

The ideal gas law and a bit of thermodynamics

The stuff that is inside the cylinder of an internal combustion engine is a gas. That is to say, it is not a liquid or a solid. During certain periods of a cycle, the gas is compressed. During other periods, it expands. The composition of the gas is not constant, either. The very act of burning changes the types of molecules in the gas. But, at all times, the contents of the cylinder are in a gaseous state.

We can imagine that the gas consists of a huge number of little particles which continuously zip around at very high speed, bouncing off each other and off the walls of the cylinder. ("Walls" includes the cylinder head and the piston head as well as the side walls of the cylinder). During the compression stroke, the particles will be the molecules which are found in air, namely, molecules of nitrogen, oxygen, argon and trace amounts of some other gases, as well as molecules of the fuel. Combustion will produce molecules of carbon dioxide and water, but the cylinder will also contain all of the ingredients of the original air which did not participate in the combustion. If the combustion was lean, there will even be some unreacted molecules of oxygen. If the combustion was rich, there will be unreacted molecules of the hydrocarbon fuel. All of these molecules can be thought of as being the little billiard balls in a zoo of activity.

Despite the compression and combustion, the gases in an internal combustion engine are at a low enough pressure that we can treat the gas like an "ideal" gas. An ideal gas is one in which the size of the individual particles is much less than the average distance they travel between collisions. Only when the pressure gets up into the hundreds of atmospheres, as it does inside the barrel of an artillery gun, does the ideal gas assumption fail. When the pressure gets so high that the ideal gas assumption fails, it is necessary to take into account the size of the particles. Fortunately, we will not have to do that.

When we talk about a closed container which holds a gas, there are four physical quantities that describe the state of the gas. There is the pressure P , the volume V and the temperature T . There is also the "amount" of gas in the container, which we could measure as its mass or as the number of particles.

What is amazing about an ideal gas is that the relationship between these four quantities does not depend on the type of molecule. A certain number of oxygen molecules inside a vessel with a given volume at a given temperature will have exactly the same pressure as the same number of xenon molecules inside a vessel of the same size at the same temperature. When you think about it, it really is amazing that the type of gas does not enter into the picture.

The reason why we are so excited about being able to model the gases inside an internal combustion engine as an ideal gas is that the governing relationship is so simple.

$$P \times V = R \times N \times T$$

Pressure \times Volume = Constant \times Quantity \times Temperature

This is the Ideal Gas Law. The pressure P multiplied by the volume V , which is the value of the left-hand side of the equation, must be equal to the temperature T multiplied by the number of particles in the container (N) and a constant R , on the right-hand side. If we know any three of the four variables P , V , T and N , we can calculate the fourth one. We never have to worry about R . It is a constant, and I will tell you its value in just a moment.

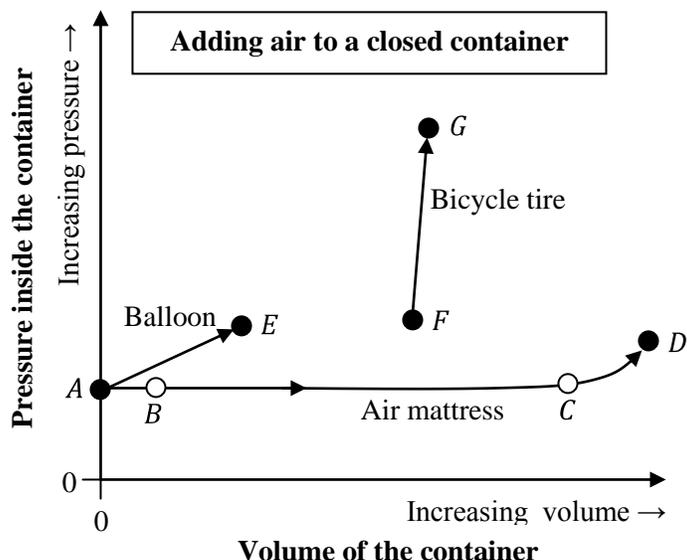
The pressure, volume and temperature which this equation uses must be measured in absolute terms. The pressure, for example, is the absolute pressure inside the vessel, not the difference between the pressure inside and outside. The temperature, too, is the absolute temperature, measured above "absolute zero". Absolute zero is such a low temperature that the atoms and molecules do not move with respect to each

other and the atoms inside a molecule do not vibrate. We measure absolute temperature in degrees Kelvin. One Kelvin degree has exactly the same size as one Celsius degree, but the two scales are shifted with respect to each other. The 0°K point is below – far below – the 0°C point. The two scales are related by $0^{\circ}\text{K} = -273.15^{\circ}\text{C}$.

So, what are N and R ? R is simply a constant. It is the constant of proportionality between the two sides of the equation. It has a known value, which is $R = 8.314 \text{ Joules}/^{\circ}\text{K-mole}$. And that word "mole" brings us to the last variable N . N is the number of particles of gas inside the vessel. The number of particles will be extremely large, of course, but it will still be just a number, like eight or thirteen, but millions of millions of times bigger. For historical reasons, the Ideal Gas Law is stated in such a way that one does not just substitute the number of particles for N . Instead, one states the number of particles in terms of moles. A mole is just a big number, 6.022×10^{23} . A cylinder which contains 12.044×10^{23} molecules therefore contains two moles and it is the number $N = 2$ which we would substitute into the Ideal Gas Law.

The Ideal Gas Law makes intuitive sense. Imagine adding more gas to the vessel (increasing N on the right-hand side) and keeping the temperature of the gas steady (leaving T unchanged on the right-hand side). Since the right-hand side increases, the left-hand side must increase, too. Either the pressure P has to increase or the volume V has to increase, or both.

I can give you three common examples which illustrate what happens when we add molecules to a closed container. We are going to blow up an air mattress, a balloon and a flat bicycle tire. As we run through the examples, follow the points on the graph below. In the graph, the horizontal direction measures the volume of the air in these containers. The left end of the horizontal axis, where $V = 0$, corresponds to no air inside. The vertical axis measures the absolute pressure of the air inside. To make things interesting, the pressure that corresponds to point A in the graph is one atmosphere.



