

A Report on Some Heat Transfer Experiments

Dale Andreatta

dandreatta@SEAlimited.com

November 14, 2004



Executive Summary

In the late summer and autumn of 2004 a series of about 100 tests was conducted relative to increasing the heat transfer efficiency of cookstoves. The goals were to come up with a simple way of virtually guaranteeing high efficiency in a cookstove, or, failing that goal, to learn about what conditions make a cookstove efficient. The first goal was not achieved, and this report covers what was learned relative to the second goal.

Almost all tests were done using a simulated wood fire with natural gas as the fuel. The gas was burned in a low velocity fully non-premixed manner, and the flame appeared to be similar to a wood fire. The firepower could be precisely controlled and measured by varying and measuring the gas flowrate. Under conditions of limited air considerable soot was formed, much like a wood flame. Most tests were done at one of 3 levels of firepower. This approach appears to be very good for studying the heat transfer process, producing repeatable results in a relatively short time.

A new technique was also developed for data reduction, using easily measured variables to estimate parameters that are important in assessing the cookstove performance. The mass flow rate through the stove, average gas temperature coming up through the riser,

air-fuel ratio, and log-mean temperature difference at the cook pot were estimated using this technique. All of these variables would be difficult to directly measure, and all are important in understanding the workings of a cookstove.

Some preliminary work was done measuring temperature distributions in the gas around the pot. The temperatures and temperature gradients near the pot are much higher on the bottom of the pot than on the sides. Some preliminary studies were also made of unconfined flames, more or less simulating a 3-stone fire. While unconfined flames can lead to very high heat transfer and efficiency, the conditions which produce high efficiency tend to be the conditions that lead to high soot production.

In the more conventional stoves, the most important variable by far in determining the heat transfer to the pot is the temperature of the gases coming up through the riser. Within a fairly narrow band of scatter, the heat transfer to the pot is a linear function of this temperature. Another way to correlate this data is to calculate a log-mean temperature difference, which is loosely defined as the effective average difference in temperature between the pot and the gases that are flowing next to the pot. The heat transfer is proportional to this log mean temperature difference.

The corollary to this finding is that changes to the stove are generally ineffective, except in how they affect the average riser temperature. For example, the common belief is that tight skirts and tight flow passages increase heat transfer by forcing hot air against the sides of the pot or the bottom of the griddle in a griddle (plancha) stove. My conclusion is that this is mostly false, that skirts and tight passages may be helpful, but mostly because they choke off the excess air flowing through the riser and keep the average riser temperature higher. The bases for this conclusion are laid out in detail.

Conditions which lead to high efficiency (relatively low ratio of air to fuel) are the conditions than can lead to higher pollution, hence the air-fuel ratio needs to be controlled within a tight band to give good efficiency without high pollution. In all cases tested, the air-fuel ratio was much higher than that required for combustion, generally at least 4 times higher, suggesting there is still room to decrease the air to fuel ratio further.

A number of schemes were tried for increasing the heat transfer regardless of the combustion conditions. None of these techniques were successful, but are reviewed here.

It is generally believed that bigger pots lead to better heat transfer because more area is available for heat transfer. This document contains a summary of a brief set of tests that was performed to test this theory. The basic conclusion was that bigger pots do lead to increased heat transfer, however the gains are not large.

Throughout this report a number of “action items” are noted, ideas or experiments that would be good to try.

Section 1: Introduction and Test Setup

Between July 2004 and November 2004 a series of about 100 experiments were conducted in the gas appliance lab at my workplace. There were two goals in this work. The ultimate aim was to come up with a method of increasing the heat transfer to the pot which could be added onto any existing stove such that good efficiency could be virtually guaranteed regardless of the type of stove, power level, operator skill, etc. Ideally, this method would not affect the operation of the stove, such that the combustion process could be kept separate from the heat transfer process. This ambitious goal was not achieved.

The second goal was to learn about the heat transfer process and what variables were important in getting good heat transfer to the pot. The bulk of this report covers what was learned along these lines.

With only a couple exceptions, all of these tests used natural gas as the fuel, burning it at low velocity in a completely non-premixed fashion. This made the flame similar to a wood flame except that the heat released in the flame could be easily controlled and measured by measuring the flow rate of natural gas. Under some conditions, the flame became very sooty, much like a wood flame, and the amount of soot left on the pot was similar to what might be left with a sooty wood flame. While no pollutant or soot measurements were made, I believe that general observations can be made about low and highly polluting fires.

Figures 1.1 and 1.2 show the burner with the flow measuring instruments in the background. The rotameters used to measure flowrate were rated for air, and a correction was used since the flowing gas was natural gas.



Fig. 1.1: The burner and flow measuring system.



Figure 1.2: The burner in operation at about the middle power level. The flame in this photo has an unusually "tongued" appearance, partly caused by the camera flash.

The pot used in most of these tests was 9 5/8 inch (25 cm) in diameter. This was the “standard pot” given to me by Aprovecho after the 2004 ETHOS Conference. In most tests 4 liters of water were used, and a lid was used on the pot for all tests.

Most tests were done at one of three power levels. High power was with a firepower of 5417 Watts, medium power was 3250 Watts, and low power was 1444 Watts. About 2/3 of the tests were done at medium power. Most tests lasted about 10 minutes and the water was heated to a temperature well below the boiling point. Under these conditions the heat input rate to the water can be assumed to be constant, no water is boiled off, and the precise control over the firepower allows a lot of experiments to be done quickly and compared to each other. All test stoves were of low mass, thus the mass of the stove and heat input to the stove body were negligible.

Of course, any promising designs would need to be tested with wood as the fuel, but the gas tests are a good research and design tool. Once the best geometry is determined using the sheet metal prototypes, the final stove could be made of better materials if needed. The test method could easily be adapted to burn liquefied petroleum gas (propane).

Section 2: Data Reduction and Calculation of Parameters

A technique was developed for taking easily measured data and estimating important stove parameters which are difficult to directly measure. In particular, the air flow through the stove can be estimated by this technique, and also the average gas temperature at the top of the riser. From these, the log-mean temperature difference can be calculated, and also the excess air ratio.

In a given test, the natural gas flow rate can be measured, which tells us the firepower. Call this P . The rate of temperature rise of the water can be measured, which allows us to calculate the rate of heat input to the water. Call this Q . Both of these quantities are typically measured in Watts.

The gases flowing through the stove are assumed to have the thermal properties of air, and to have constant specific heat. It is also assumed that little heat is lost through the double wall riser, and that this heat loss can be estimated.

Performing an energy balance on the gas flowing through the riser gives:

$$\dot{m}c_p(T_o - T_{amb}) = P(1 - f_{fo} - f_{fp})$$

where:

\dot{m} = mass flow through riser (g/sec)

c_p = specific heat, assumed to be that of air

T_o = average temperature of gas at top of riser

T_{amb} = atmospheric temperature

f_{fo} = fraction of heat leaving the flame and being lost to the environment either through radiation or through the walls of the riser, or into the walls of the riser. (The riser was double wall sheet metal, thus the heat retained by the walls should be very small, and the heat passing through the walls should not be too large.)

f_{fp} = fraction of heat leaving the flame as radiation and impinging on the bottom of the pot

One can also perform an energy balance on the gas after it leaves the riser and before it exits the stove to get

$$\dot{m}c_p(T_o - T_{out}) = Q - Pf_{fp}$$

where:

T_{out} = temperature at outlet of stove

This analysis assumes that some type of average T_{out} can be measured. In cases where gases are exiting in a vertical direction, it was found that at a given point around the circumference of the pot the exit temperature was usually fairly uniform, though this could be a different temperature from other points around the circumference. One can estimate an average outlet temperature, though this estimate may require some educated

guesses about where most of the mass is exiting. For example, if a skirt is used and the gap between the pot and skirt is not uniform, most of the mass will come out where the gap is largest. Hence, the average outlet temperature would be weighted toward the temperatures measured at the wide spots. Thus, T_{out} should be regarded as an estimate, and the calculated quantities are estimates. Still, they should be accurate enough to be useful.

