SPECIAL COLLECTION FOR WORLD REFRIGERATION DAY, JUNE, 26 2021

ASHRAE JOURNAL

THE MAGAZINE OF HVAC&R TECHNOLOGY AND APPLICATIONS ASHRAE.ORG

Refrigeration Applications

Fundamentals

An Andy Pearson Collection

Introduction to a Special Collection for World Refrigeration Day, June 26, 2021
Refrigeration Applications, Fundamentals: An Andy Pearson Collection

Start at the Very Beginning

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

The scientific principles that govern the operation of a refrigeration system are not particularly difficult to follow, but people who have been immersed in the world of mechanical cooling for too long develop a way of using language that is every bit as isolationist as the medical or legal professions. The difference is that we do not rely on Latin to keep the general public at bay; we use a Creole combination of physics, thermodynamics and Olde English. This makes perfect sense to us but can be baffling and even intimidating to a novice. Pressures are measured in units of temperature, many different words are used for the same item, efficiencies can be greater than 100% and it is even possible to have “solid liquid.”

This series of columns was commissioned by ASHRAE’s Refrigeration committee, who, in 2012, asked for a monthly article in the *ASHRAE Journal* to try to explain the wonderful world of refrigeration to the uninitiated. You, the reader, must be the judge of their success.

Andy Pearson, Ph.D., C.Eng., is the group managing director at Star Refrigeration Ltd., an industrial contractor based in Glasgow, Scotland, UK. He is an ASHRAE Fellow and has served on committees for ASHRAE, IIAR, the International Institute of Refrigeration, BSI, CEN and ISO. He is currently President of the Institution of Engineers in Scotland and has contributed over a hundred columns to the *ASHRAE Journal*.

The Temperature Lift

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This column is the first in a series that explores some of the concepts associated with refrigeration and heat pump systems. It is aimed at people who are not immersed in refrigeration technology, and who perhaps do not fully understand the jargon. I hope seasoned hands will find the series useful, too, even if only to help explain a point to their less experienced colleagues.

From "The Music Box" by Hal Roach Studios

Laurel and Hardy consider the complexities of the refrigeration cycle and temperature lift while moving a piano up the hill.

In the natural world heat flows, as a rule, between adjacent objects if one is warmer than the other. The flow of heat is always from the warmer object to the cooler one, because heat is a form of energy, and the energy is being more evenly distributed.

It is a bit like a bathtub full of water; if you push the water up one end of the tub and then release it, the water will redistribute itself so that the level is equal across the whole tub. Left to its own devices, heat, like the bathwater, always finds a level. People who want to sound smart call this entropy, but the name is not important; it is just smoothing out the bumps.

Refrigerators and heat pumps are mysterious because they seem to break this rule. They keep a temperature difference between the cold end of the system (inside the fridge) and the hot end (down the back), which is like being able to keep all the bathwater down at one end of the tub. On closer inspection, it becomes apparent that the fridge is not, in fact, breaking the rules at all. To understand this, we need to appreciate the concept of temperature lift.

To make stuff colder or to keep it cold, we need to draw heat energy out of it. This is done by placing the stuff in surroundings that are colder than the stuff. The heat energy flows from the warmer stuff to the colder surroundings. To get rid of that heat energy, we need to make the energy warmer than the outside of the refrigerator. This is done by working on the heat—putting more energy into it—so that it can flow to the outside.

How we turn the cold energy into warmer energy will be covered later in the series, so do not worry about that right now. Just consider the four temperature levels described. The stuff in the fridge is one temperature and the surroundings are colder. The heat is then lifted to a high temperature and the outside of the fridge is colder. To get the heat from the temperature of the stuff in the fridge up to the temperature outside the fridge, it is necessary first to drop the heat down to the temperature of the cold end of the fridge, and then raise it (by doing work on it) to the temperature of the hot end of the fridge.

If you do not believe that your fridge has a hot end, then feel round the back of it while it is running. The difference between the cold end and the hot end is the temperature lift for the refrigerator circuit. However, the difference between the stuff in the fridge and the outside is the temperature lift for the total system. The refrigerator works because its temperature lift is larger than the required temperature lift for the total system.

Imagine my house, which is halfway up the side of a steep hill between two terraces. There is a path sloping downhill from my basement door to the street below and another path further up the hill that slopes down from the upper street to my front door.

I have a piano in the basement that I want to place upstairs. I can roll the piano out the basement door and down to the street. I then need to do work on the piano to lift it up to the top of the hill, where I can roll it down to my front door.

Wheeling the piano down each of the paths is like heat transfer inside the fridge and from the fridge to the outside. Lifting the piano up the hill from the lower street to the upper street is like the work done on the heat within the refrigeration system.

