IMPROVE THE PERFORMANCE OF YOUR EXISTING STEAM SYSTEM

R. O. PELHAM, Merrick & Company, Aurora, Colorado

Refinery and petrochemical plant engineering staffs must understand complex-wide steam systems, and they should seek to improve the performance of these systems. Here, the adaptation of existing facilities in response to process changes in the refinery or petrochemical complex is examined.

Pursuant to the January 2013 article, "Improve your plant-wide steam network," this article is concerned with improving the performance of existing steam equipment. Several articles[1–3] have been written over the years regarding ways to improve steam systems. These articles have mainly focused on specific elements or "pieces" of the total system, and they touch on the timeless issues of boiler performance, steam turbines, steam reboilers/heaters, letdowns, process uses, condensate recovery, steam traps, deaerating, etc. Also highlighted are some aspects that generally receive less attention, or that are altogether ignored.

Steam system weaknesses and inefficiencies. Industry has not developed a specific metric that defines overall steam system performance. For the most part, it is left to the knowledge and experience of the engineers and operators to observe and develop improvements. Since the "utility system" engineer is frequently chosen from the ranks of younger engineers, and since operations personnel are responsible for their process area only and not the refinery-wide steam system, progress can be slow.

After completing all the steps discussed in Part 1 of this article, the total system is understood—but what comes next? How does one go about improving the system? There are a number of items to look for and consider. The following discussion tracks a logical material flow through the steam system, starting with the feed (i.e., cold, demineralized water), and continuing with steam generation, steam uses (and abuses) and, finally, condensate recovery.

Preheating treated makeup water to deaerator. Most energy-conservation schemes in refining come down to one of two approaches: either not using the energy in the first place, or recovering more energy from the multitude of streams being cooled in air or water coolers. In the latter case, the limitation is not the heat available, but rather having sufficient ΔT to recover the heat (i.e., second law of thermodynamics efficiencies, commonly referred to as "pinch" technology).

The colder the sink into which heat is recovered, the more ΔT is available. The two largest cold sinks in a refinery are cold crude and cold deaerator makeup water. Heating cold crude gets attention in crude preheat trains, and effort is made to practically and economically recover all heat from crude unit product streams prior to air and water cooling. The equivalent opportunity to maximize recovery of waste heat against cold makeup water rarely receives equivalent attention. Instead, the bulk of water preheating is done with low-pressure (LP) steam.

FIGS. 1 and 2 illustrate a refinery setup with letdown stations at various locations throughout the plant. Some stations were metered and some were not. Steam to the deaerator was not metered. Knowing the water rates around the deaerator allowed a simple deaerator heat balance, which gave an accurate measure of steam rate—an eye-opening occurrence.

Another attention-drawing factor was the rates through different letdown stations in the refinery. Operators discovered that the equivalent of an entire fired boiler was devoted to the sole purpose of generating high-pressure (HP) steam, which was then let down and used to preheat cold water to the deaerator. Using waste-heat streams eliminated a fired boiler.

Since deaerator feedwater still contains oxygen, heating to the 100°F–180°F range requires stainless steel plumbing for exchanger and piping to the deaerator. A new approach, which may reduce capital costs, is to deaerate the cold feedwater using membrane technology.[4] This approach is now being introduced at some US refineries.

Boiler firing and efficiency. Clearly, if fired boilers operate at 75% efficiency, it is unlikely that the steam system can be deemed "efficient." Issues include standard ones, such as recovering stack heat with economizer sections, and air preheating. A good stack temperature objective is 400°F–450°F. The aim of control furnace firing is to maintain low excess air (e.g., less than 20%).

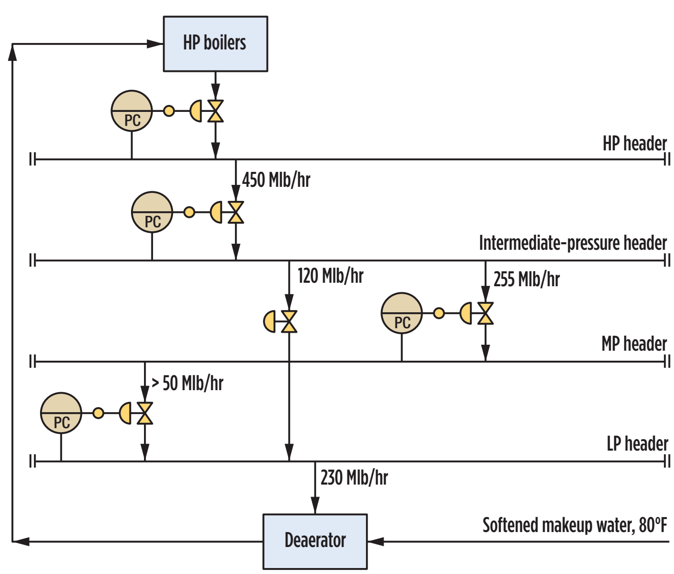

FIG. 1. Deaerator feedwater heating, no deaerator feed preheat.

