Research & Testing

Stoves, Tents and Carbon Monoxide – Deadly or Not? Part 2: Testing the Theory

Can you run your backpacking stove inside your tent, or do you risk death by Carbon Monoxide poisoning?

By Roger Caffin

February 21, 2007

Buy a small stove these days and it is likely to come covered in dire warnings about the risk of carbon monoxide (CO) poisoning and that you must not use the stove in any sort of confined space. And yet walkers have been using small stoves inside their tent vestibules in bad weather for many, many years with very few instances of trouble. What is the risk, why are all those warnings there, and how seriously should we take them?

This multi-part article explores the carbon monoxide issue. Part 1 covered the basic theory underlying how stoves work and under what conditions they emit CO. This part covers an extensive amount of laboratory work using a number of canister stoves to test the theory. Stove to pot clearance, pot diameter, burner shape, air supply and other parameters were varied to see what affected the amount of CO emitted. The results were quite clear, and validate the theory very nicely.

Parts 3, 4, 5 and 6 will examine popular canister stoves, alcohol stoves, liquid fuel stoves and solid fuels in the laboratory. Part 7 will explore what happens out in the field in a tent or under a tarp.

Recapitulation from Part 1

Part 1 examined the theory behind the combustion process in small stoves, focusing mainly on the three common hydrocarbon fuels – butane/propane, white gas and kerosene. We do know that alcohol will behave in a similar manner. The conclusions so far are:

- Carbon monoxide can be emitted by a flame under the ‘wrong’ conditions
- This carbon monoxide really can present a serious health hazard
- This hazard would seem to get worse as we go from butane/propane to white spirits to kerosene
- Some stove designs may be worse than others because the pot is placed too close to the burner, quenching the flame prematurely
- The hazard is not inevitable: there are ways to reduce it to negligible levels
- Long flames and yellow flames may indicate a CO hazard
- Ventilation is crucial under any circumstances

Outline and Scope of Part 2

The CO Test Chamber.

The first aim is to verify the theory outlined in Part 1, that CO is produced when the pot is placed too close to the burner such that the flame is cooled too soon, or ‘quenched.’ If this theory is right, we should be able to give some simple guidelines for stove operation which will minimise the chance of CO being produced.

- Effect on CO concentration as the clearance between pot and burner is changed.
- Effect on CO concentration as the stove power is varied from low to high.
- Effect on CO concentration for different models of burners all using the same fuel.
- Effect on CO concentration as the pot diameter is varied.

Most of the work in this section was done with two Kovea canister stoves and a Trekka canister stove. The Kovea stoves are made by a South Korean factory which makes stoves for many of the 'mainstream' walking companies. They are robust but a little too heavy for lightweight walkers. The Trekka is a cheap steel 'generic' stove of the sort you might find in Army Disposal or Discount stores or Wal-Mart, and is made in one of several possible factories in China. This stove was modified in several ways to test some parts of the theory. Several other sorts of stoves were used to test other parameters which could not be tested with the first two.

It would be nice to be able to make all measurements in a tent, but there are so many variables inside a tent (like the weather outside the tent!) that such measurements would not be very helpful. To separate out all the different influences it was essential that I do the testing under far more controlled conditions. So the testing described in this article was all conducted in my laboratory.

Health and Safety Guidelines

Before getting into all the details of the experiments, we should first consider what levels of CO are significant. I have scoured the web and the following table is a compilation from several sources, mainly USA and UK. It would seem that levels below about 50 parts per million for the short time required to cook dinner are not going to be very significant – especially if you do have some reasonable ventilation where ever you are.

CO level, ppm	Effect
0 – 1	Normal background
9	Max allowed for short term exposure in a living room (ASHRAE, USA)
25	Often encountered on major roads (UK)
30	Health and Safety limit for 8 hours (UK)
35	Suggested max allowable concentration for continuous exposure for 8 hours (ASHRAE, USA)
100	May be encountered on major roads during weather inversions (UK)
200	Health and Safety limit for 15 minute exposure (UK)
200	Mild headache, fatigue, nausea, dizziness. Limit for transient exposure (USA)
200	Slight headache, tiredness, dizziness, nausea after 2-3 hour exposure (ASHRAE, USA)
300	Can lead to collapse (UK MoD)
400	Frontal headache, life threatening after 3 hours

ASHRAE: American Society of Heating and Residential Air-conditioning Engineers
UK MoD: United Kingdom Ministry of Defence

There are five parameters of concern here:

- Time
- Temperature
- Gas flow rate or stove power
- Clearance
- CO Concentration

Time

A simple digital watch was used. Errors of a few seconds either way may be expected for the readings.

Temperature

Two temperature probes were used, both solid state. They were checked against a standard laboratory thermometer while immersed in water at room temperature. One was also checked in boiling water; the other did not measure that high. The former was used to measure the water in the pot: it hung down the chimney on the Test Chamber from above the pot. The latter simply monitored the temperature of the air going past the CO sensor. Measurements were made to 1 Centigrade degree.

