CONTROLLING SHELL AND TUBE EXCHANGERS

© Walter Driedger, P. Eng., 2000 May 20, 2010 April 24
First published in *Hydrocarbon Processing,* March 1998.

INTRODUCTION. Shell and tube heat exchangers are among the more confusing pieces of equipment for the process control engineer. The principle of operation is simple enough: Two fluids of different temperatures are brought into close contact but are prevented from mixing by a physical barrier. The temperature of the two fluids will tend to equalize. By arranging counter-current flow it is possible for the temperature at the outlet of each fluid to approach the temperature at the inlet of the other. The heat contents are simply exchanged from one fluid to the other and vice versa. No energy is added or removed.

Since the heat demands of the process are not constant, and the heat content of the two fluids is not constant either, the heat exchanger must be designed for the worst case and must be controlled to make it operate at the particular rate required by the process at every moment in time. The heat exchanger itself is not constant. Its characteristic changes with time. The most common change is a reduction in the heat transfer rate due to fouling of the surfaces. Exchangers are initially oversized to allow for the fouling which gradually builds up during use until the exchanger is no longer capable of performing its duty. Once it has been cleaned it is again oversized.

WHERE DO WE MEASURE? At the fundamental level, there is only one variable that can be controlled -- the amount of heat being exchanged. In practical situations it is not possible to measure heat flux. It is always the temperature of one fluid or the other which is being measured and controlled. It is not possible to control both since the heat added from one is taken from the other. Therefore the first consideration is to specify the place at which the temperature is to be kept constant. This is usually within a piece of equipment somewhere downstream of the outlet of one of the fluids. Assuming there is not much temperature change along the piping, the measurement may be anywhere between the outlet itself and the point of interest, perhaps at the base of a distillation tower. In cases where the measurement is being made downstream of a bypass valve, the further downstream, the better the mixing will be, and the more representative the measurement. On the other hand, too far down-stream may result in process dead time that can make control difficult. In cases where the "other" fluid is the one being manipulated, it is often quite sufficient to make the measurement directly downstream of the outlet nozzle of the exchanger.

WHICH STREAM DO WE MANIPULATE? The second consideration is which stream to manipulate. The complications arise from the fact that exchangers have four ports and involve two different fluids, either of which may change phase. The former feature alone allows eight different valve arrangements. Figure 3-1 allows the reader to figure them all out. The diagram assumes that it is the fluid on the shell side whose temperature is being controlled. As likely as not, it is the one on the tube side. It doesn't really make any difference to the control strategy. The real issue is which fluid is to be manipulated by the valves. For the sake of discussion we will term the two streams the "process" side and the "heat exchange medium" side. A complete tabulation of all the possibilities is:

a - Process side, outlet throttling.
b - Process side, inlet throttling.
c - Process side, bypass with outlet restriction.
d - Process side, bypass with inlet restriction.
e - Medium side, outlet throttling.
f - Medium side, inlet throttling.
g - Medium side, bypass with outlet restriction.
h - Medium side, bypass with inlet restriction.

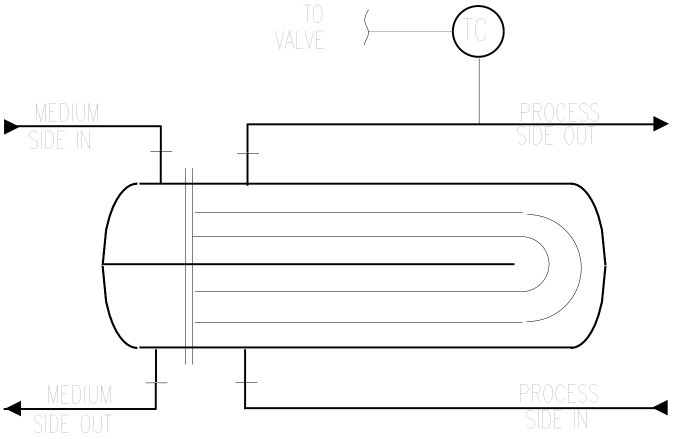

Fig. 3-1. A Shell and Tube Heat Exchanger

Among this profusion of alternatives, some must be better than others. The preferred choice depends, as always, on the particular situation.

There are a number of varieties of the basic shell and tube exchanger that can be controlled along similar lines. Plate exchangers consist of thin sheets of corrugated metal. The corrugations are formed to produce passages so that the two fluids pass in opposite directions on opposite sides of each sheet. The "shell" side and the "tube" side are essentially interchangeable.

Aerial coolers, sometimes called fin fan coolers, are similar to shell and tube exchangers except that they are all tube. The air blowing past the tubes can be considered to be in an extremely large shell.

THROTTLING THE PROCESS FLUID. It is quite meaningless to attempt to control the process temperature by throttling either the inlet or the outlet of the process fluid. The desired process flow rate is set by other requirements and these would be interfered with by manipulating the process flow. Temperature will change somewhat since flow reduction increases the residence time of the fluid and the outlet temperature will more closely approach the inlet temperature of the medium.

On the other hand, variations in process flow, caused by some external influence, is one of the major causes of temperature variation. It is often the reason why we must manipulate some other parameter to maintain constant temperature.

