The Basic Physics of Heat Pumps

Hilliard Kent Macomber
Emeritus Professor of Physics
University of Northern Iowa
June 2002

Introduction

In order to understand the basic physics of heat pumps, it is necessary to first have a clear and correct understanding of the basic concepts involved. This is all the more important as some of the key words, though familiar from daily life, are often used with a variable and hence imprecise meaning. We thus begin by carefully establishing the necessary key concepts. They are then used to describe the operation of a heat pump in an abstract way that captures just the essentials and thus is independent of whatever technology might be employed in practice. Some relevant basic laws of physics are then introduced, and limitations on heat pump performance based on these laws are derived. Finally, we briefly consider the performance of practical heat pumps.

Basic Concepts

Energy is for our purposes the capacity or potential to perform a task. An example is the energy of the working fluid in a heat pump after it is compressed.

Internal energy is the intrinsic energy a material possesses due to the motion of its atoms relative to each other and to the interactions of the atoms with each other. An example is the energy of the working fluid in a heat pump.

Heat is the transfer of energy between a system and its surroundings due to a temperature difference between the two. An example is the energy provided by a heat pump to a house. It is important to note that "heat" is not the same as "temperature," though the two words are often used interchangeably. Also note that heat, unlike internal energy, is energy *in transit* and thus is not something that a system can possess. For this reason, one should not talk about the heat in a system, as if it resides there. One may, however, and often will need to talk about the internal energy of a system.

Work is the transfer of energy between a system and its surroundings due to other than a temperature difference between the two. An example is the work done by a piston as it compresses the working fluid in a heat pump. Like heat, work is energy *in transit* and thus is not something that a system can possess.

Temperature is for our purposes what is measured with a thermometer. But here we run into some problems: (1) There are many different kinds of thermometers, and their readings generally disagree. Is there any kind that is fundamental and so can be used to do basic physics? The answer, fortunately, is yes. It uses the pressure of a gas confined in a fixed volume to tell us the temperature: Higher pressure means higher temperature, and lower pressure means lower temperature. (2) There are two common temperature scales, Fahrenheit and Celsius, but their definitions are arbitrary and do not reflect any basic characteristic of nature. Also, there is something odd about them. Suppose that the temperatures of bodies A and B are measured with a Fahrenheit thermometer, and it is found that the temperature of B is twice that of A. If the temperatures are now measured with a Celsius thermometer, this will no longer be found true. But suppose instead that the *lengths* of rods A and B are measured with an inch ruler, and it is found that the length of B is twice that of A. If the lengths are now measured with, say, a centimeter ruler, this will still be found to be true.

The defects in the common temperature scales are remedied in the "absolute" temperature scales. In them, the zero of temperature, that is, absolute zero, corresponds to zero pressure, surely a fundamental state. Above that, one can use degrees of the same size as on the Fahrenheit scale, thus producing the Rankine scale, or of the same size as on the Celsius scale, yielding the Kelvin scale. Since absolute zero corresponds to –459.67 on the Fahrenheit scale, the temperature in Rankine (R) is obtained by adding 459.67 to the temperature in Fahrenheit. Similarly, the temperature in Kelvin (K) is obtained by adding 273.15 to the temperature in Celsius.

The Heat Pump Reduced to Essentials

To understand the basic physics of the heat pump, we must focus on its essential flows of energy. The schematic diagram below shows these flows in a way that will facilitate analysis. The system of primary interest is the working fluid of the heat pump,

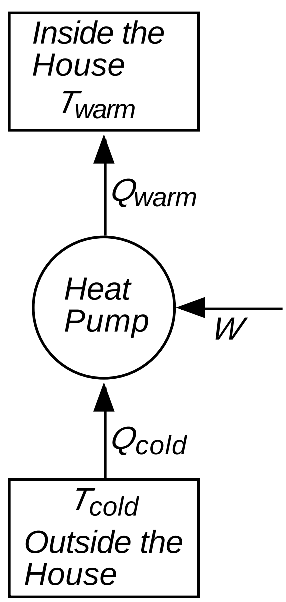

represented by the circle in the middle of the diagram. The cold environment outside the house is represented by the box at the bottom, and the *absolute* temperature of that environment is denoted by T_{cold}. The warm environment inside the house is represented by the box at the top, and the absolute temperature of that environment is denoted by T_{warm}. Heat flows from and to these environments are denoted by Q_{cold} and Q_{warm}, respectively. The work required to operate the heat pump is denoted by W. As the device operates in cycles, the heat flows and work should be understood as referring to one cycle. Let us now see what we can learn about our heat pump from some basic physics.

What the First Law of Thermodynamics Tells Us About a Heat Pump

According to the first law, the internal energy of a system can only be altered by the exchange of energy with its surroundings, that is, by heat Q or work W, and so we can write

$$\textit{Change in Internal Energy} = Q + W.$$

Heat flow *to* the system must be regarded as positive, while flow *from* the system is negative. Work performed *on* the system must be thought of as positive, and work performed *by* the system as negative.