Boiler blowdown. Over the past 20 steam system surveys, the author has observed blowdown rates of between 3% and 13% of the boiler feedwater. The major cost of a high blowdown rate is the cost of steam used to heat the cold deaerator feedwater. Additional costs are incurred from raw water cost, water-treating cost and sewage cost for the blowdown.

Two factors are in play. First is the boiler feedwater quality. Generally, in refineries, the feedwater quality is well understood, and blowdown percentage (or cycles) is set accordingly. If the blowdown target is high, the only solution is improved water treating.

The second factor is blowdown control. Ideally, blowdown is controlled automatically in ratio to boiler feedwater. In practice, however, many boilers do not have automatic ratio control. In that case, blowdown flow may be on flow control. A setpoint is frequently used to ensure that the blowdown is never less than the target percentage. Unfortunately, since steam rate varies, blowdown is usually too high when the boiler load is less than the maximum.

Worse still are boilers where no blowdown control is present; blowdown must be done manually, by outside operators. In this case, the actual blowdown rate is a guess. Given the consequences to the boiler of too little blowdown, the actual blowdown is set high as a precaution. When blowdown is high (e.g., over 5%), cost should be calculated, as explained in the preceeding article published in January 2013. That cost should be compared to the cost of installing automatic blowdown controls.

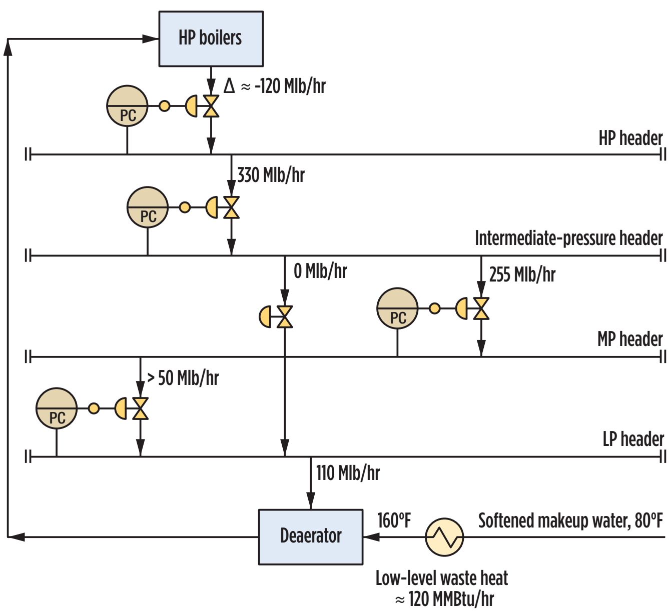

FIG. 2. Deaerator feedwater heating, with deaerator feed preheat.

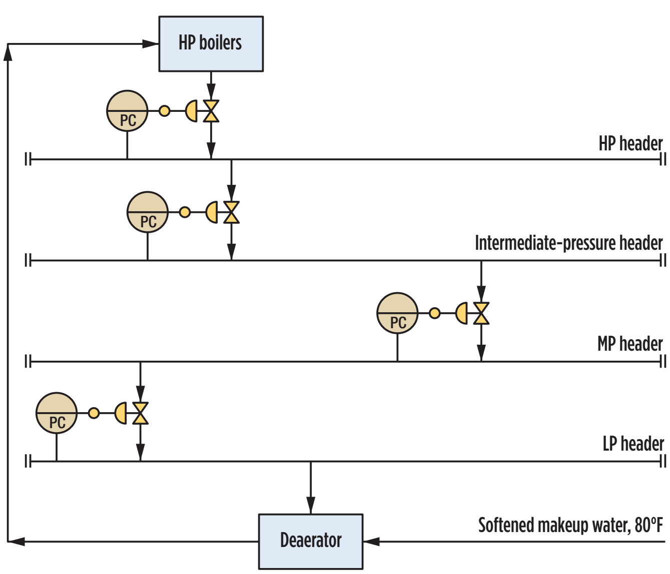

FIG. 3. Normal header pressure controls.