Gas Flow, Stove Power or Heating Rate

This is a hard one to measure at the gas end. Putting some sort of micro-flow meter into a fuel line is extremely complex. However, we don't need to measure either the fuel flow or the stove power for this project. Measuring (or calculating) the rate at which the water is heating up is enough, and this was done. All references to heating rate will be for 1 litre of water, although the graphs themselves may show heating for 1.5 litres.

In general two power settings were used during the tests: a 'Low Power' which would normally be used for simmering a stew or similar, and a 'High Power' which would be used for boiling water fairly fast or melting snow. In this context it should be repeated that really high power or flame level is usually wasteful of fuel, and with some stoves is even slightly dangerous due to 'flame lift-off.' Lift-off is seen when the gas coming out of the holes in the burner is going faster than the flame can spread and the flame lifts off the burner face: it is very visible when it happens. If lift-off happens, the flame is very susceptible to being blown out, leaving fuel going everywhere. In general the stoves were not operated with significant flame lift-off.

Clearance

This is the distance between the top of the burner and the bottom of the pot. The hypothesis in this project is that reducing the clearance will increase the production of CO when the flame is quenched by hitting the bottom of the pot too early. It was measured using a steel scale located a little bit away from the pot while the stove was running. The clearance has an experimental error of about 1 millimetre. For single measurements this will not be very significant as we will take measurements at a number of clearances and plot them out. An overall error in the 'zero' of the scale will result in the curve being moved sideways a bit. This latter does look startling at first, but may be ignored in practice.

Locating the bottom of the pot is simple. However, it is not clear what part of the burner should be used as the reference point: stoves have such a wide range of burner designs. I have used the top of the burner in most cases, but this does mean that the measurement becomes highly sensitive to the shape of the burner. Fortunately we are not concerned with the absolute value of the clearance between the pot and the top of the burner; what matters is whether the distance between the burner and the standard factory-supplied pot supports is adequate. 'Clearance' per se will usually be translated to the pot supports.

Carbon Monoxide Concentrations

Measuring CO concentration is not that easy. We are talking about parts per **million** (ppm) here: a very low level. What's more, we are talking about concentration, not an absolute amount. The stove may emit a certain amount of CO (per minute), but this gets mixed with an unspecified amount of air. How much air mixes in will depend utterly on the test conditions. This is why a fixed Test Chamber is used rather than a tent. The Test Chamber used has a large and fairly fixed opening at the bottom of the Lexan window to let the air in. The hot air goes out a chimney at the top of the Test Chamber.

Coming sideways out of the chimney is a small copper pipe. A variable-speed fan at the end of the pipe draws a continuous sample of air from the chimney, over the (blue) CO sensor which can be seen under it. As the air going up the chimney is very hot and the CO monitor is sensitive to temperature, the copper tube has a (white) water jacket around it to cool the gas inside. The second temperature sensor sits just outside the small fan, monitoring the temperature of the gas sample.

The CO monitor used here has a solid-state CO sensor, but the solid state sensor is temperature sensitive unless corrected by its own internal temperature sensor. If the air temperature changes suddenly, faster than the internal sensor can track, the readings will vary wildly until the temperature sensor catches up. Apparent spikes in the CO concentration may be seen right at the start of some test runs, when the stove was turned on. These spikes should be ignored. If the amount of cool fresh air coming into the hot flue gas stream varies, or is just a bit turbulent, the real concentration of CO will also vary up and down. So the CO measurement process can be upset by swings in temperature and air flow, and this is especially noticeable when trying to make 'real-time' measurements after sudden changes. Every time I insert my hand in the entrance to adjust the height of the stove I disturb the air flow, and I can see this in some of the measurements. I allow a delay between the adjustment and the measurement to handle this. Just don't expect perfectly smooth curves!

The heat of the stove will drive air up the chimney, and this will suck air in through the opening at the bottom of the window. I have not been able to monitor either the inlet air flow volume or the exit air flow volume. This means the CO readings I record are unique to this Test Chamber and this window opening. So what do the CO measurements mean? In themselves as absolute values, not very much. What does matter is how the CO concentration changes as we vary other things in the test. If I make a change in the test layout (such as reducing the pot to burner clearance) and the CO concentration goes up, I can reasonably infer that the increase in CO was caused by what I did. In Part 7 of this series we will look at what happens in a tent.

You also need to be aware of the differences to be expected between the confined flow in the Test Chamber and what might be encountered in a tent in the field. The extra air in the tent will dilute the CO being emitted by the stove. A stove which is recorded as emitting 50 ppm of CO in the Test Chamber might not even register in the body of a tent if there is good ventilation – not to mention when there is some seriously 'interesting' weather raging on outside!

