





Primary Metal Production

Pyrometallurgy deals with chemical reactions at high temperatures (ranging from 100°C up to 3000°C). These reactions involve numerous different solids, liquids, and gases, and are carried out using many diverse types of furnaces.





General Principles of Production of Metals by Pyrometallurgy

In pyrometallurgy, metals are extracted by converting sulfides into oxides and then reducing oxides into metals, using carbon or carbon monoxide as reducing agents. The reactions for lead, zinc, and iron are given below. The reader is cautioned that the following chemical reactions are overly simplified versions of the actual processes. In these reactions, CaCO_3 is a flux, and its purpose is discussed further in the following sections.




Lead:



Zinc:



Iron:




Heat of Reaction (Enthalpy)


The heat of reaction (ΔH), or enthalpy, determines the energy cost of the process. If the reaction is exothermic (ΔH is negative), then heat is given off by the reaction, and the process will be partially self-heating. If the reaction is endothermic (ΔH is positive), then the reaction absorbs heat, which will have to be supplied to the process.

Equilibrium

Most of the reactions used in pyrometallurgy are reversible, and so they will reach an equilibrium where the desired products are converting back into the reactants as quickly as the reactants are forming the products:



We would prefer the reaction:



and so the process will need to remove C and/or D as quickly as they are produced, so that the products cannot react to re-form A and B. A gaseous product can be removed by venting it off, and other types of product can be removed by dissolving them in slags of an appropriate composition.

Molten metals tend to dissolve impurities from the ore. For example, many ores that contain native copper have arsenic compounds associated with them. When the ore is melted, the molten copper dissolves the arsenic. As a result, the metal from the furnace can be less pure than the individual metal grains that were originally in the ground.

Gibbs Free Energy

The Gibbs free energy (ΔG) of a reaction is a measure of the thermodynamic driving force that makes a reaction occur. A negative value for ΔG indicates that a reaction can proceed spontaneously without external inputs, while a positive value indicates that it will not. The equation for Gibbs free energy is:


$$\Delta G = \Delta H - T\Delta S$$

where ΔH is the enthalpy change in the reaction, T is absolute temperature, and ΔS is the entropy change in the reaction.

The enthalpy change (ΔH) is a measure of the actual energy that is liberated when the reaction occurs (the “heat of reaction”). If it is negative, then the reaction gives off energy, while if it is positive the reaction requires energy.

The entropy change (ΔS) is a measure of the change in the possibilities for disorder in the products compared to the reactants. For example, if a solid (an ordered state) reacts with a liquid (a somewhat less ordered state) to form a gas (a highly disordered state), there is normally a large positive change in the entropy for the reaction.

Ellingham Diagrams




Construction of an Ellingham Diagram

An Ellingham diagram is a plot of ΔG versus temperature. Since ΔH and ΔS are essentially constant with temperature unless a phase change occurs, the free energy versus temperature plot can be drawn as a series of straight lines, where ΔS is the slope and ΔH is the y-intercept. The slope of the line changes when any of the materials involved melt or vaporize.


Free energy of formation is negative for most metal oxides, and so the diagram is drawn with $\Delta G=0$ at the top of the diagram, and the values of ΔG shown are all negative numbers. Temperatures where either the metal or oxide melt or vaporize are marked on the diagram.

The Ellingham diagram shown is for metals reacting to form oxides (similar diagrams can also be drawn for metals reacting with sulfur, chlorine, etc., but the oxide form of the diagram is most common). The oxygen partial pressure is taken as 1 atmosphere, and all of the reactions are normalized to consume one mole of O_2 .

The majority of the lines slope upwards, because both the metal and the oxide are present as condensed phases (solid or liquid). The reactions are therefore reacting a gas with a condensed phase to make another condensed phase, which reduces the entropy. A notable exception to this is the oxidation of solid carbon. The line for the reaction



is a solid reacting with a mole of gas to produce a mole of gas, and so there is little change in entropy and the line is nearly horizontal. For the reaction



we have a solid reacting with a gas to produce two moles of gas, and so there is a substantial increase in entropy and the line slopes rather sharply downward. Similar behavior can be seen in parts of the lines for lead and lithium, both of which have oxides that boil at slightly lower temperatures than the metal does.