The above 2 equations can be combined to solve for the mass flow and T_o . The results are:

$$\dot{m} = \frac{P(1 - f_{fo}) - Q}{c_p (T_{out} - T_{amb})}$$

and

$$T_o = T_{amb} + \frac{P(1 - f_{fo} - f_{fp})}{\dot{m}c_p}$$

Note that in the equation for mass flow, f_{fp} has cancelled out, so for estimating mass flow it is not important to estimate f_{fp} accurately. Once the above values have been calculated, the log-mean temperature difference (LMTD) can be calculated, as well as the excess air ratio. The LMTD is an average effective temperature difference between the pot surface and the gas flowing around the pot surface. The excess air ratio is the actual air to fuel ratio divided to the stoichiometric air to fuel ratio (the stoichiometric air to fuel ratio is the air to fuel ratio where, theoretically, all the oxygen will be consumed and all the fuel will be consumed). The formulas are:

$$LMTD = \frac{(T_o - T_p) - (T_{out} - T_p)}{\ln \frac{(T_o - T_p)}{(T_{out} - T_p)}}$$

where T_p is the pot surface temperature, assumed here to be 77° C for all tests.

As mentioned above, it can be difficult to measure or estimate T_{out} accurately, and in cases where T_{out} is not much above the ambient temperature, errors in T_{out} have a significant effect on the calculated results. The calculated results do not seem to be very sensitive to any other variables. The partial derivative of both T_o and LMTD with respect to T_{out} is a little over 1.

The excess air ratio is usually given by the symbol λ (lambda) and is given by the formula:

$$\lambda = \frac{\dot{m} - P / 50,000}{\frac{17.2P}{50,000}}$$

The 50,000 comes from the 50,000 J of energy in 1 g of natural gas (assumed to be methane) and the 17.2 is the stoichiometric air to fuel mass ratio for natural gas.

In conditions where the gases are exiting the stove in any direction other than vertically, the temperatures vary greatly in the exit region. Since the velocity distribution is not known, it is impossible to estimate a value of T_{out} in these cases, and this calculation technique does not work. In about half of the tests reported on here this technique was used, and the results will be given in Section 5.

One can also measure the temperature of various surfaces on a cookstove and estimate the radiative heat transfer to the pot. This was done in a number of tests, usually using an infra-red thermometer (non-contact thermometer) to measure the outside temperature of the skirt or other stove piece. Typically, only $\frac{1}{4}$ or less of the heat transfer to the pot can be accounted for by radiation from a solid surface.

Section 3: Temperature Distributions

As a preliminary study, some gas temperature distributions around the pot were measured. All measurements were made with a thermocouple, the bead of which was generally a cylinder of diameter 0.02 inches (0.8 mm). The thermocouple itself was unshielded, and a correction was made for radiative effects. This correction should be accurate for thermocouples outside the flame, though the corrected temperatures may be considerably high for locations inside the flame. All temperatures reported here are the corrected temperatures. The locations of the thermocouples were eyeball estimated, but should be accurate enough to allow general conclusions.

Figure 3.1 shows the thermocouple locations for a test with medium power, standard pot, dimensions B = 19 mm, d= 125 mm, A = 184 mm (see Fig. 5.1 for definition of dimensions).




Fig 3.1: Thermocouple locations for Test 2.

Distance from pot (mm)	T1 (° C) (location 1)	T2 (° C)	T3 (° C)	T4 (° C)
3	828	630	365	269
6	469	469	454	309
9	477	309	372	302
13	230	113	269	330
19	---	---	146	230

The table above gives the thermocouple readings for the various locations. The general conclusion from these numbers is that on the bottom of the pot the highest temperatures are closest to the surface, while on the side of the pot the highest temperatures are away from the pot. This trend held in all tests where temperature distributions were measured.

Figure 3.2 shows the locations for another test with conditions of low power, $B= 19 \text{ mm}$, $d = 125 \text{ mm}$, $A = 184 \text{ mm}$, standard pot.




Fig. 3.2: Thermocouple locations for Test 4.

Distance from pot (mm)	T1 (° C)	T2 (° C)	T3 (° C)	T4 (° C)
3	462	372	198	124
5	435	282	211	118
9	233	201	220	118
13	167	140	173	115
19	155		107	118
25				104

Again, the table above gives the thermocouple readings. Here we see temperatures generally lower than in the medium power test. Again, on the bottom of the pot the highest temperatures are closest to the pot surface, while on the side of the pot the temperatures are highest farther from the pot surface.

Another type of temperature distribution was measured, this time at a fixed distance from the pot (3mm) and varying the distance from the center of the pot. Some numbers are given below. The conditions for this test were: high power, $B = 22 \text{ mm}$, $A = 184 \text{ mm}$, $d = 126 \text{ mm}$.

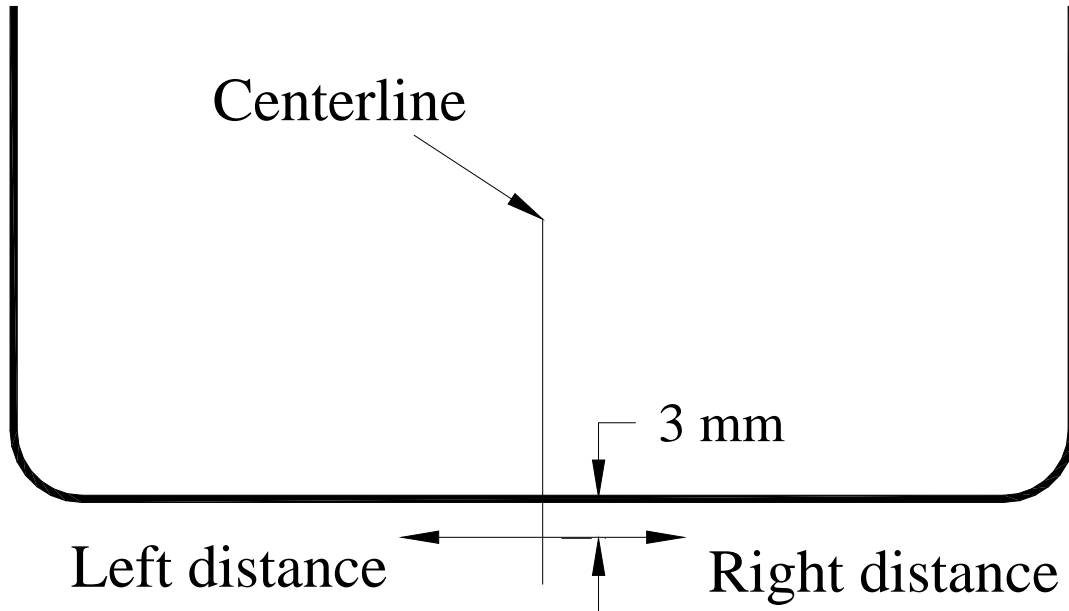


Fig. 3.3: Thermocouple locations for a high power test.

