Cheap Chills: Decarbonizing Industrial Process Cooling

Daniel Jordan, VEIC

ABSTRACT

When it comes to choosing the right system for industrial process cooling, the number of options can overwhelm any facilities engineer. At a minimum, engineers must account for process specifications like temperature and load requirements, as well as for local climate. Energy efficiency continues to play a leading role in both cost-effectiveness and in company decarbonization goals. Decision-makers sometimes use energy efficiency as a proxy for decarbonization, but are the two really interchangeable?

This paper draws on experiences in Vermont to provide a framework for distinguishing between the energy impacts and the carbon impacts of different process cooling options. Outdoor condensing units, for example, enable lower chiller condensing pressures, which can cut electricity use by 25 percent or more, but the resulting field-connected refrigerant lines become a GHG liability. In another case, an in-line dry cooler can be very effective at sharing the wintertime load in cold climates, but is it worth the extra capital cost?

The index study describing the experiences was conducted on behalf of the statewide energy efficiency utility, Efficiency Vermont. It concludes with an analysis comparing a 30-ton cooling process run on first shift vs. third shift, showing energy and carbon savings results for different technologies across each shift. The results show cost savings up to 58 percent, energy savings up to 46 percent, and carbon savings up to 49 percent.

Introduction

The industrial sector accounted for one quarter of all electricity purchased in the U.S. in 2020 (EIA 2021b). Although industrial sub-sector data for 2020 are not yet available, EIA noted that process cooling accounted for over 8 percent of all the manufacturing industry's consumed electricity in 2018 (2021a). It can therefore be estimated that industrial process cooling accounts for 1-2 percent of all electricity used in the United States.

Even with such a large share of electricity end use, the proportion of GHG emissions attributed to the generation of electricity for industrial process cooling is very likely higher than 1-2 percent. There are three reasons for this: (1) many manufacturers produce their own electricity on site, thereby eluding traditional survey data (EIA 2021a); (2) manufacturing operations may disproportionately prefer states with lower-cost and higher-carbon energy mixes (Carrick 2019); and (3) process cooling loads have high summer peak coincidence. That is, they require more energy in hot weather when grid electricity is especially carbon-intensive, because of spiking regional demand.

Electric demand aside, some industrial process cooling components are liable to emit gaseous refrigerant compounds, the most common of which have global warming potentials (GWPs) that, by mass, are thousands of times that of the carbon dioxide (CO_2). The U.S. Environmental Protection Agency (EPA) categorizes fugitive refrigerant emissions as Scope 1 emissions, the same category for on-site combustion emissions of gas or oil for heating (EPA 2020). This study shows that when considered against the electric demand of a process chiller, for example, the GWP of the refrigerant within that chiller is not trivial.

Energy efficiency and *decarbonization* refer to related energy conservation activity, but they are not synonymous. In industrial process cooling, cost-savings measures can advance building decarbonization, but not necessarily energy efficiency—and vice versa. Differentiating between these energy savings strategies requires consideration of cooling system technologies, configurations of those technologies, process requirements, and local climate. It also requires consideration of electric grid needs, dynamics, and pricing. Using consideration of these factors, together with experience in the Northeast with industrial cooling energy reduction measures, this paper addresses certain complexities of decarbonizing industrial process cooling with respect to cost and energy savings.

Cooling System Selection Criteria

Although the variety of industrial process cooling technologies is vast, they all have the same job: *Get rid of heat*. The technologies then begin to differentiate themselves in answering the follow-up questions: How much heat should they expel, and how quickly? At what temperatures? From what part of the system? To where? And, over the course of the year, how is the weather outside?

An engineer specifying a process cooling system will make decisions based on the following criteria:

- **Sizing and performance.** Small systems might not be worth the extensive plumbing and piping required for efficiency measures enabled by chilled water loops. "Crash cooling" in brewery and other such applications requires short-duration, high-power cooling that can be efficiently accomplished with a smaller unit, coupled with a thermal storage tank.
- **Process temperature requirements.** High-temperature cooling (approximately 60 degrees F or higher) might only require a simple dry cooler or a cooling tower rather than mechanical chilling (refrigeration). Dry coolers and cooling towers will be described in detail later, but they provide cooling by blowing air and sprinkling water over a heat exchanger, rather than using energy-intensive refrigeration and high-GWP refrigerants.
- **Specialized process requirements.** The cooling of plastic injection requires high flow rates and low temperature differentials (Tillou, Case, and Case 2003). In food processing, chilled water that contacts product might be subject to food industry-related restrictions.
- **Climate.** In dry climates, evaporative cooling works better than in wet climates. In climates with cold winters, cooling systems that use glycol loops can capitalize on the cold outdoor air to achieve cost, energy, and carbon savings.

Anyone advising on cooling system selection should understand that process constraints take priority. With well-articulated constraints, provided either by the facility's manufacturing and engineering staff or by a third-party engineering study, an energy consultant has a clear starting point to begin optimizing for carbon, cost, or energy savings—and ideally, all three.

Study Scope

This study draws upon experiences between VEIC's industrial energy consultants and the customers they serve through the operation of Efficiency Vermont. The systems discussed are therefore germane to Vermont and locations with similar climates: temperate, with cold winters.

Previous ACEEE study authors (Tillou, Case, and Case 2003) note that systems thinking is vital to finding energy savings in industrial process cooling. Indeed, retrofits of existing systems require critical evaluation of components as they relate to the whole. The aim of this study was different. This study has synthesized general principles in the pursuit of *categorically* decarbonizing industrial process cooling. This generalized approach might be more useful for specifying a new system than, for example, troubleshooting an existing one.