Header pressure control. On an overall basis, most steam systems control the HP header pressure by pressure control on boiler output. Intermediate and LP headers are then controlled by letdown from above to maintain the target pressure at the lower level. The general scheme is shown in FIG. 3.

A few refineries and petrochemical complexes run HP superheated boiler steam through large, double-extraction turbines. The turbines generate power. The turbine extraction rates are on pressure control at the desired intermediate header pressures. These are large, sophisticated machines. The practice is more common in European refineries than at US plants, although the difference seems to be cultural rather than technical or economic.

In terms of what to watch for in header pressure control, the following should be considered:

- The HP header
- Intermediate-pressure headers
- The LP header.

In all cases, there are two things to consider. First, how is the pressure controlled? Second, how is the target pressure determined?

Typically, the HP header pressure will be controlled by pressure control on fired-boiler output. In this setup, output on more than one boiler may be ramped up or down together to maintain target pressure. In other cases, some boilers may be on flow control, and one "swing" boiler will handle pressure control. These controls are normally in the utilities control room or console, and they are handled by the utilities operators.

However, the HP header control does not normally end there. Somewhere in the refinery, there will be one or more letdown stations. It is acceptable if these stations are configured and set to maintain only an intermediate pressure below. Frequently, however, the control may be set to maintain

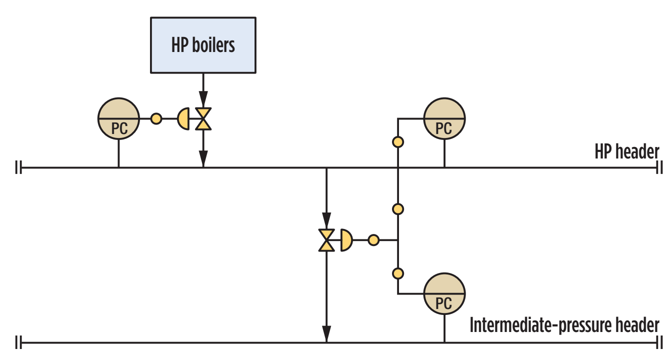

FIG. 4. Header pressure control.

HP header pressure locally. Common arguments are that the stations are too far from the boilers, that there are restrictions in the line, and that local HP waste-heat boilers are a disturbance. Regardless, there is now more than one controller performing control operations. It is key to ensure that HP boilers are not working at one end to maintain a pressure, and that letdowns on the same header are letting down to relieve pressure.

Setpoints vs. control. There are two things to keep in mind. First is the relative setting of the pressure control targets at competing controllers. The setpoints should not be configured so that the letdown relief setpoint is at a higher pressure than the boiler output pressure control (FIG. 4). In this scenario, the boiler will increase output, and the letdown will relieve the overpressure, in effect dumping excess boiler steam into the lower-pressure systems. (**Note:** The setpoints are typically configured by operators in different process departments, with no overall coordination.)

Second, the letdown station may have a split-range controller that will switch from makeup control for lower pressure to letdown control of higher pressure in response to an overpressure situation. In a recent steam system review, the author was assured that a split-range controller only let down in rare overpressure situations. An examination of a year's worth of process information data showed that the controller was in HP relief mode 46% of the time.

Another overall issue to consider is the target pressure. For example, if the nominal, 600-psi system is running at 580 psi, then the question to be answered is, "Do we have the flexibility to raise or lower the 600-psi system target pressure, and, if so, where should we set it?" In most plants, this question is rarely asked. At the HP header, the best practice is to run at the lowest pressure practical. (**Note:** This is an operating recommendation for an existing system; it may not apply in the design of a new steam system.[5])

The lower limit will be reached when, for example, a turbine does not generate enough horsepower, a reboiler does not provide enough heat, or a live steam stripper no longer meets flashpoint. This issue is worth investigation, and operation at lower pressure, with some operating margin, should be considered. The benefits of running at lower pressure include higher boiler efficiency, reduced line heat losses and reduced steam loss at steam leaks.

