Heat Engines: the Carnot Cycle

Flashlet here!

Michael Fowler Physics 142E 04/21/09

The Ultimate in Fuel Efficiency

All standard heat engines (steam, gasoline, diesel) work by supplying heat to a gas, the gas then expands in a cylinder and pushes a piston to do its work. So it's easy to see how to turn heat into work, but that's a one shot deal. The catch is that the heat and/or the gas must somehow then be dumped out of the cylinder before the next cycle: otherwise, all the work it delivered on expanding will be needed to compress it back!

Our aim in this lecture is to figure out just how efficient such a heat engine *can* be: what's the most work we can possibly get for a given amount of fuel? We'll examine here the simplest cyclical model: an ideal gas is enclosed in a cylinder, with external thermal connections to supply and take away heat, and a frictionless piston for the gas to perform (and to absorb) mechanical work:

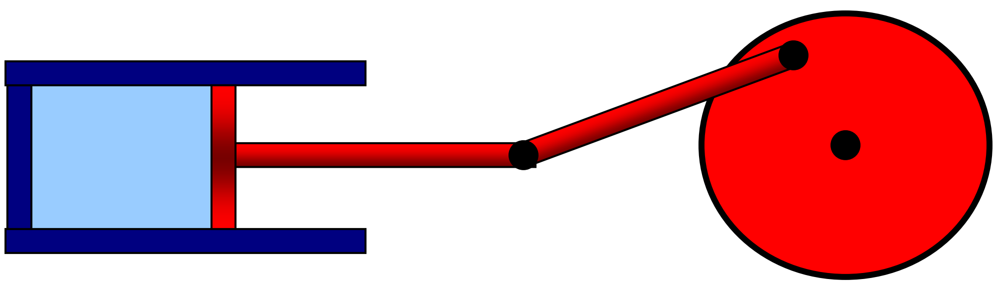

Carnot Engine

The efficiency question was first posed—and solved—by Sadi Carnot in 1820. At that time, steam engines had become efficient enough to replace more traditional sources of power in factories.

What were the sources of power before the steam engine? People, horses, windmills, watermills.

The first factory—the first organized mass production—was put together in the 1770's by Richard Arkwright. He's figured out how to automate weaving, so semiskilled labor (often children) could be used to run machines producing cloth. Rows of these machines could be powered from a long rotating rod, in the first factories this was connected to a water wheel.

Obviously, profits were maximized if the water wheel was made as efficient as possible, and a lot of engineering and thought went into figuring out how best to do this. The water loses potential energy as it is carried down by the wheel, so the most energy possible is *mgh* watts, where m is the mass of water flowing per second.

How is energy wasted? First, we need as little friction in the wheel as possible. We need smooth flow: no water splashing around. The water must flow into and off the wheel without dropping any significant height, or it loses that much potential energy without producing work.

A **perfect** water wheel would be **reversible**: it could be used to drive a copy of itself backwards, to lift up the same amount of water per second that fell.

A Modern Water Wheel in Virginia

There is in Virginia a pretty efficient water wheel: it's about 80% efficient—the Bath County hydroelectric pumped storage station. This is a water wheel, actually a turbine, but that amounts to the same thing better designed, that works both ways. Water from an upper lake falls through a pipe to a turbine and the lower lake, generating electrical power. Alternatively, electric power can be supplied to pump the water back up. Why bother? Because demand for electricity varies, and it's better to avoid if possible building power stations that are only running during peak demand. It's cheaper to store power at times of low demand.

The drop *h* is about 1200 feet, 380 meters. The flow rate is about a thousand tons a second. The plant generates about 2 Gigawatts, roughly the same as a two-unit nuclear plant.

Carnot's Idea: a "Water Wheel" for Heat

About two hundred years ago, in Carnot's time, heat was widely believed to be a fluid, called the "caloric fluid" that flowed inside materials from hot to cold, just as the electric "fluid" flowed from high potential to low potential. We now know, of course, that they were right about electric current, but wrong about heat. Nevertheless, the picture of fluid flow accurately describes heat flow as long as no work is done—all discussions of thermal conductivity, or the heat flow in cooking, etc., are still best pictures in those terms.

Anyway, following the analogy of electric current flowing from high potential to low potential, or water flowing downhill to lower gravitational potential, Carnot figured that for heat flow, the temperature was the appropriate potential, and the steam engine was like a water wheel for this caloric fluid, so the most efficient engine would have minimal friction, but also, in analogy with the water entering and leaving the wheel gently with no intermediate loss of height, the heat would enter and leave the gas in the engine isothermally (remember the temperature is analogous to the gravitational potential, thus the height). Therefore, by analogy with mgh, the drop in temperature $T_H - T_C$ measures the potential energy given up by the "heat fluid".