Let's start with an empty air mattress, at point A . We have just removed it from its original packaging, in which it was rolled up flat, so there is virtually no air inside. What little air is inside is at an absolute pressure of one atmosphere. That happens to equal the pressure on the outside, so there is no pressure differential between the outside and inside. We take a huge breath and blow in through the nipple. The air inside the mattress moves from point A in the graph to point B . The horizontal jump is the volume of air which our lungs held. The absolute pressure inside the mattress does not increase because the mattress has much greater capacity than our lungs. We will not even need to hold the nipple closed – very little of the air we blew into

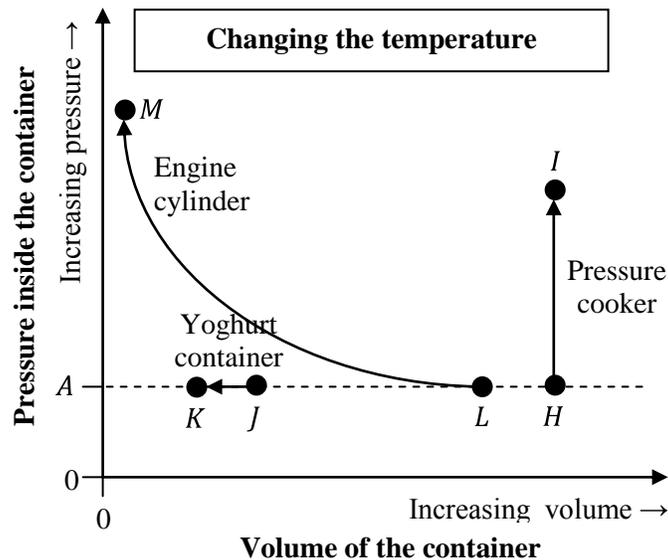
the mattress in our first blow will come back out. We will have to blow in a couple of dozen huge breaths, or perhaps even more, before the air mattress even starts to take on its final shape, somewhere near point *C* on the curve. It is only during the last few breaths that the mattress will reach its nominal physical volume. The fabric of the mattress might be slightly elastic, so our last few breaths will increase both the volume and the pressure. This is the uptick in the curve as the state of the air moves to point *D*. The pressure increases as the tension in the fabric resists the increase in volume. If we let go of the nipple now, the air inside will start to rush out.

The curve from point *A* to point *E* traces the air as we blow up a small balloon. The balloon is made of rubber. It is much more elastic than the air mattress. When there is no air inside, the balloon lies flaccid. Our first breath puts enough air inside to start to expand the balloon and increase the pressure over one atmosphere. We will have to squeeze the neck of the balloon even after the first blow to keep the air inside. Since the balloon is a lot smaller than the air mattress, we do not have to add as much air and will not move as far towards the right to complete the inflation process.

Point *F* on the graph is a flat bicycle tire. Actually, the tire is not completely empty; it is just "soft". The air pressure inside is greater than one atmosphere, so that point *F* is higher up in the graph than point *A*, but not by very much. The rubber of the tire or inner tube is quite stiff, much stiffer than the walls of the balloon. This means that the pressure inside will rise much more for each given amount of air added than the balloon did. On the graph, the tire's curve is steeper than the balloon's curve.

Since the vertical axis measures pressure and the horizontal axis measures volume, we can call this a pressure-versus-volume, or *P-V*, curve. While it can help visualize or understand the changes in a gas from one state to another, it cannot give a complete picture. The graph does not show the other two variables – temperature and quantity – which are needed to describe fully the characteristics of a gas. For example, compare the balloon curve to the bicycle curve in the graph above. The volume of the balloon changed a lot more than the volume of the bicycle tire, but we probably added three or four times as much air to the tire. The graph does not show how much air – how many molecules or moles – we added. The speed with which we worked the bicycle tire pump may have affected the temperature of the air inside the bicycle tire and that, too, is not shown on the graph.

The following *P – V* graph shows three more sample transitions, all of which are caused by changing the temperature of the gas inside a closed container. The dotted horizontal line represents one atmosphere pressure, the same vertical position as point *A* in the previous graph.



Point *H* represents the air inside a pressure cooker at the time when the cook gives the lid a half turn to seal it to the pot and places the pot on the stove. There may be potatoes and water in the pot, too, but we are interested in the air inside the chamber above the surface of the water. It will be a few minutes before the water starts boiling. During that period of time, the air will heat up. Since the volume is fixed (water is not compressible), the pressure of the air will increase. After a time, the water will start boiling. Two things happen when the water starts boiling. First, it adds more "quantity" to the gas inside the pressure cooker. When the pot was first placed on the stove, the gas consisted entirely of ambient air. Once the water boils, bubbles of steam are created. When the bubbles rise to the surface, they break and release their contents. Steam is nothing more than water vapour – molecules of H_2O in gaseous form. These molecules of water will join the molecules of whatever air was already there. Since there are now more molecules of gas in the same volume (even as it boils, the volume of water will decrease only very slowly), the pressure of the gas will go up. Secondly, the temperature of the gas will continue to rise, up to a point, and this increase in temperature will increase the pressure as before. The boiling temperature of water increases as the pressure increases, so the temperature of the gas inside the pressure cooker will be considerably higher than the 100°C we are all familiar with. The main reason for using a pressure cooker is to obtain this higher temperature, which cooks the potatoes faster.

Now, let's look at the curve from point *J* to point *K*. The object is a single serving cup of yoghurt, in a small plastic cup with a thin flexible sealed metal cover. When the product was manufactured, the yoghurt was poured into the cup at normal room temperature and pressure. The small bit of air above the surface of the yoghurt is always at one atmosphere of pressure. When you put the cup into the fridge, and the temperature decreases, the volume of the trapped air will decrease as well. The air will move to point *K* along the curve. As the trapped air shrinks in volume, the lid will become slightly depressed. The flexible lid allows the trapped volume to change and thereby equalize the pressure inside the cup with the pressure outside.

Now, look at the third curve, representing one cylinder of an internal combustion engine. The trajectory from point *L* to point *M* is similar to, but not necessarily the same as, a compression stroke. At point *L*, the piston is at the bottom of the cylinder, and the enclosed volume of the cylinder is at its maximum. Let's assume the engine does not have a turbocharger, and that the cylinder is filled full of a gasoline/air mixture at ambient pressure. We will now compress the cylinder, reducing its volume. Alert readers will observe that the volume at point *M* is about one-fifteenth of the volume at point *L*, meaning this engine has a compression ratio of 15:1. As the gas inside the cylinder is compressed, its pressure will increase. Since this is exactly what happens during the compression stroke of an internal combustion engine, why did I say this is "not necessarily the same as a normal compression stroke"? The answer to that question involves the Ideal Gas Law.