BYPASSING THE PROCESS FLUID. Process temperature can be controlled by manipulating process flow if a bypass is installed. As the outlet temperature rises (assume this is a heater), more fluid is bypassed around the ex-changer without being heated. As the two streams are blended together again, the correct temperature is achieved.

SPLIT RANGE. Bypass manipulation sounds simple but there are a few tricks to it. Firstly, there are two ways of arranging the valve controls: We can attempt to minimize pressure drop at all times, or we can attempt to keep the pressure drop constant. In neither case do we want to interrupt the total flow. If we wish to minimize pressure drop, a butterfly valve is the likeliest choice.

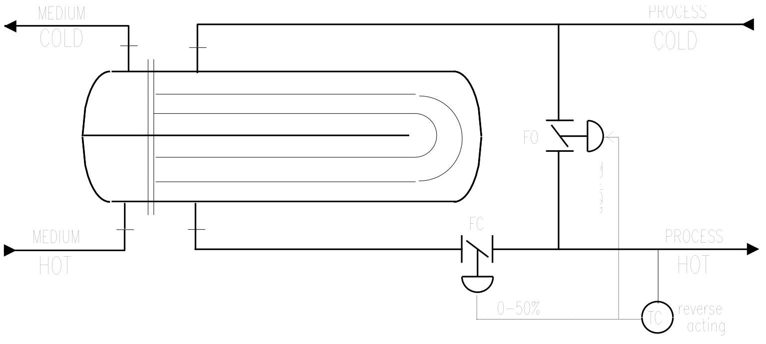

Fig. 3-2. Process Side Bypass with Restriction on the Outlet

However, even a wide open butterfly has some pressure drop. It may be greater than that of the heat exchanger itself. This means that even when the valve is wide open only half the flow, or less, will bypass the exchanger. To accomplish a greater degree of bypass, a restriction must be placed on the flow through the exchanger. The restriction should be adjustable since conditions change and we do not want more restriction than necessary. The easiest way to do this is with a hand valve. Since these valves are often in relatively inaccessible places, remote actuators may be added. Once that is done it becomes an obvious matter to arrange automatic controls so that once the bypass is fully open, the restriction valve starts to close, and vice versa.

This is, of course, a split range. The valve positioners, or I/Ps, are calibrated so that 0 → 50% signal opens the outlet and 50 → 100% signal closes the bypass valve. With this arrangement, at least one of the valves is fully open at all times and the effective C_v ranges from 100% to 200% of that of a single valve. It should be noted that by arranging for minimum pressure drop we must accept that the pressure drop, and consequently the flow, will vary as the valve positions change. Note also that there must be either a single I/P tubed to both valves, or two I/Ps (or positioners) must be wired in series. Either way, special care must be taken during construction and maintenance.

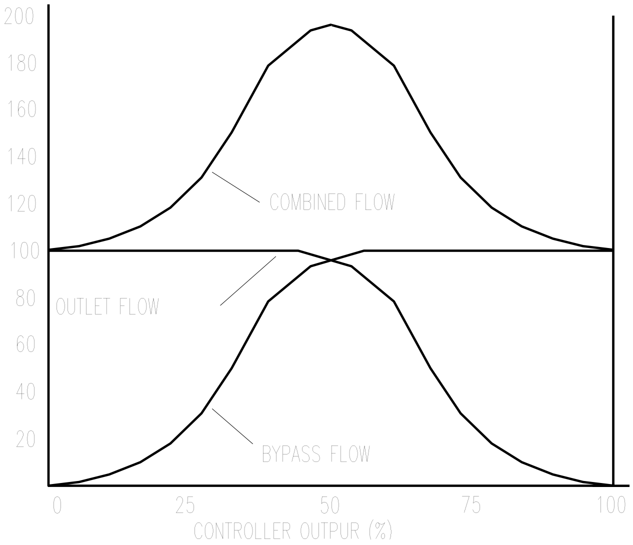

Fig. 3-3. Split Range Valves

The example shown in Figure 3-2 shows the Fail Open valve to be the one in the bypass. Let us assume that the process stream is being heated. The failure modes of the two valves is such that a signal failure to either, or both, valves will result in less heat being delivered to the process stream. Failure also means that the medium will not be cooled. The opposite failure response is easily arranged. It is a matter of choice. Once the choice has been made, the control action of the controller becomes a matter of deduction:

a) Assume that the process stream outlet is too hot. That is, it is above the setpoint.
b) Then the deviation of the controller is positive.
c) Assume the controller action is positive. This produces a rising valve output signal that will tend to open the outlet and close the bypass.
d) That would raise the temperature of the process steam. **Wrong!** The situation is getting

worse.

e) **This controller must be configured to be reverse acting.**
f) Now a rising outlet temperature will cause a falling valve signal.
g) That will open the bypass and close the outlet.
h) This will lower the temperature of the process stream thus bringing the measurement back to the setpoint.

The simplest way to carry out such a control action analysis is to trace around the loop from the measurement to the controller, to the valve, and back to the measurement. Assume the control action is positive. If the measurement is brought back to the setpoint, everything is OK. If things get worse, reverse the controller action.