There are three main uses of the Ellingham diagram:

1. Determine the relative ease of reducing a given metallic oxide to metal;
2. Determine the partial pressure of oxygen that is in equilibrium with a metal oxide at a given temperature; and
3. Determine the ratio of carbon monoxide to carbon dioxide that will be able to reduce the oxide to metal at a given temperature.

Ease of Reduction

The position of the line for a given reaction on the Ellingham diagram shows the stability of the oxide as a function of temperature. Reactions closer to the top of the diagram are the most "noble" metals (for example, gold and platinum), and their oxides are unstable and easily reduced. As we move down toward the bottom of the diagram, the metals become progressively more reactive and their oxides become harder to reduce.

A given metal can reduce the oxides of all other metals whose lines lie above theirs on the diagram. For example, the $2\text{Mg} + \text{O}_2 \rightleftharpoons 2\text{MgO}$ line lies below the $\text{Ti} + \text{O}_2 \rightleftharpoons \text{TiO}_2$ line, and so metallic magnesium can reduce titanium oxide to metallic titanium.

Since the $2\text{C} + \text{O}_2 \rightleftharpoons 2\text{CO}$ line is downward-sloping, it cuts across the lines for many of the other metals. This makes carbon unusually useful as a reducing agent, because as soon as the carbon oxidation line goes below a metal oxidation line, the carbon can then reduce the metal oxide to metal. So, for example, solid carbon can reduce chromium oxide once the temperature exceeds approximately 1225°C , and can even reduce highly-stable compounds like silicon dioxide and titanium dioxide at temperatures above about 1620°C and 1650°C , respectively. For less stable oxides, carbon monoxide is often an adequate reducing agent.

Equilibrium Partial Pressure of Oxygen

The scale on the right side of the diagram labelled " Po_2 " is used to determine what partial pressure of oxygen will be in equilibrium with the metal and metal oxide at a given temperature. The significance of this is that, if the oxygen partial pressure is higher than the equilibrium value, the metal will be oxidized, and if it is lower than the equilibrium value then the oxide will be reduced.

To use this scale, you will need a straightedge. First, find the temperature you are interested in, and find the point where the oxidation line of interest crosses that temperature. Then, line up the straightedge with both that point, and with the point labelled "0" that is marked with short radiating lines (upper left corner of the diagram). Now, with the straightedge running through these two points, read off the oxygen partial pressure (in atmospheres) where the straightedge crosses the " Po_2 " scale, and this is the equilibrium partial pressure.

It is possible to reach the equilibrium oxygen partial pressure by use of a hard vacuum, by purging with an inert gas to displace the oxygen, or by using a scavenger chemical to consume the oxygen.

Ratio of CO/CO₂ Needed for Reduction

When using carbon as a reducing agent, there will be a minimum ratio of CO to CO₂ that will be able to reduce a given oxide. The harder the oxide is to reduce, the greater the proportion of CO needed in the gases.

To determine the CO/CO₂ ratio to reduce a metal oxide at a particular temperature, use the same procedure as for determining the equilibrium pressure of oxygen, except line up the straightedge with the point marked "C" (center of the left side of the diagram), and read the ratio off of the scale marked "CO/CO₂".

Pyrometallurgical Equipment

Types of Furnaces

Shaft Furnace

These are vertical furnaces with the charge added at the top and removed at the bottom, while gas is blown into the bottom and exits the top, as shown in Figure 1. The solid charge must consist of particles coarse enough that they will not be blown out of the furnace by the gas. An iron ore blast furnace is a typical example of a shaft furnace.




Figure 1: Schematic of a shaft furnace, showing the material flows.

