

A practical treatise on the smelting and smithing of bloomery iron

Lee Sauder and Skip Williams

ABSTRACT: For several years, we have explored many aspects of the process of bloomery smelting in a shaft furnace. In contrast to most attempted reconstructions of bloomery smelting, our work has focused on the process itself, rather than its archaeological signatures. This paper describes a typical smelt of the most efficient regimen we have yet discovered. We will pay particular attention to methods that differ from those of most experimenters, especially with regard to blowing rate, slag management, and the recycling of furnace products. A bloomsmithing experiment is also described, and yields, resources, and labour requirements are quantified. We then offer a few observations where our experience differs from what we have read in the literature. Finally, we suggest that these methods, when applied to archaeological reconstructions of ancient bloomeries, may provide some missing answers for the archaeometallurgical community.

Introduction

We have been experimenting with the bloomery process since January 1998. From the beginning, our primary goal has been to smelt iron of sufficient quantity and quality for the creation of hand-forged artworks, and to explore the process for a deeper understanding of iron as an artistic medium. We strive to remain open to what the iron itself has to teach us, and to keep scientific presumptions in the background. Our interest and expertise is in iron and ironworking, not in archaeology or metallurgy. We feel that this devotion to the process and its product, rather than to furnace morphology or slag residue, has led us to uncover an approach to bloomery smelting which has the potential to provide more accurate data for historical and archaeological research than the current predominant models.

Our first eleven trials provided us with valuable experience, but produced only the most pitiable examples of blooms. These early blooms, besides being fist-sized at best, all had elevated carbon contents that made most of them unforgeable. We attempted to deal with these problems both by reducing the fuel:ore ratio

and lowering the airflow and temperature, with disappointing results. These early attempts were also hampered by particularly irreducible ores.

We built our second furnace on a modular system (Fig. 1) This allows us to explore many different furnace configurations by varying shaft heights and tuyere heights. Our first truly satisfactory bloom resulted from an attempt to make cast iron by increasing shaft height, fuel:ore ratio and, perhaps most significantly, air flow. From this serendipitous beginning, we have evolved a very efficient smelting regimen based on minimal furnace preheating, airflows from 1200–1600 litres/minute, the recharging of tapped slag, and the recycling of residue from the previous smelt.

Experiments 21–27 have all been run in a very similar manner. This paper will describe smelt 25 as typical of this series. This experiment has the benefit of especially good notes from the smelt, and the bloom is preserved and sectioned as a specimen. We then describe a smithing experiment with a similar bloom from smelt 26, a portion of which was forged to a billet and then to small ‘currency’ bars using manpower for forging and charcoal as fuel.

Smelting

Raw materials

The ore in this trial was our local 'brown ore', a dense goethite of 58% iron content, gleaned from the abandoned Longdale mines in Alleghany County, Virginia. It was roasted in a gas flame, and then broken up until the pieces ranged from 20mm to fines.

The charcoal in this experiment was commercially obtained, composed largely of oak and hickory. We broke it in a fairly cursory manner so that most pieces ranged from 30 to 80mm, and sifted out the majority of the fines. In earlier experiments using lower airflows, we found that the fines tended to clog the furnace. This is not a concern at higher blast rates, but sifting out the fines helps to keep spark damage to both the local environment and the furnace's operators to a minimum.

Furnace construction

The furnace is designed in modules to allow the investigation of many furnace types. The shaft height above the tuyere can vary from 120mm to 1.6m to recreate anything from a bowl hearth to a high bloomery or stuckofen. The diameter of 350mm was chosen intuitively, as a manageable size that would not require an onerous amount of raw material to feed it, but would provide more working room and larger blooms than our previous 300mm furnace. This furnace is similar in size, construction, and concept to those utilized so successfully by Tylecote (Tylecote *et al* 1971) and

Tholander (Tholander 1987).

The furnace as configured in this experiment had a shaft height above the tuyere of 1.0m, with the tuyere 230mm above the furnace floor. This configuration is roughly analogous to the dimensions of a Roman shaft furnace (Tylecote *et al* 1971, 343). The slag tapping arch, measuring 120x150mm, was located at the base of the furnace, opposite the tuyere. There was a single probe hole for measuring temperature 400mm above the tuyere and 100mm in from the furnace wall.