Breaking down a process cooling system into its loops is a useful (if reductionist) way to conceptualize process cooling system components. Generally, configurations of loop sequences characterize commercially available cooling system options. Thus, the study analyzed different loop configurations, as well as the implications of each one for cost, the scope of energy efficiency benefits, and GHG reductions. Moving in an increasingly pragmatic direction, the study offers examples of capital cost differences among systems with respect to their operational efficiencies and energy cost savings. It also addresses electric grid emissions reductions associated with these savings. The study also offers a modeled case study to illustrate cost and GHG differences across different time-of-day operating ranges for one industrial process cooling application in Southern Vermont.

Loops

In general, process cooling requires sequenced, closed-loop, recirculating fluids that pass heat from the industrial process all the way to its destination, which is usually outdoor air. Even for experienced engineers, distinguishing between the various loops can cause confusion. Industrial process cooling systems can comprise anywhere from one to six loops, sometimes more. The list below describes three typical loop categories that could make up a single process cooling system, moving from closest to farthest away from the process:

- **Process loop.** This loop can interface with the actual production process. The original equipment manufacturer (OEM) has built channels, ports, and possibly pumps on the process machine to control where and how this loop pulls heat from the process. After contacting the process, the loop dumps heat into the chilled-water loop through a heat exchanger. Note: It is not uncommon for the process loop and chilled-water loop to be the same, depending on requirements of the process loop fluid.
- **Chilled-water loop**. This loop is also referred to as *water loop*, *antifreeze loop*, or *glycol loop,* because common antifreeze agents are propylene glycol or ethylene glycol, mixed in some ratio with water. The loop might contain antifreeze to prevent localized freezing, either because the loop runs outdoors or because it passes through small channels likely to “ice up.” The chilled-water loop pulls heat from the process loop and dumps it either into the low-pressure side of the refrigeration loop or, directly—using fans and heat exchangers—into the outdoor air. Chilled-water loops can run through an entire plant and serve multiple processes. They are nearly as simple to plumb and pipe as a domestic water line. They also require pumps for recirculation, and they can run through rigid or flexible pipes of a few inches to one foot or more in diameter. The pipes might appear larger than their effective diameter, due to insulation, which aids in efficient cooling by preventing the loop from absorbing unwanted ambient heat; its only job is to move process heat from A to B.

- **Refrigeration loop.** Although the process loop and any chilled-water loops pass heat down a temperature gradient from hot to cold, the magic of the vapor compression refrigeration cycle is its capacity to work *against* a temperature gradient. The refrigeration cycle pulls heat from something that is already cold and dumps it into something hot, like outdoor air in the summertime. It accomplishes this by manipulating the pressure of a refrigerant. Refrigerants are highly engineered compounds that exist between liquid and gaseous states through the refrigeration loop. Although a system's chilled-water loop may contain many gallons of antifreeze or water, a refrigeration loop of the same cooling capacity would contain only a few pounds of refrigerant. Refrigeration lines might reside within flexible conduit, but they require rigid piping. This is almost always copper. The pipes will generally be less than an inch to a few inches in diameter, at most.

System Configurations

Given these loop descriptions, there are many configurations for sequencing them with respect to both the process and the outdoor air. Each configuration has its own GHG implications. Figure 1 presents schematics for three common configurations.

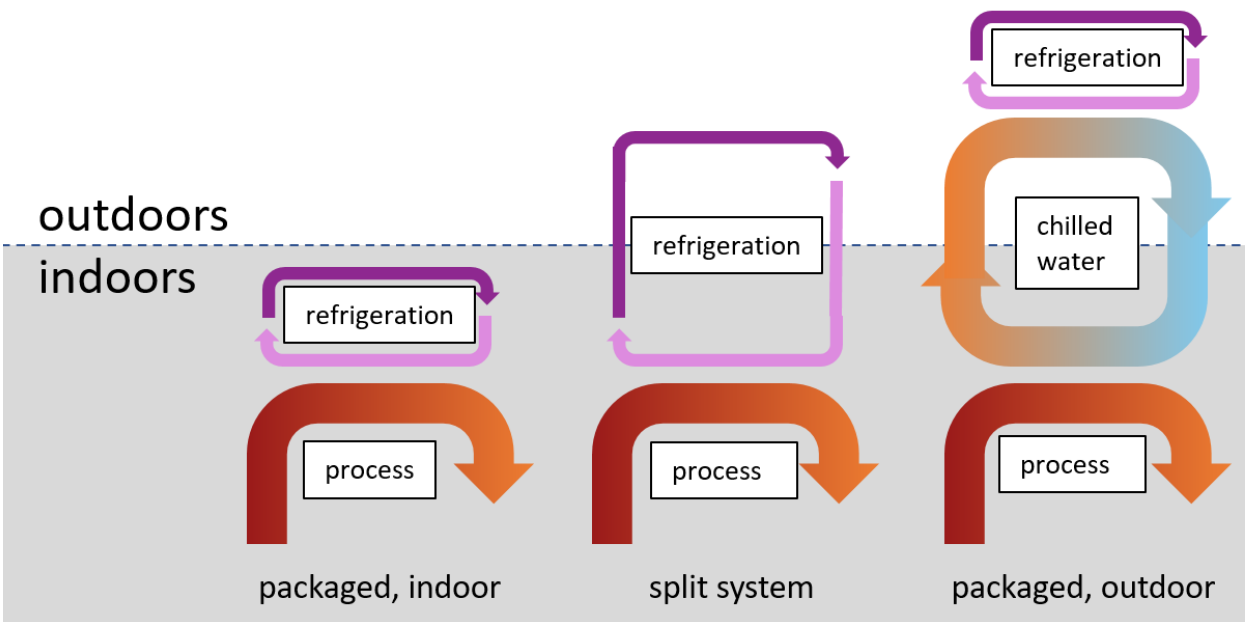

Figure 1. Three common system configurations using chillers: packaged indoor, split system, and packaged outdoor.