Finally, we need to know how long it takes for the sensor to respond to a change in CO production created by a change in test conditions. First of all, it takes time for me to change the clearance between the burner and the pot (by adjusting the lab jack): this is known to be less than 10 seconds. The burnt hot gases flow up from the stove, into the chimney, and a sample is drawn continuously from the chimney to pass over the CO sensor. It takes time for a change in CO concentration to get from the stove up the chimney and then sideways to the CO sensor. So if I start changing the clearance at zero time, it can take 5 – 10 seconds for the change to be complete, and there will be another small interval before the change in CO production is registered by the sensor. This means I have to allow a delay between making a change and reading the sensors. Typically I start making a change at 'zero' and don't take a reading until at least 30 seconds have passed. Sometimes I wait another 30 seconds (without making further changes) and take a second reading.

so the CO reading should go to zero. In general the CO reading 10 seconds after the turn-off was just starting to drop, while 15 seconds after the turn-off it was about half-way down to zero, and after 30 seconds the reading was usually very close to zero. This means it takes about 10 seconds for the change to reach the sensor, and then the change is fully sensed within another 10 seconds.

Experimental Procedure

The picture above shows the main Test Chamber used. The Test Chamber itself is a steel box with a Lexan window at the front. Air comes in at the bottom of the chamber through small holes at the side and usually through a fixed gap at the bottom of the Lexan window. This open gap below the window corresponds to an 'unlimited' air supply as in a tent. However, I have full control over how much air is available. The hot air (and steam) goes out at the top through a tall chimney which prevents any down-draft.

Most test runs were done at two heating rates: 'low' and 'high.' It will be seen that these rates are not always exactly the same between runs and stoves. The rate of heating of the water has been quoted in a number of places to allow you to see how the results relate to your practice in the field.

A number of different pots were used in the testing to explore the matter of pot diameter. They are listed below. Later on when the effects of pot diameter had been identified, most of the testing was done with the largest of the pots, the 1.5 liter GSI Bugaboo pot. The use of a large pot with a lot of water meant that the rate of rise of the water temperature was lower, so longer test runs could be used.

Pot	Volume	Diameter
GSI Bugaboo	1.4 L (1.5 qt)	175 mm (6.89 in)
MSR Titan	1.5 L (1.5 qt)	155 mm (6.10 in)
MSR Titan	1 L (1 qt)	135 mm (5.31 in)
MSR Ti Kettle30	0.75 L (3 cups)	115 mm (4.53 in)

Practical Measurement Rate

The normal regime for taking readings in a test run was governed by two factors: how much time I had before the water reached boiling, and the delay in sensing a change in CO production. The latter was determined as indicated above to be well under 30 seconds. Accordingly, I took readings every 30 or 40 seconds, depending on the conditions. At low power with a large pot containing a lot of water I was able to take two readings at 30 second intervals after every change in clearance. Ideally the first reading (at +30 seconds) should be enough, but the second reading (at +60 seconds) was a check on the system. In the early stages while determining the effect of clearance on CO production I changed the clearance in 1 millimetre steps, but when running at high power I had to use 2 millimetre steps. Otherwise the water would have boiled half-way through the series of measurements. In later testing, with the advantage of understanding the effect of clearance, I simply used larger steps. In some cases when running at low power the pot reached a temperature just under boiling and stayed there, due to steam loss. I allowed this to happen as we are not concerned with actual boil time here.

The Kovea Expedition large burner canister stove used for Clearance tests.

The Effect of Stove to Pot Clearance

This is the main focus of the theory. It was in fact extensively tested over many canister stoves, and they (almost) all gave the same result, that decreasing the clearance increased the CO level. There was one exception, but that was due to another factor, and this will be discussed in due course. For the purpose of this article we will use the Kovea Expedition as the representative stove. It is a large burner stove with a remote canister. The canister was used upright, to give a gas feed rather than a liquid feed. This stove is a quite robust and reliable unit, albeit a bit heavier than an ultralight walker would normally use. It would make quite a good base-camp stove in winter time. The Kovea Expedition has a typical large burner diameter, measuring about 59 millimetres (2.3 in) across the centre ring of holes.

The pot support arms on the Kovea Expedition are tilted (as is quite common) so the tips are higher than the centre. The standard (factory) pot clearance on this stove varies from 13 millimetre (about 0.5 in) for very small pots (110 mm or 4.3 in diameter) to 18 millimetre (0.70 in) for large pots (180+ mm or 7+ in) diameter). However, for this test the pot supports were removed so the clearance could be reduced to below the factory-set minimum. In general terms doing this is not a good idea.

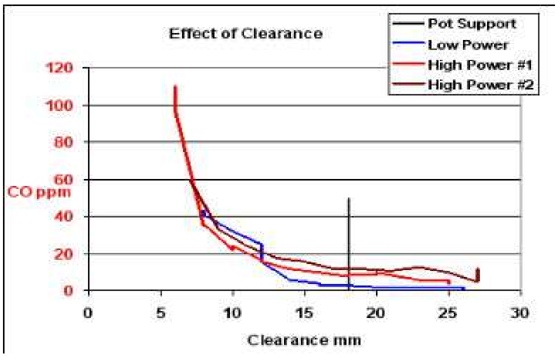

Graph 1: The effect of pot clearance for a large burner canister stove.