When the intake valve(s) of the cylinder close, a certain amount of gas (air in the case of a diesel engine and a gasoline/air mixture in the case of a gasoline engine) is trapped. No gas is added or released from the cylinder during the compression stroke. That quantity of gas, which we will say is N_0 moles, is going to be constant during the compression stroke. The Ideal Gas Law must hold true at every instant during the compression stroke. Since the quantity of gas is fixed, the Ideal Gas Law can be expressed as follows:

$$P \times V = (R \times N_0) \times T$$

The term in round brackets, being the number of moles trapped in the cylinder multiplied by the Ideal Gas Constant, is a constant during compression. The only variables are the pressure P , the volume V and the temperature T . Before we can draw a curve on the P - V graph, we have to make an assumption about what happens to the temperature. Does the temperature stay constant during the compression stroke? Does it rise?

The most obvious assumption is that the temperature stays constant, at its starting temperature, say T_0 . If so, then the relationship which holds at every instant during the compression stroke can be written as:

$$P \times V = (R \times N_0 \times T_0) = \text{a constant value}$$

The entire right-hand side of the equation is a constant. The pressure and volume are inversely related to one another. If the volume shrinks to ½, or 50%, of its starting volume, then the pressure will increase by a factor of 2. If the volume shrinks to ¼, or 25%, of its starting volume, the pressure will increase by a factor of 4. If the volume shrinks to one-fifteenth of its starting volume, the pressure will increase by a factor of 15. The consequences of this assumption are easy to understand. A process during which the temperature stays constant is called an "isothermal" process.

The problem is that the compression stroke in an internal combustion engine is not an isothermal process. Almost all isothermal processes require that the gas be able to exchange energy with the container that holds it. The container must be able to provide energy to the gas or remove it from the gas in order for the gas to be able to remain at a constant temperature.

Things happen too quickly in an internal combustion engine for there to be any meaningful exchange of energy between the trapped gas and the walls. An alternative, and much more realistic, assumption about what happens during the compression stroke is that the gas does not exchange energy with the container. Since the gas always has the same amount of energy, one might be tempted to call this an "iso-energetic" process. The powers that be prefer to call it an "adiabatic" process. (A name like that just has to have some Greek or Latin root, but I don't know what it is.)

Philosophical aside: The Ideal Gas Law holds at every instant during the compression stroke. But, since it refers to three different variables, the Law is not sufficient by itself to allow us solve for P , V and T . It is true that we know how the volume V changes as the piston head moves up the cylinder, but there remain two unknowns, the pressure P and the temperature T . The Ideal Gas Law gives us one relationship between P and T , but we still need a second relationship to be able to solve for two variables. An assumption about the type of process involved – isothermal, adiabatic, isobaric or something else – gives us the second relationship we need to get a unique solution.

Fortunately, the equation which describes the state of energy in an ideal gas as it undergoes an adiabatic transition is very simple:

$$P \times V^\gamma = \text{a constant value}$$

The arithmetic for calculating how an ideal gas behaves during the compression stroke becomes a matter of manipulating the two following equations:

Ideal Gas Law for a trapped quantity of gas	$P \times V = (R \times N_0) \times T$
Adiabatic Law for constant gas energy	$P \times V^\gamma = \text{constant}$

Although the equation describing adiabatic behaviour is easy, the thinking behind it is not. Central to the theory is recognizing that a gas has two types of energy. There is the translational kinetic energy which the individual particles (molecules) of the gas have as they travel through space. There is also internal energy inside the molecules. The atoms inside a molecule can move with respect to each other. It may help to think about the relative motions of atoms inside a molecule as vibrations or rotations around some axis. Vibrations and rotations are also kinetic energy, meaning "energy of motion". As an ideal gas is compressed, the work done by the piston head is converted into both types of kinetic energy. The speed of the molecules, treated as particles, increases. But, the speed at which the atoms inside the molecules vibrate and rotate also increases. The adiabatic assumption accounts for the effects of both types of kinetic energy.

Since the internal energy arises from atoms moving with respect to each other, it should be the case that molecules which consist of single atoms do not have this kind of internal energy. That is true. At the other end of the scale, molecules which contain lots of atoms have lots of different vibration modes and can absorb more internal energy than simpler molecules made from fewer atoms. It is this measure of complexity – the number of atoms inside the molecules – which determines the exponent γ in the Adiabatic Law. Let's look at the value of γ for different molecules. The following table shows a few.

Type of molecule	Degrees of freedom	γ (Adiabatic Index)
Monatomic – one atom: <i>He, Ar</i>	3	$(3 + 2)/3 = 5/3 = 1.667$
Diatomic – two atoms: <i>O₂, N₂</i>	5	$(5 + 2)/5 = 7/5 = 1.400$
Triatomic, collinear – <i>CO₂</i>	5	$(5 + 2)/5 = 7/5 = 1.400$
Triatomic, noncollinear – <i>H₂O</i>	6	$(6 + 2)/6 = 8/6 = 1.333$
Methane – <i>CH₄</i>	≈ 6.25	1.32
Ethane – <i>C₂H₆</i>	≈ 9.09	1.22
Propane – <i>C₃H₈</i>	≈ 15.4	1.13
Pentane – <i>C₅H₁₂</i>	≈ 25.0	1.08
Hexane – <i>C₆H₁₄</i>	≈ 33.3	1.06
Octane – <i>C₈H₁₈</i>	≈ 40	1.05

The Adiabatic Index is related to the number of "degrees of freedom" f of a molecule through the relationship:

$$\gamma = \frac{f + 2}{f}$$

What is a degree of freedom? A hockey puck sliding on ice has two degrees of freedom. It can move in the direction towards the net or at right angles to that direction. Of course, it can actually slide in any arbitrary direction on the ice, but an arbitrary direction can always be broken down into a component in the direction of the net and a component at right angles to that; hence, two degrees of freedom in its direction of movement. If we also let the players scoop the puck off the ice, then it will have three degrees of freedom, up/down as well as on the surface. A single atom in a gas has three degrees of freedom. When energy is added to that a single atom, it has three independent directions in which its kinetic energy can increase.

Every molecule has the same three degrees of translational freedom as a single atom. But, a molecule with two or more atoms has additional types of motion which can absorb energy. In a diatomic molecule, like oxygen (*O₂*), the two atoms have a shape like a dumbbell. The dumbbell can rotate around the center-of-gravity. there are two (independent) axes around which a rotation can happen. Rotation around an arbitrary axis can always be broken down into a combination of rotations around these two independent axes. As energy is added to a diatomic molecule, the two atoms will revolve around each other more and more quickly. Rotation of the molecule is another form of kinetic energy. The speed of the molecule, or its center-of-gravity, represents linear, or translational, kinetic energy; rotation around the center-of-gravity represents rotational kinetic energy. A diatomic molecule therefore has five degrees of freedom and a corresponding Adiabatic Index of $7/5 = 1.400$. This is the case at moderate temperatures.

At high temperatures, oxygen and other diatomic molecules exhibit a sixth degree of freedom. The distance between the two atoms can begin to oscillate. This motion is a vibration, another form of linear kinetic energy. At the high temperatures where this phenomenon becomes apparent, the Adiabatic Index of oxygen decreases to $8/6 = 1.333$.

Water vapour (*H₂O*) is a typical triatomic molecule – it contains three atoms. The oxygen atom sits in the middle and the two hydrogen atoms sit off to the side in a "V" shape. Because of this three-dimensional shape,

it is possible to distinguish three separate axes of rotation. So, at moderate temperatures, water molecules have six degrees of freedom and an Adiabatic Index of 1.333. At high temperatures, atom-to-atom vibrations also manifest themselves, and these atom-to-atom bonds can absorb energy, so the effective number of degrees of freedom increases and the Adiabatic Index falls.

Carbon dioxide (CO_2) is also a triatomic molecule. Unlike water vapour, however, all three atoms of carbon dioxide lie in a straight line. Mathematicians call things that lie in a straight line "collinear". That the three atoms lie in a straight line removes one of the rotational degrees of freedom. The one which is removed is the one involving rotations around the axis which runs through the three atoms. Since the atoms are so very small, rotation around this axis involves virtually no energy at all.

Bigger molecules which contain more atoms have more degrees of freedom since there are more ways the constituent atoms can revolve and vibrate with respect to each other. How much energy an individual atom-to-atom bond can absorb or release depends not just on the local shape of the molecule, but on temperature, too. An atom-to-atom bond which does not absorb much heat does not contribute a "full" degree of freedom. It contributes less to the molecule's total number of degrees of freedom than an atom-to-atom bond which easily absorbs heat. All of this means that the numbers of degrees of freedom, and the Adiabatic Indices, are not fixed numbers. Generally speaking, as the ambient temperature increases, all the atom-to-atom bonds become more receptive to absorbing energy, the effective number of degrees of freedom rises and the Adiabatic Index decreases.

Most of the changes which result from high temperatures only become significant at very high temperatures. For practical work with internal combustion engines, the variation can sometimes be ignored. For example, the following table shows how the Adiabatic Index of dry air depends on temperature. The changes are only a fraction of a percentage.

Temperature	Degrees of freedom	γ (Adiabatic Index)
0°C	≈ 4.963	1.403
20°C	5	$(5 + 2)/5 = 7/5 = 1.400$
100°C	≈ 4.988	1.401
200°C	≈ 5.025	1.398
400°C	≈ 5.089	1.393
1000°C	≈ 5.479	1.365

Working through a physical example

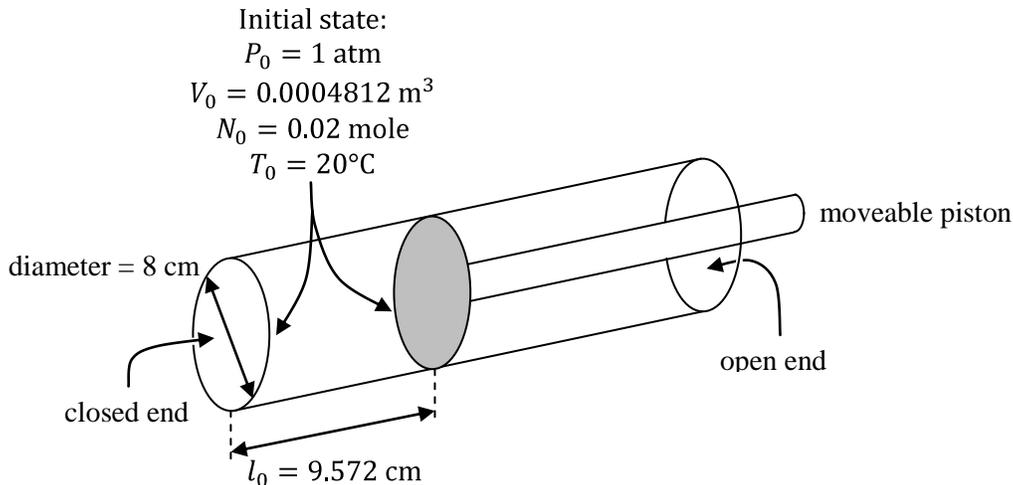
Let's attack a numerical example. We are going to take a round pneumatic cylinder, with a moveable piston at one end and some dry air inside. To add some teeth to the example, let's use realistic numbers. Let's assume the cylinder has a diameter of eight centimeters. Let's assume we add one-fiftieth of a mole of dry air ($N_0 = 0.02$ mole), and that the air is at room temperature ($T_0 = 20^\circ\text{C}$) and one standard atmosphere of pressure ($P_0 = 101,300 \text{ N/m}^2$). We can use the Ideal Gas Law to calculate the initial volume (V_0) of the air.

$$V_0 = \frac{R \times N_0 \times T_0}{P_0} = \frac{8.314 \text{ J/}^\circ\text{K-mole} \times 0.02 \text{ mole} \times (20^\circ\text{C} + 273.15^\circ\text{K})}{101,300 \text{ N/m}^2} = 0.0004812 \text{ m}^3$$

Since the cylinder has a radius of 4 cm, it has a cross-sectional area equal to $A = \pi r^2 = 3.1416 \times 0.04^2 = 0.005027 \text{ m}^2$. The volume enclosed by the the piston (V_0) is the cross-sectional area (A) multiplied by the length of the chamber, so the initial length (l_0) must be equal to:

$$l_0 = \frac{V_0}{A} = \frac{0.0004812 \text{ m}^3}{0.005027 \text{ m}^2} = 0.09572 \text{ m} = 9.572 \text{ cm}$$

The following figure shows the initial state of the dry air in the cylinder.



In the first step of playing with this apparatus, we are going to heat up the gas. Let's secure the piston in place, so the volume of the trapped gas does not change while we heat it. Using a propane torch, and heating the cylinder all round, let's heat up the air to, say, 200°C . Since the number of moles of air and the volume are fixed/constant during the heating, the only things that can change are the temperature and the pressure. For this first step, we can write the Ideal Gas Law as:

$$P \times V_0 = (R \times N_0) \times T \quad (\text{Step \#1})$$

The absolute temperature rises from $20^\circ\text{C} + 273.15^\circ\text{K} = 293.15^\circ\text{K}$ to $200^\circ\text{C} + 273.15^\circ\text{K} = 473.15^\circ\text{K}$. This is an increase of $180^\circ\text{K}/293.15^\circ\text{K} = 61.40\%$. The pressure must increase by the same percentage, from one atmosphere to 1.6140 atmospheres.