Modern distributed control systems (DCS) have a built-in option for reversing the signal to the valve. If this is used, it is not necessary to take the failure mode of the valve into account when doing a control action analysis. First think positive then select a reverse output (**not** control action) if the valve is Fail Open.

An astute observer will realize that there is a possibility for both valves to be closed. If the outlet valve fails closed, the controller will sense a low temperature. Its output will then rise and close the bypass. It is very unlikely for the valve itself to fail when the signal is still good. Nevertheless, if this remote possibility is unacceptable, both valves must be made Fail Open. How can this be done? It can't! As long as both valves are driven by the same signal, their failure modes must be opposite because their effect on the process is opposite. The only way to solve this dilemma is to have two separate outputs from the controller, one direct, the other reverse. This is not a standard feature, but it can be easily arranged with a DCS. Here is the way to do it: Send the controller output to two calculation blocks. The first doubles the signal so that the range 0 → 50% becomes 0 → 100%. Any signal beyond 50% is ignored as the output of the calculation block cannot exceed 100%. The second block subtracts the signal from 100% and then doubles it. The range 50 → 100% then becomes 100 → 0%. Any signal below 50% is ignored as the output of the calculation block cannot fall below 0%. The two outputs are then send to their respective valves. Which output goes to which valve depends on the failure mode of the valves. Note that both valves now operate on 0 → 100% signals. Some models of DCS provide for scaling and linearization in the output modules. If this feature is available, separate calculation blocks are not required.

This method of achieving split range action has additional advantages besides allowing the two valves to have the same failure mode. It costs an additional output slot, a pair of wires and an I/P but the advantage is that the two valves are "self-contained" and do not require any special treatment with respect to wiring in series, mounting of I/Ps, or split range calibration.

OPPOSITE ACTION. The second arrangement attempts to keep the pressure drop, and hence the flow, as constant as possible across the entire stroking range of the valves. To accomplish this, we would like the outlet valve to begin to close the moment the bypass begins to open. It will be fully closed the moment the bypass is fully open. This is an "opposite action" arrangement.

A three-way "flow splitter" or "diverter" valve is frequently used to combine the functions of the two valves into one body. For butterfly valves a mechanical link may be installed to join two valves to one actuator so that when one opens, the other closes. It is also possible to use two separate valves with either a single I/P, two I/Ps in series, or two separate output signals. The description in the previous section on using calculation blocks to provide opposite control action for the two valves is valid, with one difference. The signal is not doubled and only one block is needed to reverse the signal.

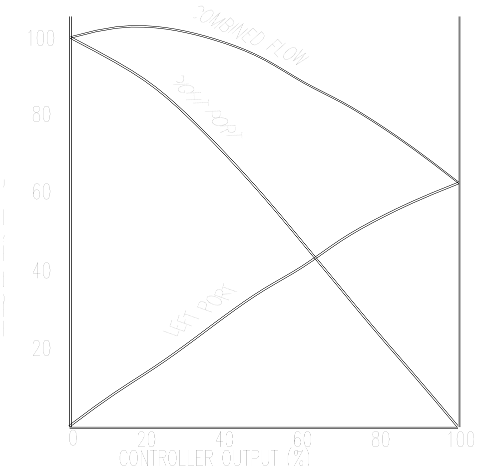

Fig. 3-4. Three Way Valve Curves

Some comments on three-way valves are in order. Figure 3-4 shows a typical characteristic. The flow on one side decreases gradually as the flow through the other increases. Ideally the flow through the inlet port is constant throughout the entire stroking range of the valve. Some three-way valves have a rather large gap between the two exit ports where both sides have more than 50% flow. Such valves should be avoided as they result in sloppy control because the loop gain near the midpoint is too small. The same consideration applies to the two-valve arrangement. The two valves should be chosen to have fairly linear characteristics so that the combined flow is approximately constant. It may be useful to plot the two characteristics on a piece of graph paper and add them up to see if the valve combination is satisfactory.

A close look at Figure 3-4 shows that the C_v of the two ports is not equal. This is almost always the case as the valve stem interferes with the flow through one of the ports. The port with the greater C_v should be open to the heat exchanger. Then the flow restriction caused by the exchanger will help to cancel out the difference.

MEDIUM SIDE THROTTLING. Avoid using a process side bypass valve with fluids that are being heated and have a tendency to break down or scorch. These include many food products and also petroleum products or other chemicals that may polymerize or coke at high temperatures. The problem is that the outlet temperature is a blend of the bypass stream and the stream through the heater. The peak temperature to which any part of the stream is exposed may considerably exceed that of the combined outlet. Over-done and half-baked don't average out! The extreme case is when the exchanger is on full bypass. The fluid trapped inside the heater will then be at the temperature of the heating medium.

The solution is to control the process temperature by throttling the heat exchange medium. In this case the heat available to the process is manipulated.