Packaged Indoor Units

The easiest industrial process cooler to install and maintain is the portable indoor chiller. The chiller sits beside the process, often on wheels, and the chilled-water loop runs to and from the chiller through flexible hoses. Within the chiller, the refrigeration loop pulls heat from the chilled water loop and, using a condenser and a fan, dumps it into the shop floor's ambient air.

The entire refrigeration cycle is contained within the single chiller frame. Hence, it is described as *packaged.*

Portable indoor chillers come in sizes up to only about 50 tons, and they are usually much smaller. The obvious drawback is that they reject heat indoors. Thus, if the indoor space is also air conditioned, then the facility must pay twice to move the heat outdoors. On the other hand, this heat will offset some of the building heating load in the winter months.

Regardless of their impacts on indoor space conditioning, indoor chillers miss an opportunity in cold climates to take advantage of heat transfer into cool outdoor air. Heat transfer happens more effectively across greater temperature differentials. This principle opens the door to efficiency gains. To this end, there are two common ways to expose the refrigeration cycle to outdoor air: (1) move the high-pressure side outdoors with an outdoor condensing unit, or (2) move the entire refrigeration cycle outdoors with a packaged outdoor system.

Split Systems (Outdoor Condensing Units)

A split system uses an outdoor condensing unit to run the high-pressure condensing side of the refrigeration cycle outdoors. In the cool weather of winter and shoulder seasons, an outdoor condensing unit can use floating head pressure to achieve high efficiency gains and capacity improvement. (See Outdoor Unit Efficiency Measures, below.)

In a split system, refrigerant lines connect the indoor evaporator unit (which pulls heat) to the outdoor condensing unit (which dumps heat). The refrigerant lines must be long enough to run from the process to the condensing unit. This is usually tens of feet, at least. The lines also must be piped through an exterior wall. This means the installer must pipe the refrigeration loop on-site, and then charge the loop with refrigerant upon commissioning, rather than having the loop built and charged at the factory. Factory charging takes place in the production of packaged units.

On-site refrigerant loop installation and charging opens the door to slight variability in installation quality. This can make the difference between a system that leaks refrigerant and one that does not. Molecules of hydrofluorocarbon (HFC) refrigerant gas are very small compared to other molecules, and they escape through joints unless the joints are nearly perfect. Even small amounts of leaked refrigerant gas can have severe greenhouse gas emission implications, as demonstrated in Table 1—even when considered next to the GHG impacts of electricity consumption.

Table 1. Common industrial process refrigerants, their global warming potentials, and the equivalent full-load hours (FLH) required for a 30-ton chiller to achieve the same electric source emissions GHG impact as a full system refrigerant leak.

Refrigerant	GWP100[1]	30-ton outdoor system charge	100% leak GHG-equivalent FLH
R-410A	2,088	75 pounds (71.1 tons CO_2e)	4,860 hours
R-407c	1,774	75 pounds (60.3 tons CO_2e)	4,140 hours
R-22	1,810	75 pounds (61.5 tons CO_2e)	4,200 hours

[1] GWP100 refers to global warming potential after 100 years, in terms of carbon equivalence; as Table 1 shows, after 100 years, one pound of R-410A refrigerant, for example, leaked into the atmosphere has the same warming effect as 2,088 pounds of CO_2.

To summarize the table: the GWP of a full leak-out of all the refrigerant in a 30-ton outdoor chiller is on par with two years' worth of emissions required to generate the electricity used by that same chiller running for one shift at full load. This metric uses a carbon intensity for electricity of .92 pounds per kilowatt-hour, which is the U.S. national average (EIA 2021c).

Although a full leak-out of refrigerant is unlikely, a survey of log books of refrigeration technicians in the United Kingdom found that "catastrophic" leaks of approximately 40 percent system charge represented the median leak rate across all leaks identified in residential and commercial heat pump systems (Eunomia 2014). Split-system heat pumps share many characteristics with split-system chillers, including field-installed refrigerant lines. Notably, the UK study also found that only 10 percent of surveyed systems leaked at all.

Lower-volume leaks might be more insidious than high-volume ones: in industrial settings, process equipment often undergoes quarterly or annual maintenance, and for chillers, this maintenance might involve topping off a chiller's refrigerant without documenting how much went in, especially for smaller systems with less stringent regulatory requirements. The Eunomia study points out that heat pumps will continue to operate even after 40 percent of the refrigerant has leaked out, albeit at reduced efficiency. If split-system chillers function the same way, then it is easy to imagine how a routine maintenance program would enable periodic recharging of a leaking chiller with no noticeable performance impacts. This would result in a significant amount of refrigerant leaked over the chiller's lifetime, far outweighing any marginal efficiency gains from a GHG perspective.

The data don't suggest that split systems are necessarily leaky, so much as that they have the potential to leak. A high-quality installation with proper vacuum testing is an effective method to mitigate preemptively the risk of refrigerant leakage of split systems.

Packaged Outdoor Units

Across the variety of applications of the refrigeration cycle in HVAC and refrigeration including industrial process cooling, packaged systems are less likely to leak than split systems. In packaged systems, the refrigerant lines are much more compact, built and charged at the factory without customized installation required case-by-case, and subject to consistent quality control measures during manufacture. And, they require less refrigerant ("lower charge") overall.

Rather than running refrigerant lines through exterior walls to the high-pressure side of the refrigeration cycle, a process cooling system can send a separate chilled water loop through the wall and to an outdoor packaged chiller. The OEM needs to equip the chillers to withstand outdoor air temperatures, but the marginal costs to this are slight. In the case of one Vermont manufacturer's installation of two small chillers—one low temp 5-ton and one very-low temp 7.5 ton, in series—the outdoor rating represented only 7.5 percent of total project cost, with the installation of this outdoor unit (mounting, plumbing, electrical) accounting for 30 percent of the total project cost. For reference, the packaged indoor chiller option would likely have had no install cost, as personnel could have installed flexible piping for the chilled water loop quite easily, and a 240-volt model could be plugged in without extensive electrical work. The split system install cost would likely have been greater than the packaged outdoor system cost due to the refrigeration piping work and, ideally, vacuum testing.

The energy trade-off with packaged outdoor units is that the chilled water loop requires additional pumping energy to recirculate the chilled water, which increases energy consumption compared to split systems. The split system, on the other hand, relies on the pressure differential already necessarily provided by the compressor to achieve the movement of refrigerant from indoor to outdoor and back again. Ideally, the additional pumping energy required by packaged outdoor systems is offset by the efficiency gains enabled by outdoor condensing, and the pumping energy can be mitigated in applications with modulating capacity through the use of variable frequency drives (VFDs) and controls.

Packaged outdoor units afford nearly the same efficiency gains as split systems with outdoor condensing—namely, the option for floating head pressure control. But in most climates, sending the chilled-water loop outdoors presents a second efficiency option: direct cooling of the chilled-water loop with dry coolers or cooling towers. Moreover, building out a chilled-water loop is the first step to inter-process heat exchange.

Efficiency Measures

Inter-process heat exchange

Facilities that require cooling might also require simultaneous heating somewhere else. For example, a mechanical room housing air compressors might be located near a room requiring hot water for washing parts. In this case, if the chilled-water loop is warmer than the water returning to the hot-water heater, then it could be advantageous for the loops to exchange heat with each other and thus boost the efficiency of both processes. With respect to the loop categories given above, inter-process heat exchange most commonly takes place between chilled-water loops.

The drawback inherent to inter-process heat exchange is that both systems then rely on each other for heat exchange. If the parts washer isn't running on a given workday, then the air compressor's entire cooling load goes to its dedicated system cooling system—which, it should be noted, otherwise runs only at fractional capacity on days when the parts washer *is* active. Most cooling systems have an option for capacity modulation, which enables efficiency measures that rely on variable heat sinks, like inter-process heat exchange.

Floating Head Pressure Controls

Also known as *low condensing,* floating head pressure on outdoor condensing units provides some of the best component-level efficiency gains in industrial process cooling. One reason for such high gains is simply that refrigeration compressors use so much energy. In the systems surveyed for this study, compressor power at baseline operating conditions[2] represented roughly half the power of the entire cooling system, and more at partial capacities.

This latest point cannot be stressed enough; it runs counter to a face-value assessment of a cooling system. For a system on the order of 10 to 100 tons, dry coolers and cooling towers comprise large metal frames the size of SUVs, with very large fans running full bore. The fans on condensing units are impressive as well, equipped with modulating electronically commutated motors (ECMs). And yet the small, unimposing compressor (a black component about the size and shape of a fire hydrant) uses approximately ten times the electric power of any one of the fans.

[2] Baseline operating conditions in this case refers to factory-set, rather than floating, head pressure.

Although floating head pressure controls are not new,[3] their implementation in the past decades has been uneven. Patenaude (2014) provides an overview of the implementation and the technology, specifically as enabled by electronic expansion valves (EEVs). EEVS operating in response to pressure control are one method for achieving floating head pressure. But an installer can achieve similar efficiency results by adjusting condenser fan cycling parameters, so long as the other system components are equipped to withstand the lower pressure differential that deeper fan cycling produces. Case by case, each system has its own "magic line"—usually determined by the expansion valve—that dictates the minimum saturated condensing temperature (SCT), which is the same thing as head pressure. It is vital to find a controls or HVAC-R technician and efficiency ally who understand the finer points of each system, since no two installations are identical. VEIC's refrigeration team has partnered with contractor networks and customers in grocery and food manufacturing to set new standards in efficient compressor / condenser controls and operation. A common measure is to work with contractors and customers to bring the factory-set minimum SCT of 110 degrees F down to 80 or 70 degrees before the condenser fans cut out. This low SCT can result in an efficiency boost of over 30 percent for the entire refrigeration system.

In industrial process cooling, where equipment specification usually requires back-and-forth between an OEM representative and a manufacturing engineer, a feature request for floating head pressure control can easily get lost. Floating head pressure controls are not the same as "flooded" head pressure controls, which are more about preventing a specific failure mode. Moreover, not all technicians or OEM representatives interpret condenser fan cycling control as floating head pressure control. If you ask about floating head pressure within high and low limits, they might think you're asking about the head pressure safety set points. This can lead to confusion because, in fact, the question might be about condenser fan cycling. Moreover, not all chiller technologies can accept SCT adjustments. Scroll compressors – which are common in refrigeration – do accept these adjustments, but some OEMs use helical rotary compressors which don't accept them.

Best-in-class head pressure controls use outdoor air temperature sensors, a pre-determined approach temperature (TD), pressure-stats, and EEVs to maintain a precise condensing temperature that's no higher than it needs to be. As Patenaude notes, low condensing enables not only more efficient chilling but also *higher chilling capacity* in the same compressor.

Dry Coolers and Cooling Towers

With a chilled-water loop, a process cooling system can send the cooling load to a dry cooler or a cooling tower, each of which operates *without* the refrigeration cycle. Instead, these components cool the chilled-water loop by blowing cold air over it (dry coolers) or by spraying water onto it (cooling towers), which evaporates and thereby pulls heat away. Under the right operating conditions, dry coolers and cooling towers are much simpler and much more efficient than chillers. Their chief drawback is that they are not suitable for making very cold chilled water (lower than 50 degrees F in temperate climates).

A process cooling system using a glycol loop may use a dry cooler and a chiller in sequence with each other. Both can run simultaneously at partial capacity, or the system might run only one or the other, depending on the load and the outdoor air temperature at that time.

[3] A paper from an early *Proceedings* from the 1988 ACEEE Summer Study conference has detailed their thermodynamics and their efficiency prospects, primarily for refrigeration systems (Wheeler and Smith 1988).

When the system switches from chiller to dry cooler, the electric load suddenly drops by 30 to 50 percent (even with floating head pressure controls on the compressor). It is advantageous to opt for running the cooler, rather than any chilling unit, as much as possible.

In Vermont, some cooling towers act as dry coolers in the wintertime. These cooling towers come equipped with electric resistance water heaters to prevent the evaporation water from freezing, but in cold climates, it makes more sense to ensure the cooling tower is specified to have the evaporation water drained in the winter so that the cooling tower can act as a dry cooler (i.e., a large fan) in cold months.

System Pricing, Operating Cost, and Operational GHG

Given the operating components and configurations, together with the known constraints of the industrial process, a facilities engineer next considers the costs of installing one of these cooling systems. These considerations involve upfront capital costs with respect to electric operating costs. In Vermont, time-of-use electric rates coincide with differences in outdoor air temperature and electric carbon intensity, which in turn affect the cooling system’s operational carbon footprint.

Initial Capital Costs

Figure 2 shows options for projects in Vermont surveyed for this study. System integrated energy efficiency ratings (IEER) are calculated as the total cooling Btu the system provides over one year, divided by the electricity watt-hours input over that same year. IEER is situationally dependent and, in this case, assumes a first-shift application in Southern Vermont. Broadly, the plot shows that more efficient systems have higher upfront costs.

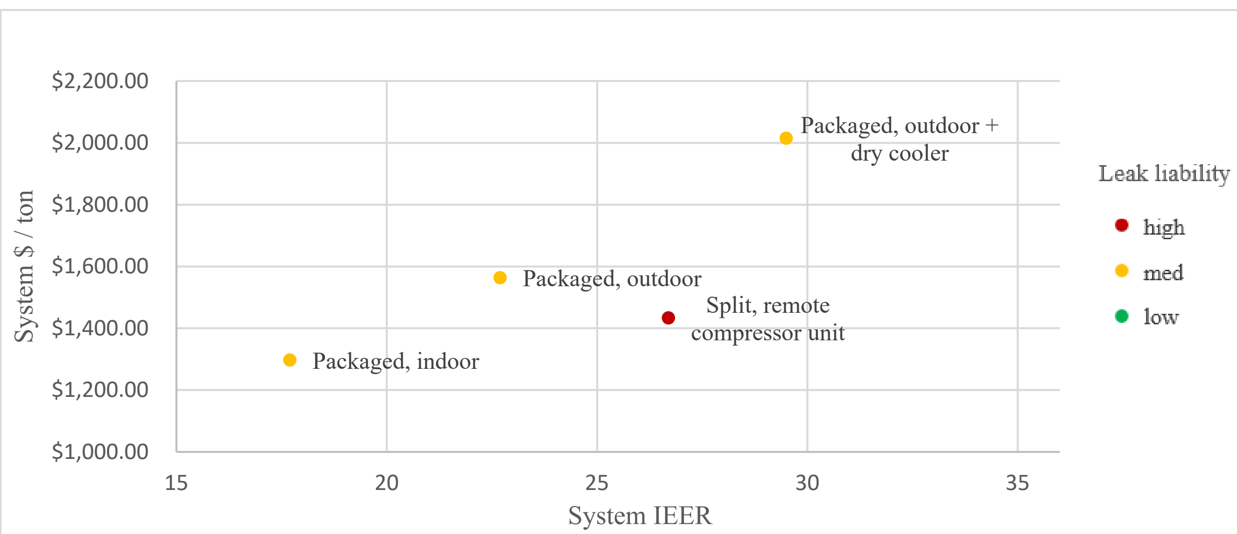

Figure 2. Process cooling system-integrated energy efficiency ratio as it related to pricing per ton of cooling. Pricing data is taken from projects between 2018 and 2020.

The extent to which an investment in a high-IEER system pays dividends will vary according to energy pricing. In Vermont, where blended energy rates for industrial customers range from \$0.07 to \$0.20 per kWh (depending on rate structure, discounts, load factor, and how many shifts a plant runs), the payback periods for the roughly linear increases in incremental gains in IEER, as depicted in Figure 2, can range from approximately one to ten years. Within that range, a specific simple payback calculation requires more details. However, the bottom line is that, from the customer's point of view, these measures nearly always pass cost-effectiveness screening over the equipment's lifetime. Expected useful life of industrial process cooling equipment is 20 years or more, and energy efficiency utility incentives may help close the gap.

Electricity Rates and Carbon Intensity

Most industrial rate structures in Vermont are subject to time-of-use billing. A unit of energy used on weekdays between 6 a.m. and 11 p.m. costs more than the same unit of energy used overnight or on weekends. An electric utility uses time-of-use billing to send pricing signals to customers to motivate electricity use in the middle of the night or on weekends, rather than in the middle of the day or the evening.

At present, the variability of a utility's cost to supply power largely correlates with the carbon intensity of the power it supplies. When electricity becomes more expensive to a utility, it becomes more carbon intensive as well, culminating in system peak periods of only a few hours. However, these few hours have an outsized effect on a utility's operating costs and carbon footprint over a given year.

Further decarbonizing industrial processes will therefore rely on considering not simply *how much* electricity the process uses, but also *when* it uses it. As utilities work to minimize the GHG impacts of their electricity resource mixes, they will push their investments into renewables as much as possible by maximizing capacity factor—not wanting to waste any of the expensive renewable energy they buy. They then can relegate the responsibility of dispatchable, on-call or "marginal" generation to fossil fuel turbines. This largely explains why carbon intensity goes up as electricity use approaches peak load.

Figure 3 demonstrates the change across one year of the carbon intensity of 1 kWh supplied by the regional grid operator, ISO New England. In Vermont, distribution utilities buy electricity from this marketplace. The emissions factors represent the aggregate (rather than marginal) emissions factors. That is, they account for the generation emissions of all the generators providing supply in the designated hour, assuming the emissions factors are spread evenly among all energy units used at that time, rather than varying by location or utility.

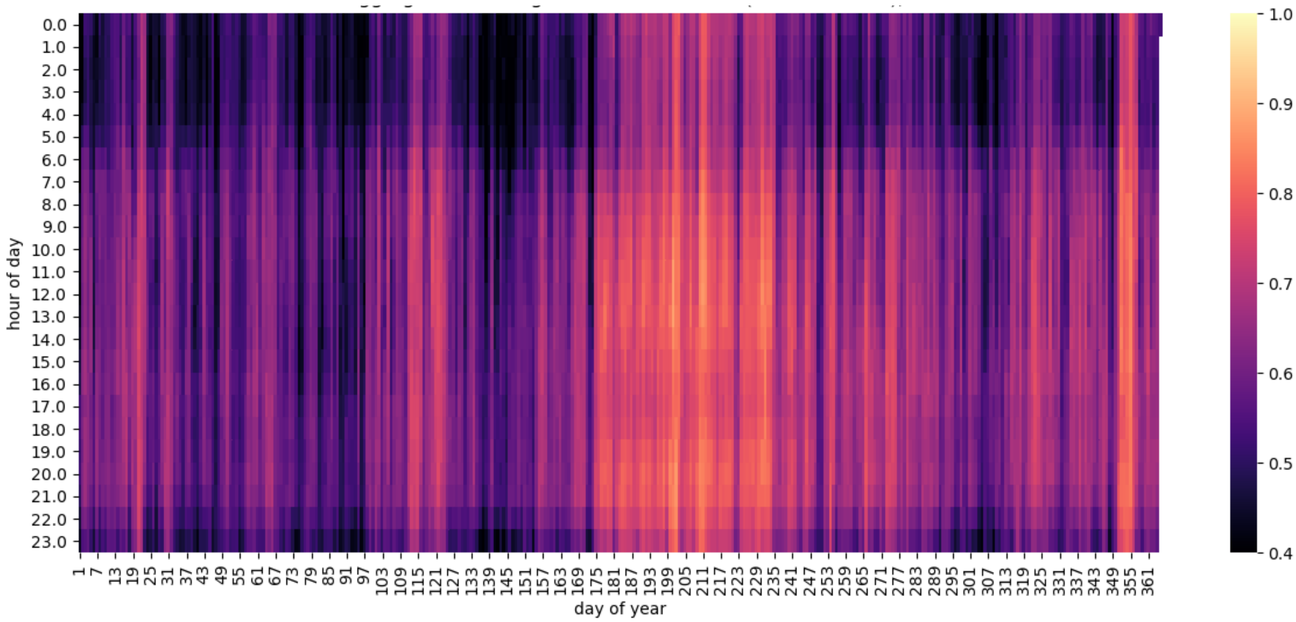

Figure 3. Emissions intensity of the New England electric grid across the 2019 calendar year. Brighter colors represent higher carbon intensity, with a maximum value of 1 pound of CO_2e per kWh, and a minimum of 0.4 pounds. Each rectangular pixel represents one hour. A vertical stripe represents a day, moving from midnight at the top to 11 p.m. at the bottom. The x-axis moves from January 1 on the left to December 31 on the right.

The intermittent dark areas along the top of this heatmap demonstrate the lower carbon intensity of overnight hours. The bright spot near the middle shows how carbon intensity is highest in the hot summer weather. The slight, intermittent dark spots across early afternoons reveal the beginning of the *duck curve*—where low residential demand, coupled with peak solar PV supply, result in low-carbon energy available during the shoulder season's early afternoons in New England.

What does this mean for industrial process cooling? Note that the off-peak period during which electricity is less expensive to customers on a time-of-use rate (11 p.m. – 6 a.m., and all day on weekends) roughly corresponds to a third-shift operation. Plants that operate only a single shift or two shifts, or even those that run lighter on third shifts or weekends, can shift load to off-peak hours when electricity is both lower carbon and less expensive, *and* cooling is more efficient because of lower outdoor air temperatures.

Moving a process from first to third shift can be a tall order, simply because it is hard to find employees willing to work third shift. However, manufacturers might be able to use thermal storage to make and store process cooling capacity at night, intending to use it the following day to offset some or all the on-peak electric cooling load. The chief limiters here are space and expense. For example, the 30-ton, single-shift process examined in the case study below is required to chill 72 gallons per minute of water / glycol from 60 to 50 degrees F. For this case, one shift of chilled-water loop storage would require a tank of approximately 30,000 gallons. Although lower-temperature and lower-volume storage might be a better option, the chiller runs less efficiently at a lower setpoint. Either option would likely double the initial cost of installation, not to mention create space constraints and service liability associated with a huge tank.

More creative solutions present themselves with phase change materials (PCM), which greatly reduce required storage volume for the same cooling capacity. The simplest PCM is ice, which draws in heat as it melts. Ice storage may be ideal for medium- and low-temperature chilled-water loops (30 to 40 degrees F). Otherwise, as in the 30-ton case study, the chiller must make colder chilled water than it would otherwise, which could help avoid on-peak costs in daytime, but otherwise hurts overall system efficiency. Thus, any manufacturer operating on a time-of-use rate and wanting to save costs and cut carbon should at least explore using thermal storage in its industrial process cooling systems.

In addition to differences in time of use, the case study model simulates the counterfactuals of different system configurations. The result offers energy, cost, and carbon savings.

Case Study Model: Load Shifting for Energy Efficiency, GHG Savings, and Cost Savings

This case study examines a newly installed process cooling system in Southern Vermont and models the differences between running the same 8-hour process on third shift (11 p.m. to 7 a.m.) and first shift (7 a.m. to 3 p.m.). The system uses a water / glycol chilled-water loop with a packaged outdoor chiller with low ambient head pressure settings running R-410A refrigerant. The system also has an in-line dry cooler set for thermostatic switchover at 40 degrees F. outdoor air temperature. However, the results in Table 2 present counterfactual systems, too (that is, what the customer *could have bought*), to model the spectrum of savings opportunities from the baseline case.

The study approximates the industrial process as a constant, 30-ton load for 8 hours. The spent water enters the cooling system at approximately 60 degrees F., and the cooling system needs to bring it down to 50 degrees F. Figure 4 shows a photograph of the system.

Figure 4. A 30-ton industrial process cooling system. The silver-colored piping is the glycol loop covered in insulation with aluminum jacketing. The two red valves in the foreground are the thermostatic switchover valves. The unit on the left is the dry cooler with 3 vertical fans. The unit on the right is the packaged outdoor chiller (whose fans aren't visible), with some of the refrigeration loop's copper piping visible. The tandem scroll compressors—the largest single energy users—are the small, grouped, black components in the bottom right.

June 20, 2020

On this warm summer afternoon, the temperature in Brattleboro reached the high 80s by noon and almost 90 degrees by the time the first shift ended. At these temperatures, the chiller's head pressure rose to an average of 100 degrees, to maintain the 10-degree approach it needs to dump heat into the outdoor air. To maintain the 30 tons of cooling required at these parameters, the compressor ran at approximately 19 kW.

Brattleboro was not the only city experiencing high temperatures on this summer afternoon; it was hot throughout New England. The carbon emissions of all generators in the ISO New England market rose to 0.82 pounds of CO_2e per kilowatt-hour, because electricity users all over New England needed air conditioning at the same time.

If the facility had decided to run this same process on the third shift the night before (11 p.m. to 7 a.m.)—when the outdoor air temperature hovered in the low 60s—the lower condensing pressure requirement would have let the compressor run at only 14 kW average. Moreover, the carbon intensity of the New England grid overnight ranged from 0.6 to 0.7 pounds per kilowatt-hour, nearly 12 percent less. Combined, the decreased electric grid carbon intensity and the increased electric efficiency, thanks to cooler temperatures, would result in net carbon savings of more than 24 percent.

The company would save on its electricity bill not only by using less energy (from running at night), but also by taking advantage of the lower overnight rates as well. The result is a 32 percent decrease in the energy costs for that one night. Table 2 shows annual electric bill savings over the course of calendar year 2020.

October 25, 2020

On this cool day in Vermont's fabled "stick season," the afternoon temperature by the end of the first shift rose to the high 40s. When the cooling process started up at 7 a.m., the chilled-water valve thermostat read 39 degrees, so the dry cooler kicked on to provide all the cooling needed for only 4.4 kW of fan power (all three fans). Around 9 a.m., the temperature rose enough above 40 that the system called for the chiller. The valves to the dry cooler closed, the dry cooler turned off, and the chiller started up. Putting out its minimum SCT of 80 degrees, the chiller's compressor used an average of 13.4 kW—lower than in the summertime, but still approximately triple that of the fans.

If the process had run during the third shift the night before when the outdoor air temperature started at just above 40 degrees F, the cooling would have begun with the chiller (13.4 kW). Then, around midnight, it would have switched over to the dry cooler (4.4 kW) to run until dawn. The carbon intensity of the New England grid that night ranged from .25 to .29 pounds of CO_2e per kWh, whereas during first shift it ranged from .29 to .33 pounds. The result is an energy savings of 46 percent, net carbon savings of 55 percent, and cost savings of 63 percent on this night alone.