In the second step of our experiment, which we will carry out as soon as the temperature reaches 200°C , we will release the piston. Since the pressure on the inside face of the piston is greater than the ambient pressure on its outside face, the piston will quickly slide down the cylinder, increasing the volume of the trapped air. The piston will slide until there is no difference in pressure between the inside and outside, that is, until the dry air returns to ambient pressure. We can assume that the expansion happens so fast that the air does not exchange any energy with the walls. That means we can treat the process which occurs during the expansion as adiabatic. Let's apply the Adiabatic Law first. At the start of the expansion, the pressure and volume of the air are 1.6140 atmospheres, or $1.6140 \times 101,300 = 163,500 \text{ N/m}^2$, and 0.0004812 m^3 , respectively. Although the Adiabatic Index will not be exactly constant during the expansion, it will be very close to 1.400. At all times during the expansion, then, the pressure P and volume V will be related by:

$$P \times V^{1.4} = 163,500 \times 0.0004812^{1.4} \quad (\text{Step \#2 - first equation})$$

At the end of the expansion, the pressure P will once again be one atmosphere. The value of $P \times V^{1.4}$ will be the same as its starting value, so the ending volume can be calculated to be:

$$101,300 \times V^{1.4} = 163,500 \times 0.0004812^{1.4}$$

$$\rightarrow V = \left(\frac{163,500 \times 0.0004812^{1.4}}{101,300} \right)^{1/1.4}$$

$$\rightarrow V = 0.0006774 \text{ m}^3$$

This volume divided by the area of the cylinder ($A = 0.005027 \text{ m}^2$) is the new length of the enclosed chamber. It is $0.0006774/0.005027 = 0.13475 \text{ m} = 13.475 \text{ cm}$. The piston will have moved a distance of $13.475 - 9.572 = 3.903 \text{ cm}$.

Although the Adiabatic Law was sufficient to enable us to calculate the new volume, we still do not have a complete description of the gas after the expansion. We do not know the ending temperature. To find the temperature we need to use a second equation, in this case the Ideal Gas Law. To do this, let's first write the Ideal Gas Law in a form where all the known and constant quantities are on the right-hand side, thus:

$$P \times V = (R \times N_0) \times T$$

$$\rightarrow T = \frac{P \times V}{R \times N_0} \quad (\text{Step \#2 - second equation})$$

$$= \frac{101,300 \times 0.0006774}{8.314 \times 0.02}$$

$$= 412.68^\circ\text{K} = 139.53^\circ\text{C}$$

The temperature of the air will have dropped from 200°C to 139.53°C .

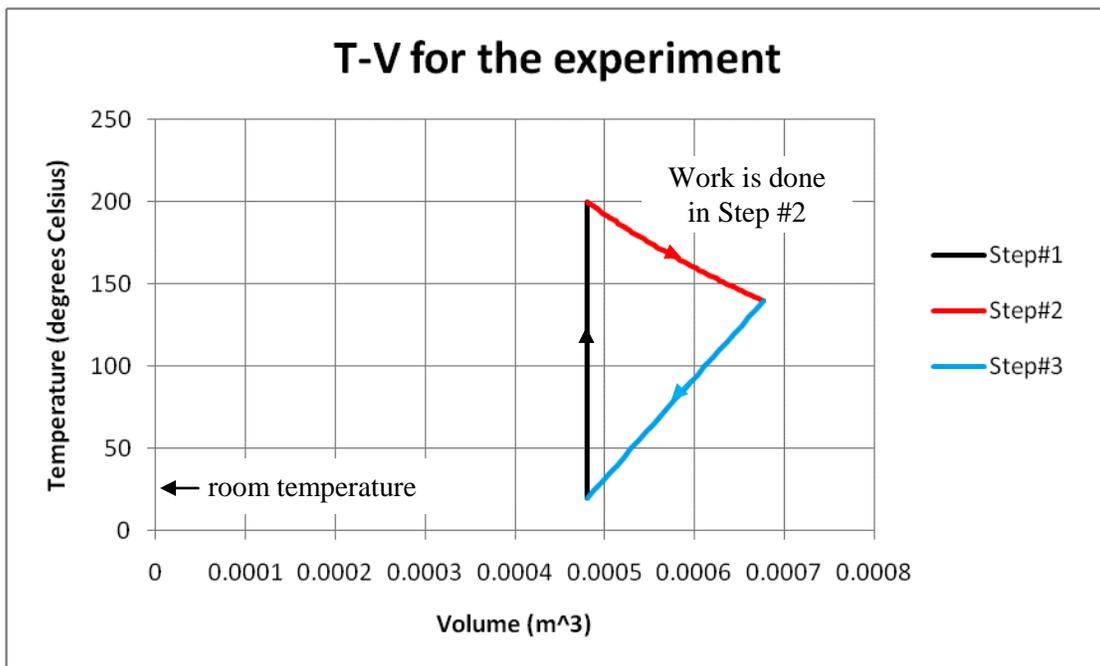
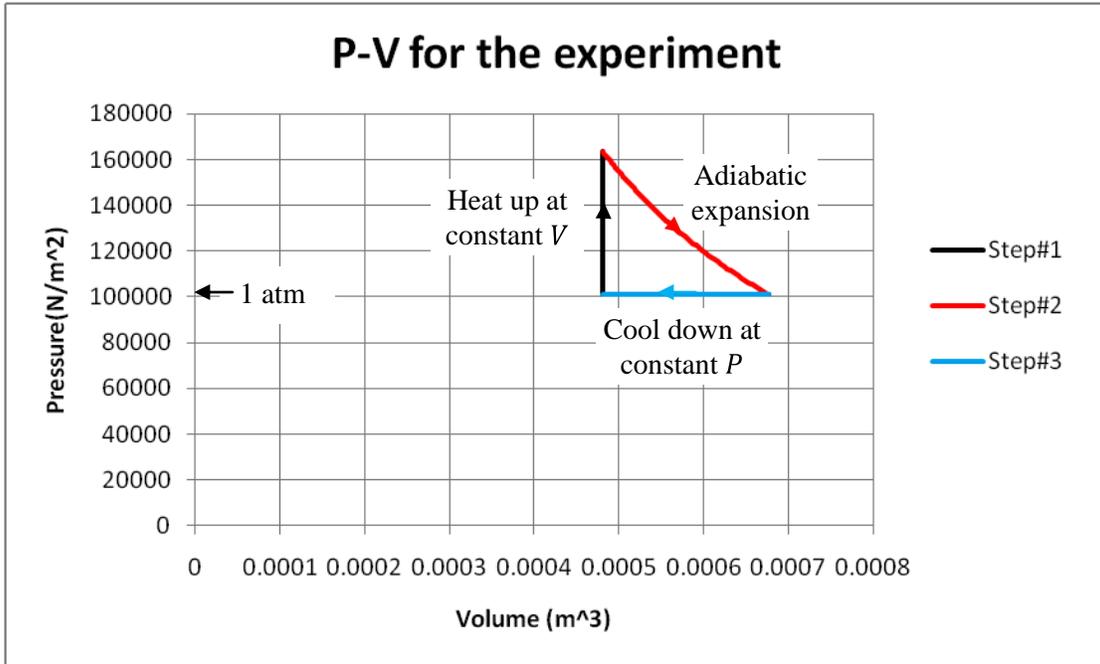
In Step #3 of the experiment, we are going to let the cylinder cool down to room temperature. As the air cools, and contracts, the piston will slide in along the cylinder, at all times keeping the pressure on the enclosed gas at one atmosphere. The cooling will stop once the temperature inside reaches room temperature. This process takes place at constant pressure. It is easiest to describe Step #3 if we re-arrange the Ideal Gas Law into the following form:

