response to production rate changes, complexity of the control strategy, the degree of interaction between controls, and the number of controllers required. These issues and others must be considered carefully by the designer, and hopefully by others involved in the subsequent operation of the process.

Once the final control scheme is selected, there are still a number of ways in which dynamic simulation can be extended to add value. By allowing links between the computer process simulation and various DCS platforms, the dynamic model used to define the process controls may be used to check-out the commissioned control strategy resident on the DCS. This exercise not only checks the integrity of the control software configuration, but also allows for preliminary tuning, and operator training. Beyond plant start-up, such a link may also be used as part of an on-line control system as well as for on-going process improvement and optimization studies.

SIMULATION CRITERIA

While the appeal of an integrated steady state and dynamic control design approach is strong, there have been no commercial simulation environments available to encourage such activity. Having examined some of the major issues involved in designing robust and workable control strategies, we can begin to see what a truly useful simulation environment should look like.

Accurate. Clearly a solid engineering foundation is critical. Without accuracy, modeling is not only a waste of time, but may make predictions that could lead to poor design decisions. Creating this solid foundation, however, is more than simply capturing the latest process engineering knowledge and presenting it together in a simulation package. It requires a judicious balance of rigor and performance that yields a tool that is at the same time useful and usable.

Integrated steady state and dynamics. As we have illustrated, the benefits of having both steady state and dynamics functionality available together are clear. While independent steady state and dynamic simulation packages are often positioned as providing “seamless” integration, many actually make

[6] Ziegler-Nichols, the IMC, and the ATV, or Auto-Tune Variation Technique have been found useful.

use of flat file exchanges and do not share a common simulation environment. For an integrated system to be truly useful, it must be able to eliminate the duplication of effort experienced when separate models are required for each mode. Ideally, one builds the model topology once, then executes steady state or dynamic solvers depending on the need.

Fast. Experience has shown that much of the process understanding that comes from simulation occurs during the model-building phase. Interaction with, and immediate feedback from the model are key elements to the effectiveness of any process modeling endeavor. Submitting runs batch-style and waiting for results is not only inconvenient and inefficient, but it removes the valuable "live" link between the engineer and the model.

Broadly Applicable. Ideally, all of the functional requirements for all types of applications should be available in one place. Conceptual design, process design, dynamic operability analysis, control strategy development, DCS interfacing and check-out, operator training, and on-going process improvement for all kinds of processes should share a common environment. This environment, however, needs to be more than a group of functional engineering tools artificially linked through file sharing or swapping. By seamlessly integrating all of the functional capabilities into a single environment, information generated in one mode is fully and immediately available for all others.

Easy to Learn and Use. In order to break the barriers that prevent its wide spread use, process simulation must be both easy to learn and to use. By using an intuitive, graphical user environment, and a comprehensive selection of *configurable* unit operation modules, simulation tools can make tremendous improvements in this area. Literally thousands of man-years of process engineering and modeling experience have been accumulated in the simulation industry. With creative ways of packaging this information available, there are very few reasons why engineers should have to write custom code or compile input files or subroutine calls in order to run simulations. By making simulation technology easy to use and available to all engineers via configurable modules, we are placing the process understanding into the hands of those who are most able to put it to effective use.

Hyprotech, Ltd. is taking a lead in this area of process simulation. Based on the criteria discussed above, Hyprotech has developed an integrated simulation environment called HYSYS® that combines steady state and dynamics functionality in one package, and also provides links to popular DCS platforms for control system check-out, operator training, and the potential for on-line dynamic model-based control. We believe that this technology will have a tremendous impact on not only the way engineers approach control strategy development, but on how we approach modeling in general.

SUMMARY

We have presented a proven steady state screening and design technique for distillation column controls. We have also highlighted some of the limitations and weaknesses of such methods when used alone. Dynamic simulation completes the analysis by providing the necessary high-frequency or initial response information that allows for proper development, evaluation, and selection of candidate control structures. Having presented an integrated design solution, we saw the need for new simulation environments that encourage such a design methodology. The features of such simulation tools include: accuracy, integration of steady state and dynamics functionality, fast execution, broad applicability, and ease of use.

REFERENCES

1. Buckley, P. S., Luyben, W. L., Shunta, J. P., *Design of Distillation Column Control Systems*. ISA, Research Triangle Park, NC, 1985
2. Finco, M. V., Luyben, W. l., Polleck, R. E., "Control of Distillation Columns With Low Relative Volatilities". *Ind Eng Chem Res*, V28, n1, Jan 1989, pp. 75-83.
3. Fruehauf, P. S., Mahoney, D. P., "Improve Distillation Column Control Design". *Chemical Engineering Progress*, March, 1994.