The example is a simple and rather straightforward application. Hot oil is being supplied to heat a process stream. It is desired to keep the process stream at a constant temperature. There is no reason to maintain the flow of oil in excess of what is needed -- it can be throttled to control the temperature. In this case the valve is placed

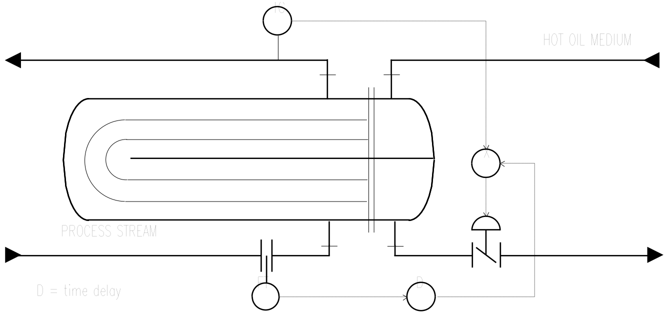

Fig. 3-5. Medium Flow Tube Side Outlet Throttling

on the outlet of the exchanger. The valve is not expected to handle a large pressure drop nor is tight shutoff of any particular value. Therefore a butterfly valve is quite acceptable. Furthermore its low pressure drop (high C_v) when wide open is an advantage.

The effects of inlet and outlet throttling are about the same, so secondary considerations come into play. In general it is a good idea to keep the pressure on a hot fluid to reduce any chance of dissolved gases bubbling out. A valve on the cooler end may be cheaper and will probably last longer. Leaks are less likely as the fluid will be more viscous than on the hot side. Thus it is best to throttle a heating medium on the outlet side.

It is rare that a heat exchanger cooling the process should have a reason not to use bypass control on the process side. Also, it is usually undesirable to throttle cooling water since it is at least mildly corrosive and is seldom clean. For this reason it is usually put through the tubes. In order to improve heat exchange and also to avoid the build-up of deposits and fouling, it is best to maintain its velocity.

If it should be necessary to throttle the cooling medium, consideration must be given to the possibility of boiling. (Remember that the cooling medium is the one being heated.) Assuming that boiling is not intended, but that the possibility exists, the valve should be placed on the outlet in order to maintain pressure on the fluid. In other words, inlet throttling is rarely used with single phase fluids.

CROSS EXCHANGERS. When the heat in one process stream is to be exchanged with another process stream, the flow on neither side may be interfered with while controlling the temperature. An example is when distillation tower bottoms are cross exchanged with the tower feed. The tower requires a high temperature at the bottom in order to function but the heat is not "consumed" by the process nor is it needed in the product. It is returned from the bottom product back to the feed. This is a common and extremely effective energy conservation measure.

As with all energy recovery arrangements, the key to success is to control the heat recovery without disturbing the process. That is, the flow of neither of the two process streams may be interfered with. The solution is to manipulate the heat transfer by bypassing one of the two streams around the exchanger. Most often control is exercised on the tube side. The failure modes of the valves are chosen to prevent overheating and flow blockage.

An interesting aside not directly related to process control: Counter-current cross-exchangers are widely applied in nature to prevent heat loss. One of these systems was first described by Herophilus in 300 BC. Arteries and veins pass very close to each other for some distance along the way to the extremities. In this way the heat in the warm blood on its way out is cross-exchanged to the returning cold blood. Biologists name this anatomical feature "rete mirabile" which means "magic net".

UNCONTROLLED HEAT EXCHANGE. In some cross exchange applications it is desired to recover all the heat (or cold) content of the product stream and to transfer it to the feed. In such cases the exchanger needs no controls at all. The feed stream usually has a second exchanger downstream of the first to boost the temperature to the required level. This exchanger is the one that is manipulated.

AERIAL COOLERS. As mentioned earlier, aerial coolers can be considered a special type of shell and tube exchanger in which the shell is the shell of the cooler. A large fan is used to blow air from below past the tubes. As with other exchangers it is possible to control the temperature by manipulating the process or the medium flows. The normal way to provide accurate temperature control is to use process flow bypass valves. In addition there are three means of manipulating the medium: Louver or damper control, fan pitch control and variable speed.

FAN PITCH CONTROL. This is an obvious means of controlling the temperature. It has the advantage of reducing horsepower as the cooling demand is reduced. As with every control technique, there are limitations. Firstly, the turndown is rather poor. This is especially important in a northern climate where it may not be possible to turn down the fans sufficiently in winter. The spinning blades still stir up the air even when the pitch is zero. Natural draft alone may provide more cooling than is required. Secondly, the pitch control mechanism can be a maintenance headache. The control system engineer must examine the equipment drawings to be sure that the mechanisms allow easy access for lubrication and repair. It is a wise idea to put separate I/Ps on each fan. Long strings of tubing with many tees make leak detection a nightmare. Each I/P requires a separate output from the control system as the current loop will not work if there are more than two in series. Note that if a single controller drives multiple fan pitch controls, the process gain of the loop is proportional to the number of fans in service. If the controller is tuned with only half the fans running, it may go unstable when the rest are turned on. A controller that is tuned with all fans running will be sloppy if some are turned off. It is, of course, possible to configure automatic gain compensation within a DCS. (Remember to check for divide by zero when no fans are running.)