But how does that relate to the energy expended producing the heat in the first place? Well, Carnot knew something else: there was an absolute zero of temperature. Therefore, he reasoned, if you cooled the fluid down to absolute zero, it would give up all its heat energy. So, the maximum possible amount of energy you can extract by cooling it from T_H to T_C is, what fraction is that of cooling it to absolute zero?

It's just

$$(T_H - T_C)/T_H.$$

Of course, the caloric fluid picture isn't right, but **this result is**! This **is** the maximum efficiency of a perfect engine.

Getting work out of a Hot Gas

Now let's turn to the details of getting work out of a heated gas. We want the process to be as close to reversible as possible: there are two ways to move the piston reversibly: isothermally, meaning heat gradually flows in or out, from a reservoir at a temperature infinitesimally different from that of the gas in the piston, and adiabatically, in which there is no heat exchange at all, the gas just acts like a spring.

So, as the heat is supplied and the gas expands, the temperature of the gas must stay the same as that of the heat supply (the "heat reservoir"): the gas is expanding *isothermally*. Similarly, it must contract isothermally later in the cycle as it sheds heat.

To figure out the efficiency, we need to track the engine through a complete cycle, finding out how much work it does, how much heat is taken in from the fuel, and how much heat is dumped in getting ready for the next cycle. You might want to look at the flashlet to get the picture at this point: the cycle has four steps, an isothermal expansion as heat is absorbed, followed by an adiabatic expansion, then an isothermal contraction as heat is shed, finally an adiabatic contraction to the original configuration. We'll take it one step at a time.

Step 1: Isothermal Expansion

So the first question is: *How much heat is supplied, and how much work is done, as the gas expands isothermally*? Taking the temperature of the heat reservoir to be T_H (*H* for hot), the expanding gas follows the isothermal path $PV = nRT_H$ in the (P, V) plane.

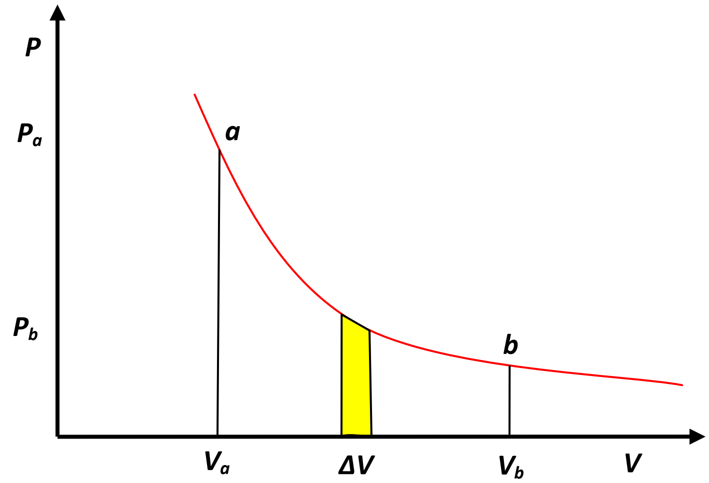

Isothermal expansion from *a* to *b* along $PV = nRT_H$

The work done by the gas in a small volume expansion ΔV is just $P\Delta V$, the area under the curve (as we proved in the last lecture).

Hence the work done in expanding isothermally from volume V_a to V_b is the total area under the curve between those values,

$$\text{work done isothermally} = \int_{V_a}^{V_b} PdV = \int_{V_a}^{V_b} \frac{nRT_H}{V} dV = nRT_H \ln \frac{V_b}{V_a}.$$

Since the gas is at constant temperature T_H, there is no change in its internal energy during this expansion, so the total heat supplied must be $nRT_H \ln \frac{V_b}{V_a}$, the same as the external work the gas has done.

In fact, this isothermal expansion is only the first step: the gas is at the temperature of the heat reservoir, hotter than its other surroundings, and will be able to continue expanding even if the heat supply is cut off. To ensure that this further expansion is also reversible, the gas must not be losing heat to the surroundings. That is, after the heat supply is cut off, there must be no further heat exchange with the surroundings, the expansion must be *adiabatic*.

Step 2: Adiabatic Expansion

By definition, no heat is supplied in adiabatic expansion, but work is done.

