Uncover Low-Cost Debottlenecking Opportunities

Don't immediately concentrate on a specific piece of equipment. Instead, look at the overall process and the interactions among units. This often can identify the most-cost-effective solution.

Dave B. Litzen,
Virtual Ideality
José L. Bravo,
Equilon Enterprises LLC

The recent strong domestic economy has presented many sectors of the chemical process industries with a problem every plant loves to have — how to make more product. Too often, though, companies do not thoroughly understand an existing plant's capacities and, so, are not prepared to challenge the existing facility to produce more. These companies take the "quick and easy" way to produce more product by spending capital to duplicate equipment or even entire processes. Or, worse yet, they do nothing and completely miss a business opportunity. Sometimes, spending major capital is the only way to achieve a business goal but, many times, the existing plant has hidden capacity that can be tapped more quickly and at less cost.

The literature is rife with articles that discuss debottlenecking specific equipment. All these references tout the traditional approach of "how can one get more hydraulic capacity through an existing piece of equipment." To this end, the classical approach for distillation, as an example, is to revamp columns with internals that allow more traffic without a loss in efficiency rather than to improve the workings of the installed equipment; some recent articles do look at squeezing more out of existing hardware in a column, as well *(1, 2)*.

This article instead concentrates on possible synergy among existing equipment in the entire plant. We have found that it is within the interaction of all the process steps and equipment where the nuggets often lie. Furthermore, we will discuss ways that allow the processing of larger feed streams without necessarily having to increase the internal throughput in equipment. When this approach is successful, it can produce more pounds of product at the end of the year with minimum capital investment.

Getting started

Understanding the existing plant capacity is a critical first step; representing it mathematically (at this point, a simple material balance) may be the easiest way to get there. Ideally, knowing the real capacity (not just what the design book said 20 years ago when the plant was built) *before* an expansion opportunity presents itself will save a lot of time. This may be the difference between a successful or a failed plant response to a potentially fleeting business opportunity. Keeping this understanding up-to-date should be an ongoing activity of the plant's technical and operating support staff.

Often, the modeling step is the most intimidating. Pure-component physical properties, binary interactions, azeotropes — scary, isn't it? The initial representation, however, does not necessarily have to be a full-blown, rigorous simulation model. For example: An existing oxygenated-hydrocarbon plant (producing four finished products in shared equipment) was to be debottlenecked several years ago. Assembling a spreadsheet to back-calculate flow rates to individual equipment or sections (mainly distillation columns and reactors) based on the product flow rates was enough definition to generate several debottlenecking plans over the years. Due to the discontinuous nature of this plant, the spreadsheet was a very handy tool to predict impacts on shared equipment.

The evolution of process simulation software has made producing a simple material balance relatively easy. At this point, it is not necessary to rigorously develop special physical properties. A total plant flowsheet simply representing all the major unit operations with flow splitters, mixers, and component separators (no flashes yet) is the best way to get started with the material balance. Using simulation software to complete the material balance will allow you to more accurately translate finished-product flow rates into equipment feed rates, which are important for the steps described below. Also, if you are interested in eventually producing a rigorous flowsheet model, working up a simple material balance in a simulator is a very good way to get the rigorous modeling activity started.

A reasonable plant material-balance flowsheet usually can be assembled in a matter of days. This exercise will force you to answer certain essential questions like what are actual stream flows and compositions — details necessary for a fundamental material balance and good for future rigorous model development.

One caveat in using a simplified model to represent the plant is that the model will not cover impacts related to vapor/liquid equilibria (VLE), such as compositional changes due to varying conditions over a range of production rates. These effects often are minor, but you must apply some engineering judgment to determine the level of model complexity required to best represent your process. Adding controls (fixing ratios and stream composition targets, for example) can compensate for some of the thermodynamic shortcomings of the simple model.

Armed with this predictive tool, you now should document each major piece of equipment's maximum-achievable feed rate and make sure it is consistent with the one in the material balance. (The term "major" has no strict interpretation. Although our example evaluation leaves out pumps, control valves, filters, and so on, these types of equipment can be included, though at the risk of increasing the complexity of the study.) Design data are always useful here but not necessary. The *real* limiting flow rates often can be determined by examining the plant's historical process data. If these data are available electronically, finding the maxima statistically is very useful. This does not necessarily determine the *true* equipment capacity, just the maximum flow that the unit has ever successfully handled.

In addition, interview the plant operators for their thoughts on maximum flows. The operators may well provide some insight on maximum feed rates to each piece of equipment, along with some ideas of what is leading to a feed limitation. For example, you may find that a distillation column is designed to handle 120-gal/min feed, but the operators will not run the column at more than 100 gal/min, often for very good reasons. The feed limit to this distillation column may not be a mechanical limit, but it is a *real* limit and needs to be understood. If no reliable data are available for a unit's maximum operating capacity, you will have to calculate the capacity. Estimating the capacity of a distillation column, for instance, is done fairly easily with a commercial simulator (which usually gives a conservative value) or a standalone tray-capacity program.

To illustrate the process of assembling a debottlenecking plan, let's use an imaginary process that produces a product called ABC. Figure 1 depicts the simplified process schematic. At this point, we will assume that we have a simple process model representing this process, and that we've determined all the major unit operations' maximum feed rates (marked by an "X" on the streams in Figure 1). The debottlenecking target for this example is 255 million lb/y of finished product ABC.

Once the maximum equipment feed rates have been determined, combining these data with the predictive tool completes the "evergreen" process of maintaining plant capacity knowledge. We've observed that using the sensitivity analysis feature found in some commercial simulators works very well for this activity. By generating plots for each individual unit operation (unit-operation feed rates vs. plant finished-product rate shown in Figure 2), you can transpose the equipment feed-rate limits on each plot to determine the overall plant product rate at which that equipment will become a bottleneck. In this example, the maximum feed capacity for the stripper before flooding occurred (or will occur, if this is an estimated limit vs. a limit based on plant data) is 117 gal/min, and for the extractor is 5.8 gal/min. These feed rates, provided all other bottlenecks were removed, would translate to an annual ABC production rate of 210 million and 218 million lb/y, respectively. Obviously, this plant will have to implement a capital improvement or, better yet, a low-cost process change to remove these bottlenecks to achieve the target rate of 255 million lb/y.

The stair-step chart

Once you have translated all the major-equipment-limiting feed rates to plant product rates, summarizing the bottlenecks on a "stair-step" chart by arranging them from lowest to highest appears to be the most effective way to communicate the data (see Figure 3). Technical and nontechnical people alike seem to quickly grasp a presentation of data in this single-page form.

To enhance the usefulness of the stair-step chart, add cost estimates

■ *Figure 1. Simple schematic of the ABC process.*

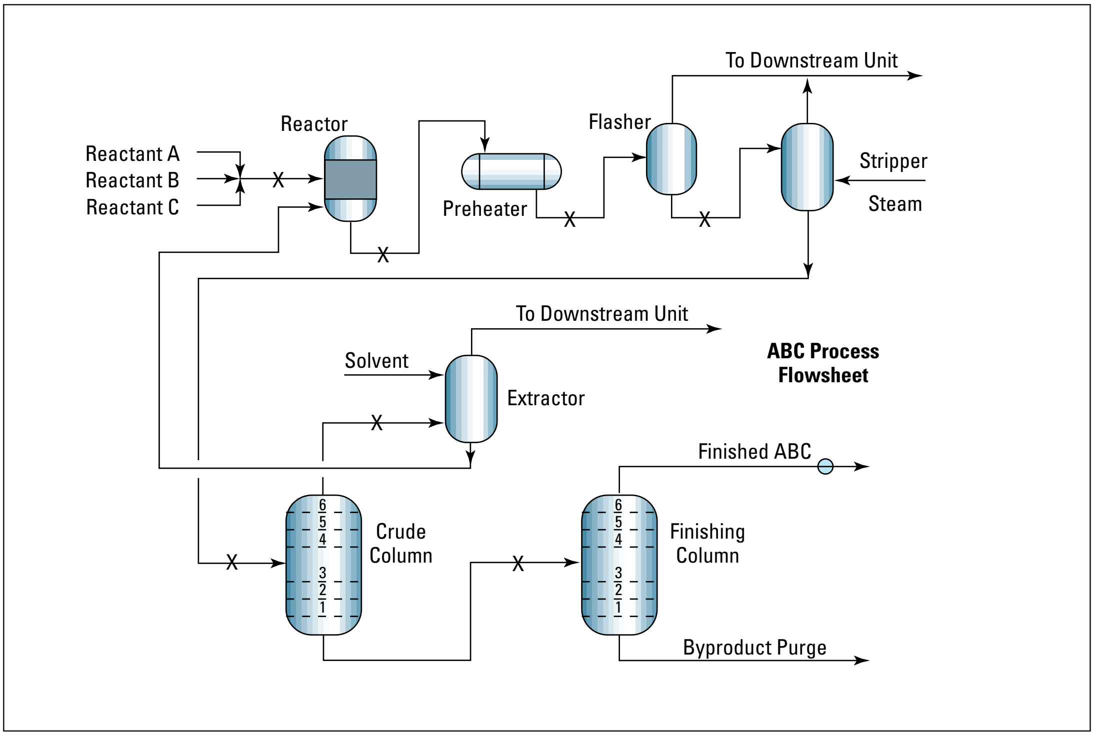

for debottlenecking each piece of equipment; this allows easy visualization of the cost/benefit of each "step" as you progress toward your debottlenecking goal. For this cost estimate, assuming that the equipment bottleneck cannot be removed by a process change will result in the most conservative estimate. This chart also will identify the equipment that has more-than-adequate capacity to achieve your goal. Having this information gives you the opportunity to creatively remove an equipment constraint below the target production rate by taking advantage of extra equipment capacity elsewhere — that is, equipment synergy.

The chart in Figure 3 indicates that the stripper, as mentioned previously, is the primary plant bottleneck. The example presumes that for $2 million the plant can raise the capacity of the stripper to support plant rates of 255 million lb/y or more. Figure 2 can be used to determine the minimum feed rate to the stripper at the debottlenecking target (see the dotted extrapolation line). After the stripper capacity is expanded, the new stripper-capacity bar can be moved toward the right on the chart and reranked against all the other unit operations. As long as the capacity is above the target line, the stripper no longer will come into play for this evaluation (unless its new capacity can be utilized somehow to remove other bottlenecks with simple low-cost process changes). Now, the extractor becomes the plant bottleneck. In essence, the $2-million expenditure

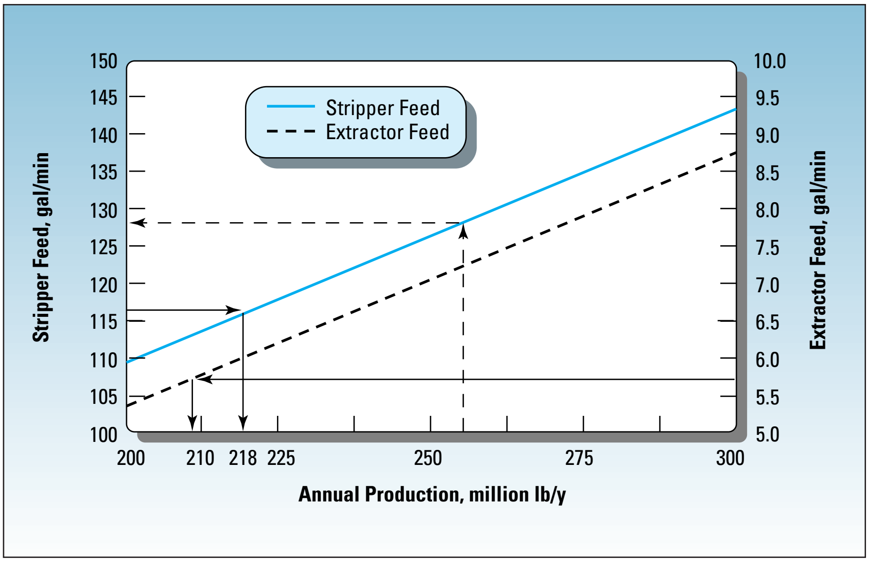

■ *Figure 2. Plotting individual-equipment feed rates vs. annual production can pinpoint the amount of additional feed required to meet the output target.*

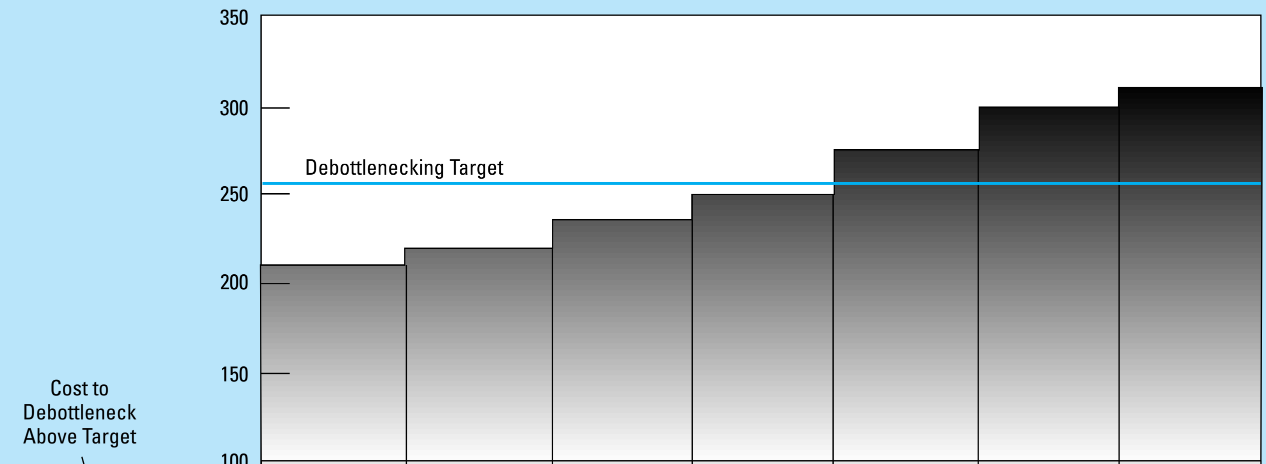

	Stripper	Extractor	Finishing Column	Reactor	Crude Column	Flasher	Preheater
Capital Cost, $ million	2	0.5	1	0.5	0	N/A	N/A
ABC Production, million lb/y	210	218	235	251	274	300	310

Figure 3. A single stair-step diagram clearly shows which major equipment need to be debottlenecked, and which have excess capacity that may provide debottlenecking opportunities.

for the stripper increased overall plant capacity by 8 million lb/y, which as such may be a poor investment if that were the only step taken toward the 255-million-lb/y capacity goal.

Generally, all the unit operations that fall below the target line need to be expanded to achieve the goal. In some situations, however, the stepwise removal of *all* bottlenecks beneath the target line might result in a case of diminishing returns. The data on the stair-step chart provide all the basics for a cost-benefit analysis. For example, if the cost to raise the reactor capacity were $5 million instead of $0.5 million, it might make sense to lower the target to the current reactor capacity. If the target is far above the capacity of most unit operations in the plant, the cumulative debottlenecking cost from the stair-step chart readily might make apparent that building a new plant is the only cost-effective solution.

Our experience suggests that if the debottlenecking target is no more than 20–30% above the nameplate capacity, you should have a pretty good chance of utilizing the existing plant (due to safety factors and conservative assumptions used in its design) to achieve your debottlenecking goal.

So, how can an evaluation like this save money? By pinpointing which equipment require debottlenecking and which do not, it allows you to focus on possible synergism. For instance, Figure 3 shows that the crude column, the flasher, and the preheater have capacities greater than the debottlenecking target, while the stripper is the primary bottleneck fixing the present plant capacity. These observations should stimulate some potentially-low-cost debottlenecking solutions. One idea: increasing the duty of the preheater and removing more volatile compounds in the flasher (as both have extra capacity) would reduce the required steam-stripping load in the stripper, which might boost stripper capacity enough to reach the removal target. Or, taking advantage of the extra capacity in the crude column by raising reflux to improve separation might allow decreasing the solvent rate to the extractor, reducing reflux on the finishing column and total flow to the reactor. This should result in increased capacity for all three unit operations with one small process change in the crude column.

This systematized approach to understanding plant capacity is most useful if done *before* the need to debottleneck presents itself. Ideally, once the stair-step chart has been assembled, you should keep it evergreen by updating it once per year or whenever a major piece of equipment is replaced or added.

Generating, evaluating, and ranking ideas

Once the overall plant capacity has been assessed, you can turn to the fun and exciting part of being creative to remove bottlenecks. The game (okay, objective) is to come up with ideas that are based on process changes that require no capital. Use the stair-step chart to pinpoint where the priorities are and where the synergism might lie. It's almost a necessity at this point to develop for the major process steps rigorous simulation models (ones that use fundamental physical properties and possibly additional interaction data to determine the VLE and VLLE) to help evaluate your creativity. Having validated, rigorous models not only will improve the accuracy of the results but also the speed at which they are generated. If process models do not already exist, generating new ones may be the most-time-consuming step in the entire evaluation process — particularly if component physical properties are not immediately available for your process. Having a proven set of physical properties for process modeling is important for all processes, but extremely important for those that may contain many nonideal interactions. Indeed, your success may well depend on an accurate representation of these physical properties.

Let's assume the component physical properties are available for your process. Of course, having process models that mimic the plant is the desired objective. Model convergence does not guarantee this *(3)*. To properly validate a process model, you must compare the model input and results against stable plant data. Analytical and process data that are normally recorded to operate the plant often are adequate to use for model validation. Many modern plants have these data electronically archived for easy access and evaluation.

To generate a basis for the validation, rely on data taken during stable plant operation — ideally 30–50 consecutive days with no upsets or interruptions — to determine process means (or medians, if data are not normally distributed). Then, use the resulting means (or medians) for fixing model input and comparing calculated results. Validating the model at different rates and conditions would result in a more robust model, but is not really necessary for evaluating debottlenecking options.

Using any simulator, models can be run manually to adjust process parameters (column tray efficiencies and heat-transfer coefficients, to name two) to match plant data. To aid the validation process, some simulators have tools to regress process parameters using techniques similar to those for physical-property regression. Additionally, these tools usually can help to reconcile plant flow-rate data (as raw plant flow-rate data rarely do material balance) to improve validation.

Use the stair-step chart to pinpoint where the priorities are and where the synergism might lie.

The validated models now can be used to evaluate debottlenecking options case by case. The simulator's sensitivity analysis again can be a time-saving feature for the case evaluations.

The evaluations hopefully should result in many projects, both capital and noncapital. These should be ranked, with the noncapital projects given immediate consideration for implementation. In addition, estimating the installed capital cost for the remaining projects is necessary. Following that, you must choose an economic index and index target consistent with business management objectives. Examples of an economic index are simple payout (capital cost divided by annual revenue), earning power (maximum cost of capital at breakeven point), or capital efficiency (5-y net present value divided by capital cost). Process risk and staff time consumption also may require consideration in the ranking process.

Some real-life examples

We already have used this strategy successfully on numerous projects, including:

Reactor/column combination. In the oxygenated hydrocarbon process mentioned previously, the plant was asked to exceed its understood capacity for product. The stair-step chart indicated that the reactor was limiting the manufacturing process. The initial impulse was to install an additional reactor. But, with the knowledge we had of the equipment constraints, we knew that the purification column downstream of the reactor had *plenty* of capacity. We also learned from operator interviews that the reactant concentration target in the reactor product, set at 0.1 wt. % to meet finished product specifications, was what established the maximum feed rate to the reactor. As an alternative to installing a new reactor, the plant instead relaxed the self-imposed reactant concentration in the reactor exit to 1 wt. %, allowing an increase in feed to the reactor. Using the excess capacity of the purification column, the plant boosted reflux, top product recycle, and control tray temperature (consistent with model predictions) to ensure finished product specifications were met. Sounds simple — but the plant was prepared to spend \$1–2 million to achieve the same results. The alternative was free, except for the minor variable cost due to higher recycle and reflux. More importantly, the results were immediate.

Column pressure leveraging. In another oxygenated-hydrocarbon manufacturing plant, the finishing column was designed to remove and recycle reactant (a higher boiling component than the finished product).

The 100-tray column required a relatively high reflux to successfully purify the finished product to its required specification. This, in turn, consumed much of the column capacity. The column, for convenience, was operated at atmospheric pressure.

Very often, changing the column pressure can squeeze a few percent more capacity from distillation columns. Decreasing the system pressure usually results in increased relative volatility of the key separation components, reducing the required reflux (or number of stages) to meet a separation target. On the other hand, reducing the pressure decreases the vapor density, increasing vapor velocity and cutting overall column capacity. Because of these two counteractive effects, a natural pressure optimum will occur *(1)*. This effect is easily observed by assembling a sensitivity analysis in a column simulation.

For the oxygenated-hydrocarbon finishing column, we can fix the composition targets of the key components (the product and the reactant in this case) in the top and the bottom of the column, vary the column operating pressure independently, and change the reflux and purge rate to meet the target compositions. Have the simulator estimate the maximum flood factor for each pressure case. Plotting flood factor vs. pressure across the range of evaluated pressures allows you to observe the optimum (see Figure 4).

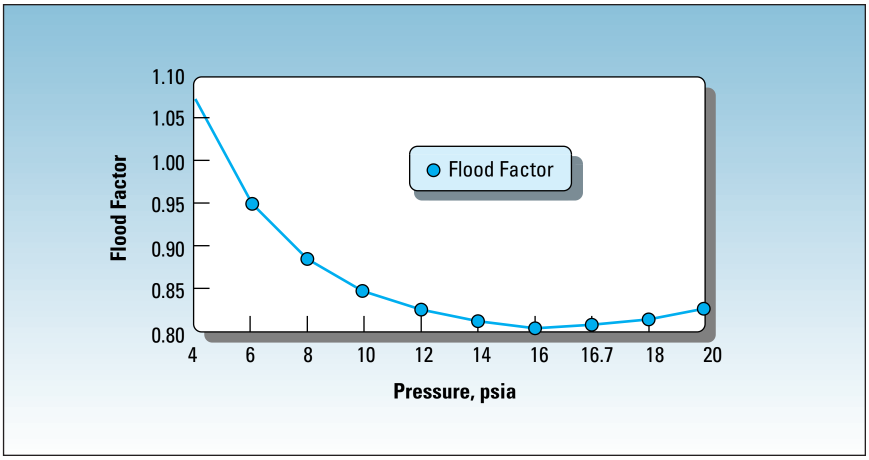

Figure 4. With fixed top and bottom compositions, a distillation column model can exhibit a minima when plotting flooding factor vs. operating pressure.

For the oxygenated-hydrocarbon finishing column, a 1–2% capacity increase should be achieved by moving the operating pressure up slightly from atmospheric to 16 psia, the optimum. In the actual case, more improvement was required; so, the column was refitted with higher-capacity, more-efficient trays to obtain the target capacity.

Column pressure leveraging — the other way around. Sometimes, decreasing the pressure of a column can remove a bottleneck. Indeed, a classic way to increase feed capacity is by reducing pressure and lowering reflux. In the retrofit of a debutanizer and deisobutanizer (DIB) combination in a gas plant, we went even a step further by decreasing tray spacing (a step traditionally associated with lower capacity) to obtain more stages and lower reflux.

Three issues turned out to be most important for the integration and debottlenecking of these two towers:

1. Increasing debutanizer pressure — to provide hotter overhead that can be used to reboil DIB.

2. Decreasing DIB pressure — to give cooler bottoms that can be reboiled by the debutanizer overhead, raise relative volatility of *iso*butane (thus enabling lower reflux), and cut hydraulic capacity available in tower shell.

3. Adding to DIB tray count — to boost theoretical-stage count and reduce reflux required, at the expense of tray spacing and hydraulic capacity.

The combination of items 2 and 3 results in large energy savings and increased capacity due to lower internal column traffic. Additional hydraulic capacity in the DIB at reduced pressure and tray spacing only can be achieved via high-capacity sieve trays. Interestingly, pressure and tray-spacing reduction normally are considered to cut column capacity. In the case of the DIB, however, the gains from improved separation performance from more trays at lower pressure overcome the loss of hydraulic capacity by requiring lower reflux. When this is combined with the superior flooding capacity of the improved sieve trays, the DIB can be debottlenecked by moving in the seemingly counterintuitive direction of lower tray spacing (more trays) and lower pressure (better volatility).

Feed-port/sidedraw optimization. Many distillation specifications target very low concentrations of single, identifiable compounds. These stringent targets often can lead to opportunities to optimize feed-port and, if applicable, sidedraw locations on purification columns. For example, a chlorinated-hydrocarbon finishing column (with a sidedraw product) was designed to remove both heavy and light chlorinated impurities from the finished product. Reflux ratio already was maximized to meet purity specifications and reduce yield loss. The plant was unable to incrementally increase feed rate without sacrificing either purity or yield loss. Using a sensitivity analysis similar to that described above, optimum feed-port and sidedraw locations can be determined simultaneously by maximizing vapor traffic

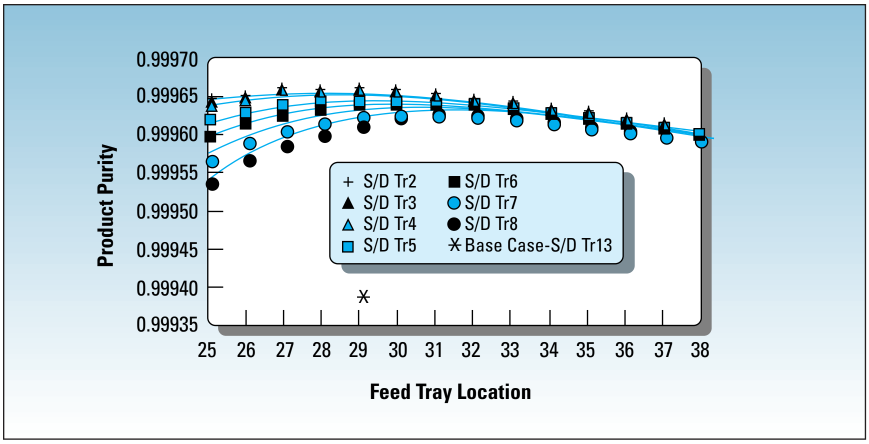

Figure 5. A sensitivity analysis can be used to determine optimum feed-port and sidedraw locations.

(by maximizing reflux at a given feed rate), fixing the product loss in the purge stream, and varying feed-port and sidedraw locations to locate the point of highest purity (see Figure 5). You can use other combinations of variables (fixing purity and yield loss to determine minimum reflux, for instance) to gain the same results.

In both this and the previous pressure-optimum case, a flowsheet optimizer (found in many flowsheet simulators) also could be used. This optimizer should come up with the same answer as the sensitivity analysis, but only will give the optimum point (if one exists). We prefer the sensitivity analysis, because it helps visualize the optimum. This can be very valuable in determining whether the independent variable optimum has a flat or sharp response and, thus, how sophisticated solutions need to be to maintain the optimum. In this case, the true optimum is in a position (feed port on Tray 27 and sidedraw on Tray 2) that does not exist. The shape of the curves, though, led to the best choice of existing feed-port and sidedraw location, a change that was made for no capital cost and minimal loss of production.

Column realignment. In the chlorinated hydrocarbon system mentioned above, the existing plant used two columns in series to remove light ends to a very low concentration to meet the finished product specification. The second light-ends column was the smaller of the two, and limited the capacity of the overall finishing process. Reevaluating the column capacities individually indicated that each column on its own could meet the light-ends specification. So, the plant has been modified (via piping changes) to operate the light-ends columns in parallel instead of in series, providing a quick, 40% increase in light-ends removal capacity for minimal cost.

Compressor/preheater balancing. In a plant with a light olefin feedstock, the process requires vapor recompression and preheating of the olefin before it is reintroduced to the reactor system. The stair-step diagram indicated that the compressor was the bottleneck to increasing production capacity at the present byproduct yield, but the fired preheater was the bottleneck if the plant was targeting reduction of yield loss at a fixed capacity. Simply understanding the limitations, tradeoffs, and synergies between these two pieces of equipment allowed the plant to increase product rates by 15% by operating closer to the compressor capacity threshold. The deeper understanding of the individual equipment capacities stimulated replacing the furnace burner tip with a new design (about $1,000 installed) that enabled the preheater to operate with 30% more heating duty. During the occasional low-product-rate periods where the fired preheater becomes the bottleneck, yield loss has been reduced, saving the plant up to $1 million/y.

Use of flooding detection mechanisms. One way to maximize capacity in a process is to maximize the effective onstream factor. Too often, though, plants push their equipment without considering the impact on onstream factor. Yet, when processes are operated very close to their limits, excedences can mean lengthy recovery periods that end up reducing the real onstream factor and output. This is particularly true in distillation steps — a column that floods may not produce valuable product until it completely recovers from the flooding episode.

If pushing a column to 98% of its capacity from 90% lowers reliability to, say, 90% from 100%, you actually wind up worse off, because output will have dropped to 88% of nameplate from 90%. The trick is to operate at, say, 98% of capacity 100% of the time. That would represent removal of an 8% bottleneck in the operation.

For columns, the use of reliable indicators to warn the operator that a flood episode is approaching may prove very valuable. If the column is at 98% of maximum capacity and an incipient-flooding indicator turns on, the operator may choose to back off the feed to avoid the costly flood episode. Then, once the danger is averted, the capacity of the column can be challenged again, using the flooding indicator as a warning.

By using such indicators, we succeeded in debottlenecking a large distillation column that had some reliability issues at high rates. The effective capacity of the column was increased by more than 5% by operating the column at a rate equivalent to 98% of calculated capacity for most of the time and backing off

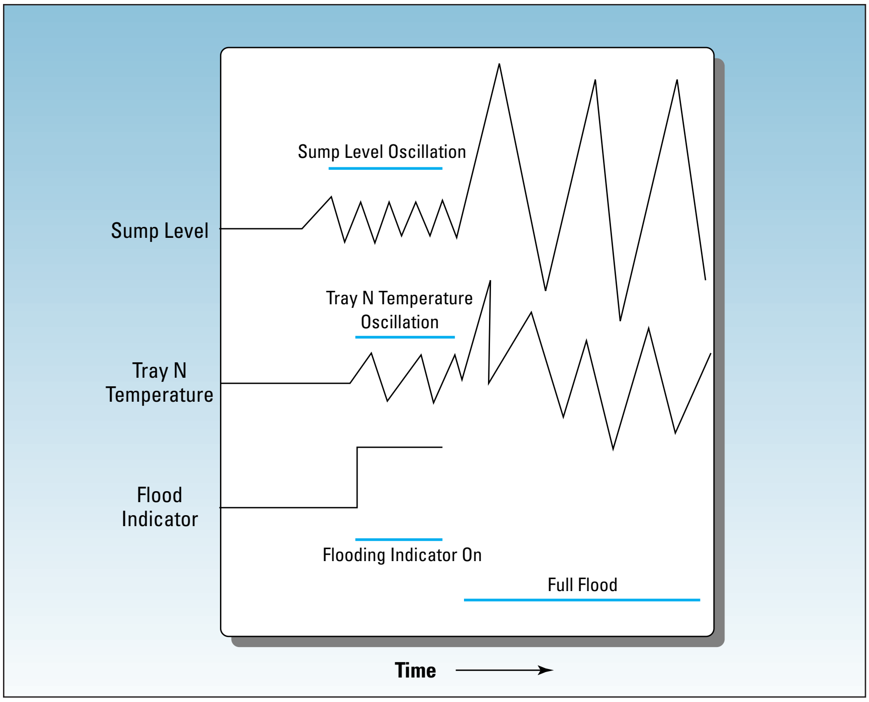

Figure 6. The response of two variables reliably predicts impending flooding and, so, can trigger a warning indicator.

when the incipient-flooding indicators turned on. Flooding episodes in a typical year were reduced from an average of 12 to 2. In this case, two variables provided a tipoff of incipient flooding:

- liquid temperature 10 trays above the bottom; and
- liquid level in the sump of the column.

These two variables showed very repeatable oscillations with a fixed amplitude and period several minutes before every recorded flooding episode (see Figure 6). The control system was programmed to detect these particular oscillations and turn on the incipient-flooding indicator. The operator then would back off feed until the indicator turned itself off. After a small period of time, the operator would ramp up the feed again. Because actual flood point may change with time or variations in factors such as feed composition, pressure, or pressure drop, the operator then may be able to get to the specified 98% point without the incipient-flooding indicator relighting. If conditions have not changed sufficiently, the indicator will alert to operator to stop the ramp up.

Traditional indicators such as temperature profile and pressure drop can alert that flooding already is occurring, but cannot reliably warn of the approach of flood. In contrast, the dynamic behavior of temperatures and levels can be a clear indication of impeding flooding.

Literature Cited

1. **Fair, J. R., and A. F. Seibert,** "Understand Distillation-Column Debottlenecking Options," *Chem. Eng. Progress*, **92** (6), p. 42 (June 1996).
2. **Sloley, A. W.,** "Should You Switch to High Capacity Trays?," *Chem. Eng. Progress*, **95** (1), p. 23 (Jan. 1999).
3. **Horwitz, B. A.,** "Hardware, Software, Nowhere," *Chem. Eng. Progress*, **94** (9), p. 69 (Sept. 1998).

Uncover the opportunities

The more thoroughly you understand the plant capacity, the more ideas you should be able to come up with to debottleneck the plant. Simulation can help you to more quickly and accurately evaluate these ideas. The more ideas you evaluate, the more likely you should be to identify a low-cost, low-impact debottlenecking plan. And, that means money.

Debottlenecking a plant, in many instances, may involve changes in the reliability of continued operation at higher rates so as to attain a higher effective onstream factor. In other cases, first principles may allow you judiciously to adjust operating conditions, such as pressure, to achieve maximum feed throughput without exceeding the internal hydraulic capability of devices such as distillation columns. CEP

D. B. LITZEN is principal of Virtual Ideality, Rapid City, SD (Phone: (605) 355–9342; Fax: (605) 355–0980; E-mail:Dlitzen@compuserve.com), an engineering consulting firm. Before starting that firm, he spent nearly 16 years in various process engineering and manufacturing assignments for Shell Chemical Co. He specializes in building and extending process models, particularly for nonideal chemical systems. He received a BS in chemical engineering from South Dakota School of Mines and Technology. A registered Professional Engineer in Texas, he is a member of AIChE.

J. L. BRAVO is a separations specialist with Equilon Enterprises LLC (a Shell/Texaco/Star Alliance joint venture), Houston (Phone: (281) 544–6721; Fax: (281) 544–8123; E-mail: jlbravo@equilon.com). Previously, he worked for Jaeger Products, and was manager of the Separations Research Program at the Univ. of Texas at Austin, where he received an MS in chemical engineering. He has extensive experience in the design and troubleshooting of mass-transfer equipment, and has published more than 30 articles in the areas of distillation, absorption, extraction, and stripping. In addition, he is the senior author of a book on separations technology, and an instructor at AIChE's Distillation in Practice course. A member of AIChE, he is a registered Professional Engineer in Texas.