The Kovea Expedition was tested at both low and high power settings. Typical low power heating rates (normalised to 1 litre) were around 3.5 C/minute (6.3 F/minute). This is a 'simmer' power. High power heating rates of 9.1 C/minute (16.4 F/min) and 13.8 C/minute (24.8 F/min) were used for runs #1 and #2.

Typical results are shown to the left, but I emphasise that these curves are representative of a very large number of tests. It should be remembered that the CO readings will always contain some 'noise.' In my opinion, these curves all effectively overlap. Low power performance and high power performance are essentially the same as far as CO emissions go. That said, we will show later that some stoves do emit more CO at high power – but following the same sort of curve.

What we can clearly see is that the CO levels are low when the clearance is good, but start to rise as the clearance shrinks below a certain point. Where is the threshold for the deterioration? It's hard to say, but for any clearance over about 12 millimetres (0.47 in, on this stove) the CO readings seemed to have 'bottomed.' However, as the clearance drops, the CO emission rises – quite spectacularly at low clearances.

> ***The closer the pot is to the flame, the more the flame is quenched and the more CO is emitted.***

It is very important here to note the clearance in the graph which corresponds to the factory default position for the pot supports. As discussed previously, the absolute value for clearance is quite arbitrary, usually being the distance between the top of the burner and the bottom of the pot. But there is no reason to believe that this definition is anything special. What does matter is where the pot supports are **on this scale**. This is where the pot will be when you use the stove in the field. In this case we can see that the pot supports are well clear of the point where the CO emission starts to rise. This is good!

The Effect of Pot Diameter

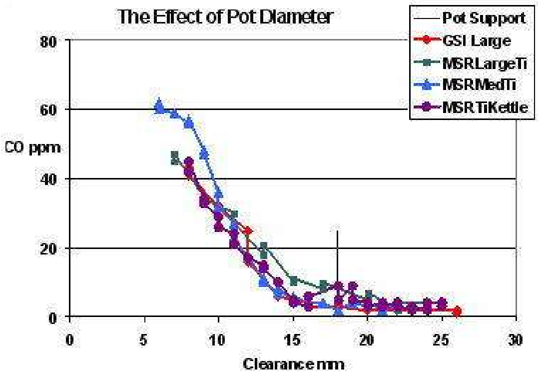

Graph 2: The effect of pot diameter for a large burner canister stove at low power.

There have been rumours in some magazines and medical journals that pot diameter may influence the amount of CO being emitted, with smaller diameter pots giving lower levels of CO. The theory presented in Part 1 of this Series does not have a mechanism to support this idea, but it was tested carefully anyhow. The four pots listed above range in diameter from a large 175 millimetre (6.89 in) down to a fairly small 115 millimetre (4.53 in). They were all tested on the Kovea Expedition at various clearances. The graph to the right shows representative results for a low power setting; the high power results from this stove are very similar, as shown in the previous graph.

There is the usual noise or scatter in this graph, but the overall results are fairly clear:

> ***Pot diameter does not have any significant effect on CO emission.***

If you think about it, this makes sense. The flames go up and hit the pot. The places where they hit will have the most influence on whether the flame is quenched. Once the flame bounces off the impact point on the pot the outer diameters are not going to interfere all that much. This is very good, as we know from other work that the larger diameter pots are more efficient in heating water than the small diameter pots. When the pots are too small a lot of the heat from the flames goes up the sides of the pot and is wasted.

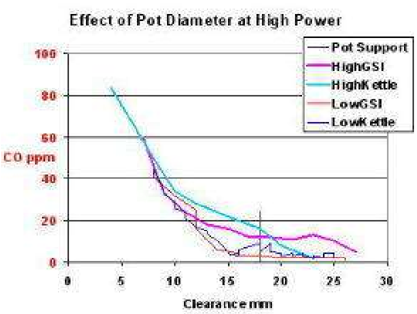

high power.
Note that the Pot Support line shown in the graph at 18 millimetres is for the large GSI pot. The equivalent pot support line for the small MSR Ti Kettle would be at 13 millimetres. At this lower clearance it is clear that the CO emission is starting to rise slightly – due to the shape or tilt of the pot supports. But this means that a smaller pot on this stove would cause a slightly **higher** CO emission, rather than the rumoured lower emission. However, the effect would not be very large, and later on we will see far higher levels of CO being emitted by some other stoves.

The Effect of Burner Diameter

The small burner Kovea Moonwalker.

The theory allows little obvious reason for the results from a small burner head to be very different from the results from an equivalent large burner head: the flame chemistry is exactly the same. This was examined by comparing the results obtained for a Kovea Moonwalker with the results for the Kovea Expedition. The Moonwalker is a remote canister stove with a burner diameter of 34 millimetres (1.3 in) to the middle ring of holes, while the Expedition has a burner diameter of 59 millimetres (2.3 in).

