



Burning Desire for Efficiency

Posted by [JoulesBurn](#) on June 28, 2012 - 3:32pm


Topic: [Demand/Consumption](#)

Tags: [energy efficiency](#) [[list all tags](#)]

This is a guest post by Tom Murphy. Tom is an associate professor of physics at the University of California, San Diego. This post originally appeared on Tom's blog [Do the Math](#).

Ever wonder how efficient it is to heat water? Of course you have! Ever measured it? *Whoa, mister, now you've gone too far!*

I recently devised a laser-phototransistor gauge to monitor my natural gas meter dial—like ya do. As a side benefit, I acquired good data on how much energy goes into various domestic uses of natural gas. Using this, I was able to figure out how much energy it takes to heat water on the stove, cook something in the oven, or heat water for a shower. Together with the knowledge of the heat capacity of water, I can compute heating efficiency from my measurements. What could be more fun? I'll share the results here, some of which surprised me.



Heating Basics

The amount of energy it takes to heat water is so well-established, that it is the *basis* for several prominent units of energy. For instance, the calorie is the amount of energy it takes to heat one gram (1 ml) of water by 1°C. As a straightforward extension, 1 [kcal](#) = 4184 J (often Calorie with capital C) is how much energy it takes to heat one kilogram (or liter) of water by 1°C. Likewise, 1 [Btu](#) = 1055 J is the amount of energy it takes to heat one pound of water by 1°F.

So if I want to take 500 ml of water from 18°C to boiling, I need to expend 82×0.5 kcal to get the job done, or 171.6 kJ.

Measuring the Gas


My natural gas meter usually receives little attention from me—which is saying something for a person as measurement/data crazed as myself. The reason is that the meter does not provide sufficient information to track small expenditures without vigilant monitoring. As explained in the [pilot lights post](#), there are dials that make revolutions once every half-cubic-foot and two-cubic-feet, then a jump to 1000 cubic feet. The jump is so large that one cannot walk up to the meter and know at a glance how many wraps the high-resolution dials have made since the last look.

Rather than parking myself outside for days on end to keep track of my gas gauge, I bought a cat-toy laser pointer and modified it to be powered by a constant 5-volt power supply. I aimed the laser at the 1/2-cf dial so that the black needle would interrupt the beam once per revolution. Then I set up a photo-transistor to “watch” the laser spot, and tuned the sensitivity so that the needle would make a robust change in the photocurrent. The laser and detector were shoved into

a crudely drilled block of plywood to point at the same spot on the dial face, and then clamped to the meter. A dishtowel provided ambient light baffling, and a plastic shopping bag gave it some modest protection from rain and a certain trashy look. Thankfully, our gas meter information is now digitally transmitted. I can't think what the meter reader's reaction would be to walk up on this clap-trap arrangement (glowing red at night, no less).




Contraption attached to gas meter. A red glow can be seen from the laser. The green/black wires go to the phototransistor.



Close-up showing the laser spot illuminating the half-cubic-foot dial, ready to sense the passage of the dark needle. See also a picture of the same meter face-on in the pilot light post.

I kept the electronics in the garage (only the laser and sensor were outside), accompanied by a data acquisition unit and a computer to log the continuously sampled series (3 second samples were sufficient). Not the most elegant way to get the data, but I used stuff I had on hand. And it worked. And the plots are nice.



Here we see the hot water heater pilot light causing the sensing circuit to respond in a periodic fashion. Every time the needle crosses in front of the laser, the sensor voltage goes way up. The needle can be a little jerky, so it sticks and slips—causing some spikes to be stunted and steppy. In this plot, ten revolutions span 11.2 hours. At 510 Btu per revolution, this corresponds to 133 W of continuous power in the pilot.

Most of the time, the water heater pilot light creates our only gas demand. In the summertime, even our shower activity—albeit modest—is satisfied by the pilot. We can see in the plot above the periodic signature of the needle interrupting the laser, and also a sinusoidal behavior between needle crossings. Unintentionally, the sinusoid gives a useful cue as to exactly when a high-flow demand (stove, oven, heater) is initiated, so that one does not have to wait until the needle event for the first indication. The sinusoidal pattern is presumably from scattered light off of irregularities in the needle hub.

Even without the extra benefit from the hub, one can determine how much gas was used by noting how many cycles appear between the languid pilot spikes, compared to how many revolutions the pilot would have produced in the same period.


Each half-cubic-foot of gas is equivalent to 510 Btu, since 100 cubic feet delivers 1.02 Therms, and a Therm is 100,000 Btu. Therefore, each turn amounts to 538 kJ, or 0.149 kWh.