Since the heat pump operates in cycles and at the end of a cycle returns to the same state as at the beginning, we can see that the change in internal energy in one cycle must be zero. Then, with careful attention to what is positive and what is negative, the first law tells us that

$$0 = Q_{cold} - Q_{warm} + W,$$

or

$$Q_{warm} = Q_{cold} + W.$$

The last equation shows us immediately what is so attractive about heat pumps: Part of the heat that it delivers is coming from outside the house and is free! Of course, to make this happen, you must put work into the heat pump, and the work costs money. This brings us to the question of the performance of heat pumps.

Defining the Performance

What we probably care about most is how much heat the pump provides per dollar spent on work. The heat provided is measured by Q_{warm}, while the cost of the work is proportional to W. This suggests that we measure the performance by the ratio Q_{warm}/W. It will be useful to have a name for this ratio, and the word "efficiency" might come to mind. This word has, however, long been associated with heat engines. In them, there is a *single* energy flow into the device, in the form of heat, and two energy flows out of the device, one in the form of work and the other in the form of heat. The efficiency expresses the work as a fraction of the heat that flows into the device and is necessarily less than 1. In the case of a heat pump, however, there are *two* energy flows into the device, in the forms of heat and work, and a single energy flow out of the device, in the form of heat. Thus, we get out more heat than we put in work, and this makes our ratio greater than 1. An "efficiency" that is greater than 1 seems impossible, given the context in which the term is defined. Worse, it may suggest some sort of deep mystery or swindle where there is neither. For these reasons, a different name, "coefficient of performance," is used. Thus, for a heat pump, we can say that the

$$\text{Coefficient of Performance} = Q_{warm}/W.$$

What the Second Law of Thermodynamics Tells Us About a Heat Pump

For a given Q_{warm}, better performance means a smaller W, and the smallest W that we can imagine is none at all, that is, zero. But is that possible? If it were, then heat would flow from the cold outside to the warm inside without any help at all. The second law of thermodynamics rules this out because it states, in one form, something that you probably already know:

Heat will not spontaneously flow from a lower temperature to a higher temperature.

Nature must then impose an upper limit, but how do we find it?

We begin with the observation that all practical heat pumps have wasteful imperfections, such as friction and turbulence in the working fluid. An ideal heat pump that lacks these imperfections would, we expect, have better performance and thus might help us find the theoretical upper limit to heat pump performance. It will not be a practical heat pump because, among other things, it will have to operate very slowly to avoid turbulence. Still, its existence will not violate any law of nature, and so we may consider it.

An important aspect of the ideal heat pump is that it will be *reversible*, meaning that it can equally well go the other way around its cycle. All energy flows will then be reversed but unchanged in magnitude. The device will produce work rather than requiring it and so will act as a heat engine. We can then imagine connecting it to a hypothetical heat pump so that it supplies the work required to run the pump. If we now assume that the hypothetical heat pump has a higher coefficient of performance than the heat engine did when acting as an ideal heat pump, we find that the combination of the two devices will violate the second law of thermodynamics in the form stated above. Therefore, the coefficient of performance of the hypothetical heat pump must be less than or equal to that of the ideal one. If the hypothetical pump is reversible, it is easy to show that its coefficient must be the same. If it is practical and hence not reversible, a more general version of the second law (employing entropy considerations) shows that the coefficient of performance of the hypothetical heat pump must be less than that of the ideal one. These are very important conclusions, so let us state them again:

All ideal reversible heat pumps operating under the same conditions have the same *coefficient of performance.*

All practical and hence irreversible heat pumps have lower coefficients of performance than an ideal reversible heat pump operating under the same conditions.

Now all that we have to do is calculate the coefficient of performance of some ideal reversible heat pump and we will have our upper limit for all heat pumps. That is next.

The Carnot Cycle

The Carnot (pronounced "car-no") cycle was originally devised to study the upper limit of the efficiency of a heat engine but will work equally well for a heat pump if we just execute the cycle in reverse. We imagine that the working fluid of the pump is a gas of a particularly simple sort, an ideal gas. Though no real gas is actually ideal, most gases are nearly so under fairly typical conditions. The gas is confined in a cylinder with a movable piston, as shown in the diagram below.

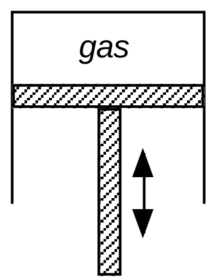

When the piston moves up, it compresses the gas and does work on it. When it moves down, the reverse happens, and the gas does work on the piston (and whatever is connected to it). The piston is insulated so that no heat can flow through it. Any heat flow between the gas and its environment then takes place solely through the walls of the cylinder, which can be insulated or not as required. When not insulated, the walls are held at whatever temperature is necessary, and heat can then flow into or out of the gas while its temperature remains fixed at the temperature of the walls or nearly so.