Table 2. Differential savings for the case study across calendar year 2020

	First shift					
	Incre-mental capital cost	Annual energy cost savings	Energy simple payback period, years	Annual energy savings, kWh / %	Annual carbon savings, tCO2e	Carbon savings + 5% fugitive penalty
Packaged indoor*	-	-	-	-	-	-
Split system	$6,715	$1,431 [23%]	4.69	12,961 [23%]	4.67 [23%]	-3.22 [-16%]
Packaged outdoor	$14,620	$480 [8%]	30.5	4,350 [8%]	1.48 [7%]	-
Outdoor + dry cooler	$28,150	$1,871 [30%]	15.04	16,955 [30%]	5.94 [29%]	-

	Third shift					
	Incre-mental capital cost	Annual energy cost savings	Energy simple payback period, years	Annual energy savings, kWh / %	Annual carbon savings, tCO2e	Carbon savings + 5% fugitive penalty
Packaged indoor	-	$1,360 [22%]	n/a	10 [<1%]	1.79 [9%]	-
Split system	$6,715	$2,642 [43%]	2.54	14,924 [27%]	6.80 [33%]	-1.09 [-6%]
Packaged outdoor	$14,620	$1,902 [31%]	7.69	6,313 [11%]	3.89 [19%]	-
Outdoor + dry cooler	$28,150	$3,570 [58%]	7.89	25,711 [46%]	10.16 [49%]	-

The reference (*baseline) case in Table 2 is a packaged indoor chiller on first shift. The analysis excludes any additional air conditioning costs. What drives the savings in moving away from the packaged indoor chiller in both the first- and third-shift cases is the low head pressure enabled by outdoor operation. The deeper cost and carbon savings on third shift are due to the time-of-use electric rate and the lower carbon intensity of energy at nighttime. The column on the right reveals that with as much as a 5% annual leak rate, the carbon gains through energy efficiency of outdoor condensing are negated by the small leak rate of the field-installed split system.

Conclusion

The results show that when comparing system options across first shift and over the course of a year, savings in energy, energy cost, and operational carbon closely track one another. It may then be appropriate for an engineer to use energy savings as a proxy for carbon savings when evaluating options restricted to first shift operation. However, comparing first shift to third shift operations appears to decouple savings across the three categories. For example, when compared to a packaged indoor unit running first shift, the packaged outdoor unit running on third shift reduces operational energy by 11%, but it reduces operational carbon by 19%. So, the packaged outdoor unit runs somewhat more efficiently during the cool weather overnight, but the additional drop in carbon intensity of the electric grid during these hours only compounds the operational carbon savings.

The study also reveals that refrigerant should be understood as a GHG liability. Even small leaks can offset the carbon savings of energy efficiency.

More generally, the results show savings varying from a 16% carbon *penalty* to 58% cost savings across a rather straightforward small process cooling application. Across the spectrum of process cooling options, careful evaluation and accounting will reveal differences in cost-effectiveness which are not trivial.

References

Carrick, Alex. 2019. "Rankings of States by Industrial Subsector Jobs – Manufacturing." *Construct Connect, Economic News.* Cincinnati, OH. https://www.constructconnect.com/blog/economy/series-1-7-rankings-states-industrial-sub-sector-jobs-weight-concentration-maps-manufacturing

EIA (U. S. Energy Information Administration). 2021a. *Annual Energy Outlook,* Table 4, February. Washington, DC: EIA. https://www.eia.gov/consumption/manufacturing/data/2018/pdf/Table5_1.pdf

EIA (U. S. Energy Information Administration). 2021b. *Monthly Energy Review,* Table 7.6, March. Washington, DC: EIA. https://www.eia.gov/energyexplained/electricity/use-of-electricity.php

EIA (U. S. Energy Information Administration). 2021c. *Frequently Asked Questions (FAQs)..* Washington, DC: EIA. https://www.eia.gov/tools/faqs/

EPA (U. S. Environmental Protection Agency). 2020. "Direct Fugitive Emissions from Refrigeration, Air Conditioning, Fire Suppression, and Industrial Gases." December. https://www.epa.gov/sites/production/files/2020-12/documents/fugitiveemissions.pdf

Eunomia Research & Consulting Ltd and the Centre for Air Conditioning and Refrigeration Research (London Southbank University), 2014. *Impacts of Leakage from Refrigerants in Heat Pumps*. https://www.ammonia21.com/files/decc-refrigerants-heat-pumps.pdf

Tillou, M., M. Case, and P. Case. 2003. "Energy Efficiency in Industrial Process Cooling Systems: A Systems Approach." *Proceedings of the 2003 Summer Study of Efficiency in Industry.* Washington, DC: ACEEE: 4-290 – 4-300. https://www.aceee.org/files/proceedings/2003/data/papers/SS03_Panel4_Paper_26.pdf

Patenaude, A. 2014. "EEVs Enabling Low Condensing Refrigeration." *RSES Journal*, April. Arlington Heights, IL: 18-21. https://www.rses.org/assets/rses_journal/0414_Condensers.pdf.

Wheeler, G., and G. Smith. 1988. "Refrigeration Energy Savings with Floating Head Pressure." *Proceedings of the 1988 ACEEE Summer Study on Energy Efficiency in Buildings.* Washington, DC: ACEEE: 4.123 – 4.134. https://www.aceee.org/files/proceedings/1988/data/papers/1988_V4_015.pdf