The furnace is constructed of an outer layer of steel sheet, lined with 50mm of insulating board, and then with a 70–80mm thickness of castable refractory as the furnace lining. In this experiment, the final 300mm of shaft height was a hollow steel section which functioned as an air preheater. The preheater provided 100–200°C of preheat. The only major effect of this preheating is the conservation of charcoal. The preheater does raise the maximum temperature we can achieve, but this capacity is rarely used. Although our earlier smelts utilized ceramic tuyeres, we used a very simple water-cooled tuyere in this experiment.

In general, the details of furnace construction were chosen for durability, practicality and economy. The use of preheated air and a durable water-cooled tuyere (innovations of the 19th century) have no significant impact on the process other than improved economy in fuel consumption and furnace maintenance. These minor anachronisms of furnace construction have been deliberately chosen to allow us to concentrate on the smelting process and its variables in a series of repeatable experiments, eliminating the variables of changing furnace shapes, sizes, clay compositions etc. Thus our observations can apply to shaft furnaces of any era or locale.

Blowing apparatus

Our earlier trials used squirrel-cage fans for an air supply. This type of fan does not overcome back pressure very well, so we found it difficult to control the blast as the smelt progressed. The back pressure in the furnace increases as the reduction reactions escalate. The air source we are now using is a vacuum fan, which generates better pressure, and delivers approximately 1600 l/min at full blast through the 50mm tuyere orifice. Airflow was regulated by a simple valve which dumps excess air, rather than constricting its flow.

Air flow was measured and calculated using a Kestrel

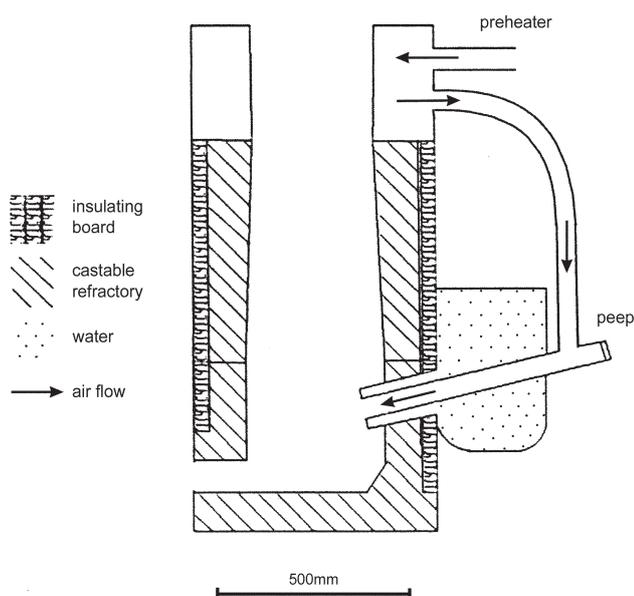


Figure 1: Cross-section of the modular smelting furnace

1000 anemometer. We measured the air speed 10mm from the tuyere orifice of the (obviously cold and empty) furnace, and multiplied this speed by the tuyere area to arrive at a volume measurement. The air flow values thus arrived at, for various valve settings, are reported below. This may not be a very accurate measurement, but serves as a reasonable approximation, which has been roughly corroborated with more accurate air flow measurements, using an air flow meter, at lower rates of 200-300 l/min.

Airflow at these rates can be supplied by the vigorous application of carefully counter-weighted blacksmith's bellows, or more easily by a pair of bellows.

Description of smelt 25

This experiment took a total of 5 hours and 30 minutes from lighting the fire until the removal of the bloom. We think of the smelt as breaking down into four general phases: preheating, charging of ore, recycling and recharging, decarburization and burndown.

Preheating

We kindled a fire, and preheated the furnace with wood strips, utilizing natural draught through the open tap arch. After 30 minutes, we loosely blocked the tap arch, added charcoal, and began a blast of 1275 l/min. Preheating with charcoal continued for another hour, consuming 18.5kg of charcoal.

Charging of ore

After 1¹/₂ hrs of preheating, we added our first charge consisting of 6.8kg of charcoal followed by 6.8kg of ore. At this time the temperature at the probe hole had reached 850°C. Each of the succeeding charges were identical in charcoal and ore weight to the first. Each charge was added as there was room to do so in the top of the furnace, at about 20-minute intervals. By the time of the second charge, the temperature had risen to 980°C. The exhaust gases from the furnace did not ignite until just before the third charge. By the time of the third charge, the temperature at the probe hole had fallen to 930°C. The blast was then increased to 1500 l/min.

The rationale for air rate and temperature changes needs to be explained. We don't often bother to measure the temperature at the probe hole, but when we do, we usually expect to record 1000° to 1050°C. This measurement is not the major factor in our decision-making. We most often base our temperature and airflow decisions on three other major factors.