VARIABLE SPEED. Fully variable fan speed control is becoming more common on aerial coolers. The fan motors are often quite numerous but not very large. However, the cost of the electronics has come down considerably in recent years. One way to cut costs is to connect both fans of one bay to the same set of VFD electronics. On the other hand, two-speed fans have always been quite common. This is especially true in climates with extreme seasonal temperature swings. Reducing the fan to half speed results in an 85% reduction in electric power demand. (Remember that electrical power varies as the cube of fan speed.) Cutting the speed of an electric motor in half requires only a reconnection of the wiring to a multipole stator. This can be accomplished by having two electrical starters wired in different ways. The increase in cost is not very large.

A variation of the two-speed motor is to arrange for reverse flow. This can be extremely useful in climates where icing is a problem. Reversing the flow blows warm air through the inlet louvers and serves to melt any accumulated ice. This should be done before ice build-up is too large or large chunks of ice may be sent crashing down onto other pieces of equipment or even personnel.

LOUVER CONTROL. Automatic louver control has similar problems as fan pitch control. There is an additional problem of hysteresis. The louvers seldom move smoothly for long and it becomes very difficult to maintain stable control as dirt and wear accumulate over time, especially in sandy or dusty environments.

In climates with strong seasonal temperature swings, it is possible to stabilise the air temperature and prevent icing by controlling internal recirculation. The exchanger is fitted with a duct leading from the top outlet to the bottom inlet of the unit. Dampers are placed in this duct and at the air intake. A temperature controller senses the air above the fan and controls it by opening the recirculation duct and simultaneously closing the intake. Note that an opposite action arrangement, as described above, is appropriate. Figure 3-6 shows a possible arrangement using both outlet louver control from the process and recirculation control off the internal air temperature.

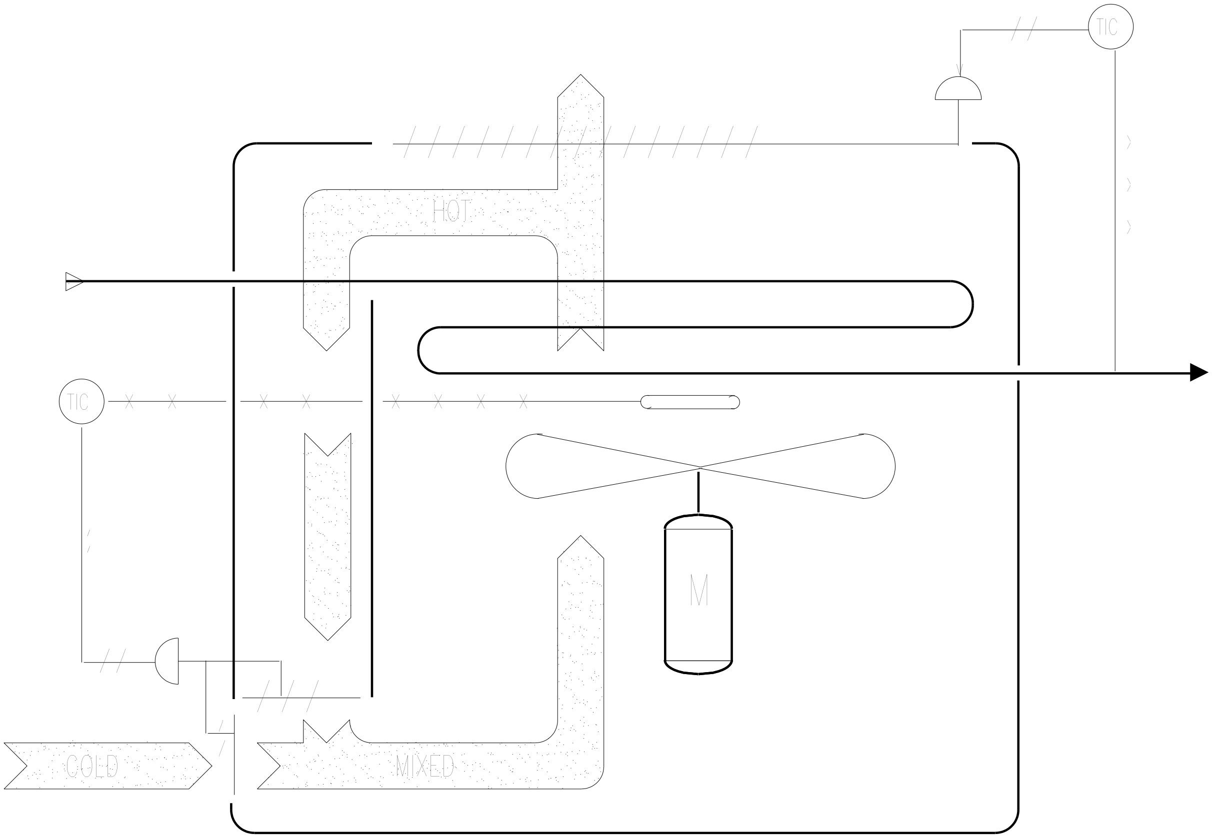

Fig. 3-6. Aerial Cooler Louver Controls

ADVANCED TRICKS – FEEDFORWARD. Large heat exchangers have both dead time and considerable thermal inertia. These two factors can make control difficult. Feedforward can be usefully applied if load changes are a problem. Since the heat demand is proportional to the process flow rate, other things being equal, a flow rate measurement can be used. Figure 3-7 shows a typical arrangement. Note that the output of the TC is multiplied by the flow signal. That is because the heating medium flow rate must be roughly proportional to the feed flow. This works best if the installed characteristic of the valve is linear. If the exchanger is very large, it may be necessary to insert a lag or