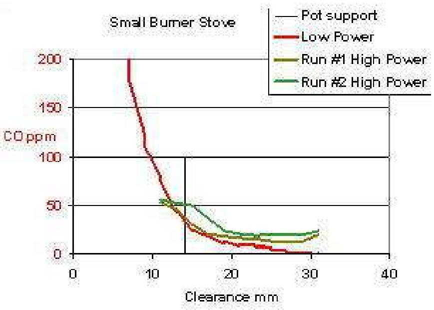

Graph 4: The effect of pot clearance on a small burner canister stove.

First we check that this stove shows the same sort of clearance effects that the Kovea Expedition shows. This is dealt with by the graph to the left, which shows almost exactly the same sort of curves for low and high power operation. We note here however that the pot supports are located a bit into the rising part of the curve, so this stove would in practice give off a bit more CO. This is not so good.

Now we compare the two different burner heads, with just the low power performance of the Moonwalker compared to the low and high performance of the Expedition. Immediately we see that doing so is not that easy. The blue curve for the Moonwalker *seems* to be seriously different from the curves for the Expedition. However, we must remember that the zero of the clearance scale means very little in practice, and if we shift the zero for the Moonwalker clearance by just 4 millimetres the curve for it slides sideways, and can be made to overlap the other curves. This is shown by the dashed blue curve. So the overall behaviour as a function of clearance is as before.

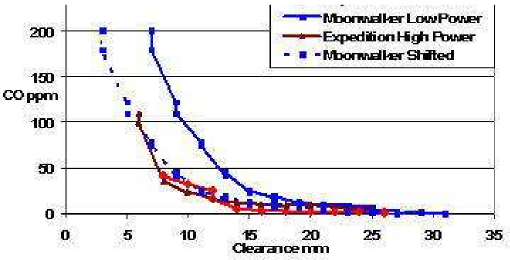

Graph 5: The effect of burner diameter.

It should be noted that the results for the Moonwalker above 100 ppm of CO were taken with the pot in actual contact with the tip of the preheat tube, and the pot was tilting slightly. The readings are valid, but one would not normally get a pot this close to this burner. The data points were taken just out of interest, because they do show the natural progress of the curve.

I will anticipate things a bit here and say that some other small-diameter burners (e.g. Snow Peak GS100) do exhibit very low CO emissions, but the reason for this is almost certainly tied up with the design of the burner head rather than its diameter.

What do we make of this? I draw two linked conclusions, using this and other data. The first conclusion is that comparing burner head diameters for CO emissions is a futile exercise. The second is that burner head design is vastly more important than diameter, although we have not yet come to the data for this.

The Effect of Burner Hole Size


The generic Trekka stove. (The blue arrow indicates the windshield.)

I have mentioned that some stoves display a flame lift-off when the gas velocity out of the burner holes exceeds the maximum flame velocity. This is of course very dangerous: one puff of wind and the whole flame could go out, leaving fuel going everywhere. But it may also affect the combustion cycle and influence how much CO is emitted. It is hard to test the effect of flame velocity independent of stove power under most circumstances, but there was an opportunity to do so with a cheap steel Trekka stove. This is shown to the right. The blue arrow points to what many vendors call a windshield, but which could equally be called a radiation shield as it masks the canister from the heat of the flames.

I had previously found that this stove was severely underpowered compared with most canister stoves, and had done some experiments with it. I had decided that the holes in the burner face were too small and were limiting the flow of fuel/air mix. So I had drilled out two rows out of three of those holes from 1.2 millimetres to 1.7 millimetres: a 100% increase in area for the two main rows of holes. Testing the modified burner head showed that I had successfully increased the safe (no flame lift-off) peak power from the vendor-quoted 5,600 BTU/hr to about 10,000 BTU/hr. Fortunately I still had an identical but unmodified Trekka stove which could be used for the comparison.

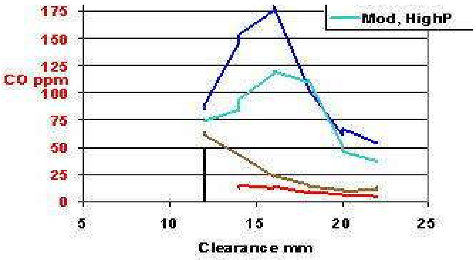

Graph 6: The effect of burner hole size.

The performance of this commodity stove is not as good as for one of the more expensive 'brand-name' stoves. Its CO emission curve at high power has a very strange shape as well – we will examine that later. For the present we will simply note that this Trekka stove does emit a lot more CO than the Kovea stoves. This actually allows us to do some testing.

What effect on the CO emission should slowing down the gas flow have? At high power we would expect that the lower gas flow rate out of the holes in the burner head should mean that it takes longer for the burning gas to reach the bottom of the pot. Not much longer, it is true, but long enough perhaps to reduce the CO levels. This was found to be so when I compared the stock (un-modified) Trekka burner head (dark blue line) with the modified Trekka burner head (light blue line) at high power: the results of this test are shown to the right. Clearly the lower gas velocity has reduced the CO emission at high power over some range of clearance. However, it would seem that the CO emissions at low power (red and light brown lines) actually rose slightly, and this is harder to explain. I believe the confusion here is due to the effect of the radiation shield on the stove performance, and I will return to this later.