Never assume that the nominal nameplate pressure is the actual pressure, however. The author was recently involved in a process and instrumentation (P&I) meeting between engineering company designers and refining company operators. In the meeting, it was revealed that the designers had assumed that the 150-psi system was operating at 150 psi and had designed on that basis. Refinery personnel knew that the 150-psi system ran at 130 psi, but no one had ever pointed out that fact.

Intermediate-pressure headers raise similar questions to those above, but they are more complex. There is a mixture of steam supplies (e.g., letdown stations, process steam generators, condensate flashing, back-pressure turbine exhaust), as well as a mixture of uses (e.g., letdowns, heating/reboiling, turbines and process steam).

The most critical question is, "Are we sure that controllers are not set somewhere that cause simultaneous letdown from above (to maintain pressure) and letdown to below (to relieve pressure)?" Each option will look normal in the particular process area in which it occurs, but it will not be logical from a total system viewpoint.

It is always good practice to raise the question, "What has determined the actual pressure at which the pressure is controlled?" Is it arbitrary, or has it been lowered to maximize power from a critical turbine that exhausts into this header? Or has it been raised to achieve maximum power from a turbine using this header as inlet steam, or from some reboiler where a few more degrees were needed to bring sufficient heat into some column? If there is no known constraint, then the benefits of raising or dropping it by some Δpsi can be brought into question. The answer can be in either direction; it is specific to the steam suppliers and to the users connected to that header. Chances are good that the pressure is suboptimal.

The LP header will have steam supply from letdowns, turbine exhaust, waste-heat boilers and condensate flash. Uses are limited to those things that can be achieved with steam having a temperature of 240°F–260°F (10 psi–20 psi of steam). Every steam system is different, but, overall, refineries and petrochemical complexes tend to be long on LP steam. If there is excess, it will vent to atmosphere.

Venting steam—which ultimately comes from a fired boiler—is expensive and frequently noisy; it is a visible signal of imbalance and should be avoided if possible. However, some refineries (in the author's experience, approximately 20%) treat continuous venting of LP steam as inevitable, and it has become accepted as a simple cost of doing business.

One important thing to look for is whether different letdown and atmospheric vent pressure controllers in different parts of the complex are configured with the inlet setpoint higher than the atmospheric vent setpoint. The result is continuous venting (FIG. 5). The author has observed this phenomenon on several occasions. Usually, the engineer attempts to find the causes of the steam imbalance that is causing excess LP steam to vent. This can be frustrating, as it is not a material balance issue; it is a control problem. The solution is to understand the normal setpoints of the letdown and vent controls. Keep in mind that these controls are typically in different sections of the refinery, controlled by operators in different control rooms, and there is no overall coordination.

Another important thing to look for is whether excess LP steam is being hidden by excessive use. Examples discovered include:

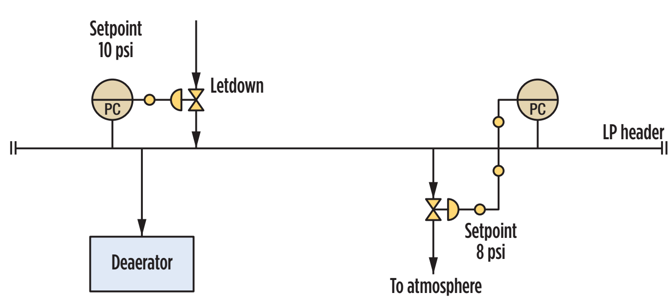

FIG. 5. Low-level header pressure control.

- Excess amine regeneration
- Very high iC_4/olefin ratio in an alkylation unit to consume excess LP steam in the deisobutanizer tower reboiler
- Very high steam stripping in a sour water stripper

The ultimate question to be answered is, "How many letdowns are really needed for system control?" Generally, the number is less than the number of existing letdowns.

- Hidden vents (i.e., condensing excess steam in air or water coolers, or even venting steam in a cooling tower).

The solutions in these cases are:

- Examine all sources of LP steam and evaluate methods to reduce or eliminate that supply
- Consider alternate, legitimate uses for LP steam
- Check options to thermally recompress the LP steam (occasionally, there may be a large letdown of 600 psi–150 psi of steam that could be used in an eductor to thermally compress the LP steam to a more useable pressure).

Number of letdown stations. Discerning the real number of letdown stations in a refinery is always an interesting challenge. There are the "official" stations, and there are also the "forgotten" ones. Conversations with the boiler house operators will typically elicit information about two to three letdowns from the HP system. These may include a few that are under boiler house control, and perhaps another significant one out in the refinery. Beyond that, further investigation is generally required.