Let us begin our cycle with the piston in its lowermost position. The following four steps (which bear careful reading) are then carried out: (1) With the cylinder insulated and the gas initially at temperature T_{cold}, the piston is slowly pushed part way up. This requires work, compresses the gas, and causes its temperature to increase to T_{warm}. (2) With the cylinder not insulated but held at the temperature T_{warm}, the piston is slowly pushed to its uppermost position. This requires more work and further compresses the gas. The compression causes the temperature of the gas to start to rise above T_{warm}, but this in turn causes heat Q_{warm} to flow out of the gas, thereby keeping its temperature slightly above T_{warm}. (3) With the cylinder insulated, the piston is slowly pulled part way down. The gas does work on the piston and expands as it does so. This causes the temperature of the gas to decrease to T_{cold}. (4) With the cylinder not insulated but held at the temperature T_{cold}, the piston is slowly pulled down to its lowermost position. The gas does more work on the piston and expands further. This causes the temperature of the gas to start to decrease below T_{cold}, but this in turn causes heat Q_{cold} to flow into the gas, thereby keeping its temperature slightly below T_{cold}. At this point, we are back to where we started, and the cycle is complete. Work was done on the gas when the piston moved up and by the gas when the piston moved down. It is not hard to see, however, that during each upward movement, the gas pressure was higher than during the counterpart downward movement. Thus, more work was done on the gas than by the gas, and so the net effect is that work W was done on the gas.

The cycle just described shows how a heat pump, though admittedly not a practical one, could work. An analysis of the energy flows then yields the following simple but very important result:

$$Q_{cold} = (T_{cold}/T_{warm})Q_{warm}.$$

This result, we should be reminded, would be found for *any* ideal reversible heat pump operating between the same temperatures. Let us now make use of it to find the upper limit of the coefficient of performance of a heat pump.

The Upper Limit of the Coefficient of Performance

Earlier we learned that the first law of thermodynamics told us that

$$0 = Q_{cold} - Q_{warm} + W.$$

We may rewrite this in the form

$$W = Q_{warm} - Q_{cold},$$

and then substitute the result given at the end of the previous section for Q_{cold} to obtain

$$W = Q_{warm} - (T_{cold}/T_{warm})Q_{warm} = [1 - (T_{cold}/T_{warm})]Q_{warm}.$$

This then yields

$$Q_{warm}/W = 1/[1 - (T_{cold}/T_{warm})].$$

But the ratio on the left is the coefficient of performance, and the expression on the right is its upper limit, the maximum permitted by nature. Therefore, we can can now say that

The upper limit of the coefficient of performance of a heat pump operating between the temperatures T_{cold} and T_{warm} is $1/[1 - (T_{cold}/T_{warm})]$.

We note again that the temperatures must be *absolute* temperatures.

It may be useful now to see some values for the upper limit of the coefficient of performance in various climates. The table below gives these values for the month of January in three cities: Saint Louis, Missouri; Des Moines, Iowa; and Minneapolis, Minnesota. We assume that the temperature inside the house is 70ºF, corresponding to an absolute temperature of about 70 + 460 = 530 R. The heat pump is assumed to extract its heat from the air outside the house (an "air-source" heat pump), and for its temperature we use the monthly mean air temperature for each city.

City	Upper Limit of Coefficient of Performance
Saint Louis	*12.0*
Des Moines	*9.5*
Minneapolis	*8.1*

It should be kept in mind that the performance of a heat pump actually varies as the temperature outside a house varies during the day, from day to day, and from season to season. Nonetheless, we can draw the useful and correct conclusion from the table that the colder the climate, the less well does a heat pump perform. The reason, of course, is that it becomes harder and harder to move heat from a lower and lower temperature outside the house to the warm temperature inside. This has given rise to heat pumps that extract their heat from underground sources ("ground-source" heat pumps). The temperature underground below the frost line is roughly 50ºF everywhere. For that temperature, we find that the upper limit of the coefficient of performance is about 25, surely a great improvement over the values shown in the table. Finally, we note that the imperfections, some of them unavoidable, in presently available air-source heat pumps might reduce their coefficients of performance to something of the order of 4 in a climate like that of Des Moines. Still, even a coefficient of performance of 4 means that for each unit of work that goes into the heat pump, 3 units of free heat from outside the house combine with the 1 unit of work to provide 4 units of heat to the house.

Further Information and Issues

The best widely available source of further information is the Internet. A search by the author using "heat pump" and the Google search engine produced about 645,000 references. Some of them would undoubtedly be of interest, particularly in connection with very important practical issues not addressed here. Such an issue, for example, is the problem that as the temperature outside a house drops, the need for heat rises but the performance of a heat pump drops. This means that unless the heat pump is large enough, supplementary heating may be needed. But which is better, a large heat pump or supplementary heating? And then there is the important but changing issue of competing methods of heating. Which is best in a given climate, and what exactly does "best" mean?