Stove-top Boil

Enough with the preliminaries. I lured you in with promises of efficiency measurements, and went off instead into some mad scientist tangent.

If I place 500 ml of water in a copper-bottomed 4-quart (approx. 4 ℓ) pot, and set this pot—without lid—on top of the largest burner and let-er rip, what efficiency do I get? The flames are

mostly hitting the bottom of the pot, but maybe a smidgeon of curling around the sides. Any guesses?



Two languid pilot light needle sweeps frame this series. Shortly after the first pilot spike, the gas stove was turned on, producing two quick spikes and a sped-up periodic signature in the baseline. It is 65.3 minutes between pilot spikes, during which time the needle made three turns. Ordinarily, it would take 67 minutes between pilot spikes (one turn). We can therefore attribute 2.03 turns to the gas stove, amounting to 1033 Btu, or 1.09 MJ. The 1.95 minutes between stove spikes translates to a burn rate of 15,700 Btu/hr, about 460 of which are from the pilot light.

Starting from an ambient temperature of 18°C (water at equilibrium), and heating to a boil, the gas meter reveals 1.09 MJ of energy was used on top of the steady burn from the hot water heater pilot light, at a rate of 4500 W (about 15,000 Btu/hr). Meanwhile, it was supposed to take only 171.5 kJ of energy to boil this amount of water. So the efficiency was a mere 16%.

I was surprised. Are you surprised? I mean, burning the fuel is itself highly efficient. The pot sits right atop the flame. Dirt simple. Some heat escapes around the side, some goes to heating the pot itself, but come-on—16%?!

In another test on another night, I boiled 500 ml of water starting at 16.5°C on a smaller burner (1670 W, or 5,700 Btu/hr). I used the same pot, but this time with its lid. I put the burner on full power, but its smaller diameter meant that no flames visibly licked around the bottom of the pot. Want to guess as to the efficiency in this mode? I deem this mode to be about as efficient as I am likely to get on the stove-top using standard practices. The answer: I used 640 kJ of gas for something that should have taken 175 kJ for an efficiency of 27%. An enclosed kettle on the same burner delivered the same result.

Okay, so a 68% improvement over the full-throttle approach is a significant gain (although it took almost twice as long). But still, I was surprised that I top out at less than 30% on a gas stove burner.

Incidentally, from the rate of cooling after the burner was turned off, I gather that the water

remaining in the open pot lost heat at a rate of about 300 W through convection, radiation, and possibly some conduction to the grate. These mechanisms are roughly proportional to ΔT between the water and the environment, so on average come to about 150 W during the linear ramp-up of temperature from ambient to boiling. This would have cost 63 kJ in the slow case (about 10% of the total), or about 35 kJ in the full-blast case (3% of the total). But the lidded case would actually lose heat less quickly than these numbers indicate, as the loss rate is based on an *open* pot after the boiling test was complete. In either case, loss from the pot to the environment appears to play a relatively minor role.

How About Microwaves?

Okay, the stove top delivered disappointing efficiency results. But right above the stove, I have a microwave. Naively, I expected the microwave to be something like 80% efficient at delivering energy to the water. The thinking goes like this: the interior of the microwave oven is a good reflector for microwaves, so they rattle around until they find an absorber; namely, the food/water. Of course there will be some loss in generation at the magnetron tube (I imagine it gets hot). And the walls may not be perfect reflectors, so that after twenty bounces, there may not be much energy left. But still, 80%—right?

This time, I put 500 ml of 18°C water into a plastic measuring cup (63 g) and heated on full power for 90 seconds, measuring the temperature afterwards. The water reached 50°C, demanding 67 kJ of energy. Meanwhile, both a Kill-A-Watt and TED (The Energy Detective) told me that I used 157 kJ of utility energy. That's 43%, folks. Yes, I'm disappointed too. A later test at 120 seconds produced 60°C water (this time from 22°C) for a computed 40% efficiency.

Considering that electricity often comes from fossil fuels at 30–40% efficiency, the microwave (as tested by me) is only 15% efficient at transforming fossil fuel heat into heated water. That's worse than the stove top—especially the slow-and-steady version.

Electric Kettles

What about electric kettles, where a heating element is immersed into the water, directly heating the thing you *care* about? To boot, these kettles often have insulated sides, keeping the heat where you want it.

I borrowed a kettle from work, and heated 500 ml of water from 18°C. I was able to suspend a thermometer in the water in a way that did not compromise the lid or spout. After 152 s, the water was making plenty of happy bubbling noise and the thermometer read 100°C—although apparently it was not as well-calibrated as I would have expected, since it sailed up to 105°C after 173 s. The water was obviously in a vigorous boil by then, and I shut the kettle off at 210 s (I should have let it turn off on its own, but lost patience with its inefficient lethargy).

Both the Kill-A-Watt and TED agreed that the kettle consumed about 1440 W. So while I needed $0.5 \times 4184 \times 82 = 171.5$ kJ to raise the water to boiling from 18°C, I spent 219 kJ by the time the (erroneous) thermometer read 100°C; 249 kJ by the time a full, roiling boil was obvious; and 302 kJ by the time I shut the device off. Using the middle number, I calculate 69% efficiency. I get 78% if I use the shortest time, when I *thought* I was at boil. But by the time I shut it off, we were down to 57%, and falling. Who knows how far it would have sunk before shutting itself off (kicking myself that I don't *know*).

If the shutoff is sufficiently sensitive, it seems electric kettles are capable of something like 70% efficiency or better. It would seem to beat the pants off the other methods—except for two caveats.

The first is the one raised for the microwave. If your electricity comes from fossil fuels, 70% efficiency becomes 25% in turning fossil fuel energy into hot water. The stove top is competitive.

The second caveat is that kettles suffer a sometimes rather large inefficiency due to lax filling practices. A kettle is considered to be a reservoir. Few measure the amount of water they put in. As long as it's *enough*, game on. So the tendency is to overfill. All the water gets heated. Only some is used. This practice is probably even more pervasive in communally-shared kettles, and could easily cut a deep hole into the net efficiency. By contrast, microwave practices typically heat exactly as much as is consumed. A kettle with 50% over-fill turning off on its own sweet time can easily sink to the levels achieved by the microwave.

Hot Water Heater

While we're talking about the efficiency of heating water, how well does my gas-fired hot water heater perform in transferring combustion energy into the water tank? Seems like it should be pretty good: few places for the heat to go except into the water (although the flue does get super-hot).

Practically speaking, efficiency suffers from slow heat loss during idle times, and from loss through the pipe walls during delivery. Tankless applications can avoid these two loss mechanisms. But let's separate these out and ask only about the direct efficiency with which combustion energy gets into the water in the first place. To do this, I poured hot water into a bathtub until the water heater came on. I let the heater complete its business, then waited about an hour, after which I drew another batch of hot water until the heater activated again. This way, the second heating did not have to compensate for long-term loss of heat from the reservoir.

Mind you, this experiment represented unusually excessive hot water use in our domicile, but I did it for the people. And we did manage to enjoy baths in the bargain.

I measured the temperature of the water emerging from the (hot only) tap, and from a tap right at the inlet from the street (adjacent to water heater, too). I used the water meter by the street to gauge volumetric use to a precision of 0.01 cubic feet (0.3 ℓ). And, of course, the gas meter told me how much energy was consumed.

On the first wave, I used 1.39 ft³, or 39 l, heated from 22°C to 53°C (remind me to turn the heat down; I should also note that the water came *out* at 50°C on the first draw, but it had cooled in the tank for some unknown time). Multiplying these fine numbers by 4184 J/l/K, I compute an energy demand of 5.1 MJ. Meanwhile, my gas gauge made 17.15 revolutions at 510 Btu/rev for a total of 9.2 MJ. The efficiency computes to 55%.

On the second withdrawal—this time based on a recently heated tank, it took 51 l before the heater engaged, after which the gas dial made 19.15 turns. This time, the demand was 6.6 MJ, while the gas cranked out 10.3 MJ. The efficiency is 64%.

Again, I find myself disappointed. I need to re-adjust my intuition about how straightforward it is to channel heat from a flame into water on the other side of a metal wall.

Bonus Round: Gas Oven

I seriously doubt I'll get around to another post detailing the things I learned from my laser-gauge gas meter. So I'll stick in here what I learned about my gas oven. How much energy does it take to "charge" up? How much power to keep steady? How much "on" time equivalent does it take for the pre-heat phase?

I heated my oven to 425°F (218°C) from an ambient 20°C temperature. It took 9.5 minutes to arrive at the setpoint, during which time the half-cf dial on the gas meter made five turns. Correcting for the water heater pilot light rate, this action took 2480 Btu, or 2.6 MJ (0.72 kWh). The burner operated at about 4800 W (16,300 Btu/hr). Thereafter, I found that it took a steady 1500 W to maintain temperature: I left the oven unopened and undisturbed for the better part of an hour to make sure I reached equilibrium—all for the sake of experiment. This means that the preheat phase uses the same amount of energy as 30 minutes of steady operation. You don't need fancy gauges to tell you this if you note the preheat time is 10 minutes and observe that the burner is on one third of the time during steady operation.

If you're heating a pizza in the oven that requires ten minutes of cook time, then only 25% of the total energy is spent in cooking mode, the other 75% in preheat. If cooking something for an hour, the preheat surcharge drops to 15%.

Although I am embarrassed to reveal the efficiency of cooking pizza in the oven (since this is not unknown in my household), I owe it to myself to carry out the calculation. Let's say I heat a 383 g pizza by 200°C (but that it's heat capacity is in between that of water (at the high end) and more typical materials—say 2000 J/kg/K. So I need to inject $0.383 \cdot 2000 \cdot 200 = 153$ kJ. Meanwhile, my oven heats up for ten minutes and then the pizza spends ten minutes in the oven. I count 2.6 MJ for preheat, and an additional 0.9 MJ for the cook time. In the end, I manage to get 4.4% of the expended energy into the pizza. If I instead tried 6 minutes of full-power equivalent in the microwave oven (at 1750 W), I might get near 25% efficiency—and a flaccid crust.

Summary Table

Lots of numbers thrown around in this post. Here's a table of my results.

Activity	Efficiency	Adjusted Efficiency
Boiling water on gas stove, full blast, no lid	16%	—
Boiling water in same pot, smaller burner, with lid	27%	—
Boiling water in kettle on small gas burner	27%	—
Heating water in microwave oven	43%	15%
Boiling water in electric kettle	50–80%	18–28%
Hot water heater, including tank loss	55%	—
Hot water heater, without tank loss	64%	—

The adjusted efficiency is the fossil fuel equivalent if electricity is derived from fossil resources (coal, natural gas) at 35% efficiency. The range on the electric kettle depends on how quickly the kettle shuts off in response to boiling water.

Testing Without Lasers

If you wanted to replicate or extend these experiments, is it hopeless without the laser-sensor gauge I created? Not at all—although slightly less convenient. In fact, I carried out the hot water heater test after I had already dismantled the gauge. Here's one trick. Once you've characterized the burn rate of various devices (e.g., stove burners on max flame; oven while burner is on; water heater; furnace; etc.) then all you need is a way to measure time when the (audible) burner is on.

As alluded to above, keeping track of the fraction of time the oven burner is on in steady operation is enough to tell you how much energy goes into preheat vs. holding temperature. Even without calibrating a stove burner, different configurations (lid, different pots, etc.) can be compared against each other just by timing.

To measure the rate of gas usage of various devices, make sure that device is the only thing on, and time how long it takes the half-cubic-foot (or 2 cf) dial to make one revolution. Together with knowledge that each half-cf translates to 510 Btu (538 kJ), you're pretty much set. Having a way to measure volumes and temperatures also came in handy for me.

What of It?

Heating water is less efficient than I originally thought. All the same, it will always take 1 kcal to heat 1 kg of water 1°C, and heating water is something we will always be interested in doing. Yet, efficiencies are middling-enough that we can't expect gigantic improvements. Heat pumps could break the 100% efficiency barrier (by a factor of several), but these are impractical for small-scale applications.

Should you react to the numbers above by rushing out to buy an electric kettle, with the promise of tripling the efficiency over the stove-top solution? In part, this depends on your source of electricity. If your electricity is derived from fossil fuels, and you otherwise have a gas stove, then it's probably not worth it. But even if going from an electric stove to an electric kettle, we must consider the embodied energy of the kettle. I discussed two methods for estimating embodied energy in another post, which for this case results in something like 50 kWh of investment. Each cup (250 ml) of water takes 0.03 kWh of energy at 70% efficiency. If you're tripling efficiency, then you save 0.06 kWh per cup, and need to boil over 800 cups before the thing pays for itself, energetically. Maybe this makes sense. But it's not a burning imperative.

As with many things, a far more effective strategy is to first recognize how and why you use sources of energy. After developing an awareness, you are far less likely to heat excess water in a kettle that you don't plan to consume. Shorter, less frequent showers can have a far bigger effect on your energy use than how you heat water for tea. It comes back to behaviors.

In the meantime, it's kind-of nice to have some numbers for water heating efficiencies—as disappointing as the numbers themselves are.



This work is licensed under a [Creative Commons Attribution-Share Alike 3.0 United States License](https://creativecommons.org/licenses/by-sa/3.0/).