The discovery of additional letdown stations requires conversations with process personnel in other areas of the refinery or complex. These discussions will typically affirm several more stations. This process typically requires some debate and clarification as to exactly at which pressure levels one or more of the letdown stations function.

The final effort is to make another investigative round, this time including board operators in different control rooms. It is advisable to ask about the letdown stations that are unused or that have been shut in or removed in the past. This process typically turns up one or two more candidates requiring field inspection to find out what stations are still in use. After this step is completed, all of the letdown stations can be added to the steam system drawing.

The next step is to find out how much steam is actually being let down. This is frequently a challenge. Some of the main letdown stations may be metered; others may need to be estimated by valve characteristic and ΔP across the headers. Further insight might be gleaned by the steam system balances.

The ultimate question to be answered is, "How many letdowns are really needed for system control?" Generally, the number is less than the number of existing letdowns. The required approach is to ask why each letdown is needed. Ideally, there is just one letdown between each pressure level; i.e., a steam system with four pressure levels requires three letdown stations. Justification for additional letdowns then occurs. Typical justifications include:

- One letdown is needed to control pressure locally
- Two letdowns are needed if each one's line size is too small
- Letdowns are needed to protect local critical equipment.

These claims must be analyzed carefully. Reasonable concerns are legitimate, although letdowns should not be added without a thorough analysis of where and how many are needed.

In a recent steam system review, 11 letdown stations were uncovered—three in just one process unit. The recommendation was to reduce the number of letdowns from 11 to three, which required proper sizing of the stations.

What are reasonable letdown flowrates? Once the number of letdown stations has been determined and everything possible has been done to identify total steam flows, two new issues emerge. First, the total aggregate amount of steam being let down may be larger than realized. (**Note:** This scenario is acceptable if it is legitimately required for low-level uses.) Often, however, it inspires ways to use LP steam that are much less efficient than simply not making the steam in the first place. A telltale indication is a large use for preheating cold deaerator feedwater. That flow is generally not metered and is given little recognition, and it is very easily calculated by deaerator heat balance. The overall objectives are to minimize LP steam consumption and, in turn, back up the system to reduce boiler-fired steam production cost.

The second issue is that, when letdown rates are minimized, the question then becomes, "How much letdown is legitimately required for system control?" The author has yet to establish a general guideline or best practice in this area. Typically, a value of around 10% of steam demand at the outlet pressure header is reasonable to maintain control. In reality, it becomes a question of analyzing size, rate and frequency of demand swings at the particular header, and using some judgment as to the letdown capacity reasonably needed to handle that level of demand variability.

Desuperheating. It is not uncommon to discover at least one desuperheating station in a complex steam system. It is typically found where there is a large letdown steam flowrate, or where steam is superheated to begin with (e.g., typical fired-boiler steam that has a superheater section and becomes more superheated across the letdown station).

There are two points to understand. First, highly superheated steam is not beneficial to most refinery users. It is not advantageous for reboilers and steam heaters, since some surface area is used up in cooling steam (low heat-transfer coefficient and duty) as opposed to the surface area used up in condensing steam (high heat-transfer coefficient and duty). Second, highly superheated steam is typically not efficient in hydrocarbon stripping, where moles of steam per mole of hydrocarbons is the target, and fewer (but superheated) steam moles may not have value. Superheat is typically beneficial to back-pressure steam turbines. However, if steam is let down in parallel to a back-pressure turbine, then turbine efficiency and steam superheat is not important. Steam superheat is favorable in condensing turbines; however, most refineries have already eliminated these turbines.

Expanding steam across a valve is an adiabatic process; no heat is gained or lost. Instead, the outgoing steam is more superheated than the incoming steam. Since steam letdown superheats steam, and since superheat is technically a disadvantage in many applications, desuperheaters are not uncommon. Desuperheating simply injects water into the letdown station. The effect is to increase pounds of steam, with each pound being of a lower quality or a lower superheat.

A short dialogue on wet steam vs. dry steam may be useful at this point. Steam engineers like to design so that the whole system is dry. There is some opinion that the whole system must be dry for safe operation. The author was once almost asked to leave a European refinery when he proposed an idea that would have created some wet steam. The belief was that the whole system needed to be dry everywhere. Ironically, at the next system on which the author worked, most of the steam was generated at saturation, with no superheat, and the whole system was wet.