Left temperature (° C)	Distance from centerline (mm)	Right temperature (° C)
1900	0	1900
2008	25	2008
1700	51	1517
1275	76	1275
1132	101	1065

As expected, the highest temperatures are in the central zone of the pot, with the temperatures falling sharply at the outer part of the pot.

Various other temperature measurements were made under other conditions. The gas temperatures around a pot on a modern gas kitchen stove were measured, as well as temperatures around the pot when a laboratory burner (Fischer burner) was used. With the Fischer burner 35 mm below the bottom of the pot very high temperatures were measured at a location 3 mm from the pot and near the outer edge, with temperatures decreasing as one moved farther from the surface. On a gas stove a similar trend was noted, with the hottest gas closer to the pot surface and decreasing temperatures farther away. The peak temperatures in the gas stove test were generally lower than in the Fischer burner test.

The gas flowing below the pot seems to segregate itself with the hotter gas staying up close to the bottom of the pot with the cooler gases forming a hot but not quite as hot layer below the hottest layer. Very close to the pot, less than 3 mm, the gas would be cooler as heat is transferred to the pot surface.

It would be expected from these temperature distributions that the bulk of the heat transfer goes through the bottom of the pot. As an independent confirmation of this observation, all tests were done with chlorinated city water in the pot. As the water heats up, it will start to form bubbles first in the hottest portions of the water. In all cases, the vast majority of the bubbles was on the bottom of the pot, further suggesting that the bulk of the heat transfer goes through the bottom of the pot.

When a skirt or other tight channel is used, temperature measurements generally can not be taken in the tight channel, hence a greater portion of the heat transfer could come through the pot sides. In the tight skirt tests, it was still observed that the majority of the bubbles were on the bottom of the pot. More will be said about the effectiveness of skirts and tight passages in Section 5.

Action Item: It would be interesting to use a pot with 2 or 3 separate chambers, such that the temperature rise of each chamber could be measured separately and the temperature rise could be related to the heat entering that chamber. One chamber would cover the center bottom of the pot, another chamber would cover the outer bottom region, and the third chamber would cover the sides of the pot. With such a device, one could directly measure the heat entering the pot from the 3 separate regions. I have done some preliminary design calculations which show, however, that it would be difficult to design separators that would prevent excessive heat from being transmitted from region to region. Perhaps some other clever means of achieving the same result could be designed.

Action Item: It would be good to take similar measurements on a plancha stove, and to see how these measurements vary with the size of the channel under the plancha. Conventional wisdom is that this channel needs to be kept as narrow as possible to get good heat transfer, but I expect that the temperature distribution and heat transfer will not depend on the channel depth. If this is true, it would allow plancha stoves to be made with much looser tolerances on some of the dimensions, which would make them easier to build and maintain.

Section 4: A Brief Study of Unconfined Flames

A brief study was conducted of unconfined flames, using the natural gas burner described previously without any sort of duct. These experiments more or less simulate a 3-stone fire.

The results are summarized in the graph below. The distance dimension is the vertical distance from the top of the burner pipe to the bottom of the pot. A standard pot was used in all tests, and 3 power levels were used.

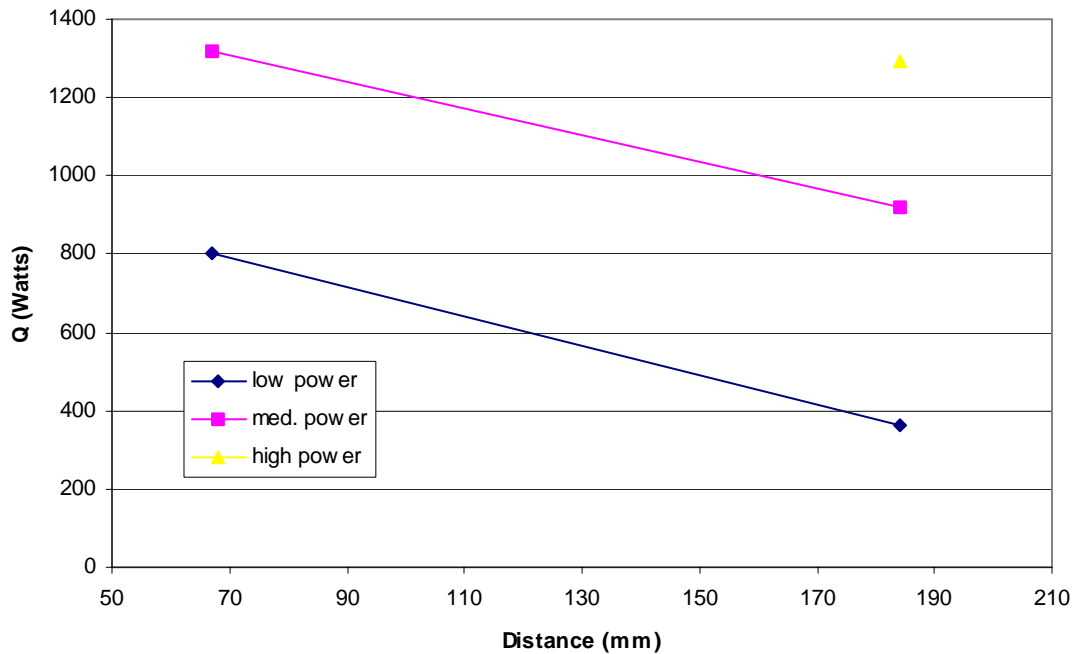


Figure 4.1: Heat transfer to the pot vs. distance from the burner to the bottom of the pot.

It can be seen that putting the pot closer to the burner increased heat transfer effectiveness greatly. The higher low power point represents an efficiency of 0.56, quite a good number. This would be expected, since it has long been known that cool air is entrained into a fire, such that the average gas temperature decreases rapidly above the flame (Cox and Chitty, 1980). One way of getting consistently good efficiency would be to put the pot directly above the burning wood, however this leads to high soot formation at the higher power levels.

It is known that open fires are highly variable in efficiency, sometimes attaining good efficiency, sometimes poor efficiency. This study suggests that the distance from fire to pot is an important variable, as well as the firepower, and also the pot size. If the test is performed outdoors any crosswind would be important. These 4 factors together could explain why a wide range of efficiency numbers are recorded for a similar test setup.

Section 5: How to Achieve High Efficiency

5.1 Test Results

A general drawing of a stove is shown below for reference purposes. Not all part of the stove were used in all tests. The skirts and bases were of sheet metal with shiny outsides, except for small sections painted black so that the infra-red thermometer could be used to measure surface temperature. (See Fig. 1.2, for example.) The shiny surface will have low emissivity, and will loose heat by radiation much less than a non-shiny surface. Preliminary calculations show that the bulk of the heat loss from such a surface is through radiation rather than convection. Thus, while the skirt and base were not insulated, they did have considerable resistance to heat loss.

The riser was double walled furnace duct, which should also have considerable resistance to heat loss.