The first factor is our view through the 'glory hole', the sight glass that allows us to judge the temperature by looking through the tuyere into the bloomery's hot zone. We look for a very brilliant white, greatly in excess of forge-welding temperature. At the time of this experiment, we had no instrumentation to measure this temperature, but in subsequent trials it has ranged from 1450° to 1650°C.

The second factor in our temperature decisions is the condition of the slag bath. We can monitor this by careful probing through the tap arch. The first slag to work its way to the base of the furnace tends to be a frothy, high-iron slag that quickly cools to a bowl-shaped insulating layer on the furnace bottom. The pool of slag that fills this bowl is a high-iron, free running 'tap slag' in which the bloom forms. We want this pool to stay very hot and liquid, and to gradually melt into and enlarge the bowl of solidified slag.

The third decisive factor is intuition. Some of our decisions might look like whim to an observer, but we have come to know what the furnace looks, sounds, smells and feels like when it's making iron.

Shortly after the third charge, the blast was increased to 1625 l/min, and then reduced back to 1500 l/min at charge four. Shortly before the fifth charge, we saw a few sparks of burning iron at the tuyere, and reduced the blast again to 1275 l/min. The blast remained between 1275 and 1500 l/min for the remainder of the smelt.

Recycling and recharging

The sixth charge consisted of 5.8kg of 'gromps' and magnetic slag recovered from the previous smelt, along with 1kg of ore to make up a 6.8kg charge. By 'gromps' we mean bits of unconsolidated iron and magnetic material that either failed to adhere to the bloom in the furnace, or were removed from the bloom during its initial cleaning.

The next cycle of the smelt is the tapping and recharging of slag. First we scraped as much of the early semi-solid sponge slag from under the incipient bloom as possible, cooled it, broke it up, and recharged it into the top of the furnace with an equal weight of charcoal. This sponge slag contains reduced iron that never had a chance to adhere to the bloom, as it was the first material through the furnace. Thereafter, liquid slag is tapped from the furnace by poking through the solidified slag at the tap arch. In order to keep the incipient bloom covered, and to maintain the heat reservoir of the slag, we try not to tap more than 4-5kg of slag at one time.

We cool the tapped slag in water, break it up, and return it to the furnace with an equal amount of charcoal.

We tend to continue this recharging sequence until we detect a change in the quality of the slag that indicates a lessening iron content, such as increasing viscosity or a decrease in the rapidity of freezing. Sooner or later, the slag becomes too difficult to tap, and we proceed to the burndown phase. In this smelt, a total of 14.5kg of slag was re-charged with an equal amount of charcoal, followed by a further 9kg of charcoal.

Decarburization and burndown

Finally, we added charge of 5.5kg of ore and 4.5kg of charcoal. This final charge of ore has a pronounced decarburizing influence on the bloom. In our usual practice, another 5–10kg of charcoal would be added for burndown, but this time we did not do so. A little over an hour later, 5¹/₂ hours after kindling the fire, the furnace burden had burned down enough to allow us to disassemble the furnace and remove the bloom.

In this case, the upper sections of the furnace were removed as the furnace burden burned down. To remove

the bloom, the lower section was tipped on to its side, and the bloom knocked out from below. In more recent smelts, we have removed the bloom through an enlarged tap arch. In either case, by the time we remove the bloom, the slag above it is still molten and the slag below the bloom has solidified, leaving a 'plano-convex' slag mass, which is usually destroyed during the process of bloom removal.

The bloom

Immediately after the bloom's removal from the furnace, we hammered all loosely consolidated material off it. The bloom thus cleaned was roughly 250mm in diameter, 120mm thick, and weighed 14kg. We later sawed the bloom into approximate halves along the vertical plane axial to the tuyere. As can be seen in Figure 2, the bloom is very dense, with little 'spongy' character except at the periphery. Blooms such as this one do not appear to be simply an agglomeration of particles that have fallen from above. Rather it appears that in the oldest section of the bloom, in its centre, the interstices of the sponge iron have been filled by iron particles reducing in situ. The carbon content, judged

Table 1: Data for smelt 25

Time (hh:mm)	Charcoal (kg)	Ore (kg)	Slag* (kg)	Gromp (kg)	Air (l/min)	Temperature ^{oo} (° C)	Operational notes
0:00	-	-	-	-	-	-	start preheat with wood
0:30	22.7	-	-	-	1275	-	switch to charcoal, blast on
1:15	-	-	-	-	-	850	
1:28	6.8	6.8	-	-	-	-	charge 1
1:48	6.8	6.8	-	-	-	980	charge 2
2:02	-	-	-	-	1500	935	first liquid slag in front of tuyere
2:08	6.8	6.8	-	-	-	-	charge 3
2:28	6.8	6.8	-	-	1625	930	charge 4
2:40	6.8	6.8	-	-	1275	-	charge 5, furnace now hot enough
2:58	6.8	1.0	-	5.8	-	1045	gromp charge
3:15	5.0	-	5.0	-	-	1025	recharge 1
3:35	4.5	-	4.5	-	-	-	recharge 2
3:49	4.5	-	5.0	-	-	-	recharge 3
3:54	4.5	-	-	-	-	-	
4:10	4.5	-	-	-	-	-	
4:14	4.5	5.5	-	-	-	-	decarburizing charge and burn down
4:45	-	-	-	-	-	-	remove preheat section
5:10	-	-	-	-	-	-	remove shaft section, consolidate visible bits onto bloom
5:25	-	-	-	-	0	-	remove tuyere, blower off
5:30	-	-	-	-	-	-	remove bloom
Totals	91.0	40.5	14.5	5.8			

Notes: Furnace configuration: 1m shaft height above the the tuyere, air preheating on, tuyere 230mm above the floor. Fuel type: hardwood charcoal. Ore type: goethite (Victoria mine). * slag recharged. ^{oo} temperature at point 400mm above tuyere, 100mm in.

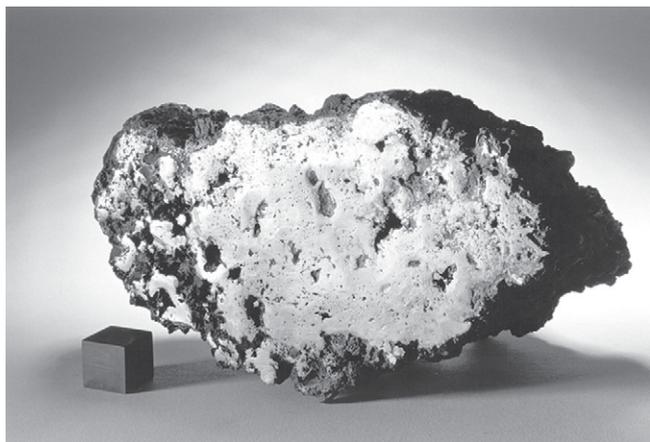


Figure 2: Sectioned bloom with 22mm cube for scale

by hardness and spark test, indicates that this bloom is of a low carbon content, similar to that of modern mild steel. This bloom has a density of roughly 4500kg/m³. The 14kg bloom represents a yield of 60% of the iron available in the ore, and required 94kg of charcoal and 5 hours 30 minutes to produce (Table 2).

Bloom smithing

In order to quantify the useable iron in a bloom of this type, as well as to supply a comparison with Peter Crew's work with a non-slag tapping furnace, we forged a similar small Llyn Cerrig Bach currency bar. Unlike in Crew's experiment, all heating was done with charcoal and all forging by manpower alone (Crew 1991).

From bloom to billet

The bloom forged in this experiment was the 13.5kg bloom produced from smelt 26, which was very similar to the bloom from smelt 25 described above. The entire bloom was first heated in a gas furnace, to replicate the hot bloom removed from the bloomery. It was easily split in two in this single heat, using a splitting maul as a hot cut, driven by sledges. We selected the larger of the two sections, weighing 7.7kg, to forge into a 50 x 50mm billet similar to Roman billets (Cleere and Crossley 1985, 48; Sim 1998, 55).

The half-bloom was heated for forging in a modern cast-iron bottom-blast blacksmith's hearth with several bricks

Table 2: Yield from smelt 25

Ore	Input		Bloom	
	Available iron	Weight	Yield	
40.5kg	23.5kg	14.0kg	60%	

stacked around it to increase the depth of the hearth. This approach to heating with charcoal was found to be fairly unsatisfactory, as we were unable to bring the entire piece to a welding heat all at once. We also briefly attempted to convert this forge configuration to a side blast, with poor success. All forging at this stage was done as near to a welding heat as possible. The lead smith wielded a 2kg hand hammer, and the two strikers used 4.5 and 5.5kg sledges. With a bloom of this type, there is no need for gentle compression before welding and forging. Despite our heating difficulties, in 29 heats we managed to forge the half-bloom to a 50x50x300mm billet in 3hr 40m (11 man/hours labour). The finished billet weighed 5kg. This represents 65% of the starting weight of the bloom section, and a yield at this stage of 36% of the iron available from the ore.

Although we use the above numbers in the time and yield analyses below, we feel that this procedure could be greatly improved, and thus made more accurate. Our lack of experience in forging with charcoal slowed the entire process, and our difficulty in attaining a welding heat left cracks that had to be rewelded in subsequent smithing. Also, modern strikers are trained more for precision than power, since modern smiths have power hammers for heavy work. A striking team trained in heavy forging, such as the anvil or anchor smiths of more recent centuries, would be vastly more efficient. Given a more efficient furnace design and a better trained set of hammermen, it is conceivable that these billets could be forged in much less than an hour.

From billet to bar

We forged two currency bars of similar dimensions from portions of this billet. The first trial was a continuation of the morning's billet smithing, using the same forge set-up and personnel. A roughly cubical section was isolated from the billet by fullering, and then hammered to a flat bar of approximately 75mm wide by 20mm thick. A major transverse crack appeared where the billet had been most drastically deformed during the fullering. We welded this crack as best we could, and then folded the flat section in two, and faggot-welded the entire length. We then cut this flat section from the billet and continued to draw out the bar, welding up any cracks as they appeared. We made the mistake, when faggot-welding, of forging the bar to a fairly thin flat cross-section, which slowed the drawing out process considerably. A more efficient forging technique would have been to forge it to a square cross-section that would then be forged to the correct flat stock in the final heats. The latter portion of the work was carried out with a single striker. The

Table 3: Smelting and smithing yields for smelt 26.

Input		Bloom		Billet		Currency bars	
Ore	Available iron	Weight	Yield	Weight	Yield	Weight	Yield
40.9kg	23.7kg	13.5kg	57%	8.55kg	36%	4.95kg	21%

final result, after 2hr 30m (representing 6.5 man/hours of work), was a socketed currency bar 25x60x500mm, weighing 623g. Forging from bloom to billet to this final bar required 73kg of charcoal.

The following afternoon, Sauder forged another currency bar, utilizing the lessons of the previous day's work. The forge was re-configured as a shallow side-blast forge by filling the firepot with refractory insulation. A more traditional approach to this would have been to create an insulating base of packed cinders (Fitzgibbon 1990). We set our tuyere from the bloomery on the forge table, and stacked up firebricks around it to create a hearth. This hearth set-up functioned much more satisfactorily than either of the previous day's attempts. The side tuyere was not troubled by slag blocking and choking by charcoal dust as the bottom blast had been. Also, by having a shallow, well insulated base, the bottom of the hearth reflected heat to the fire. Somewhat counter to expectation, the hottest zone of the fire was not in front of the tuyere but below it, halfway between the tuyere and this reflective base.

A 30mm chunk of the 50x50mm billet, weighing 538g, was hot-cut completely from the bar. All forging was done with a single 2kg hand hammer. This smaller chunk was easy to bring to a welding heat. Four heats were used to weld from all three directions, bringing the billet chunk to a cube. After a total of 17 heats in 1h 35m, we had a short bar 16mm square and 210mm long. A further hour of forging and 16 heats produced a socketed currency bar 4.5mm thick, 22.5mm wide and 375mm long, weighing 312g. This represents 58% of the starting billet weight, and a yield at this stage of 21% of the iron available from the ore. Forging this bar from the billet consumed 20.5kg of charcoal. Total labour required for this bar was 2h 35m.

This bar was smaller than we had intended. If we had started with a 50mm length of billet, it is reasonable to assume that a 500g bar could have been forged in three hours.

Spark testing of both currency bars indicated a low overall carbon content. One small zone in the large bar was of a medium carbon content, similar to spring steel,

but the rest of both bars were of a low carbon content, between modern mild steel and wrought iron.

Analysis

In the following analyses, we have not included the gromp charge in our yield calculations, as at least as much gromp is produced as is consumed in each smelt. The billet and currency bar weights are extrapolated as if the entire bloom had been worked to these forms. If the entire bloom had been worked to small bars, it would have produced 10 bars of 500g each.

Yield analysis

Table 3 shows the smelting and smithing yields for smelt 26.

Fuel consumption

Table 4 demonstrates the consumption of charcoal in smelting and smithing. As with our estimate of yields, we used our results based on smithing one section of the bloom from smelt 26 to estimate the fuel needed to smith the entire bloom .

Labour analysis

Table 5 shows our labour analysis for the whole series of operations from mining to forging of currency bars, based on smelt 25. It should be noted that 9.76 man-hours/kg of the above figure represent the smithing process, which our experiments replicated inefficiently, so some further reduction of the total figure is likely.

Table 4: Fuel consumption.

Process	Charcoal used (kg)
Smelting	94
Billet smithing	78
Bar smithing	238
Total fuel	410
Fuel consumption/kg of bar	82

Table 5: Labour analysis

Process	Time (man-hours)
Mine and roast ore, cut wood, make charcoal (assumed as per Crew 1991)	50
Smelting—assuming a three-man team (as per Crew 1991), Smelt 25: 5hr 30min duration	16.5
Billet smithing—extrapolated to full bloom	18.8
Currency bar smithing—extrapolated to full bloom @ 3 man-hours per 0.5kg bar	30
Total time to smelt and forge 10 bars, each of 0.5kg	115.3
Total labour time/kg	23.06

Discussion

Our smelting method differs from that of other experimenters in three aspects, (or at least three that can be easily communicated by written language). These aspects are: the air rate, the management and manipulation of slag, and the recycling of furnace products.

Air rate (the myth of the overblown bloomery)

Early archaeological experimenters in the bloomery process used air rates in the neighbourhood of 0.4 l/min/cm² of hearth cross section (Tylecote *et al* 1971, 362). This air rate seems to have been arrived at due to fairly theoretical criteria (Tylecote *et al* 1971, 348).

It is understandable that later experimenters stayed within this range. Our earlier experiences with blasts of these lower rates indicates that as the blast approaches 0.6 l/min/cm², the carbon content of the bloom increases, and the slag near the bloom turns to a drab green low-iron slag. Others have noted this phenomenon (Crew 1991, 27; Harvey 1988, 36). Further increases above this blast rate produce copious incandescent sparks at the tuyere, indicating the re-oxidation by the blast of any iron which has reduced in the stack above, as well as burning of the incipient bloom, which adheres to the wall just below the tuyere. A furnace run on blasts of 0.4 to 0.8 l/min/cm² will resemble Figure 3(a).

If the blast is increased still further, in the neighbourhood of 1.2 to 1.5 l/min/cm², conditions in the furnace change drastically. The hot zone of the furnace enlarges to encompass most of the hearth's cross section. The burden will burn down much more evenly across the furnace, rather than in a narrow cone which funnels all material directly in front of the tuyere. Iron particles which have reduced in the stack do not have to pass directly in front of the tuyere on their way to the slag bath below, and those which do are protected by the more copious molten slag above the tuyere level. As

the hot zone is also expanded downwards, the bloom forms much lower in the furnace, and is thus much more easily protected by the molten slag bath. A furnace that is run on higher blasts will resemble Figure 3(b).

At these higher blast rates, the bloom does not adhere to the furnace wall, and so is easily removed from the furnace. In this experiment, due to the small slag-tapping arch, the furnace was disassembled to remove the bloom. In subsequent trials, with a tap arch almost as wide as the interior diameter of the furnace, the bloom was pulled out through the tap arch.

At yet higher blasts for prolonged periods, in the neighbourhood of 1.6 l/min/cm², carbon content again seems to elevate, and incandescent sparks indicate re-oxidation.

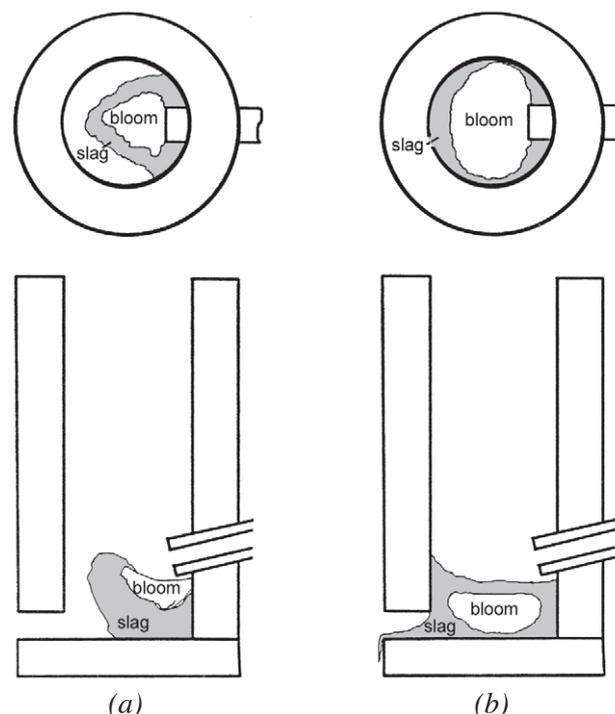


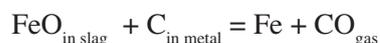
Figure 3: Bloomery run with an air blast of: (a) 0.4–0.8 l/min/cm² and (b) 1.2–1.5 l/min/cm².

Slag as a physical, chemical and thermal resource

Slag fulfills two physical functions in a furnace: protection and transportation. Molten slag coats and protects reduced iron particles from re-oxidation. After the bloom begins to form, we also try to keep it physically covered by molten slag at all times, to protect it from re-oxidation. Slag flow also serves to transport reduced iron particles to the locale of bloom formation. The re-charging of the first slag to reach the bottom of the furnace utilizes the transportational function of the slag. Iron particles that did not have a chance to coalesce into a bloom are thus carried back to the active zone of the hearth.

Both protection and transportation require a liquid, free-running slag. The fluidity of the slag is a function both of its chemistry and its temperature. Higher temperatures facilitate slag flow through all parts of the hearth.

High-iron slag also serves two chemical functions: reduction and decarburization. These two functions are often simultaneous: wustite in the slag is reduced by carbon in any iron with elevated carbon content, decarburizing the metal even as it produces more. This mechanism was perhaps described most clearly and succinctly by Espelund (1997, 54) as



Note that the product of this reaction is not only more iron but also more reducing agent. We think the lovely chain reaction thus initiated is the real workhorse of bloom formation, and that reduction within the stack merely provides a seed for reduction in the slag bath below. Reduction of wustite by direct contact with bits of charcoal that survive to hearth level also contributes to bloom formation. Admittedly, we are unable to quantify or directly observe these reactions; remember that these are artisans' working hypotheses.

Low temperature, small slag baths, and low fuel:ore ratios only serve to inhibit these hearth-level reactions. A small low carbon bloom, composed of loosely accumulated stack-reduced particles, is like an ungerminated seed. This type of bloom accounts for the difficulty reported by many researchers (*eg* Crew, Sim) of consolidating the bloom without breaking it apart. Immersing this ungerminated seed into the fertile environment of a hot and active slag bath produces a very dense bloom that is in no way fragile, and may be hammered vigorously from the start.

Finally, the slag performs vital thermal functions. The growing slag bath, as well as the incipient bloom itself, provides both a reservoir of heat and a source of radiant energy that keeps the temperature of the furnace from falling with the addition of each fresh charge of ore. This heat reservoir, along with the exothermic nature of the reduction reactions taking place, provides the not-so-gradual increase in furnace temperature in the latter stages of the smelt. This is another reason for restraint in the tapping of slag, which removes heat from the furnace.

The hot, fluid slag also tends to carry heat down to the lower part of the hearth, allowing the hot zone, and the bloom itself, to sink lower in the furnace as the smelt progresses. The hot bloom also tends to melt its way down towards the bottom, leaving room for more bloom formation above (we think we have yet to approach the limits of charge weight or bloom size in this furnace).

The tapping and recharging of slag ensures a constant flow of this physical, chemical, and thermal resource through and around the growing bloom. The use of these slag manipulation techniques is not necessarily limited to slag-tapping furnaces, however. Many of the same goals could be accomplished through the recycling of slags from previous smelts. Also, any pit furnace with a bottom of combustible material, like the grass used by the Haya of Tanzania, and in Scandinavian slag pit furnaces (Schmidt 1997; Mikkelsen 1997), could provide a slow subsidence that ensures a constant flow of fresh slag across the bloom.

The recycling of gromps

We use here the term 'gromps' to refer to the leftover bits of mixed slag and reduced iron that litter the furnace site after a smelt. Like most experimenters, we initially refrained from the recycling of iron and slag from previous smelts, in hopes of making each smelt an isolated, measurable event. But ancient smiths would rarely have carried out any smelt as an isolated event, and would surely have recycled all material possible. An accurate reconstruction of an ancient bloomery process should therefore use recycled material.

In our yield analysis above, we have not included recycled material as an input in our calculations. Theoretically, each smelt will produce as much gromp as was put into it. In reality, each smelt produces more gromps than it consumes, as a glance at the ground around our furnace proves.

So why use it if there is no net gain from its use? Because the use of gromps is a variable that significantly alters the process. When we use gromps in our last charge, we get blooms of greater density. We can't posit a mechanism for this effect, only report our empirical observations. This recycling charge seems almost to have a catalytic effect. We perceived no impact on bloom density when we used the gromps as the initial charge, only when added last, before the slag tapping and recharging phase of the smelt.

A few miscellaneous observations

Preheating: The conventional practice of attempting to bring the furnace to full operating temperature before charging is simply a waste of charcoal. The heat reservoir of molten slag, and the exothermic reduction reactions, are what really bring the furnace to full heat.

Control of carbon content: The production of high-carbon steel in a bloomery is often presented as a special accomplishment of a particular technology. In our experience it is difficult not to produce high carbon steel in a bloomery, and the challenge is, rather, to create a soft, low carbon iron. We find the most important technique for the control of carbon content is ensuring a constant flow of high-iron slag across the bloom.

The addition of a late charge of fresh ore also has a pronounced decarburizing effect on the bloom. We stumbled on this technique by accident, but have since found that Evenstad prescribed this final charge (Espelund 1997, 54). This procedure could apply to any furnace type.

Conclusions and suggestions

As can be seen from the analyses above, the three inter-related techniques of increased air flow, constant monitoring and manipulation of the slag bath, and the recycling of furnace materials, lead to very different results in terms of yields and labour requirements than those published by most other experimenters. Most significantly, the labour requirements to produce a kilogram of fully smithed iron drop from the 25 man-days suggested by Peter Crew to a mere 23 man-hours in our experiment, a difference of a full order of magnitude.

These improved yields and efficiencies are a result of the constant improvement of working practice. We have smelted and forged over half a ton of blooms in this furnace, with an average bloom weight of over 13kg.

Having experienced this much success, as well as suffering many failures, we feel certain that the techniques and conditions we have described here are essential to the creation of high quality bloom iron.

We think that the techniques we have described can, and should be, applied to any form of bloomery. We also think that the application of these methods can lead to revision of the conclusions of earlier experiments, and that the resulting iron more closely matches the archaeological record. After all, if you are not reproducing the product, you are certainly not reproducing the process.

We offer this information in hopes that those with more scientific and archaeological training, not to mention institutional support, may be able to apply what we have learned to experimental archaeology.

Acknowledgements

Many thanks are due to: Martha Goodway, Pati Wattenmaker, and Elizabeth Sauder for encouragement and advice, Wayne Raynal for his thorough note-taking, Frank Settle and Doug Metcalf for ore analysis, Yates Spencer for mighty wielding of the sledge, and Washington & Lee University and Virginia Military Institute for various assistance and access to resources. This paper is dedicated to Dr William C Sauder, a great scientist and a great father.

References

- Cleere H and Crossley D 1985, *The Iron Industry of the Weald* (Leicester).
- Crew P 1991, 'The Experimental Production of Bar Iron', *Historical Metallurgy* 25, 21–36.
- Espelund A 1997, 'The Evenstad Process — Description, Excavation, Experiment and Metallurgical Evaluation' in L Norbach (ed), *Early Iron Production — Archaeology, Technology and Experiments* (Lejre), 47–58.
- Fitzgibbon M 1990, 'Designing and Fabricating a Firebox', *The Anvil's Ring* 18, 40–42.
- Harvey D 1988, 'Reconstructing the American Bloomery Process', *The Colonial Williamsburg Historic Trades Annual*, 19–37.
- Mikkelsen P 1997, 'Straw in Slag-Pit Furnaces' in L Norbach (ed), *Early Iron Production — Archaeology, Technology and Experiments* (Lejre), 63–66.
- Sim D 1998, *Beyond The Bloom- Bloom refining and iron artifact production in the Roman world* (BAR IS 725: Oxford).
- Schmidt P 1997, *Iron Technology in East Africa: Symbolism, Science and Archaeology* (Bloomington).
- Tholander E 1987, *Experimental Studies on Early Iron Making* (Stockholm).
- Tylecote R F, Austin J N and Wraith A E 1971, 'The Mechanism of the Bloomery Process in Shaft Furnaces', *Journal of the Iron and Steel Institute*, 209, 342–364.

Authors

Lee Sauder is a professional artist-blacksmith and sculptor. He began his apprenticeship in 1973 at the tender age of 12, and has supported himself solely by his art since 1981.

Address: 229B McLaughlin Street, Lexington VA 24450 USA. e-mail: thesauders@rockbridge.net

Skip Williams is an electrical engineer and computer systems specialist employed by Washington & Lee University. He maintains his perspective through the study of older, saner technologies.

Address: University Computing, Washington & Lee University, Lexington VA 24450 USA.

e-mail: swilliams@wlu.edu