My conclusion here is that the speed of gas flow out of the burner holes is very important for safety (avoiding flame lift-off) but only slightly important for CO emissions. In short, it is not a major factor by itself.

The Effect of Burner Design

Snow Peak GST100 burner head.

We have dealt with the simple measurable factors; now we must look at the more complex ones. In particular we must look at what effect the design of the burner head has on the emission of CO. If we compare the burner heads for the Trekka and the Kovea Expedition we see that the burner face (where the gas comes out) on the Expedition is at a steeper angle, while the face on the Trekka is nearly horizontal. To go to an extreme, the burner face on the esteemed Snow Peak Gigapower is nearly vertical. The tilt of the burner face controls the angle the flames come out at. Why is this important?

Consider a gentle flame coming out of the nearly horizontal Trekka burner face. It will go straight up to hit the pot directly above the burner: a short distance. Raise the power (gas flow) and the flame will simply hit the pot harder and sooner. The flames on the Kovea Expedition will come out at a bit of an angle and will tend to splay outwards as they reach the pot. As the gas flow is increased the flames tend to splay out even more. It is no coincidence that the Expedition has a lower CO emission. Now consider the flames from the Snow Peak GST100: they come out almost horizontal and don't really get to 'impact' on the base of the pot at all. There is almost nothing to impede the completion of the combustion cycle here. We find that the GST100 records a CO emission level of about 5 ppm over a wide range of clearances and power levels. This is a very good design.

The radiation shield may not seem like part of the burner design, but experiment shows it is. Having that steel cup sitting just under the burner affects the flow of the flames in the very small space between the burner, the pot and the wall of the radiation shield. Consider once again the dark blue performance curve for the Trekka at high power shown in the graph above. There is a certain amount of CO being given off at the minimum clearance – the factory standard clearance. Normally we would expect the level of CO emission to falls as the clearance is increased, but in this case it actually increases for a while. Why is this so?

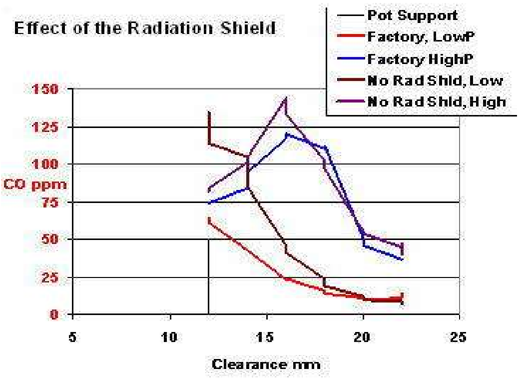

Graph 7: The effect of removing the radiation shield.

At the factory clearance the radiation shield is bottling up the flames a bit and this keeps them hot and allows them to recirculate. This actually keeps them burning for longer. Users of small alcohol stoves are familiar with the use of a windshield to do the same thing. As this clearance increases the flames are less bottled up, and can cool faster. The cooling of the flame allows more CO to be emitted. However, when the clearance is increased sufficiently the flames don't hit the pot in time to quench the combustion so much, and the CO emission starts to drop as originally expected. This explains the hump in the dark blue curve. The light blue curve shows the performance when the gas velocity is reduced by the larger holes. A lower velocity should mean the flames stay bottled up better and stay hotter and emit less CO. The hump in the CO emissions is still there, but it is lower as predicted. So far so good; but what should happen if the radiation shield is removed completely?

Since the radiation shield is bottling the flames up a bit and keeping them hot, removing the radiation shield should remove this effect. The flames should cool down sooner. The CO emissions should therefore rise compared to the factory state. Precisely this effect is seen in the graph to the right. The blue and red curves are for high and low power operation in the factory state, with the radiation shield in place. The purple and brown curves are for high and low power operation with no radiation shield in place. Both the no-shield curves show higher levels of CO emissions. It is interesting that in this case the removal of the radiation shield has a greater effect on the low power operation, where the flames go straight up to hit the pot. The CO emission level reaches the same value at minimum clearance as for the high power operation at its worst clearance. In my opinion the very flat angle of the burner face is responsible for this poor performance.

The Effect of Air Flow and Pressure

All of the testing described here was done at an altitude of about 188 metres (617 ft) above sea level. This is fine, but what happens when you take one of these stoves up to the top of a mountain, where the air pressure is so much lower? It would seem reasonable to suggest that a lower air pressure would mean a lower supply of air into the burner and into the flames, and that would mean a higher level of CO emissions. It is widely accepted that something like this does happen in the mountains. Can we demonstrate this?

There are two ways of trying to demonstrate such behaviour. One is to reduce the air pressure, but this is difficult with a Test Chamber like this. You can try to vary the air flow into the burner column through the inlet holes by the jet. This is easier to do, and some tests were conducted this way. Varying the available air supply was also done, and is discussed below, but this does not really approximate a reduced air pressure.