So, how much superheat is appropriate? That is a tough question, as there is no measure of superheat (or wetness) in the system, and steam quality is not always known at the end of the line, or in distant areas, or following certain weather events, such as a thunderstorm. On the basis that more pounds of steam of lower quality are more useful than fewer pounds of high-superheat steam, there are a couple of general guidelines:

- Desuperheat large letdown flows of already superheated steam
- Aim for enough superheat to avoid condensing in steam lines (around 30°F–40°F superheat), assuming lines are properly insulated.

The local process area 'protection' racket. There is an old adage that says, "If you don't ask the right questions, you won't get the right answers." In studying existing steam systems, this adage applies. It is the author's custom, when talking with process staff and board operators in various refinery process areas and control centers, to ask what steam system control practices are applied in local areas to protect their domain.

Typically, in any process area, there are equipment inadequacies, such as a back-pressure turbine that struggles to put up enough pressure for the reflux pump, or a reboiler that does not quite achieve the duty a column needs. As a result, these items are "protected." A spare letdown station is used to maintain a specific local pressure or to tweak steam supply when needed, with no measurement and no record.

Some of these practices are appropriate. However, some of them may have consequences elsewhere in the steam system. A local control practice may impact other parts of the system in a different process area—i.e., Area A may be normal, but Area B may experience an upset or a consequence for reasons not understood. The net result is negative from a total-system viewpoint.

Steam vs. electric-driver sparing. Refineries and chemical complexes tend to have a large number of back-pressure steam vs. electric-driver sparing options. Local circumstances generally dictate which driver is normally run and which is normally used as the spare. However, switching to or from a back-pressure turbine affects the steam letdown between the same pressure headers. Turning on a 150-psi–20-psi turbine will reduce the existing letdown by the amount of steam the turbine uses. Conversely, taking a steam turbine out of service will increase letdown.

A more expensive issue occurs when the LP header is out of balance and excess steam from too many turbines online is vented to atmosphere. A more common issue is to use a steam-sparing or electric-sparing option to provide some rough control on reducing excessive letdown rates. Actual practices vary widely, depending on the severity of the variation and the ability to switch turbines.

A good practice is to provide a central source of information and control on the following management and procedural issues:

- Monitoring the actual letdown rates
- Targeting the minimum and maximum desirable letdown rates
- Considering the steam-sparing and electric-sparing options available
- Having knowledge of which driver is running for each sparing option
- Following established operator instructions on the sequence of which drivers are to be added or removed
- Adhering to procedures to request switching if the next combination to be switched is in another control room.

TABLE 1. Amount of HP condensate flashed to generate LP steam

Condensate pressure, psig	Flash (steam) pressure, psig	Condensate flashing to steam, %
600	150	16
600	50	23
600	20	26
150	50	8
150	20	12
50	20	4

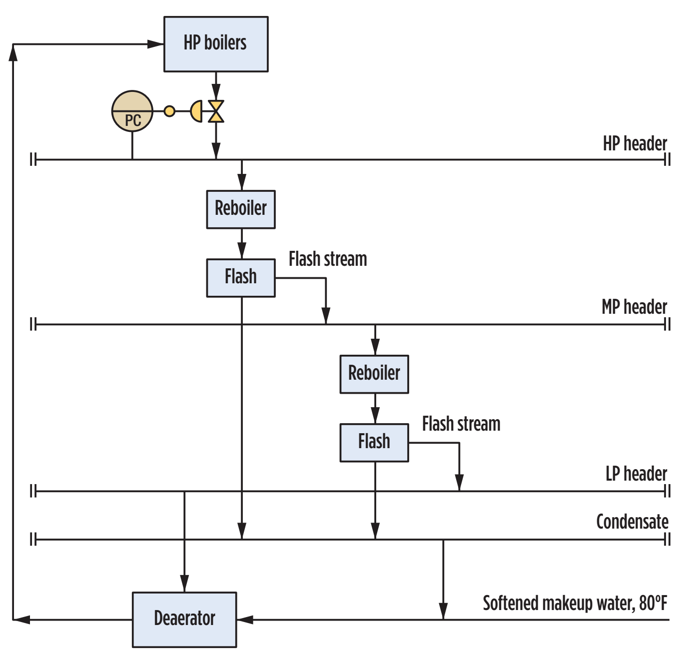

FIG. 6. Condensate flashing.

Condensate flashing and condensate recovery. Condensate is produced from three main sources:

1. Reboilers and steam heat exchangers
2. Condensing turbines
3. Steam tracing/tank heating.

The most critical question is, "Are we sure that controllers are not set somewhere that cause simultaneous letdown from above and letdown to below?" Each option will look normal in the particular process area in which it occurs, but it will not be logical from a total system viewpoint.

The biggest supply comes from reboiling/heating. This typically happens at some high or intermediate pressure level and is "clean," as it has not been in contact with hydrocarbons. In a perfect steam system, this condensate will be flashed to produce steam for use in the LP system, and the LP condensate is collected for return to the boiler feed system (FIG. 6). Wherever possible, condensate should be flashed in condensate flash pots to generate additional LP steam and help prevent dangerous water hammer in condensate return piping. TABLE 1 provides a perspective on the amount of HP condensate that can be flashed to generate LP steam.

The values in TABLE 1 assume that condensate is available at its saturation temperature. The percentage of condensate actually recovered as steam will be less, depending on system heat losses. This is largely a function of insulation quality and geography. For a large, well-insulated steam reboiler, with condensate collected and flashed close by, losses will be small. At the other extreme, condensate recovered in a distant tank farm and brought back onsite may experience significant cooling before being flashed.

Condensate should be tracked carefully as part of the steam system definition. A well-designed and well-operated steam system should recover up to 70% of available condensate. Losses tend to build from the accumulation of small quantities of condensate that do not justify lines and pumps for recovery. These quantities are instead dumped to local sewers, or even to the ground.

For steam systems that recover less than 50% of condensate, effort should be made to locate points where condensate is lost, to determine the size of this loss, and to consider the cost of recovering the condensate vs. the value of the condensate recovered.

Note that some complexes define the percentage of condensate recovery as:

$$\text{Percentage of condensate recovery} = (\text{condensate recovered} \div \text{steam production}) \times 100$$

In any one plant, this may have merit as a measure of whether the local situation is improving or regressing. However, this criterion is useless when comparing different refining and petrochemical complexes.

A refinery or petrochemical complex with little demand for heating/reboiling steam (e.g., a plant where an extensive hot oil system is used), and with high usage of steam for live stripping or process consumption (e.g., a hydrogen plant), will expect low recovery of condensate relative to steam produced. A more meaningful measure of condensate recovery is:

$$\text{Percentage of condensate recovery} = (\text{condensate recovered} \div \text{condensate produced}) \times 100$$

In the author's experience, different complexes range from a low (20%) recovery to a high (75%) recovery.

Vents to atmosphere. Continuous venting of LP steam to atmosphere has occurred in approximately 25% of the steam systems reviewed by the author. This is a costly practice, since every pound of steam vented represents the full cost of steam. Continuous venting is caused by either a fundamental steam mass balance issue, or by improper control practice.

Commonly, continuous venting is a mass balance "imbalance"; i.e., more LP steam is produced than used. The solution is to either find effective ways to use the LP steam, or to reduce the supply of LP steam. Supply reduction includes checking all letdown sources, the replacement of back-pressure turbines exhausting to the LP steam system, or finding other uses for waste heat used to generate LP steam.

While LP venting is typically assumed to be a steam imbalance problem, it can also be caused by an improper and unrecognized control issue. A shortlist of things to investigate includes:

- Actual setpoints for the various letdown and vent controllers
- Board operating practices regarding those setpoints
- Control system configuration
- Header pressure drops in and between different process areas
- Transient steam demand or supply
- Conflicting local process issues being protected.

It is important to recognize that LP steam controls may be scattered throughout the refinery. Local area practices are dictated by local area concerns. Commonly, there is no overall coordination; a "systems view" is critical to solving the problem. HP

LITERATURE CITED

1. Pelham, R. O. and R. D. Moriarty, "Survey plants for energy savings," *Hydrocarbon Processing*, July 1985.
2. *Improving Steam System Performance: A Sourcebook for Industry*, 2nd ed., February 2012.
3. Tanthapanichakoon, W., "Saving energy in multilevel steam systems," *Chemical Engineering Progress*, January 2012.
4. Dogan, A., "Using cold boiler feedwater for energy recovery," *Petroleum Technology Quarterly*, Q1, 2012.
5. Peterson, J. F. and W. L. Mann, "Steam system design—how it evolves," *Chemical Engineering*, October 14, 1985.