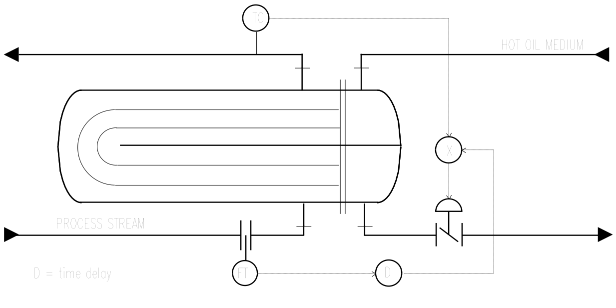

Fig. 3-7. Feed Forward Control

some other form of delay into the flow signal to prevent it from acting too soon and causing a reverse spike to appear in the temperature. Note that the dead time is inversely proportional to the flow rate and some "typical" value must be used. Some brands of DCS have the option of a variable delay time. This allows delay to be inversely proportional to flow rate.

TEMPERATURE OPTIMIZATION. Another "Advanced Trick" involves optimization of a fired heater. Heat is being supplied to the reboiler of a deethanizer as shown in Figure 3-8. It is required to keep the temperature at the bottom of the tower constant. The heating medium is hot oil which is being heated by a fired heater and circulated by a pair of pumps. Since the tower bottoms is being boiled, and is also very clean, it goes on the shell side. The oil goes through the tube side where the outlet is throttled by a butterfly valve. A position transmitter has been added to the valve. Its output goes to a Position Controller with a setpoint of about 80% open. The output of the Position Controller is cascaded to the setpoint of the Temperature Controller of the furnace. The effect is to maintain the furnace, and the hot oil, at the lowest temperature consistent with the heat demand of the tower. It works as follows:

a) As the heat demand rises, the valve opens further.
b) When the valve is open beyond 80%, the setpoint to the furnace Temperature Controller is raised.
c) As the temperature of the hot oil rises, the valve closes to near the 80% value.
d) As the heat demand of the tower falls, the valve closes below 80%.
e) As the valve closes, the setpoint to the furnace is lowered until the valve is once again at its 80% target.

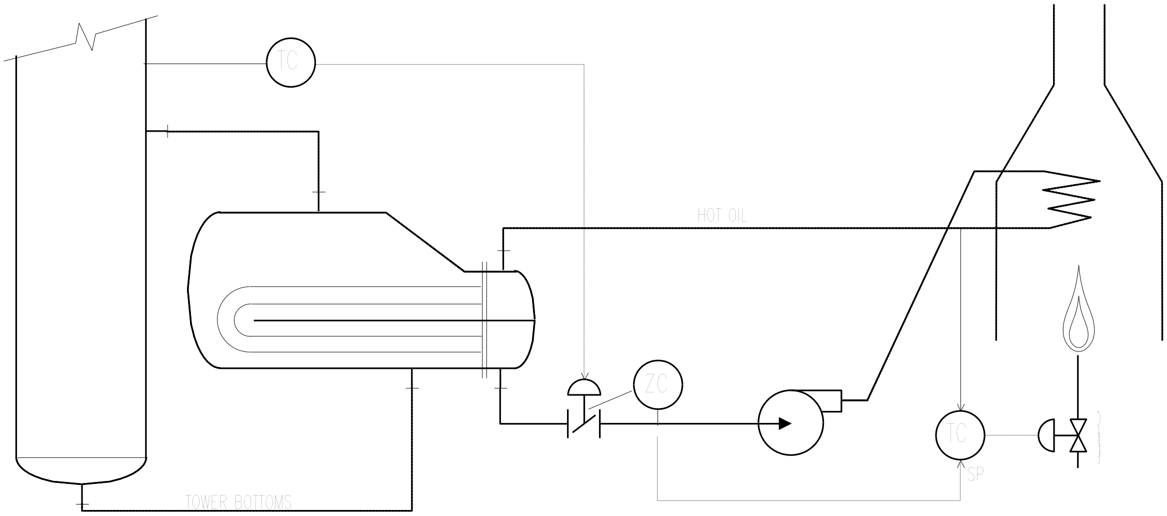

Fig. 3-8. Furnace Optimization

In this way the temperature of the hot oil system is kept at its lowest acceptable value and a minimum of heat is lost by the furnace or the piping.