The work the gas does in adiabatic expansion is like that of a compressed spring expanding against a force—equal to the work needed to compress it in the first place, for an ideal (and perfectly insulated) gas. So adiabatic expansion is *reversible*.

$W_{\text{adiabat}} = nC_V(T_c - T_b)$—this is the loss of internal energy by the gas on expanding against the external pressure.

Steps 3 and 4: Completing the Cycle

We've looked in detail at the work a gas does in expanding as heat is supplied (isothermally) and when there is no heat exchange (adiabatically). These are the two initial steps in a heat engine, but it is necessary for the engine to get back to where it began, for the next cycle. The general idea is that the piston drives a wheel (as in the diagram at the beginning of this lecture), which continues to turn and pushes the gas back to the original volume.

But it is also **essential for the gas to be as cold as possible on this return leg**, because the **wheel** is now having to expend work on the **gas**, and we want that to be as little work as possible—it's costing us. The colder the gas, the less pressure the wheel is pushing against.

To ensure that the engine is as efficient as possible, this return path to the starting point (P_a, V_a) must also be reversible. We can't just retrace the path taken in the first two legs, that would take all the work the engine did along those legs, and leave us with no net output. Now the gas cooled during the adiabatic expansion from *b* to *c*, from T_H to T_C, say, so we can go some distance back along the reversible colder isotherm T_C. But this won't get us back to (P_a, V_a), because that's on the T_H isotherm. The simplest option—the one chosen by Carnot—is to proceed back along the cold isotherm to the point where it intersects the adiabat through *a*, then follow that adiabat back to *a*. That route keeps the gas as cold as possible for as long as possible, minimizing the external work needed to get it back to the original state.

To picture the Carnot cycle in the (*P*, *V*) plane, recall from the previous lecture the graph showing two isotherms and two adiabats:

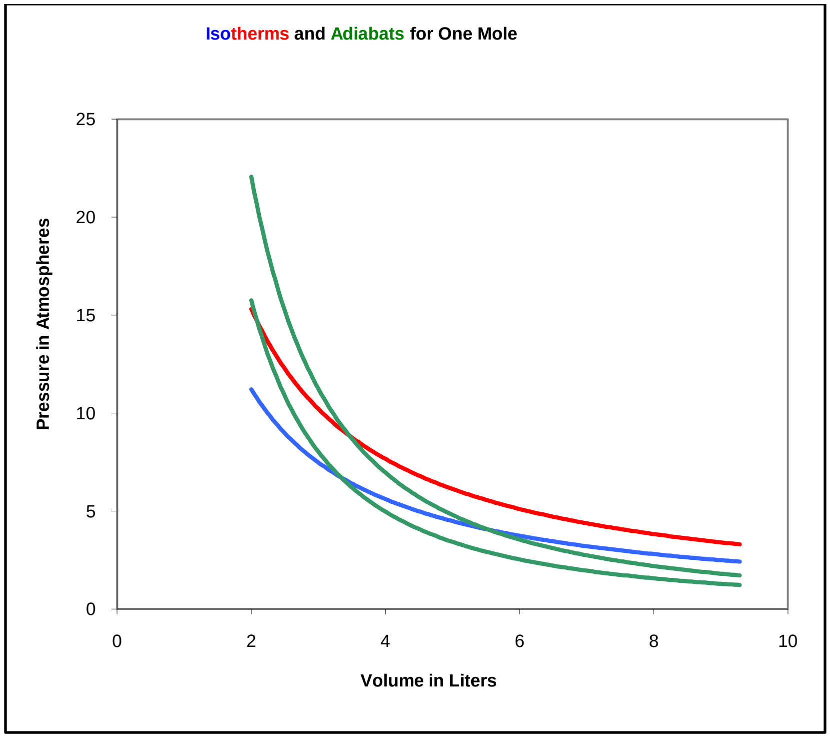

Carnot's cycle is around that curved quadrilateral having these four curves as its sides.