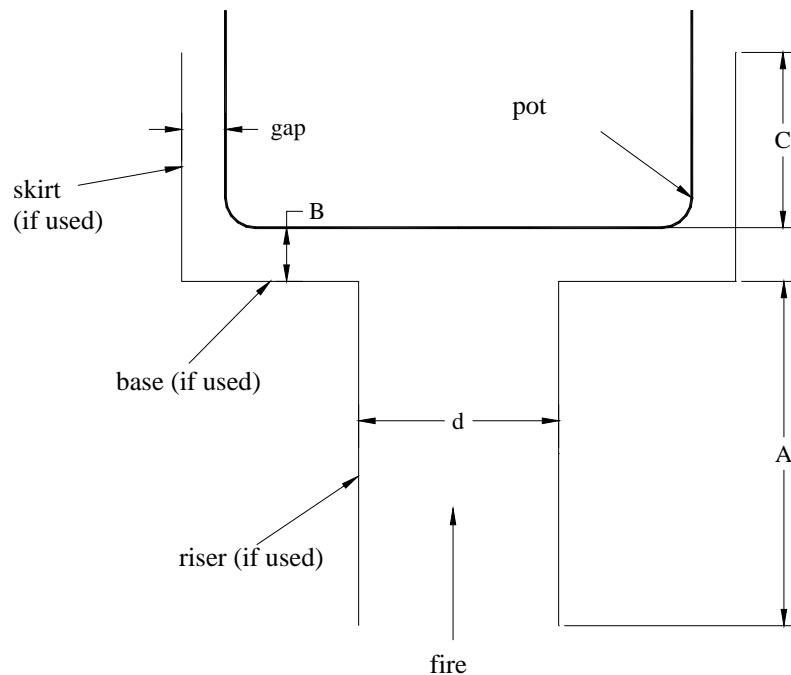


Fig. 5.1: General stove dimensions. Not all dimensions will apply to each case.

One interesting way to look at the results of these tests is to use the analysis presented in Section 2, where the average riser temperature and log-mean temperature difference (LMTD) are used. If one uses simplistic reasoning, one would expect the heat transfer to the pot for a given pot to be proportional to the LMTD. It was noted in Section 2 that the methods of Section 2 can not be applied to all tests. In particular, there must be a well defined output temperature, T_{out} , for this method to work. Of the roughly 100 tests performed, 42 were selected for further study using the methods of Section 2.

Figures 5.2 and 5.3 begin to explore these results. Figure 5.2 shows the heat transfer against the average riser temperature, T_o . The best fit line is given through the data. It appears that if this line were extended to the horizontal axis (where the heat transfer would be zero) the value of T_o would be about 70°C . This would approximately be the pot temperature, which suggests that this method is reasonable.




Figure 5.2: Heat transfer vs. average riser temperature for 42 data points.

One can generally see three groups of data points corresponding to low power, medium power, and high power tests. Most of the tests were at medium power, so the bulk of the data points fall in this middle group.

Figure 5.3 shows the heat transfer for the same 42 test points plotted against the estimated log-mean temperature difference. The linear curve fit in Fig. 5.3 is forced to go through the origin. The most interesting aspect of these 2 charts is the fairly narrow scatter band of the data. The 42 tests given include tests with tight, loose, and no skirts. Some stove prototypes had a base, others did not. Some stove prototypes had a riser, some did not. As one can see there is a wide range of temperature and heat transfer conditions.

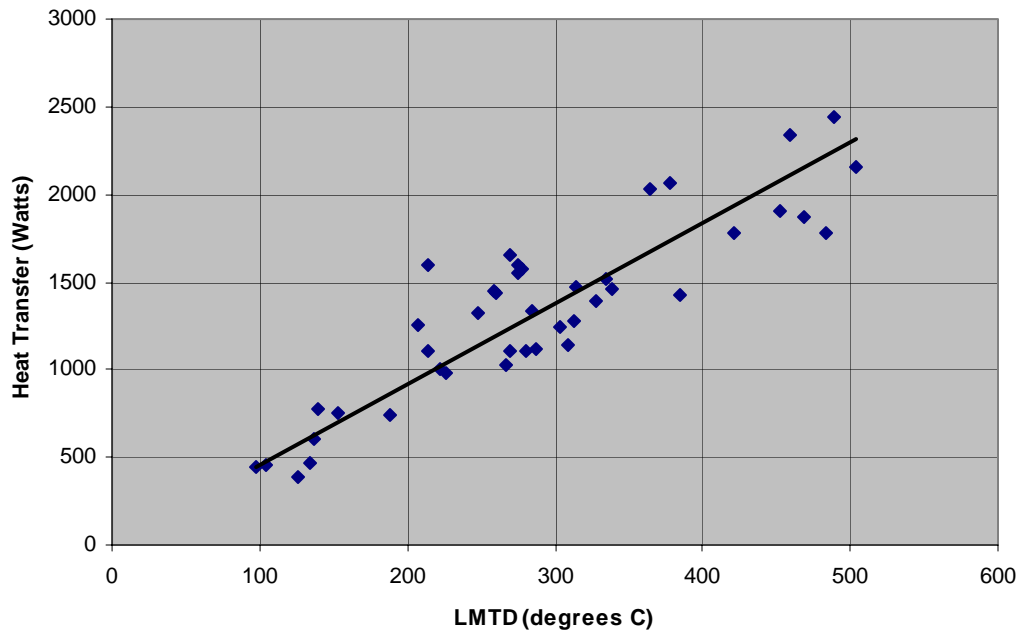


Figure 5.3: Heat transfer vs. log-mean temperature difference for 42 data points.

One common belief about stoves is that the key to high heat transfer efficiency is tight passages around the pot. Skirts are the most common way to achieve this. The above data suggests that the heat transfer is largely a function of the gas temperatures and not a strong function of skirt or stove geometry. In other words, adding a skirt to an existing stove might slow down the gas flow through the stove, giving less excess air to the fire, keeping the gas temperatures high and leading to high heat transfer, but the main cause of the increased heat transfer is the increased gas temperatures, not the presence of the skirt. If increased gas temperatures could be achieved in some other way, high heat transfer would also result.

As further evidence that a skirt has limited value, one can compare the results of 4 tests with similar conditions. The conditions for these tests were: medium power, $B = 51$ mm, $d = 126$ mm, $A = 184$ mm.

Conditions	Min. Flow Area (cm ²)	Heat Transfer (Watts)
No skirt	127	1107
Skirt, gap = 19 mm, highly non-uniform	127	1025
Skirt, gap = 19 mm, fairly uniform	127	1109
Skirt, gap = 9 mm, highly non-uniform	73	1465

We see that the skirt has little or no effect until it is tight enough to close down the minimum flow area, whereupon it makes a significant difference.

This is not to say however that skirts are completely ineffective. If one breaks the above tests down into 2 groups, one with tall skirts and one either with no skirt or with a minimal skirt one gets the results shown in Figs. 5.4 and 5.5.