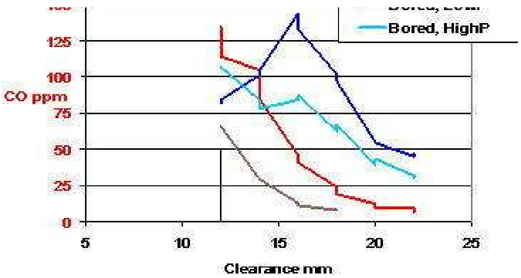

Graph 8: The effect of increasing the air supply.

The Trekka burner seems poorly designed. In addition to the problems already described, it seemed to me that the flames from the Trekka burner were lacking slightly in air supply, and that possibly the air supply inlet holes might be too small. It is well known that if you block the air inlet holes in the burner column the flames will become very long as it takes time for the air to diffuse into the flames above the burner. This is what happens at high altitude. But what would happen if the air supply inside the burner was increased – which can be done by increasing the size of the air inlets around the jet in the column?

The four air inlet holes in the column of the Trekka burner are about 4.8 millimetres (0.19 in) diameter in the factory state. They were bored out to 5.5 millimetres (0.22 in) diameter for an increase in area of about 31%, and then to 6.0 millimetres (0.24 in) diameter, for an increase in area of 56% over the factory value. This should increase the air supply inside the burner head to the flames. The results for one test run are shown to the left, for the 31% increase in the inlet hole area. The bold red and blue curves are for the stock burner, while the muted pale blue and brown are for the burner with the larger inlet hole area. Clearly, increasing the air supply into the central burner column has decreased the CO emission – as expected.

Increasing the hole sizes to 6.0 millimetres (0.24 in) diameter gave a corresponding (extra) decrease in the CO emissions. Blocking one of the four 6.0 millimetre inlet holes gave a corresponding increase in CO emissions, again as might be expected. Both changes were in accord with the percentage changes in burner air-hole area. They have not been shown in the graph here to avoid confusion over what the curves are showing.

In an associated experiment I tried injecting extra air into the air inlet holes in the burner column. This was expected to bear some similarity to the effect of increasing the flow of oxygen to an oxy-acetylene torch: the oxy-flame goes from long, wavering, orange and rather ineffectual at low oxygen flow to short, pointed and very blue at higher oxygen flow – and it's a lot hotter too. To make this experiment I rigged up a length of capillary tube to point upwards through one of the holes in the burner column. Turning on a gentle air flow pushed extra air into the burner column and converted the somewhat lazy flames with their orange tips into short and quite fierce looking flames. The CO level dropped as well. Sadly, doing this in the field is not so easy, but the point has been made.

Having demonstrated the effect of increasing the air supply at the burner air inlet holes, an experiment was done to decrease the air supply over all. This is not so easy to do for several reasons. It is not sufficient just to close off one of the holes in the central burner column, as air is also brought in around the flames above the burner. It is necessary to reduce the overall air supply. This would correspond, very roughly, to operating the stove in a confined space with inadequate ventilation. We know from both history and these experiments that doing so is dangerous.

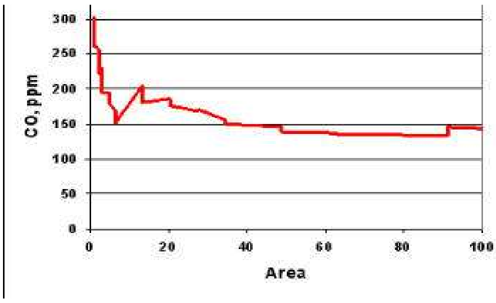

Graph 9: The effect of decreasing the air inlet area or air supply.

Doing this with the Test Chamber is not without difficulties. Provision was made for this in the design, with the inlet vents seen at the bottom right hand side of the Chamber. However, closing these off does not necessarily reduce the available air in a consistent manner. It is possible to get some air coming down the chimney when the air flow at the bottom is restricted. This can be compounded by the effect of the fan drawing air over the CO monitor: some of this air can be sucked down the chimney if the updraft is not strong enough. A test was run using the generic stove as this is known to be a CO emitter anyhow, but some of these interfering effects were seen.

In the graph to the right the effect of steadily closing off the air vents is shown. The horizontal axis is square centimetres of air vent. The red curve shows the results for a medium power. It may be seen that as the area of the air vents is decreased (moving left), the CO level rises – in general. The recorded data does actually extend way out to the right to an effectively unlimited air supply, but that part of the curve is flat. However, it is clear that the amount of CO being emitted does start to noticeably increase as the air vent area drops below about 60 square centimetres or 9 square inches. In Part 1 of this Series I mentioned that the Coleman Fyrestorm stove carries a warning that the user should ensure an air outlet of at least 10 square inches (65 square centimetres): this is close to the finding here for the air inlet.