While the total removal only lifts the piano from the basement to the front door, the work done in lifting it takes it from the foot of the lower path to the top of the upper path. If that height is more than necessary, for example if I take the piano farther up the hill than I have to, then I need to do more work to lift the piano and the removal is less efficient than necessary.

In a refrigerator the total system lift is usually set by the system specification: from the desired temperature to the ambient. The temperature lift within the refrigeration system is a function of the performance of the heat exchangers. The key to refrigeration system efficiency is ensuring that the temperature lift within the refrigeration system is as small as possible. If the total system temperature lift also can be reduced, that is even better.

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■

Temperatures, Pressures, & Refrigerants

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This column is the second in a series exploring refrigeration and heat pump concepts without using jargon.

Most stuff we come across on a daily basis has a boiling point and a freezing point. For example, water at sea level will boil at 212°F (100°C) and freeze at 32°F (0°C). The subtle disclaimer "at sea level" is important because the boiling temperature for all stuff is dependent on pressure, so it varies at different conditions.

This is how pressure cookers are able to cook food much faster than in an open saucepan. The water inside the pressure cooker is boiling at a higher temperature and so less time is required for cooking. Conversely, it also explains why it is difficult to make a decent cup of tea up a high mountain; the water boils, but it is not at a high enough temperature to properly infuse the tea leaves. To put some numbers against this, if the pressure in the cooker is double the normal atmospheric pressure then the water will boil at about 250°F (121°C). At the top of Mount Everest (29,029 ft [8848 m]), water boils at about 160°F (71°C). This is definitely not hot enough to make tea.

The substances that are used as working fluids inside refrigerators and heat pumps all have a fixed relationship between their pressure and their boiling point. Some of them, such as R-11, operate at very low pressures and others, for example, carbon dioxide, run at much higher pressures, for the same boiling point. For most other substances used as refrigerants, the pressure at a given temperature is somewhere between these two extremes.

Photo courtesy of Royal Geographical Society (with IBG)

Edmund Hillary wonders if he will ever have a good cup of tea again.

In last month's column, the temperature lift for the refrigerator circuit was described. The way in which work can be done on the heat taken out of the stuff in a fridge is by pressurizing the working fluid so that its boiling point rises. In this way, the same working fluid can be boiled at low pressure (when it is colder than the stuff in the fridge) and reliquefied at high pressure (when it is warmer than the outside of the refrigerator). While the working fluid is boiling or reliquefying it does not change temperature or pressure but the amount of liquid in the total fluid content changes. If liquid is at its boiling point, for example just at the beginning of the boiling process or just at the end of the liquefying process, then this is called *saturated liquid*. At the opposite end of these processes, when the liquid is all turned to gas, or the gas is just on the point of liquefying, this is called *saturated vapor*. Calling it *vapor* instead of *gas* is just another example of refrigeration people wanting to sound smarter than they are. Maybe they don't like the thought of admitting that they have saturated gas.

For most situations within the refrigeration system, to know the condition of the working fluid, it is sufficient to know the temperature and the pressure. If these are known, then everything else about the fluid can be calculated, or looked up in tables, or most likely these days read off a computer screen or an app. However, when the working fluid is boiling or liquefying, the boiling point and pressure are locked together by their fixed relationship. Knowing the temperature and pressure of the fluid is not enough to determine the other properties. This is a bit like when you drive your car into a tunnel and the satnav system loses contact with the satellite. You know the car is in the tunnel but until it emerges from the other end and can be seen again it is not possible for the satellite to tell exactly where in the tunnel it is.

Unfortunately, for the interested operator who wants to know what's happening all around the refrigeration circuit, the two "tunnels" in the system are where most of the heat transfer action takes place. In the evaporator the working fluid is boiled from liquid to gas, and in the condenser it is reliquefied. Most of the refrigerant in the circuit will be in one or the other of these places, and its other properties are hidden from view by the tunnel effect. The normal way of describing the condition of the fluid is as a percentage quality. If the fluid has a vapor quality of 10%, then 90% of it is liquid, and 10% is gas. If it has a vapor quality of 80%, then only 20% of it is liquid. The vapor quality is the percentage of the weight of working fluid in the vapor state. However, because the density of liquid is much higher than gas, the liquid only takes up a very small amount of the space.

Andy Pearson, Ph.D., C.Eng. is group managing director at Star Refrigeration in Glasgow, UK. ■

Sometimes It's the Simple Things

Improving Efficiency

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This column is the third in a series exploring refrigeration and heat pump concepts without using jargon.

There are many ways that the efficiency of a refrigeration system might be compromised. Turning this around: it is fair to say that there are many ways in which the efficiency of a refrigeration system might be improved. Funnily enough, sometimes the simplest things are the least likely to be done in practice.