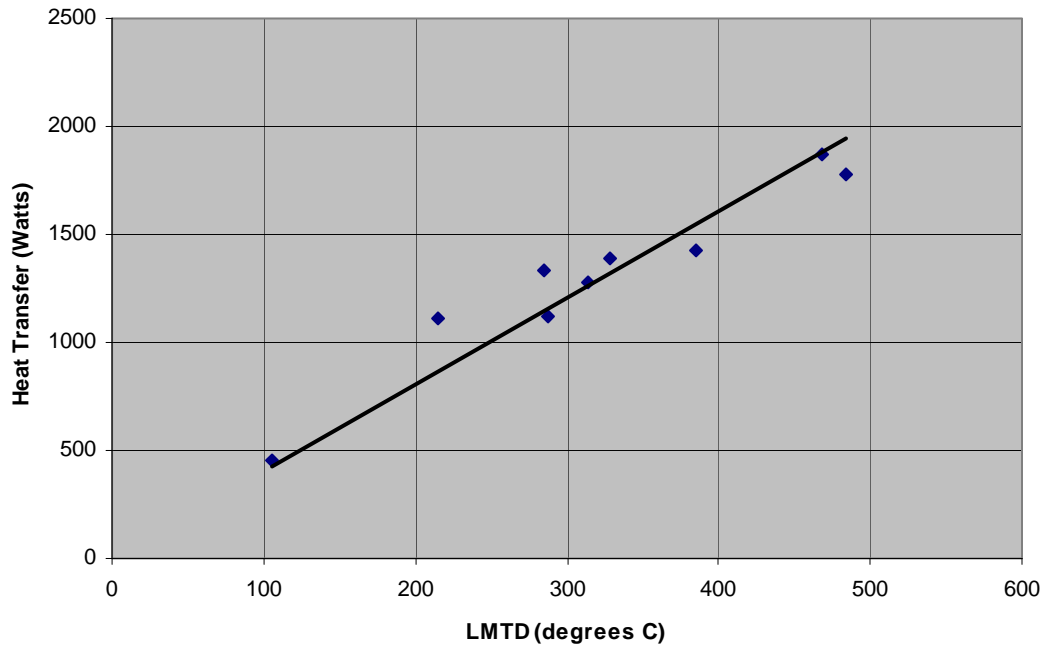


Figure 5.4: Heat transfer vs. log-mean temperature difference for unskirted stoves.

The best fit line again goes through the origin and has a slope of about 4.0 Watts/° C.




Figure 5.5: Heat transfer vs. log-mean temperature difference for skirted stoves.

The best fit line goes through the origin and has a slope of about 4.75 Watts/° C. This is 19% higher than the slope of the line through the non-skirted stove data, and suggests that a skirted stove will tend to have an efficiency somewhat greater than a non-skirted stove, once the effects of increased gas temperature have been accounted for. One can again see the narrow bands of scatter for the stoves. Obviously, more skirted stoves were tested than unskirted, leading a greater variety of conditions and greater scatter.

A number of tests were done in which the outside temperature of the skirt was measured with a non-contact thermometer (infra-red thermometer). Since the skirt was a single layer of sheet metal, the inside temperature would be essentially the same as the outside temperature, and the radiative heat transfer can be estimated from the skirt to the pot. Typically, this heat transfer is on the order of 10-25% of the total heat transfer, which suggests that the main benefit of having a skirt is the increased radiative heat transfer to the pot. Note that without a skirt, the hot gases still flow along the sides of the pot, but not as tightly.

One might investigate whether a base (as shown in Fig. 5.1) helps. Two pairs of tests can be used for a comparison. The conditions here were: medium power, A = 184 mm, d = 126 mm.

Dimension B (mm)	Min. Flow Area (cm ²)	Base (yes/no)	Heat Transfer (Watts)
19	76	No	1124
19	76	Yes	1333 (19% more than without base)
57	127	No	797
51	127	Yes	1107 (39% more than without base)

It appears that a base makes a considerable difference. Some basic temperature measurements of the base were taken with the infra-red thermometer, and these can be used to estimate the radiative heat transfer. In both cases in the above table, the estimated heat transfer from base radiation is on the order of 20-30 Watts, not nearly enough to account for the difference between base and no-base conditions. It appears that a base is a good investment in increasing stove efficiency, especially since a flat base need not be sized to fit a certain pot. In other words, it's easy to make a base, while it is more difficult to make a skirt, and much more difficult to make a tight fitting uniform gap skirt.

To further examine this question of the effect of excess air on heat transfer, a series of 13 tests were done at various values of A, B, and d (see Fig. 5.1 for definition of these variables). Figure 5.6 shows the heat transfer vs. minimum hot flow area. The minimum hot flow area is the smallest area through which hot gases will flow. This may be thought of as the “choke point” of the flow. (Since cold input gases are much more dense, they can flow through a much smaller area without being choked.) If dimension B is less than

$d/4$, the choke point will be where the flow turns the corner at the top of the riser, and the minimum flow area will be $Bd/4$. If dimension B is greater than $d/4$, the minimum flow area will be in the riser itself, and the area will be $\pi d^2/4$.

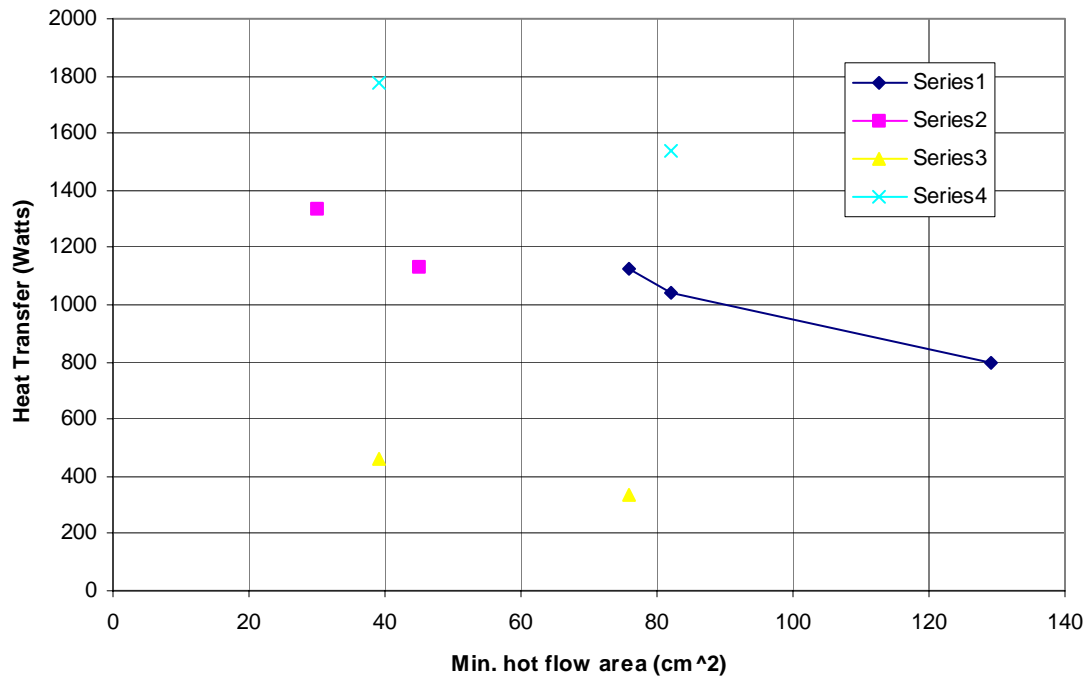


Figure 5.6: Heat transfer to pot vs. minimum hot flow area.

Series 1-Medium power, $d=126$ mm, $A=184$ mm.

Series 2-Medium power, $d=76$ mm, $A=305$ mm.

Series 3-Low power, $d=76$ mm, $A=305$ mm.

Series 4-High power, $d=126$ mm, $A = 279$ mm.

It can be seen above that reducing the flow area increases the heat transfer. This is presumably a result of less excess air being pulled up through the riser to mix with the hot gas from the flame. Similar results were seen in the various tests with bases and skirts.

Figure 5.7 lends further evidence to the excess air theory. Here are 3 pairs of tests with different riser heights but the same power and flow area. A taller riser would tend to produce more chimney effect, hence more excess air and lower temperatures and heat transfer. This is indeed what we see.

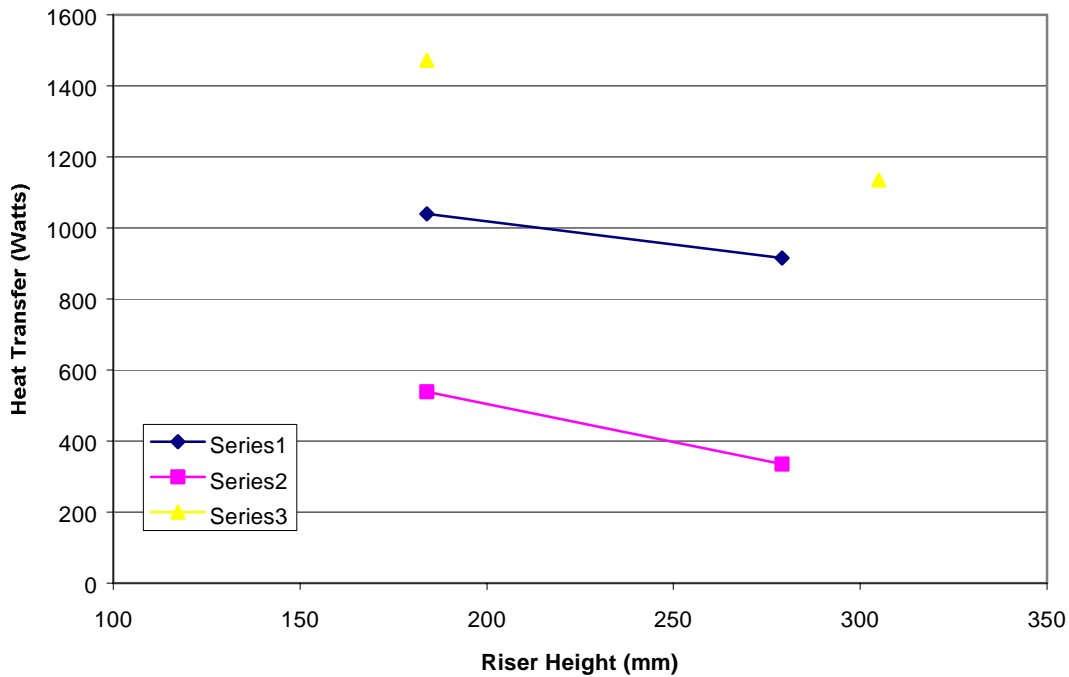


Fig. 5.7: Heat transfer to pot vs. riser height.
 Series 1-Medium power, $d=126$ mm, min. flow area = 82 cm^2 .
 Series 2-Low power, $d=126$ mm, min. flow area = 76 cm^2 .
 Series 3-Medium power, $d=76$ mm, min. flow area = 45 cm^2 .

Under some conditions, considerable soot was formed in the above tests. This was particularly when the firepower was high and the flow area was small. Thus, there appears to be a fundamental tradeoff between efficiency and pollution. However it might be possible to find a region with both good efficiency and good pollution, or it might be possible to find a way of getting lower soot at lower excess air ratios.

The mass flow estimation technique from Section 2 says that even under highly sooting conditions the excess air ratio (ratio of actual air flow to the minimum airflow required for complete combustion) is still at least 3. It is likely that there is plenty of excess air, but the mixing of air with fuel is not right to achieve clean combustion. Other studies have achieved clean combustion in a downdraft burner with excess air ratios of around 1.7 (Khan, 1991). It could be that there is not enough mixing and combustion finishes with not enough air in some zones, or it could be that there is too much mixing, quenching the combustion gases before the reactions are complete.

Action Item: It would be good to repeat the above tests with a quantitative measurement of soot and CO, as all of the above soot conditions were qualitatively observed only.

Action Item: It would be good to repeat the above tests with wood, and to vary the mixing conditions to find the right amount of mixing to give clean combustion with high temperatures.

5.2 Comparison with previous results

In June and July of 2000, Ken Goyer built and tested a number of stoves at the Aprovecho Center in Oregon. While he tried to achieve steady combustion at nearly constant power during each test, his tests used wood as the fuel, and thus there could be variation in power between tests. Regardless, it is interesting to compare his results with the ones presented here.

His results showed that chimneys (risers) with smaller diameters and shorter heights were more efficient. This agrees with the results here, and agrees with the theory that more excess air leads to lower heat transfer efficiency. He also did a number of tests showing that dimension B (Fig. 5.1) has a large effect on efficiency, with smaller values of B leading to higher efficiency. This also agrees with the results here. He also showed that the addition of a skirt led to a large increase in efficiency, however, one can calculate from his reported dimensions that the skirt cut the minimum hot flow area by over 50%. Thus, one would expect the decrease in air flow to be the major source of the increase in efficiency.

One significant difference between Goyer's results and the results here are his conclusion that adding insulation to a skirt increases the efficiency. My results showed that insulating the skirt led to only a small increase in efficiency. One reason for this discrepancy might be that my uninsulated skirts were made of fresh sheet metal which is shiny on the outside, and thus will lose a lot of heat. Thus, my uninsulated skirts lost less heat than his uninsulated skirts, while my insulated skirt probably lost more heat than his insulated skirt.

Perhaps the largest difference between Goyer's results and mine was that he measured the oxygen in the outlet stream. One can use this to calculate an excess air ratio, and if one does this one gets numbers much smaller than mine. His average excess air ratio was 2.92, he occasionally saw numbers as low as 1.29. My excess air ratios were rarely less than 3.

5.3 Conclusions

If one were to draw up a preliminary list of "steps" that one could take to get good heat transfer efficiency, it would look something like the following.

1. Match the size of the fire to the size of the cooking task. A 3 kW fire at 50% efficiency delivers the same power to the pot as a 5 kW fire at 30% efficiency, while burning only 60% as much fuel. It is easier to get high efficiency out of a small fire than a large one.
2. Choke off the excess air as much as possible without sooting.

3. If possible, use skirts and large pots.
4. For the simmering phase, use as small a fire as possible, and a lid on the pot.
Note that the amount of power required to keep a pot simmering with a lid is very small, hence it is likely that the fuel consumption during the simmering phase will depend more on how small a fire can be used and maintained, preferably without operator intervention, than it will on the heat transfer efficiency.

Action Item: It would be good to design a stove that could keep a small simmering fire going without operator intervention (in other words, it would be good to have a stove with a high turndown ratio). If a lid were used on the pot, the wood consumption could be made very small, even if the heat transfer efficiency were not great.

Section 6: The Effect of Pot Size

It is generally believed that the pot size has a strong effect on the stove efficiency, with larger pots leading to better heat transfer. A brief study was conducted along these lines, using 3 different size pots at the same 3 power levels as used in the main part of this study. The middle pot size was the standard size used in the majority of these tests. Tests were conducted using the unconfined burner case, which is roughly like a 3-stone fire. (See the figure on the cover page of this document.) The results are given below.

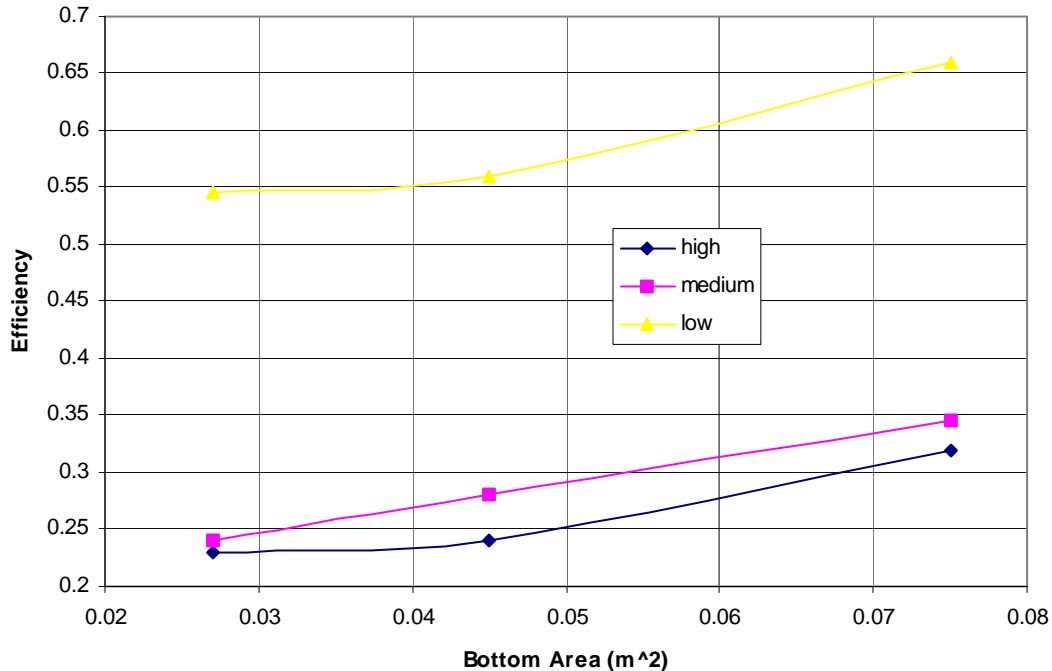


Fig. 6.1: The effects of pot size and power on efficiency.

For these tests, at each power level the distance from the burner to the pot was fixed for each power level, but was not the same for all power levels. (See Section 4 for a discussion of the effect of burner to pot distance.) The distance was always such that little or no soot was formed. The results are shown vs. bottom area only, as it was previously shown that most of the heat transfer comes through the bottom.

It can be seen that increasing the size of the pot increases efficiency, but the effect is not terribly large at any power level. There may be several reasons for this. First, it was shown earlier that if one looks at the gas close to (3mm) from the bottom of the pot, the hottest gases by far are close to the center of the pot. The heat transfer should be largely driven by this temperature difference. Second, a certain amount of heat will go directly from the flame to the pot by radiation, and this will be concentrated at the center of the pot, and a larger pot doesn't intercept more radiation. This is true whether a riser is used or whether the flame is unconfined. Finally, according to theory the boundary layers will get thicker as one moves away from the center of the pot, leading to lower temperature gradients and lower heat transfer even if the temperature is high.

Section 7: Various Methods to Increase Heat Transfer

This section relates the results of a number of tests that were attempted in order to come up with a way to increase the heat transfer. In most cases these methods are not promising, but are included here for completeness.

7.1 A sloped-sided pot

It was previously described (see Section 3) that most of the heat transfer seems to be through the bottom of the pot rather than the sides. One might wonder, does the pot surface need to be horizontal for this effect to take place?

A bowl with sides about 49° from vertical was used as a pot. The area on the sides of the pot in contact with water was 2.9 times the bottom area of the bowl. If the heat transfer to the sloping sides was similar to the bottom, this represents a way of increasing the effective area of the pot by 2.9 times.

The bottom area of the bowl was about 87% as much as the small pot used in the tests described in Section 6. The heat transfer was 13% more than the small pot at the same power level (medium power). The heat transfer per unit area of the bowl that was in contact with water was only slightly greater on the bowl than on the small pot. Thus, it appears that a pot with sloped sides is slightly more effective at transferring heat than a vertical-sided pot, but only slightly.

7.2 Finned pot

A prototype finned pot was built using aluminum tape as the fins. While the tape is very thin (0.005 inches or 0.127 mm) it can be doubled over to double the thickness. The aluminum has such high conductivity that the calculated fin efficiency is on the order of 90%, and the high heat transfer efficiency keeps the fin relatively cool. The total surface area of the fins was a little over twice the area of the bottom of the pot. A photo is given in Fig. 7.1.



Fig. 7.1: The prototype finned pot used with the Fischer burner. The fins protrude about 25 mm from the bottom of the pot.

Two tests on the finned pot were performed. The first was done prior to the development of the pseudo-wood flame described in Section 1. This test was done with a Fischer burner in a duct. The same standard pot was used as in most of these tests. The heat transfer was about 65% greater with fins than without. The increase in heat transfer was much smaller than the increase in area. One reason for this could be that the fins were about 25 mm long, and thus extended well beyond the zone where the hottest gases were present.

In both the finned and unfinned pot tests the efficiency was very low, about 9.5% with the unfinned pot and 16% with the finned pot. The probable reason for this low efficiency was that the Fischer burner created a jet and sucked a very large amount of excess air up the duct, giving low temperatures and low efficiency. The point to this test was to compare efficiency in finned and unfinned pots, and these tests should not be compared with later tests done with the pseudo-wood burner and restricted excess air.

The second test of a finned pot used the same pot with the same fins, but bent the fins over. Recall that in Section 3 it was shown that the hottest gases are always close to the pot, hence having fins that stick out 25 mm from the pot penetrates far through the layer of hottest gas and most of the fin surface is exposed to much cooler gases. In this second finned-pot test, the pseudo wood flame was used. The finned pot only absorbed 31% more heat than an unfinned pot under the same conditions, but efficiencies were about 0.36 with the fins and about 0.275 without.

While fins did increase the heat transfer significantly, the question of how to make a finned pot inexpensively remains a difficult question.

7.3 Vortex generation

A delta-wing aircraft will develop 2 strong roll vortices, one over each wing (Sforza, 1975). A stove prototype was designed that had a triangular piece of sheet metal in the riser with one point down. This formed a configuration similar to that of a delta-winged aircraft, and it was observed that strong roll vortices appeared in the flame similar to those above the wings of such aircraft. It was hoped that this rolling motion would stir up turbulence and provide increased convective heat transfer on the bottom of the pot.

Comparing similar tests with and without the delta-wing showed that it provided no improvement in heat transfer. If anything, heat transfer was slightly worse with the delta-wing. It is possible however, that the delta-wing would provide the kind of large scale mixing that is often useful to reduce the pollution output of a stove.

Action item: Try the delta-wing concept in a wood burning stove while measuring pollutants to see if it improves the completeness of the combustion and lowers the pollutants. The results may depend on whether the delta-wing is the choke point of the flow or not. The test should be done with the delta-wing as the choke point, and with the delta-wing not as the choke point.

7.4 Radiators

If a solid surface is at about 551°C it will radiate approximately $25,000\text{ W/m}^2$. If the area of the pot is 0.04 m^2 this will produce 1 kW of heat transfer. The idea here was to have a system that absorbed heat from the flue gas a significant distance from the pot bottom, then this system would get hot and radiate heat to the pot. While the surface area of the bottom of the pot is limited, the solid radiator could have greater area. Also, after the flue gas passes through the radiator it still has substantial heat left to transfer to the pot by standard convection.

A typical setup was to use several layers of wire mesh above the riser and below the pot. The hot gases flowed through the wire mesh, heating the wire mesh, which radiated heat to the bottom of the pot. In one typical test, the central portion of the wire mesh heated to red heat, however, the process of pulling heat out of the flue gas cooled the gas so rapidly that considerable soot was formed. While the hot central section of the mesh presumably transferred considerable heat to the pot by radiation, only the central section of the mesh was hot, and the overall heat transfer to the pot was essentially the same as without the wire mesh.

Similar experiments using metal plates of various geometries and beds of stones instead of a wire mesh produced similar results. In no case was there significant gain in efficiency. In most cases there was a considerable increase in soot production.

7.5 Bottom-choked combustors

It was noted in Section 5 that the main key to achieving good efficiency is to limit the excess air such that the average riser temperature is high. This can be done by limiting the flow at either the top or bottom of the combustor. It was noted that the limiting factor in choking the flow is the point where the flame starts to be sooty, however it was also noted that even under sooty conditions the excess air ratio is 3 or more. One might suspect that insufficient mixing is the cause of the high soot production in the presence of considerable excess air. If the flow is choked at the bottom, there will be considerable turbulence downstream of the flow constriction, one might expect that this will improve mixing and lead to greater ability to reduce excess air, leading to higher temperatures and greater efficiency, while still providing a clean flame.

The table below shows the comparison between 2 prototype stoves. In both cases the riser was 126 mm diameter and 305 mm tall. In each case the flow was reduced until a minimal amount of soot was produced, either by constricting the flow at the top or at the bottom of the riser. (In the low power tests the flow was restricted as much as practically possible and no soot was produced.) In the bottom-limited tests, a tall skirt was used, while in the top-limited tests, only a short skirt was used. One would expect the bottom-limited stove to be more efficient, all else equal, due to the presence of the tall skirt.

Power level	Efficiency (top limited)	Efficiency (bottom limited)
High (5417 W)	0.39	0.38
Medium (3250 W)	0.43	0.45
Low (1444 W)	0.56	0.52

One finds, however, that the top limited stove without a skirt was about the same as a bottom limited stove. The mass flow rates for the 2 stoves are comparable at a given power level. Thus, while bottom limiting is a workable idea and gives good efficiency, top limiting is even more effective and probably a lot easier.

7.6 Turbulence Promotion

When a gas flow from a small passage in to a larger passage the speed of flow slows down and the kinetic energy decreases. In theory, the lost kinetic energy shows up as increased turbulence, and it might be reasonable to expect that this turbulence would lead to better heat transfer.

A pair of tests were done, one with the skirt gap increasing in the direction of flow, and the other with the gap decreasing. By the above theory, the one with the diverging gap should have better heat transfer. The minimum gap in both cases was about the same, such the flow area should be about the same. Contrary to expectations, the converging gap stove had about 10% more overall heat transfer than the diverging gap.

Section 8: A Promising Idea?

While most of the work here was an investigation rather than a design project, the lessons learned were applied to design a new type of stove and perform some preliminary tests. I call this stove a “skirted campfire” or perhaps a “skirted 3-stone fire”. The stove would fit into the category of stoves called refugee stoves, as it is built from a single sheet of metal, could probably be made a very low cost, but would not be expected to last for a terribly long time. The skirted campfire is shown in one form below.

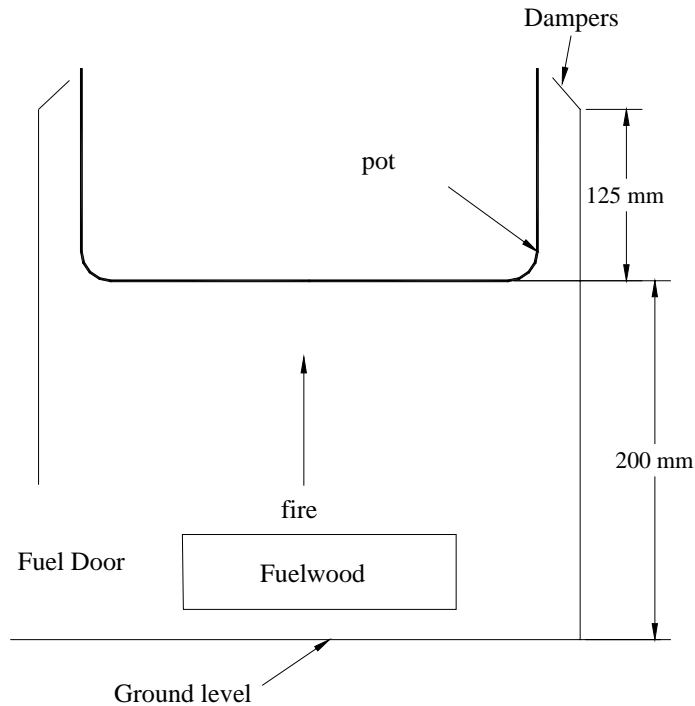


Fig. 8.1: The proposed skirted campfire stove.

The stove is basically a sheet metal ring with approximately the dimensions shown. The gap between the pot and the skirt is intentionally not made too tight, perhaps 12-15 mm. A fire would be started as normally would be with a 3-stone fire. Once the fire was started the stove would be placed over the fire and the pot placed in the stove. Fuel would be added as necessary through the fuel door.

At the top of the skirt could be flexible metal dampers as in the prototype, or some other variable method of restricting air flow. The basic idea is that for a given power level, the dampers would be opened or closed so that clean combustion was barely attained, thus giving a good balance between having enough excess air for clean combustion, but not having so much excess air as to cool the gas flow and give low efficiency. The user could set up the stove and close off the dampers until soot started to form, then open them slightly. If the stove were used outdoor, the dampers could be closed off more in order to get better efficiency at the expense of more pollution. At low power levels the dampers would be closed off more than at high power levels.

A variation of the above would include a skirt with a slight taper, wider at the top and narrower at the bottom. There would be no dampers, and the minimum flow area would be controlled by raising the pot to open the flow area and lowering the pot to close off the flow area. Some type of adjustable pot support could be used.

Some preliminary tests were performed with the gas burner. In these tests, the dampers were adjusted such that soot formation was minimal. The average gap around the pot was about 15 mm, but was not uniform. At the high power setting, the dampers were adjusted so that very little soot was formed.

Power level	Damper setting	Efficiency	Best efficiency at that power level while still keeping soot minimal
High	Medium	0.375	0.39
Medium	Nearly closed	0.51	0.51
Low	Very closed	0.52	0.56

At medium power the dampers were nearly closed and again very little soot was formed. At low power the dampers were closed even more and no soot was formed. The table above shows the efficiency achieved with the skirted campfire, and also the best efficiency of any design at that power level while still keeping soot minimal. The highest efficiency was achieved with stoves similar to Fig. 5.1 with the dimensions optimized. The skirted campfire gives good efficiency over the entire power range while still producing little or no soot. While the efficiency was not quite as high as with other designs, the simplicity of this stove is much greater.

Action item: Obviously this stove would need to be tested further. Tests should be done with full pollution monitoring, and of course with wood instead of gas as the fuel. The design alternative described above with the tapered sides should be investigated.

References and Acknowledgements

Thanks go to my company, SEA Ltd. for allowing me to do most of this work in their gas appliance lab. Also thanks go to my co-worker, Pete Susey, for coming up with the idea of a diffusion flame gas burner, without which this work would never have gotten started. Thanks also go to my mom for lending me her large pot and bowl for these tests. Thanks also go to my two volunteer lab assistants, Lisa Holt and Sue Fitzgerald.

Cox, G., and Chitty, R., A Study of the Deterministic Properties of Unbounded Fire Plumes, 1980, Combustion and Flame, V. 39, pp. 191-209.

Khan, Hasan R., Woodcombustion Studies, Part II Edited by E. Schutte, K. Krishna Prasad, and E. Moerman, Chapter1, 1991.

Sforza, Pasquale M., Vortex Augmentor Concepts for Wind Energy Conversion, Proceedings of the Second Workshop on Wind Energy Conversion Systems, 1975, NSF-RA-N-75-050.