There is a glitch in the curve below 20 square centimetres, where the air flow in the chimney changed somewhat. I was not able to fully control the air flow inside the Test Chamber under these conditions. Despite this, it can be seen that the CO levels do continue to rise overall as the air flow decreases. At a very low air supply the CO level above the stove has reached well into the dangerous zone. The UK Ministry of Defence classifies 300 ppm as leading to a person's collapse. Around here I decided enough testing had been done!

The test was repeated at a higher power and the behaviour was essentially the same: the CO level goes up as the air supply (inlet area) decreases.

These experiments do not match the effect of a windshield placed around a stove, although they do come quite close. They certainly do suggest however that a generous air inlet area is required. Any stove which needs a really tightly closed windshield to operate may be a CO hazard. This will be examined in a future part of this Series.

Other Factors

There are other factors which may affect the amount of CO being emitted by a canister stove but which have not been adequately analysed here. I mention them because they do have relevance to other sorts of stoves, albeit in a less direct manner.

The interior of the burner head on a canister stove is largely unknown on most stoves: you simply cannot see inside. Is the burner head simply a hollow shell, or does it have any structure? Certainly, some burner heads have quite convoluted shapes. Some burners have mesh inside the burner head, and this is visible at the burner holes. The holes are quite small, so the mesh is not there solely to prevent flash-back inside the burner head. Why is it used on some burners and not others? Looking inside the head of the Vargo Jet-Ti it can be seen that the burner column continues up inside towards the top. The burner head on the Snow Peak GS100 looks quite similar to the Jet-Ti, but it seems to have an internal baffle coming out of the holes in the burner head.

I have described the effect of enlargening the holes in the burner column. It increases the amount of air sucked in and so improves the fuel/air ratio, making for a better burn with lower CO levels. But what about the internal diameter of the burner column? Does this influence the amount of air sucked in? Many burner columns are about 12 millimetres (0.5 in) on the inside bore. Would increasing this help? I have not been able to properly test this, but logic would suggest a small increase should do no harm. Some other testing I did with a smaller bore seemed to imply that a smaller bore could limit the overall flow, and that this would reduce slightly the amount of air being sucked in.

The position of the top of the jet with respect to the air holes in the burner column has not been explored here. If the jet was positioned a significant distance higher inside the burner column, would this alter the amount of air sucked in? In fact it does, but this may be partly due to the bulk of the jet obstructing the air flow. If the jet is significantly lower than the air holes and the flow of fuel out of the jet is low, there is a risk that the fuel may come out of the air holes to burn there instead of up at the burner head. This would not be good.

The size of the jet will affect the speed of gas flow, and this will affect the amount of air sucked into the burner column. Changing the factory jet for one with a larger jet hole reduced the speed of flow and this radically changed the nature of the flames. They lost their short 'fierce' look and became long and 'lazy' like the flames on an alcohol burner – and the CO emission went way up. Obviously a high jet speed is critical to achieving low CO emissions. But we also know that different fuels (kerosene, white gas, butane/propane) require different-sized jets, and some manufacturers supply a set of different jets for their multi-fuel stoves. I do know from experience that using the wrong jet can be a big mistake!

Summary

We have tested the effect of various factors on the emission of CO from a canister stove, with the following results:

- **Burner to pot clearance**
  When this is too low the amount of CO emitted can become dangerous. The flame is quenched too soon. This is the core hypothesis outlined in Part 1.
- **Pot diameter**
  This has little or no effect on the amount of CO emitted, despite some rumours otherwise.
- **Burner diameter**
  By itself this has little effect on the amount of CO emitted.
- **Burner hole sizes**
  This has little effect on CO emission, but small holes can lead to the 'flame liftoff' hazard.
- **Burner design**
  The angle the flame comes out of the burner interacts with the burner to pot clearance and does influence the amount of CO emitted. Burners with tilted faces are better.
- **Radiation Shield**
  This can interact with burner design to affect the amount of CO emitted, but this is not a simple factor.
- **Air supply**
  It would be surprising if this had no effect, and it has been shown that the effect is real, and the results are similar to

matching area of outlet. Stoves which require a really tight windshield to operate may be a CO hazard.

Subsequent parts of this Series will look at the performance of actual stoves: canister, alcohol, white gas, kerosene and solid fuel.

Conclusion – so far

The opinion of the author is thus:

> ***There should be no CO hazard from operating a good canister stove inside a tent provided care is taken to ensure adequate ventilation. Highway levels of CO may be worse.***
>
> ***Other hazards exist of course – you could set the tent alight!***

Categories:
Research & Testing

By Roger Caffin

Research Scientist. Been walking all my life, mainly off-track - we don't have 'trails' here, and I always go with my wife. Summer and winter, lowland and highland, Australia and Europe. Forced into UL gear by heavy packs and increasing age. :-)

Discussion

Forums › Stoves, Tents and Carbon Monoxide – Deadly or Not?

Viewing 4 posts - 1 through 4 (of 4 total)

Feb 20, 2007 at 9:14 pm Reply #1221956

Benjamin Smith BPL MEMBER

Companion forum thread to: