

Update on the "Practical Treatise"

Lee Sauder, Winter 2010-2011

The experiment described above was intended to demonstrate some of the ways our working practice has evolved in the years since we wrote the HMS paper.

In general, our work has tended to go forward by going backward, utilizing more traditional furnace materials and construction. We did all our early work in a rather large furnace built of steel and modern refractory, with a water-cooled tuyere, and often with preheated blast. I think this may have led some to discount the applicability of our ideas to ancient smelting practice. But I still stand by all those ideas, and any changes and recantations that may follow are driven by my greater understanding of the working qualities of the bloom iron, rather than the use of more traditional materials.

Our current work tends to use a smaller clay furnace, with a copper tuyere, and occasionally using bellows for air supply. This is not just the result of curiosity or an urge for authenticity, but practicality as well. Clay has the tremendous advantages of being inexpensive, easily alterable, and almost endlessly recyclable and reusable. Compared to modern refractories, clay's combination of insulative, reflective, and emissive qualities also seem much more suited to the bloomery furnace. All of which I suppose should have been self-evident, since the technology evolved in clay furnaces! But we're all a bit susceptible to the hubris of modern man's technical prowess, aren't we? Most importantly, the use of clay has allowed us to easily and cheaply explore many different furnace designs.

The size of the furnaces have shrunk, with the usual goal of producing a bloom of 8 to 12 kilos. This simply reflects a retreat from the usual Yankee madness of "bigger is better". An 8-12 kilo bloom is just more pleasant to work on than a 20 kilo one.

The early furnaces were designed to be dismantled to remove the bloom, or to have the bloom removed through the top. Once we began using clay, though, we quickly converted to removing the bloom through the bottom of the furnace, by digging away enough below the tap arch to allow the bloom to be pried free.

Our use of the copper tuyere was first inspired by descriptions of its use in the Catalan furnace. We have stuck with it as it is extremely durable, convenient, and safe (the water-cooled tuyere caused our worst injury to date, knock on wood). The key to its durability is its conductivity. It simply needs to be long enough that the cooling effect of the air blast can keep the temperature of the tip from rising to an orange heat.

Though most of my own smelts continue to use a vacuum blower for an air source, I have successfully used hand-powered bellows on occasion. All of Shel and Steve's smelts at Colonial Williamsburg have relied on bellows, and all their successful smelts have been conducted along the lines of the experiments described above.

In the discussion section of our '02 HMS paper, we argued for a different view on three aspects of bloom smelting: 1) air rate, 2) slag management, and 3) the recycling of furnace products.

In the early days of our bloom smelting, our goals tended to focus on the yield of the smelt. In the years since, our focus has shifted towards the quality of the resulting metal. This is a much more delicate, difficult, and complicated realm, and that has altered my view a bit on some of those issues.

1) Air rate: In the '02 paper, Skip and I proposed an air rate of 1.2 to 1.5 liters/min/cm² of furnace area. This may be excessively quantitative and specific, but I still maintain that if you're not getting the results you desire, blow harder.

Over the years Skip and I have helped a number of frustrated people get their first good bloom, and find that people commonly fail for one or more of three reasons: 1) they're not blowing hard and hot enough, 2) they're using bad ore, or 3) they're stopping the smelt too soon (i.e. not charging enough total ore).

Another aspect of air rate we addressed in '02, that was perhaps overlooked by being poorly expressed, is that low air rates tend to produce a high carbon product, and high air rates tend to produce low carbon product. Many people still think of this the other way around, but my experience has shown this time and again. There are, of course, many other factors that affect carbon content, which I'll address briefly below.

Since reasonably accurate measurement of air rate is a slippery business, we realized some time ago that it's more sensible to measure the rate of charcoal consumption, since this is really a much easier, direct and more accurate measure of the air entering the furnace. As noted above, we usually do this as a time needed to burn a 2 kg charcoal charge (or actually, at home, a 4 lb. charge). As I dig back in my notes from several very different furnace types, the optimum rate of charcoal consumption works out to a range of from .3 to .4 grams of fuel/minute/ cm² of furnace area. I am struck by the fact that the optimum rates we've come up with in each furnace, for high quality blooms, seem to all work out very closely to .35g/min/cm².

2) Slag as a physical, chemical, and thermal resource: Ah, now we're getting down to the nitty-gritty! I still stand by our assertions in the '02 paper, but we have learned so much more here.

I see slag chemistry as the key to the working qualities of the metal, both directly and indirectly. Directly, in that if you end up with an iron-depleted slag in your bloom, you have a metal that welds poorly and fractures easily, and is in general, just difficult to work. Indirectly, in that the chemistry of the slag bath during all stages of the smelt is the key to the bloom's carbon content.

In short, I don't think it's too bald a statement to say that low-iron slag = bad metal, and high-iron slag = good metal.

We judge the iron content of the slag by observing its behavior upon tapping, and by examining its fracture. We strive for a free-running slag that on freezing, piles up on itself nicely. On fracture, we want to see a grey metallic sheen we refer to as "wustitey", and a certain shape of fracture surface I don't believe I could describe meaningfully.

I often see comments in the literature equating very high iron content in archaeological slags with inefficiency on the smelter's part. Rather, from my point of view, a very high iron content in the slag would point to mastery of the process.

3) Recycling gromps: Ok, time to recant. I have found that the recycling of gromps has a deleterious effect on metal quality, so I have largely abandoned the practice.

In several smelts I have made entire blooms largely of gromp, to make a second (or even third) bloom in a hot furnace. These gromps were charged, or even placed directly in the hearth, with fresh ore to make a nice iron-rich slag bath. Though these blooms were dense and lovely to look at, they invariably have resulted in a nasty and recalcitrant metal.

Though there may be a good method of reprocessing this metal, simply adding it to the ore charge of a bloomery smelt is probably not it.

On Quality Control:

I have a long way to go to be able to claim any mastery here, but I may as well take this opportunity to share some things we've found so far. It seems to me that the working quality of my blooms is largely determined by the control of three (interrelated) factors: slag chemistry, carbon content, and phosphorus content.

Slag chemistry: I've covered this pretty well above, but in addition to the iron content of the slag, there are also things we definitely want to avoid getting into the slag.

Too much alumina, from melting furnace wall, can cause a slag that makes for poor welding. To avoid this, we try to build our furnaces with thin walls to shed enough heat to keep the clay from melting and slagging much, especially trying to keep it thin in the hot zone around the tuyere.

The addition of too much calcium, added in an attempt either to increase yield or to control phosphorus, also seems to cause welding problems. I think I perceive trouble from this even if the slag viscosity during the smelt seemed fine, and the carbon content of the bloom has stayed low.

Carbon control: Some of the more important factors we've identified as affecting carbon content are air rate, slag chemistry, fuel:ore ratio, tuyere angle, type of ore.

Air rate and slag chemistry we've covered above.

The effect of fuel:ore ratio is pretty straightforward and widely understood. Our general charging practice stays pretty close to a 1:1 ratio, and we'll push it as heavy on the ore as we can, especially during the latter stages of the smelt. We're rarely able to get much higher than 4 fuel to 5 ore without knocking down the heat in the furnace.

Steeper tuyere angles promote higher carbon contents. Tuyere angle has a lot to do with the shape of the bloom, and the depth at which it forms, both of which affect carbon content. The best angle varies with each furnace, but generally seems to hover in the area of 15 to 22°.

My experience with ores beyond my usual limonite are limited. But it seems that very rich ores really want to make steel. I also suspect ore of a very regular size, such as magnetite sand ore hematite sandblasting grit, tends to make steel. I think a variety of sizes in the charge helps ensure adequate slag formation, since the different sizes travel through the stack at different rates, some reducing in the stack more than others.

Phosphorus control: It took Skip and I a long time to realize what a significant factor the phosphorus content of our ore is. And it's taken us even longer to start to get some degree of control over it by any procedural means beyond ore selection. Since we have basically no access to any modern, professional analysis, we have been judging our phosphorus content by working properties and a bit of hillbilly metallography. We are still very much in the midst of working through this, but here's what we think so far. Much of this work has been through collaborative effort with our "smeltfest" crew.

Once again, the air rate/temperature/ tempo of the smelt seems to have a large effect. Too slow a tempo produces a higher carbon, or worse, a carbon/phosphorus combination that doesn't want to work up easily. Too high an air rate/tempo seems to raise the P content of the metal enough to cause cold-shortness. So a medium tempo seems best. Not adagio, not presto, but andante.

In terms of manipulating P content through slag chemistry, addition of calcium seems to be a poor choice, for the reasons mentioned above. Addition of manganese ore to the mix shows more promise in this regard. Don't ask me why, but manganese seems to make a pretty radical improvement in metal quality.

Another thing I'd like to note is that from a smithing perspective, phosphoric iron is a really good thing. It makes for an iron that is dense and tough but forges beautifully. Some of the highest phosphorus blooms forge so beautifully it's hard to believe there could be anything wrong with them when they're cold!

The Goldilocks principle: All the foregoing points to a consistent theme. The key to making quality iron seems to be a matter of finding balance. Not too hot, not too cold. Not too fast, not too slow. Not too much, not too little. Just Right.

A related observation is that success in smelting is built of careful attention to small incremental steps. It's easy for two or three small overlooked details to compound into a big problem. So it pays to attend to every aspect of the job as carefully and thoroughly as possible.

This requires plenty of time. Work up your ore, charcoal, and repair clay in advance, make any tools you need, and generally give yourself the time and space to pay attention.

This really emphasizes the importance of a long, sustained working practice applied to any attempt to understand early ironmaking technologies. Short and scattered experimental campaigns, with a different furnace and different ore every time, can cause more confusion than enlightenment, if they don't find and explore that sweet spot every furnace has.

Forge it now!: One extremely simple (and therefore easily overlooked) key to producing good iron is to forge it immediately, fresh and hot from the bloomery.

No matter how long you try to reheat a cold bloom in a hearth, it will never be as nice and hot through and through as it was when it came out of the furnace. And all that long slow time that you try to get a thorough heat on a cold raw bloom, you are altering its slag chemistry. So the metal you get from starting with a cold bloom will not be the same as metal starting from a hot one.

Even a simple compaction and splitting of the bloom really helps reduce the time and fuel needed to reheat.

New questions and directions:

Well, there's enough of these to last several lifetimes, aren't there?

Part of what I've found so fascinating about working with bloom iron over the years is its unpredictability and irregularity. Lots of the sculpture I've made from it has been about exploring its failures and funkiness. But as I come to understand it better, I'm becoming most excited about the possibilities of really good bloom iron. I can see the potential to produce iron that is superior to currently available industrial steels for artistic forging.

A few of the directions that seem particularly fruitful at the moment:

Phosphorus: We've still got lots to do here, not just in getting some control over the metal's phosphorus content, but how best to exploit P iron's characteristics: its forgeability, its hardness, its corrosion resistance.

Manganese: As I mentioned above, addition of manganese ore to our charge has a pretty dramatic effect on the resulting iron. So much so, that this is the only thing I've ever considered keeping as a trade secret. But I'm just not that kind of guy.

I assume this is a matter of slag chemistry only, but I have no idea what's really going on here.

Tapping regime: Immediately after the conference, I had the opportunity to help begin an excavation of a medieval bloomery site that Jake Keen located. Most of the bloomery slag I've seen in my life was made by me. Here I had a week to ponder the slag produced by someone with real mastery of the craft. I was struck by its density, the size of the slag flows, and the incredible consistency of it from one piece to the next.

Though for years I have been trying to ensure a constant flow of high iron slag across the bloom, this medieval slag seemed to be telling me that I was being timid in this regard. On returning home, I tried to smelt to reproduce the slag I saw by tapping heavily and constantly, and continuing with fresh ore charges rather than recycling slag.. This smelt produced one of the most remarkable blooms yet, in terms of iron quality. As of this writing, I only have been able to do one smelt along these lines, but you can be sure there's more to follow.

Remelting: To date, we have mostly refined our bloom iron by means of folding and welding. I am currently very curious about the possibility of refining the iron by remelting, not to produce steel but to make soft iron. This is described explicitly by Evenstad (Evenstad,1968), and more obliquely by others. Are we overlooking an entire step in the ironmaking process?

I guess it's time to quit writing and go find out.

Gerry McDonnell, Roger Doonan, and Hector Cole. And very special thanks to my brother Jake Keen for being Jake Keen.

References:

Evenstad, Ole, *Journal of the Historical Metallurgy Society*,25(,1968).