$$P \times V = R \times N_0 \times T$$

$$\rightarrow V = \left(\frac{R \times N_0}{P} \right) \times T \quad (\text{Step \#3})$$

Since the term in parentheses is constant during the cooling, the change in volume will be proportional to the change in temperature. The temperature will decrease by $139.53^\circ\text{C} - 20^\circ\text{C} = 119.53^\circ\text{C}$. To calculate the percentage decrease in temperature, we cannot use 139.53°C . We have to use its Kelvin equivalent, 412.68°K . The percentage decrease in temperature is $119.53^\circ\text{K}/412.68^\circ\text{K} = 28.964\%$. The volume must decrease by the same percentage, from 0.0006774 m^3 to $0.0006774 \times (1 - 0.28964) = 0.0004812 \text{ m}^3$. This is the original volume of the trapped gas, before we started Step #1. That we have returned to the starting point should not be a surprise. The pressure and the temperature after Step #3 are the same as they were at the outset. Since we did not add or remove any gas, the volume must be the same, too.

The following two graphs show how the state of the air changed during the three steps of the experiment. The first graph is the same kind of pressure-verses-volume graph that I described above. The second graph complements the P - V graph by showing what happens to the temperature. In it, the temperature is plotted on the vertical scale and the volume is (again) plotted along the horizontal



Since the sequence of three steps ends up where it started, it is a cycle we could repeat over and over. Unlike the cycle in an internal combustion engine, the same gas, or "working fluid" as they call it, is used throughout the cycle. During Step #2, when the gas inside is expanding, the piston is forced to travel down the cylinder. With the right mechanism, we could get the moving piston to do some real work. In fact, we have the tools we need to calculate how much work this apparatus could do during Step #2. In the next section, we will calculate the work which the expanding gas could do.

Tracking the energy through the cycle

As a preliminary step, let's calculate how much dry air the cylinder contains. We put in 0.02 moles, but moles are not a unit we use in our day-to-day lives. The following table lists the constituent gases in dry air. Five of these gases – argon, neon, helium, krypton and xenon – are the "inert" gases we encountered in the earlier paper, on stoichiometry. They are the atoms which have exactly the right number of electrons to complete their outer-most shells. Since they are complete, they do not tend to react with other substances. That is why they appear in air; they simply do not team up with other atoms. The percentages which these component gases make up of dry air are shown in the second column. Nitrogen and oxygen are the main ones, with argon in third place and representing only one percent. The percentages given in the second column are comparisons by volume, which is the same as comparing the numbers of individual molecules which are present in a given quantity of dry air. But, the molecules of the different gases have different weights, so the composition of dry air by weight is slightly different. The third column gives the weight of each component gas, in grams per mole. (The weight expressed in grams per mole happens to be the weight of an individual molecule expressed in AMU.) The fourth column gives the weight of each component gas in one mole of dry air. The total of the fourth column, 28.97 grams per mole, is the weight of dry air. We can therefore calculate that the 0.02 moles of dry air we put into the cylinder weighs $0.02 \text{ moles} \times 28.97 \text{ g/mole} = 0.5794 \text{ grams}$. Half a gram is about one-thousandth of a pound. The last column in the table compares the constituent gases by mass.

Component	% by volume	Mass per molecule (g/mole)	Mass per average molecule (g/mole)	% by weight
Nitrogen	78.09%	28.02	21.88	75.52%
Oxygen	20.95%	32.00	6.704	23.14%
Argon	0.933%	39.94	0.373	1.286%
Carbon dioxide	0.030%	44.01	0.013	0.046%
Neon	0.0018%	20.18	0	0.001%
Helium	0.0005%	4.00	0	0.00007%
Krypton	0.0001%	83.8	0	0.0003%
Hydrogen	0.00005%	2.02	0	0.0003%
Xenon	0.000009%	131.29	0	0.004%
Total	100%		28.97	100%

Let's return to the Adiabatic Index γ , which is a letter in the Greek alphabet, pronounced "gamma". We treated it above as a mathematical thing related to the number of degrees of freedom of molecules in a gas. While it is that, it first arose in thermodynamics in connection with the "specific heat" of substances. Recall from school that one BTU (British Heat Unit) is the amount of heat which will raise the temperature of one pound of pure water by one degree Fahrenheit. This is a concept like the calorie, which is the amount of energy needed to raise the temperature of one gram of pure water by one degree Celsius. Pure water is used as the reference substance, with which all other substances are compared. The amount of heat required to raise the temperature of substances other than pure water is expressed in multiples of the amount of heat needed to do the same job for water. The multiple is called the other substance's specific heat. The beauty of talking about the specific heat of an arbitrary substance is that it does not depend on the units used to measure the energy, or the amount of substance heated or even by how much the temperature changed.

The specific heats of liquids and solids are pretty straight-forward to measure. Their volumes change by only small amounts when heat is added. Gases are another matter. They have to be confined inside a closed container. Adding heat can increase the pressure or the volume or both, depending on what the container does. It happens, therefore, that a gas has two different specific heats. One of the specific heats is measured with the pressure held constant. The other is measured with the volume held constant. Unsurprisingly, they are called

the specific heat for constant pressure and the specific heat for constant volume, respectively. If we divide the former by the latter, we get:

$$\gamma = \frac{\text{specific heat for constant pressure}}{\text{specific heat for constant volume}}$$

Yes, this ratio is our old friend, the Adiabatic Index. The following table sets out typical values for dry air. Like the Adiabatic Index itself, the two specific heats vary slightly with temperature. The specific heats are reported in units of kiloJoules per kilogram-°K. A kiloJoule is, as usual for that prefix, just a thousand Joules.

Temperature	Specific heat at constant pressure (C_p)	Specific heat at constant volume (C_v)	Adiabatic Index $\gamma = \frac{C_p}{C_v}$
0°C	1.005 kJ/kg°K	0.716 kJ/kg°K	1.403
20°C	1.005 kJ/kg°K	0.718 kJ/kg°K	1.400
100°C	1.009 kJ/kg°K	0.720 kJ/kg°K	1.401
200°C	1.026 kJ/kg°K	0.734 kJ/kg°K	1.398
300°C	1.047 kJ/kg°K	0.751 kJ/kg°K	1.395
400°C	1.068 kJ/kg°K	0.767 kJ/kg°K	1.393

We can use the specific heats in this table to determine how much heat was added to or removed from the dry air in the cylinder during the steps in our experiment.

The heating we did in Step #1 was done with the volume of the trapped air held constant. Therefore, it is the specific heat for constant volume which applies during this step. We already calculated the weight of the air in the cylinder. It is 0.5794 grams, or 0.0005794 kilograms. In Step #1, we raised the temperature by 180°C, from 20°C to 200°C. Let's use the specific heat at 100°C as representative of this range of temperatures. We calculate the energy added as follows:

$$\begin{aligned} \text{Energy added} &= \text{mass} \times C_v \times \Delta T \\ &= 0.0005794 \text{ kg} \times 0.720 \text{ kJ/kg}^\circ\text{K} \times 180^\circ\text{K} \\ &= 0.07509 \text{ kJ} \\ &= 75.09 \text{ Joules} \end{aligned}$$

We can make the same kind of calculation for Step #3. In that step, the gas cooled and energy was released. Since the cooling during Step #3 was done at constant pressure, we will use the specific heat at constant pressure instead of the specific heat at constant volume. The same mass of air – 0.5794 grams – was involved. This time, the temperature decreased by 139.53°C – 20°C = 119.53°C. From the table of specific heats, it looks like the average of C_p during Step #3 is probably near 1.007 kJ/kg°K. The energy added is given by:

$$\begin{aligned} \text{Energy added} &= \text{mass} \times C_p \times \Delta T \\ &= 0.0005794 \text{ kg} \times 1.007 \text{ kJ/kg}^\circ\text{K} \times (-119.53^\circ\text{K}) \\ &= -0.06974 \text{ kJ} \\ &= -69.74 \text{ Joules} \end{aligned}$$

The temperature change in Step #3 is algebraically negative because the ending temperature was less than the starting temperature. This results in a negative amount of energy "added". All the minus sign means is that energy was removed during the cooling step, not added.

The difference between the energy we added in Step #1 and the energy released during Step #3 is 75.09 – 69.74 = 5.35 Joules. This is the work the piston could do during Step #2. We know this because that is the

only possible use of energy during Step #2. We assumed the expansion in Step #2 was adiabatic which, by definition, means that the gas did not exchange any thermal energy with its surroundings.

In principle, we could repeat the cycle over and over again. Each time, we would secure the piston and heat up the cylinder with the propane torch. We would then release the piston, which would do 5.35 Joules of work if it was attached to some mechanical load. Then, we would let the cylinder cool down. Doing things this way would not be very efficient. Some of the heat from our propane torch will be diverted and not be absorbed by the gas. It will heat the outside of the cylinder, the supports, the torch itself, and so on. Letting the cylinder cool down in Step #3 is a complete waste. We did not make any effort to recover any of the energy from the heat released or, as an even better alternative, to skip the cooling step entirely. However crude, though, the apparatus is still an engine of sorts. We can calculate its theoretical efficiency as follows:

$$\text{Efficiency} = \frac{\text{Useful work done}}{\text{Energy input}} = \frac{5.35 \text{ Joules}}{75.09 \text{ Joules}} = 7.12\%$$

This is a very low efficiency compared with that of internal combustion engines. The point is that the theoretical efficiency of any engine whose working fluid is a gas can be calculated using exactly the steps we have followed here – calculating the thermal energy added to and removed from the gas, and then taking the difference.

Reconciling the work done during the expansion stroke

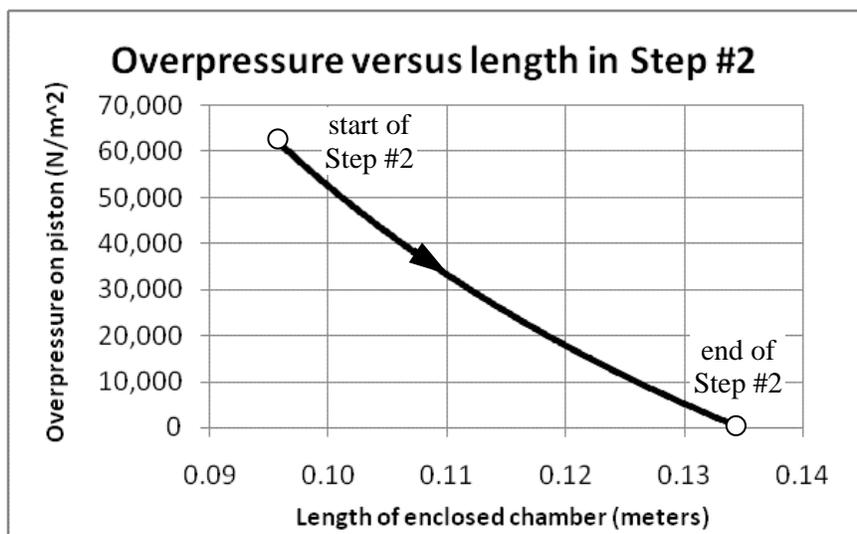
We can make a direct calculation of the mechanical work the piston could do during Step #2. The area of the piston head is 0.005027 m^2 . This is the area A of the cross-section of the cylinder, which we calculated above from its 8 cm diameter. If the pressure of the air inside the cylinder at any instant in time is P , then the force F which the air inside exerts on the piston head is the pressure multiplied by the area, $F = P \times A$. We must not forget about the air outside. The pressure we used in all the calculations above is the absolute pressure. And, indeed, the force $F = P \times A$ is the absolute force which the air inside exerts on the inside face of the piston head. The air outside presses on the outside of the piston head, and the absolute force which the outside air exerts resists the force exerted by the inside air. The net force acting on the piston head, for which I will use the symbol F_{net} , is the difference in pressure from the inside to the outside, multiplied by the area.

$$F_{net} = (P_{inside} - P_{outside}) \times A$$

At the start of Step #1, the absolute pressure inside was $163,500 \text{ N/m}^2$. The air outside was at standard atmospheric pressure, $101,300 \text{ N/m}^2$. The net force acting on the piston head at the start of Step #2 is equal to:

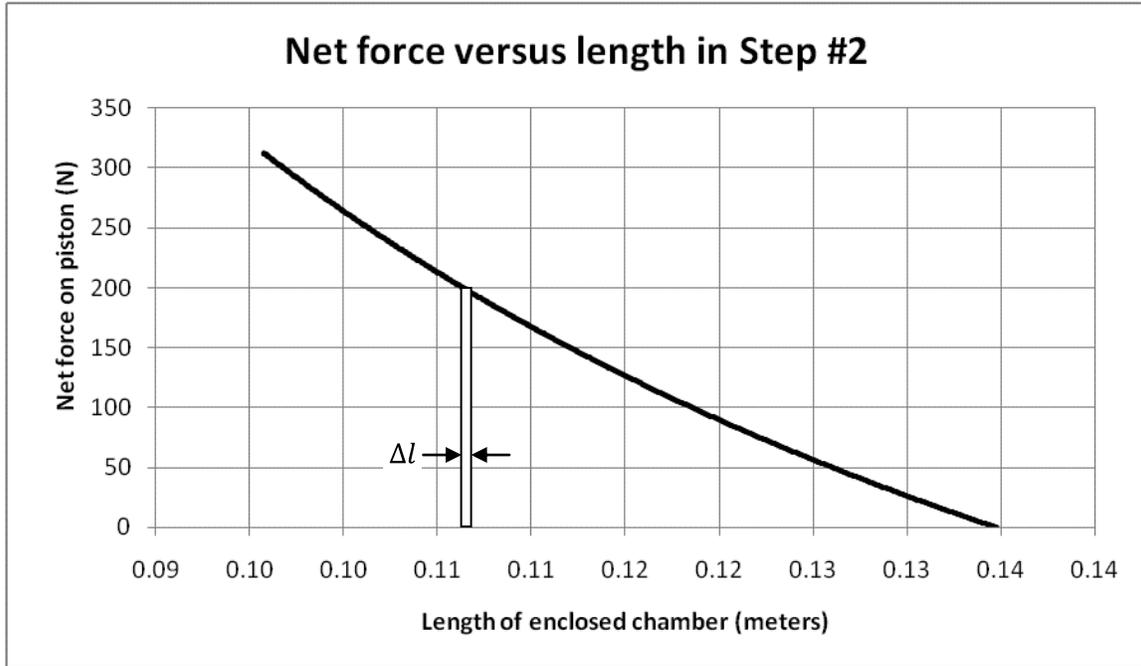
$$F_{net, start} = (163,500 - 101,300) \times 0.005027 = 312.7 \text{ Newtons} \approx 75 \text{ pounds}$$

At the end of Step #2, the gas inside has expanded to atmospheric pressure, so the net force on the piston head is reduced to zero. Throughout Step #2, as the pressure inside decreases, the net force on the piston head decreases. The graph to the right shows how the pressure decreases during Step #2. The horizontal axis is the instantaneous length of the enclosed part of the cylinder. At the start of Step #2, the



enclosed length is 9.572 centimeters; at the end of Step #2, the enclosed length has increased to 13.475 centimeters. The vertical axis is the overpressure, being the excess of the pressure inside over the pressure outside.

The overpressures on the vertical axis ($P_{inside} - P_{outside}$) of the graph above can be converted directly into net forces by multiplying them by the area. The corresponding graph, of the net force versus the cylinder's enclosed length, is as follows:



A force pushing on something can do work. If the pushed object moves, then the mechanical work done is equal to the magnitude of the force multiplied by the distance through which the object moves. (If the pushed object does not move, then no mechanical work is done.) The graph, which shows the net force versus the length of the enclosed part of the cylinder, contains the information we need to calculate how much mechanical work was done.

Consider the thin rectangle shown in the graph. I picked this particular rectangle so that its height would be 200 Newtons. As the piston moves from the cylinder length corresponding to the left edge of the rectangle to the cylinder length corresponding to the right edge, it moves a distance Δl . The net force during this little move is approximately 200 Newtons, so the mechanical work done is approximately equal to the product $200 \text{ N} \times \Delta l$. Extending this thought process will show that the mechanical work done by the piston as it moves down the black curve is equal to the area underneath the black curve. To see this, imagine a whole sequence of thin rectangles covering up the area underneath the black curve.

Making a rough approximation of the area To the extent that the curve is almost a straight line, we can approximate the area underneath it as the area of triangle. The height of this triangle is 312.7 Newtons, being the net pressure on the piston at the start of Step #2. The width of the triangle is the total distance through which the piston moves, taking distances from above, $13.475 - 9.572 = 3.903 \text{ cm}$. Since the area of a triangle is one-half the base times the height, the mechanical work done is approximately equal to:

$$W = \frac{1}{2} \times 312.7 \text{ N} \times 0.03903 \text{ m} = 6.10 \text{ Nm} = 6.10 \text{ Joules}$$

A Joule of energy is equal to one Newton multiplied by one meter. In fact, this is the relationship by which the Joule is defined.

Because the black curve is concave upwards, one would predict that the triangular estimate of the area would overstate the true area. That is the case. The triangular approximation gives 6.10 Joules versus the 5.35 Joules we calculated from the balance of heat energy.

Making a better approximation of the area A better approximation of the area can be found by explicitly dividing the black curve into a large number of short segments, and constructing a thin rectangle for each one. I did this using an Excel spreadsheet. I divided the black curve into segments by calculating the volume, and thus the cylinder length, for each pressure between the starting pressure $163,400 \text{ N/m}^2$ and the ending pressure $101,300 \text{ N/m}^2$, in steps of 100 N/m^2 . This resulted in 623 segments and 623 corresponding thin rectangles. The sum of the areas of all these little rectangles is 5.27 Joules. This is much closer to the 5.35 Joules calculated using the heat balance. It should be noted, though, that even a perfect integration of the area underneath the curve will not be exactly equal to the calculation using heat energy (except by coincidence). The reason is this. We calculated the heat absorbed by the gas during Step #1 using a constant value for the specific heat at constant volume. Similarly, we calculated the heat given up by the gas in Step #3 using a constant value for the specific heat at constant pressure. In reality, both specific heats vary with temperature. By using an average value over a range of temperatures, we gave up some accuracy.

Jim Hawley

© October 2014

If you found this description helpful, please let me know. If you spot any errors or omissions, please send an e-mail. Thank you.