The definition of efficiency I like best is that it is the ratio of "what you want to do" to "what you need to do." Fridge guys think in terms of the cooling effect and the work done on the gas, but this is a bit blinkered and so sometimes misses the point.

For a pie maker what he wants to do is make pies; for a brewer it is beer. In a distribution warehouse, "what you want to do" might be measured in pallet movements and in a data center, it may be uptime or processing speed.

"What you need to do" is a bit more universal; generally it is work input, measured in kWh, which after all is what you pay for, so not an unreasonable thing for an operator to be interested in. A factory that does not measure its output is unthinkable, and yet few of them measure the electrical input to their refrigeration system in a meaningful way; and even fewer correlate that input to their output in order to track their overall efficiency.

Cold stores give a very good example of the way in which fridge guys can be too blinkered. If the cold store doors do not fit very well and the room temperature is regularly a few degrees higher than it is supposed to be, the temperature lift required of the refrigeration plant will be reduced. That is supposed to be a good thing, right? Yes, but if the result of this additional heat load is that the compressor runs all day and all night then the kWh required will ultimately be higher than if the doors were fixed, the store got down to temperature and the compressor switched off.


© CBS Studios

Captain Kirk hears that his pies per kWh metric has gone down 10%.

Unless the store operator is measuring what he wants to do (hold pallets at the required temperature) and comparing it with what he needs to do (use electricity), he might be fooled into thinking he has a very efficient system because the temperature lift is reduced and the compressor COP is good.

Efficiency improvement is therefore a very simple three-step process.

Step 1 is to measure what you want to do and divide it by what you need to do (expressed in kWh electrical use for the whole refrigeration system). This will give a rather odd metric, which might be pies per kWh or beers per kWh or pallets per kWh, but it will mean something to the business.

Step 2 is to find things to do that increase that metric, preferably by reducing the number under the line: the "per kWh." The biggest user of kWh is the compressor so that is a good place to start. Clean the suction filter in order to reduce the temperature lift that the compressor sees. If the machine is old and worn, give it a thorough overhaul to reduce internal inefficiencies, perhaps even replacing the rotors of a screw compressor to reduce tip seal losses.

Step 3 is to compare the metric before and after the reduction and then take the difference to the boss and ask for a pay raise. It is essential to build up a record over time of the results of these measurements. This will help to spot when something goes wrong whether it happens quickly or is a gradual decline.

Beyond the compressors, there are many more users of kWh, mainly fans, pumps and heaters. They can all be targets for reducing the kWh figure.

A more subtle way to reduce kWh is to look at the run time of the plant and figure out ways to reduce it without compromising on the number over the line. Systems with hot gas defrost are a favorite of mine for this. All will be revealed in a later article.

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■

Load Matching and Capacity Control

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This column is the fourth in a series exploring refrigeration and heat pump concepts without using jargon.

In the May column we considered ways to measure refrigeration system efficiency and skirted round the issue of how to improve it. The simple answer to this is to find ways to reduce the electrical input required to achieve the necessary throughput of pies, beer or pallets without reducing the quality of the product. Since pretty much everything in the system is variable (including the quantity of product being processed, the cleanliness of the heat input and output devices and of course the weather), it follows that the cooling demand required of the plant will not always be the same.

Most systems very rarely run at maximum capacity. They are designed for a hot summer day when the overall temperature lift will be at a maximum. If the cooling load is a function of ambient temperature, then warm weather is a double whammy: the load is high and the temperature lift is also high.

Things that freeze stuff in batches (including ice rinks and thermal storage systems as well as food factories) have a high heat load at the start of the freezing process, but this quickly reduces as it becomes more difficult to take out the remaining heat from the middle of the stuff. Most refrigerating systems, including water chillers, operate most of the time on less than 70% load, and some of them may usually operate on less than 50% load.

To decide on an efficiency improvement strategy, it is necessary first to determine what is expected of the plant as the load varies. Many compressors become progressively less efficient as they reduce capacity, particularly screw compressors with slide valve control. Others will tend to reduce electrical consumption in proportion to load reduction, but might only be able to change the load in quite large steps.

With some system designs it is even conceivable that the plant becomes more efficient as it offloads, particularly if the compressor speed can be changed to match the requirement, or if the full surface area of the heat input and output devices can be used when only a portion of the full cooling capacity is required.

Tonto explains how to loosen the reins to improve efficiency.

If your efficiency improvement strategy is not based upon a right understanding of this key behavior of the plant, then it is likely to give poor results.

If the compressor is most efficient when on full load, then the cooling requirement can still be matched by arranging the system so the compressor runs at full load for a portion of the time and then switches off.

The Earl of Essex, commander of the Parliamentary forces at the start of the English Civil War in 1642, is said to have remarked, "stone dead hath no fellow," when calling for the death penalty for the Earl of Strafford, principal advisor to King Charles. Certainly, the compressor cannot consume any less energy than when it is not running, so a system that is most efficient on full load will give the highest pies per kWh if it alternates between running flat out and switching off. However, too many stops and starts per hour will also hurt the efficiency (and perhaps the reliability) so some folk prefer their plant to match the heat load exactly in order to avoid stops and starts.

The key to combining both strategies is to make sure that the temperature control band is not too narrow. A wider band will allow the compressor to load up to its peak efficiency and run for a reasonable length of time. It will then switch off, and will stay off until the control temperature rises up through the band and triggers a restart. For example, if the control band is only 1°F (0.6°C) wide and the plant is 50% loaded, then it could run all day on 50% of full load (which would typically consume as much as 70% of full power). If it was loaded up and then switched off, it might cycle on and off again quite rapidly.

For example, if one on-off cycle takes five minutes, there will be 12 starts per hour, which would be inefficient and potentially damaging. However, simply changing the control band to make it ±3°F (±1.7°C) would cause the compressor to start, load up to 100% and run like that for 15 minutes before switching off for a further 15 minutes—only two starts per hour. It is a common mistake to think that tight temperature control equals good practice in all respects.

If it is okay to loosen the reins a little, plant efficiency may be much improved, provided product quality is not affected. With chilled water systems, fitting a larger tank can have the same effect as widening the control band, provided it does not push the operating temperature of the chiller lower than it needs to be.

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■

Reducing Leakage

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This article is the fifth in a series exploring refrigeration and heat pump concepts without using jargon.

System designers go to great lengths to ensure their equipment is capable of dealing with the sorts of variation in load that were described in the June column. This includes the changes in the amount of working fluid that is contained in the cooling part of the circuit. If the cooling requirement is large then lots of liquid will be boiling and most of the space in the cooler will be taken up with gas. However, when the load reduces and the boiling is less vigorous then more of the space is filled with liquid, which has to come from somewhere else in the system. This is why systems need a tank with a lot of liquid in it, so that the plant does not fail to perform when the load is light and more liquid is required in the cooler.

This leads to two specific difficulties for the operator of the system. First, it is very difficult to say what is the right liquid level in the tank because that depends on what is happening in the rest of the system. Second, if some of the working fluid leaks out of the system its disappearance might not be immediately obvious. However, just at the time it is most needed it will not be there, and the system is likely to misfire in one way or another. A shortage of working fluid might cause the system to run less efficiently than it should or to stop working altogether.

Some people are more passionate about leakage reduction than others.

Leakage of working fluid causes many other problems. The fluid might be toxic, or flammable, or smelly or perhaps harmful to the environment in some less obvious way. Even if it is non-toxic it might suffocate people who are working near the leak or who walk into a room in which the oxygen has been displaced by colorless, odorless refrigerant. If a system is "short of gas" then someone has to purchase some more, someone has to get in their truck and drive to the site and someone has to return the empty cylinders once the job is done. However, if the leak is not found and fixed then they will be back next month, and the month after, and the month after that. In some countries "topping up" a system that is known to be leaking is illegal and carries a heavy fine if convicted.

Leaks are therefore bad news and should not be allowed to happen. The good news is that most of them can be prevented, with a bit of care, some forward planning and a dose of common sense. Most leaks come from broken pipes, worn or faulty seals and loose fittings. Pipes can be broken by the effects of excessive vibration, either through fatigue failure, through abrasion of the pipe against another object or through work-hardening. Pipes also fail due to corrosion, particularly under insulation if the vapor seal is not maintained. Many of these causes can be eliminated at the design stage of the system by using suitable materials, ensuring adequate clearances around pipes and avoiding screwed fittings wherever possible.

For small-diameter pipes stainless steel is a much more durable material than copper, and is not significantly more expensive when the total cost of the system is considered. Anyone who says they care about leaks but is still using copper gage lines on their plant is not credible.

Maintenance also has a big part to play in the reduction of leaks, including checking that the vapor seal on insulated pipe is in good order. Inspection of uninsulated pipe for signs of wear and corrosion—and treating them as soon as they are found—is probably the single most effective anti-leak measure. Spotting and eliminating excessive vibration is another key element of ongoing leakage prevention.

In this case it is useful to take benchmark vibration readings from time to time because gradual increases might not be noticed in a day-to-day inspection routine. It is also important to take the readings under a variety of operating conditions, particularly because vibration levels might be higher on part load than full load, especially if the system is speed controlled.

Ultimately, the key to leakage reduction lies in your mental attitude. If you lived in the superintendant's house at the hydroelectric plant in the lee of the Hoover Dam, you would be passionate about leakage prevention. There are two tricks that help keep plants in a leak-free condition. The first is to tell yourself that there is definitely a leak and your job is to find it. It is amazing what a difference this makes in comparison to wondering whether the system is leaking or not. The second trick, when a leak has duly been found and repaired, is to tell yourself that the system is still leaking, and you still have to find the leak. Developing this kind of passion could save thousands of dollars over the life of the plant.

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■

Reducing Load

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This column is the sixth in a series exploring refrigeration and heat pump concepts without using jargon.

Previously in this series, we thought about efficiency being defined as "what you want to do" divided by "what you need to do." For example, a beer maker would measure his efficiency in "beers per kWh." A corollary is that inefficiency occurs when, in addition to doing what you want to do, you are unwittingly doing a whole bunch of other things as well. These other things will consume additional resources and so will make the plant less efficient than it should be.

Some of the other things that get in the way of making beer efficiently might be pumping more glycol than necessary, or alternately heating and cooling parts of the system unnecessarily. Any time that one bit needs to be heated, for example during a cleaning operation, care should be taken to ensure that as little heat as possible spreads to the rest of the system.

Fans and pumps that move the cooled stuff around are all energy users, and the energy put into them ends up heating them up. This was proved conclusively by James Joule in England in the mid-nineteenth century. Joule was the son of a brewery owner and spent lots of time mucking about with tubs of water and mechanically driven paddles, trying to prove that heat has a mechanical equivalent. This kind of thing is taught in elementary school now, so it is difficult to imagine what engineering would be like before this was common knowledge.

To this day, the English drink their beer warm in honor of master-stirrer James Joule.

Nevertheless, many people seem to forget elementary school when it comes to plant operation, and the number of extra fans and pumps left running in complicated systems for no apparent reason is substantial. Like the weather, described in the fourth article (*ASHRAE Journal*, June 2012), fans and pumps are a double whammy. They use electricity and so increase the "what you need to do" part of the efficiency equation, but as Joule found out, the electricity turns into heat, which then needs to be removed to keep the stuff cold. In this case you are paying the electric bill twice, once to put the energy into the stuff, and once again to take it back out again. It makes sense to aim to keep the amount of pumping and blowing to a minimum while you are making beer (or pies or pallet movements for that matter).

Other kinds of load may accidentally be added to cooling systems but are less obvious than additional temperature rise. In air conditioning (and also in cold storage) extra wetness can be a heavy burden. The amount of water, in the form of gas, held in the air is invisible although it has a big influence on how we feel. It's important to control it if what you want to do is to keep the folks indoors comfortable, but it is difficult to tell just by looking at whether the control is being done wisely or stupidly.

A classic example of stupid control (which might give very comfortable indoor conditions) is when water is taken out of the air during the cooling process and then other water has to be put into the air to stop it from becoming too dry. If the substance providing the cooling (whether it is water, glycol or refrigerant) is colder than 50°F (10°C), then it's likely that condensation will form on the cooling coil. That water, taken from the airstream, will reduce the humidity of the air, and may make it too dry for comfort. Rather than over-drying the air, which requires a lower cooling temperature, it is far better to set the cooling temperature so that it delivers the right humidity straight off the cooling coil. This reduces both the heat load and the temperature lift on the cooling plant.

Where stuff is cooled through a wide temperature range, it is sometimes possible to use higher temperature cooling for the first stage and so reduce the load on the cooling plant. This might be done with a higher temperature plant, or with cooling tower water in an air cooler, or even with fresh air (cleanliness permitting). In a bakery, just making the conveyor belt between the ovens and the chiller take a detour around the building can knock a significant lump out of the total cooling requirement.

In a pasteurizing plant, where the stuff being processed needs to be heated quickly and then chilled again, it is possible to arrange for it to meet itself coming the other way. The incoming stuff is preheated by the hot stuff further down the line, leaving the heating and cooling equipment providing just the last few degrees in each direction. This principle could be applied to many other processes but is often thought to be too complicated. That's a shame because, in the right circumstances, both the cooling and heating loads can be reduced to about one-third of the total requirement.

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■

Saving Energy in Refrigeration Systems

Parasitic Loads

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This article is the seventh in a series exploring refrigeration and heat pump concepts without using jargon.

A parasite is an organism that derives some benefit from a host by modifying the host's behavior or bodily chemistry to gain some advantage. This is a more sophisticated meaning than our subject this month: for cooling systems a parasitic load is simply an additional demand on resources that could be smaller, or perhaps need not be there at all. They do not necessarily affect the rest of the system, and they may even help to make control systems more stable by adding a base load to an otherwise diminishing demand.

There are two types of parasitic loss: deliberate and inadvertent. Deliberate ones include the things we discussed last month such as fans and pumps. They are necessary for successful operation of the plant, but could perhaps be better managed so their impact on the total energy use of the system is reduced and performance is improved.

A common characteristic of deliberate parasitic loads is they remain constant when the main cooling demand is reduced. For example, the oil pump on a screw compressor may only be 5% of the main drive motor power, but as the compressor unloads that percentage increases. Likewise, in some chilled water systems total kWh consumption of the water pumps is more than the compressor motors on the chillers because the pumps run 24/7, even when the compressors are switched off. In these cases, there are significant benefits to be gained by reducing the liquid flow, for example, through the use of variable speed pumps. This applies to condenser pumps in water-cooled systems, too, and the savings can be very significant.

There are many methods for chilled water and condenser water flow control, but the heart of the problem lies with the configuration of the heat loads, not the water pumps. If a system has several heat loads, they all run all the time and they each have a modulating bypass valve to control the amount of cooling done; then the flow in the main cooling loop will be constant and variations in load will appear as changes in the temperature difference between cold water flowing to the load and warm water returning from it. When the cooling demand is low in all units, they will still draw the maximum flow, but will not make good use of it.

©anonw.wordpress.com

You wait ages for a decent chiller control algorithm and then several come at once.

This is a bit like a bus company sending the maximum rush hour number of buses all through the day. It gives great service to the traveller at lunchtime but every bus is nearly empty, and it is clearly not a very efficient way to operate. The simplest way to avoid this problem is to reduce flow to each load when less cooling is required. This drops the flow rate within the cooling coil, but since the load is light and water is supplied at a constant temperature, this should not be a problem. It also changes the operating point for the pumps, so the water system design guys do not like it because they have to think a bit more, but the pump will be happy to find a new balance point further up the curve.

If the cooler needs to have a fixed flow, then the variable main flow can be achieved by recirculating water through the cooler until it warms up enough to be released back to the chiller. This requires the extra complication of a recirculating pump and three-port valve at the cooling load. The main chilled water pumps still have to run up and down the curve like mice on a clock.

Greater pump power savings can be achieved by reducing pump speed. If the speed is one-half then the power is one-eighth, but be careful because pressure capability is only one-quarter of what it was. Generally, speed control only operates between 70% and 90% of full load speed but the power savings are still very attractive.

Now the water system design guys really have to scratch their heads to make sure each cooling load gets its fair share of water under all conditions. If the loads are on different floors of the building, this can be tricky because convection in the rising headers can disrupt flow. In all cases, the real name of the game is to make sure the return water temperature coming back to the chillers is as warm as possible.

Inadvertent parasitic losses are easier to deal with. These are things like open windows when the air conditioning is on, fresh air dampers stuck wide open, computers left on all the time, and so on. In a refrigeration plant, leaks from the high-pressure side to the low-pressure side of the plant (which are like leaving an internal window open) fall into this category. To paraphrase a well known sportswear company, "Just don't do it!"

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■

Downsides of Refrigeration

Defrost Cycle

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This article is the eighth in a series exploring refrigeration and heat pump concepts without using jargon. We start to look at some of the unsavory things that creep into refrigeration plants. The first of these is the defrost cycle.

Any air-cooling equipment that runs at a temperature below 32°F (0°C) will tend to form ice on the heat exchange surfaces, even if the air itself is above freezing. The physical science behind frost building is extremely complicated, but the result is very simple: a ball of ice that fails to deliver the required cooling effect. Building ice to this extent can cause structural damage if the steel was not designed to cope with the excess weight—the ice can weigh as much as the cooler. It can also create health and safety hazards in the store including slippery floors and chunks of falling ice.

When to defrost is a classic optimization problem. Defrosting puts heat into the cold store, and as always in these situations you pay for it twice; once to put it in, and once to take it back out again. Running the plant with too much frost on the coils is also inefficient. It widens the temperature difference (rolls the piano further down the hill than necessary [*ASHRAE Journal* March 2012]) and increases the temperature lift. So the right time to defrost is not too soon after the last one, but before the system performance falls off a cliff.


Dr. Zhivago would have liked a few more defrosts.

The simplest defrost is using ambient air—either the air in the room, or in some tropical climates, ducting air from outdoors into a closed chamber. This is almost free. You are only paying for the fan power to blow the air through the coil, and the electricity required to bring the metalwork back down to operating temperature when refrigeration restarts. The latter cost may be substantial, depending on the structure of the air cooler, and it is well-nigh impossible to avoid, since the frost has to be melted, and the melt water then has to run off the metal surfaces and down the drain.

Another simple defrost system is the use of electric heater elements built into the air cooler. This is generally very effective and easy, but it requires a hefty electrical supply and might cause extensive damage if a heater element fails. In larger systems so-called "hot gas defrost" is often preferred to electric elements. This is a bit of a misnomer because the temperature of the gas at inlet to the coil is not important and it does not need to be hot. The important thing is that the pressure is high enough to make the gas condense at a temperature higher than the freezing point of water. It is the latent heat that transfers to the surroundings when the gas condenses (exactly the same as when the gas is reliquefied in the condenser) that heats up the frost and melts it. This can be done very effectively at a condensing temperature of 50°F (10°C), which is colder than the minimum condensing condition for most systems.

There are many ways in which a poorly maintained defrost system can kill the efficiency of the system. Defrosting too often is every bit as bad as not defrosting enough. Forcing the compressor discharge pressure up during defrost (perhaps to make sure that the gas really is "hot") wastes compressor power for the whole system and overheats the metalwork in the air cooler, increasing the cooling required to restart the system.

Gas that passes through the air cooler without condensing (perhaps because the flow control is based only on pressure) can put a false load on the compressors. Even worse, gas that leaks through the control valve when the air cooler is not being defrosted (perhaps when one of the other units is on defrost) puts a huge load on the compressors and can increase the energy used by the system by more than 10% without giving any obvious symptom.

This type of internal leakage can be detected by looking for signs of gas flow at the control valves when coolers are not supposed to be defrosting. Typical signs include higher temperatures in the pipes leading to the valve, lack of frost on the valve and the sound of gas whistling through the valve. Apart from the damage this does to the electricity bill, it can also erode the valve, and in severe cases it can lead to liquid hammer and condensate-induced shock which can cause major plant failure. Defrost efficiency is also affected by convention currents through the air cooler, which can be reduced by fitting fans socks or hoods.

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■

Good, Bad or Ugly

Lubricating Oil

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

This article is the ninth in a series exploring refrigeration and heat pump concepts without using jargon. Last month we started looking at the unsavoury side of refrigeration. This continues with a sideways look at lubricating oil.

About seven years ago we started building chillers with centrifugal compressors that have magnetic bearings and, therefore, do not need any lubricating oil. Until then I knew, in theory, that oil is a good thing in a compressor and a bad thing in the rest of the plant, but I had never truly appreciated quite how ugly it can be.

Oil in a refrigeration system is a necessary evil; it is critical to the continued operation of the compressor, but when it gets to the low temperature part of the circuit it becomes a troublesome nuisance. It coats the inside surface of the heat exchangers, acting as a layer of insulation. It collects in pools in low spots and can cause sudden changes in operation that destabilize plant operation and can damage components.

Last year I saw a horrific video made by Garden City Community College that filmed the behavior of refrigeration oil in ammonia at low temperature. This stuff was disgusting! It had the consistency of two-part epoxy resin adhesive shortly before it hardens, and it flowed with the speed and agility of a dead possum. The mental picture I had cherished for years—of the oil in low temperature receiver vessels willingly flowing down small pipes into oil recovery pots—was shattered.

The problem with oil is that we need it to be reasonably stiff (but not too thick) in the compressor bearings, which are hot, and yet to be free-flowing and easy to recover in the low-temperature part of the plant. However, the oil is an uncooperative player in the game; it gets thicker and stickier as it gets colder. It is even possible to have an oil that is too thin and runny when hot to be any good in the bearings and at the same time is too thick and sticky when cold to be any use in the evaporator. If the oil is capable of absorbing high-pressure refrigerant, it

There must be better ways of choosing refrigeration oils than this.

might be of even less use in the bearings.

This problem becomes even trickier if the oil is susceptible to changes in its properties over time. If it is made from long strings of carbon atoms it might become thinner as time goes by; if it gets contaminated with dirt, water or other chemicals, it can turn to sludge and become impossibly thick.

The "grade" of oil is important. This single number gives a measure of how runny the oil is at about 100°F (38°C). Typical grades are 32, 68, 100 and 220. However, since neither of the two parts of the system that are of interest, from an oily point of view, operate at this temperature, it is a kind of an unhelpful measure. The key temperature in the compressor bearings, where the oil must be thick enough to work, is likely to be about 200°F (93°C). In the evaporator, where the oil has to be thin enough to flow, the temperature might be 0°F (–17°C) or lower.

Traditional refrigeration oil was refined from crude, as a by-product of gasoline production. Refrigeration systems required highly refined, stable product usually based on cyclic carbon-based molecules. These oils get runny fast as they are heated, but get thick quick when they are cooled. They were very common in CFC and HCFC systems, but they do not dissolve in modern refrigerants.

New types of oil were developed to go with the new HFC refrigerants. They all seem to have names that are three-letter acronyms starting with P and combining any two from A, E, G and O. However, don't be fooled. They are significantly different and putting the wrong one into your system could have disastrous results.

For example, oil suitable for HFC refrigerant put into an ammonia system might turn to expanded foam: not the best way to look after the compressor. Some of the synthetic oils show much less variation in their thickness as the temperature changes, and remain relatively free-flowing even down to –50°F (–46°C). This also means they are thicker where they need to be—in the compressor.

Until magnetic bearings or other oil-free technologies are developed for more traditional compressors—screws, pistons and scrolls—the best advice you can follow with regard to lubricants is "pay close attention to the compressor manufacturer's recommendation for the refrigerant you are using and follow it *to the letter*."

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■

Jargon Buster

By Andy Pearson, Ph.D., C.Eng., Fellow ASHRAE

It is time to dig a little deeper into the jargon, although you should bear in mind that it is theoretically possible to talk about refrigeration systems without using any jargon. Fridge guys are generally not good at this; they have a language of their own and they are not ashamed to use it. The following lists terms you might encounter.

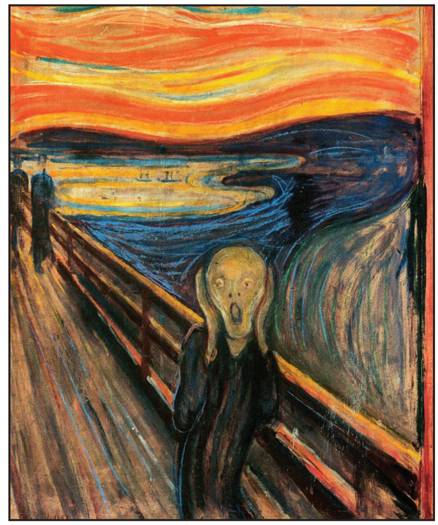

Munch knew what it felt like to listen to two fridge guys chatting.

Condensing Turning from gas to liquid. This transfers heat to the surroundings without changing temperature.

Evaporating Turning from liquid to gas, absorbing heat from the surroundings without changing temperature.

Compressing Reducing the volume of gas to raise its pressure, so it can be condensed at higher temperature.

Expanding Dropping the pressure of liquid so it can be evaporated at lower temperature. Sometimes some of the liquid boils in the process, hence the expansion tag.

Cycle The combination of the four processes of condensing, expanding, evaporating and compressing which creates a closed loop in order to keep the cold stuff cold.

Enthalpy The heat energy content of the fluid used in the cycle.

Entropy The total internal energy of the fluid. For most practical purposes you don't need to worry about entropy unless you are looking in depth at compressor performance.

Heat A form of energy that can be moved from one substance to another.

Work Another form of energy that can provide a useful effect on a substance but in the process is turned into heat.

Exergy Don't go there! For most practical purposes, like entropy, an understanding of exergy is not necessary. It might just create unhelpful confusion. From the Latin for "Get me out of here!"

Adiabatic A change of refrigerant condition with no change of enthalpy, usually the pressure reduction in the expansion process. See also *Enthalpy*.

Isentropic A change of refrigerant condition with no change of entropy, only relevant for compressor analysis. See also *Entropy*.

Liquid When a fridge guy says "liquid," he usually is specifically referring to condensed refrigerant. He may seem oblivious to the fact that there are other liquids in the system. For example, oil in the compressor and perhaps water in the condenser are liquid too, but they are not "liquid."

Vapor The fancy pants term for evaporated refrigerant, i.e., gas.

Gas See *vapor*.

Flash Gas The gas created during the expansion process as the refrigerant pressure is reduced.

Solid Liquid Another one of the fridge guy's funny expressions. Solid liquid is not solid; it means "liquid" that has no gas in it.

Saturated Temperature The temperature at which the refrigerant would boil or reliquefy at a given pressure.

Suction Pressure The pressure at which the gas is sucked into the compressor. This is slightly lower than the evaporating pressure. Fridge guys often use the saturated temperature when talking about suction pressure, so if the "pressure" is in Fahrenheit, that's why.

Discharge Pressure The pressure at which the gas is blown out of the compressor. This is slightly higher than the condensing pressure. Again, the saturated temperature is commonly used to indicate the pressure.

Superheat The difference between the actual temperature, and the saturated temperature for a gas.

Subcooling The difference between the saturated temperature, and the actual temperature for a liquid.

Heat Rejection Taking heat from the discharge gas to condense it to liquid, but not using it in any other way.

Waste Heat The amount of heat that is rejected, that is to say thrown away from the high pressure side of the system.

Heat Recovery Taking heat from the discharge gas in order to condense it to liquid and using the heat in some other process.

Refrigerant The evaporating/condensing stuff in the system, even if the system is a heat pump.

Non-condensible Other gas (usually air or nitrogen) that gets into the system in addition to the refrigerant. It cannot reliquefy and so it collects in the condenser and makes the temperature lift higher than it needs to be, which makes the system less efficient.

Andy Pearson, Ph.D., C.Eng., is group managing director at Star Refrigeration in Glasgow, UK. ■