COMBINATION CONTROL. Sometimes a heat exchanger is used to heat, or cool, a fluid whose total flow is being controlled by some other parameter. The most straightforward way of controlling this is to use a three-way valve, or two butterflies, to control the heat exchange and to use another valve to control the total flow. The flow control valve must be on the common line either upstream or downstream of the exchanger. This arrangement has two valves in series and cries out for a way of eliminating one of them. If the positions of the inlet and bypass valves are controlled separately so that the total C_v is controlled by the Flow Controller and the difference between the C_vs is controlled by the Temperature Controller, complete control can be achieved with only two valves. Figure 3-9 shows how this can be done.

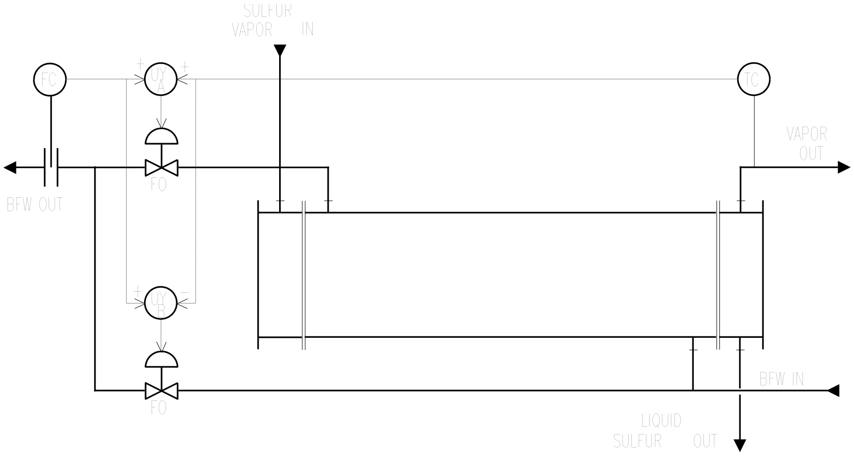

Fig. 3-9. Combined Flow and Temperature Control

The example uses boiler feedwater to cool a sulphur condenser at the same time the water is being preheated. The sulphur vapour, being the more difficult fluid, is in the tubes. The water is in the shell. Since we want to make certain that the water does not boil, we will put the valves on the outlet side. The valve controlling the outlet of the exchanger receives a signal equal to half the **sum** of the two controller outputs. The valve controlling the bypass receives half of the **difference** between the two controller outputs. Assuming that the installed characteristic of both valves is linear, the combined flow of the two valves is then dependent entirely on the Flow Controller. The difference between the two flows is dependent on the Temperature Controller. In this particular situation it is desirable that both valves are fail open. If the failure mode of either, or both, valves is fail closed, the signs of the summing/scalers UY-A and UY-B will have to be changed to give the proper result.

EQUIPMENT PROTECTION. The usual shell and tube exchanger has no moving parts nor any in-put of external energy. There are few machinery protection issues. Severe corrosion is sometimes a problem. If so, corrosion detection devices may be installed. These consist of a thin wire or film of the same material as the exchanger. The wire is held in a holder that is inserted through a nozzle into the exchanger. Two electrical connections are accessible from the outside. When the resistance is measured, the extent of corrosion can be determined directly. These devices are not normally connected into a data logging network. The usual practice is to make the measurements with a portable monitor on a regular basis. Intrinsically safe monitors are available for hazardous locations.

Aerial coolers require protection from the energy introduced by the electric motors. The most serious hazard is a thrown blade. The resulting vibration is quite severe and can cause extensive damage. A simple seismic vibration switch mounted on the structure that holds the lower bearing of each fan is quite sufficient. It works by having a small weight held in place by a magnet against the force of a spring. A "bump" dislodges the weight from the magnet and allows it to open the shutdown contact. The usual method of "calibration" is a light whack with a hammer. A button allows the operator to reset the switch by pushing the weight back against the magnet. Switches with remote electrical reset can be bought but it is always best for an operator to look at the machine and determine the cause of the shutdown before restarting the equipment.

Precautions must be taken when reversing a motor that has been running. Such a change is a considerable shock to the machinery. The usual approach is to provide a time delay interlock so that sufficient time has elapsed to be certain that the fan has stopped rotating before the motor can be started in the opposite direction. If this is not done the fan will most likely trip on vibration. The nuisance of resetting locally mounted vibration switches will encourage the operators to be more careful in the future.

SAFETY. Overpressure is the only common safety issue affecting shell and tube heat exchangers. They are pressure vessels and as such are subject to the same codes and practices as other pressure vessels. That means the ASME Boiler and Pressure Vessel Code, Section VIII, Pressure Vessels, Parts UG-125 to 136[1] dealing with pressure relief devices. This specification gives very clear guidelines concerning all aspects of pressure relief requirements and application.

Pressure relief must be provided for both the shell and tube sides. If the source of overpressure is from upstream, the relief valve for that stream is best placed on the inlet. Otherwise it does not matter much whether it is on the inlet or outlet so long as they are inside any control or isolation valves. It is not sufficient to put minimum stops on the valves as these are easily altered. Even if the stops are welded in place, the valve may be replaced at some future date and the modification forgotten. If careful analysis shows that there are no process, fire, or failure conditions that could possibly require relief valves, it is still strongly advised to install thermal reliefs on both sides of any exchanger that is capable of being blocked in. It may be argued that the fluid is gas or that the process is not capable of adding heat to the blocked in exchanger. This argument overlooks the various unanticipated conditions that may arise during testing and maintenance. A worst case scenario: A cooler was taken out of service and steam cleaned. No one had drained the cooling water which expanded in the tubes and ruptured the joints. True, good maintenance practice would have prevented this incident. But then an NPS ¾ relief valve would have provided a permanent solution and would have cost a lot less than the damage caused by its absence.

ACCESSORY INSTRUMENTS. Since the purpose of a heat exchanger is to transfer heat from one fluid to another, instrumentation must be provided to check that this is happening. A thermometer is required at each inlet and outlet. TEMA[2] recommends NPS ½ nozzles on each of the four major nozzles. In practice they are not always useful. Firstly, in the process industries, a NPS 1 threaded connection, or even an NPS 1½ flange is the minimum allowed for thermowell connections to piping or vessels. Secondly, exchangers are often installed in such a way that thermometers on the nozzles are inaccessible without ladders or platforms. This is especially true if they are stacked. A more useful approach is to cancel the TEMA connections and to provide appropriate connections in the piping. (I almost broke my neck once trying to read a thermometer at the top of a stacked heat exchanger. My own fault, of course. I should have put it in a better place, I should have insisted on a platform, I should have used a proper ladder, I should not have crawled around on equipment with snow on it.) "Every angle" thermometers are the only kind to get. You never know exactly where they are going to end up.

It is not unusual to have a slip-in butterfly valve installed directly on an exchanger nozzle. A thermowell located in the nozzle may jam the valve.

On every project there is someone trying to save money. Eliminating the thermometers from the thermowells is often a candidate for dubious cost cutting. Consider the installed cost of the piping connection and the thermowell along with all their associated documentation. Then consider the cost of a thermometer. Remember that a plug must be provided for the empty well. The savings are negligible. Readily accessible temperature readings can eliminate long discussions, hypothesizing and delays when a process mysteriously doesn't work right.

The primary ailment of heat exchangers is plugging and fouling. The diagnosis is based on differential pressure. For this reason TEMA[2] recommends an NPS ½ pressure connection on all four nozzles. These connections suffer from the same deficiencies as the ones for thermometers. The solution is the same: Put the right connections in the best place in the piping. A shutoff manifold with a spare port that can be used for a differential pressure indicator should be used. The problem is that the differential pressure may be only a small fraction of the static pressure. A 10 psi difference cannot be read from two pressure gauges with 1000 psi scales. (War story: During commissioning, a cooler was found to have an inlet pressure of 595 psi and an outlet pressure of 586 psi. The solution to this strange violation of the pressure/flow relationship was to switch the two gauges.) As with thermometers, it is poor economy to eliminate pressure gauges in favour of pressure connections only. Just see how long it takes to find two pressure gauges of the right range when you need them.

When two exchangers are stacked in series, they are often connected flange to flange. In this case, it is not necessary to have PIs and TIs on both mating nozzles. A single set is quite sufficient. Unfortunately, it may not be possible to put them in an easily accessible location. It may also be necessary to modify the nozzles to conform to project standards.

PARALLEL HEAT EXCHANGERS. Aerial coolers may be viewed as a number of coolers in parallel. A single thermometer at the inlet is sufficient but a separate one at each outlet to the header is essential. There is absolutely no other way to identify individual plugged or fouled sections without taking the whole thing apart.

Heat exchangers in parallel do not share the flow equally. Symmetrical piping is like the perfect life: at best a pious intention. (Some people think it is more like Santa Claus.) The problem is that an initial flow imbalance can grow. If a reduction in flow causes fouling which further restricts the flow, a positive feed-back loop is set up which can cut one exchanger entirely out of circulation. Some way must be found to force the flows to balance. Unfortunately, it is meaningless to attempt to control any variable without measurement, and flow measurement is expensive. Not only are the instruments expensive in terms of installed cost and maintenance, the required piping arrangement is also expensive. Thus automatic flow balancing is rarely installed except in extremely critical service such as a furnace with multiple tube passes.

If a bank of parallel exchangers is controlled using a bypass it is only meaningful to have a single bypass for the entire bank. If the temperature of stream A is being controlled by throttling stream B, every exchanger must have its own valve or imbalance is sure to result. In fact, a separate control loop on every exchanger is probably a good idea. Care must be taken in the location of the temperature sensors so that each senses only the contribution of the exchanger it is controlling. An example is a pair of reboilers at the bottom of a tower.

REFERENCES

1. ASME Boilers and Pressure Vessels Code, Section VIII, Pressure Vessels, Division 1, Unfired Pressure Vessels, Parts UG-125 to 136, Pressure Relief Devices. (http://www.asme.org/catalog/)

2. Standards of the Tubular Exchanger Manufacturers Association. http://www.tema.org

3. API STD 661, Air-Cooled Heat Exchangers for General Refinery Service. (http://www.cssinfo.com/apigate.html)

4. API STD 660, Shell and Tube Heat Exchangers for General Refinery Service. (http://www.cssinfo.com/apigate.html)