Let us redraw this, slightly less realistically but more conveniently:

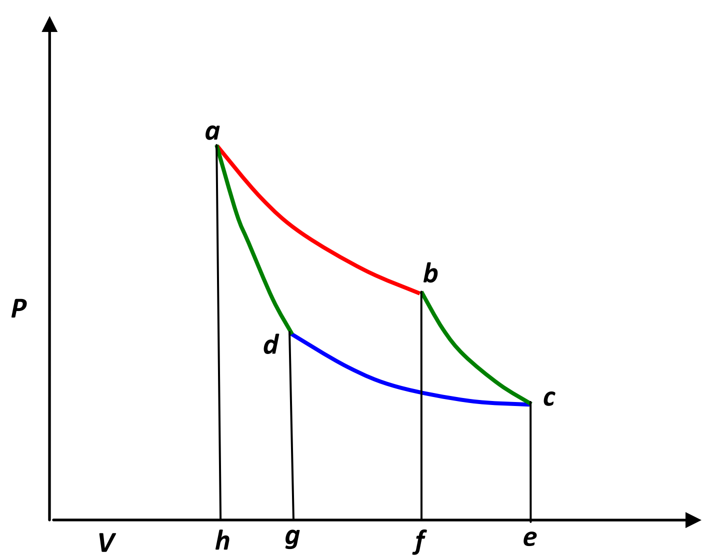

Efficiency of the Carnot Engine

In a complete cycle of Carnot's heat engine, the gas traces the path *abcd*. The important question is: what fraction of the heat supplied from the hot reservoir (along the red top isotherm), let's call it Q_H, is turned into mechanical work? This fraction is called the *efficiency* of the engine.

Since the internal energy of the gas is the same at the end of the cycle as it was at the beginning—it's back to the same P and V—it must be that the work done equals the net heat supplied,

$$W = Q_H - Q_C$$

Q_C being the heat dumped as the gas is compressed along the cold isotherm.

The efficiency is the fraction of the heat input that is actually converted to work, so

$$\text{Efficiency} = W/Q_H = (Q_H - Q_C)/Q_H$$

This is the answer, but it's not particularly useful: measuring heat flow, especially the waste heat, is quite difficult. In fact, it was long believed that the heat flow out was equal to that flowing in, and this was quite plausible because the efficiency of early engines was very low.

But there's a better way to express this.

Now the heat supplied along the initial hot isothermal path *ab* is equal to the work done along that leg, (from the paragraph above on isothermal expansion):

$$Q_H = nRT_H \ln \frac{V_b}{V_a}$$

and the heat dumped into the cold reservoir along *cd* is

$$Q_C = nRT_C \ln \frac{V_c}{V_d}.$$

$Q_H - Q_C$ looks complicated, but actually it isn't!

The expression can be greatly simplified using the adiabatic equations for the other two sides of the cycle:

$$T_H V_b^{\gamma-1} = T_C V_c^{\gamma-1}$$
$$T_H V_a^{\gamma-1} = T_C V_d^{\gamma-1}.$$

Dividing the first of these equations by the second,

$$\left(\frac{V_b}{V_a}\right)=\left(\frac{V_c}{V_d}\right)$$

and using that in the preceding equation for Q_C,

$$Q_C = nRT_C \ln\frac{V_a}{V_b} = \frac{T_C}{T_H} Q_H.$$

So for the Carnot cycle the ratio of heat supplied to heat dumped is just the ratio of the absolute temperatures!

$$\frac{Q_H}{Q_C}=\frac{T_H}{T_C}, \quad \text{or} \quad \frac{Q_H}{T_H}=\frac{Q_C}{T_C}.$$

Remember this: it'll be *important in developing the concept of **entropy***.

The work done can now be written simply:

$$W = Q_H - Q_C = \left(1-\frac{T_C}{T_H}\right)Q_H.$$

Therefore **the efficiency of the engine, defined as the fraction of the ingoing heat energy that is converted to available work, is**

$$\text{efficiency} = \frac{W}{Q_H} = 1-\frac{T_C}{T_H}.$$

These temperatures are of course in degrees Kelvin, so for example the efficiency of a Carnot engine having a hot reservoir of boiling water and a cold reservoir ice cold water will be $1-(273/373)=0.27$, just over a quarter of the heat energy is transformed into useful work. This is the very same expression Carnot found from his water wheel analogy.

After all the effort to construct an efficient heat engine, making it reversible to eliminate "friction" losses, etc., it is perhaps somewhat disappointing to find this figure of 27% efficiency when operating between 0 and 100 degrees Celsius. Surely we can do better than that? After all, the heat energy of hot water is the kinetic energy of the moving molecules, can't we find some device to channel all that energy into useful work? Well, we can do better than 27%, by having a colder cold reservoir, or a hotter hot one. But there's a limit: we can never reach 100% efficiency, because we cannot have a cold reservoir at $T_C = 0K$, and, even if we did, after the first cycle the heat dumped into it would warm it up!