4. Fruehauf, P. S., Mahoney, D. P., “Distillation Column Control Design Using Steady State Models: Usefulness and Limitations”. *ISA Transactions*, 1993.
5. Jacobsen, E. W., Skogestad, S., “Multiple Steady States in Ideal Two-Product Distillation”. *AIChe Journal*, V37, n4, Apr 1991, pp. 499-511.
6. Papastathopoulou, H. S., Luyben, W. L., “Tuning Controllers on Distillation Columns with the Distillate-Bottoms Structure”. *Ind Eng Chem Res*, V29, n9, Sept 1990, pp. 1859-1868.
7. Skogestad, S., Jacobsen, E. W., Morari, M., “Inadequacy of Steady State Analysis for Feedback Control”. *Ind Eng Chem Res*, V29, n12, Dec 1990, pp. 2339-2346.
8. Tolliver, T. L., McCune, L. C., “Distillation Column Control Design Based on Steady State Simulation”. *ISA Transactions*, 1978.
9. Tyréus B. D., Mahoney D. P., “Applications of Dynamic Simulation”. *Proceedings of the Chemical Engineering Chemputers II Conference*, March 1994.

APPENDIX A

Development of the multiple steady state condition for binary distillation with fixed mass inputs [5].

$$L_w = LM \qquad M = y_D MW_1 + (1 - y_D) MW_2$$

where L_w = mass reflux
L = molar reflux
M = reflux molecular weight
y_D = mole fraction of light key
MW_1 = molecular weight of light key
MW_2 = molecular weight of heavy key

If we take the partial derivative of L_w with respect to L, we have:

$$\frac{\delta L_w}{\delta L} = M + L\frac{\delta M}{\delta L}$$

$$\frac{\delta L_w}{\delta L} = M + L(MW_1 - MW_2)\frac{\delta y_D}{\delta L}.$$

If $MW_1 > MW_2$, as is often the case, then $\frac{\delta L_w}{\delta L}$ may be negative depending on the difference in molecular weights, the magnitude of $\frac{dy_D}{\delta L}$, and the magnitude of L. A negative sign here suggests that under some conditions, increasing mass reflux will actually *decrease* molar reflux. Figure A1 shows the relationship between mass reflux and molar reflux for a system exhibiting multiple steady state solutions. Figure A2 shows how overhead purity changes with changes in molar reflux (i.e., $\frac{dy_D}{\delta L}$). Recall that this relationship, along with the difference in molecular weights and the value of L determines the sign of $\frac{\delta L_w}{\delta L}$. At very low and very high overhead impurities, $\frac{dy_D}{\delta L}$ is small (see Figure A2), and wee see that molar reflux increases monotonically with increasing mass reflux. However, between about 0.001% and 0.35% overhead impurity, $\frac{dy_D}{\delta L}$ becomes very large (steep slope in Figure A2), thus driving $\frac{\delta L_w}{\delta L}$ negative such that molar reflux actually decreases with increasing mass reflux.

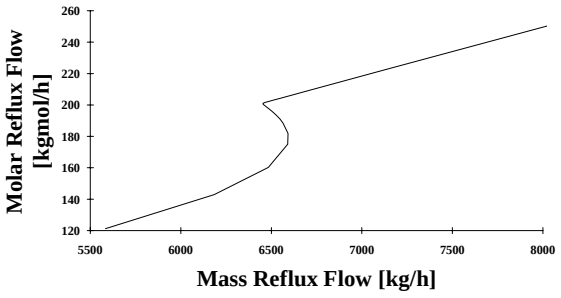

Figure A1 - Molar Reflux vs. Mass Reflux

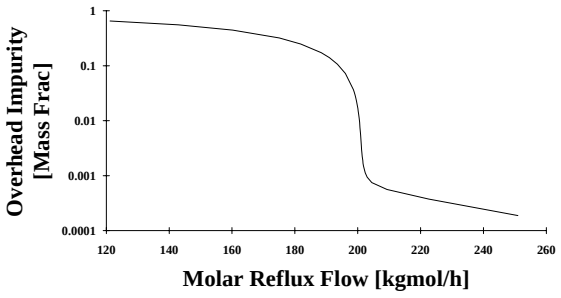

Figure A2 - Overhead Impurity vs. Molar Flow