Muffle Furnace

This type of furnace is used when the material being heated should not be contaminated by the heating fuel. This is accomplished by enclosing the material in a chamber, with the fuel burned outside of the chamber, as shown in Figure 2

Hearth Furnace

Hearth furnaces allow the burning fuel to come in contact with the material being heated. This allows very high temperatures to be reached. This type of furnace includes reverberatory furnaces and rotary kilns (Figure 3).

Electric Furnace

Electric furnaces heat the charge by running a massive electrical current through it (the larger furnaces use approximately 20,000 amps at 50 to 500 volts). These work by immersing electrodes in the charge, as shown in Figure 4.




Figure 2: Schematic of a muffle furnace




Figure 3: Hearth furnaces. (A) Reverberatory furnace; (B) Rotary Kiln




Figure 4: Schematic of an electric furnace. The electrodes can be connected to DC, AC, or 3-phase AC, depending on the application.


Drying and Calcination

Drying


Drying is the removal of water from the ore, using a moderate amount of heat (temperatures on the order of 100°C). Only the mechanically bound water is removed (water filling pores and cracks, or that is adhering to the particle surfaces). Chemically-bound water, such as water of hydration in the ore crystal structure, is not removed by drying processes.

Calcination

Originally, calcination referred to the heating of limestone above 900°C to drive off the CO₂ and produce lime:




In current practice, calcination refers to any process where the material is heated to drive off volatile organics, CO₂, chemically bound water, or similar compounds. For example:




Roasting

Roasting involves not only heating, but also reaction with a gas. It is typically used to convert sulfides to oxides by reaction with air (air is usually used as an oxidizing agent, because it is free). For example:



Roasting can also use other gases, such as chlorine (to produce volatile chlorides):



If the reaction is exothermic enough (strongly negative ΔH), then autogenous roasting may be possible. This is where the heat for the roast is provided by the roasting reaction, and fuel is only needed to get the reaction started. For this to work, the furnace must be designed to capture and use the heat produced to bring fresh ore up to the roasting temperature. Not all ores will “burn” in this way, and so in many cases supplemental fuel will be needed to maintain the roasting temperature.

Basic Steps in Roasting:

- Particles are heated.
- Reactive gas (air, oxygen, chlorine, etc.) contacts the particles.
- Particles react with the gas.
- Gaseous reaction products are carried away.

Since the particles do not melt, the reaction starts on the particle surface and gradually works in to the particle core, as shown in the Shrinking-Core reaction model (Figure 5)




Figure 5: The Shrinking-Core model. As the shell of oxidized ore becomes thicker, it becomes more difficult for fresh gas to reach the unreacted ore, and so the roasting rate slows down. It is often difficult to react the last bit of material in the center of the particles.

Basic Roasting Terms:

- Dead Roast: the ore is completely reacted, and leaves the process cold.
- Sweet Roast: the ore is completely reacted, but leaves the furnace still hot.
- Sour Roast: the roasting reaction is not run to completion.

Types of Roasting Furnaces:

Hearth Roaster (Figure 6): This type of furnace is suitable for coarse particles, and is common in




Figure 6: Hearth Roaster. Typically, these have 6 to 12 hearths, and most of the roasting occurs as the particles drop from one hearth to the next.

older facilities. The capacity is on the order of 100-200 tons/day.

Flash or Suspension Roaster (Figure 7): These furnaces process very fine particles, and take advantage of the rapid reaction rate to provide autogenous heating. The capacity is approximately 3 to 4 times that of a similar-sized hearth roaster.

Fluidized Bed Roaster (Figure 8): Fluidized bed reactors are used not only for roasting, but also for drying and calcination. Gas is bubbled up through a bed of particles, with the particles large enough that they are not swept out of the furnace, but small enough that the gas can expand the bed so that it will behave as a fluid. These reactors provide excellent gas/solid contact, without requiring extremely fine particles.

Smelting

In smelting, the ore is brought to a high enough temperature that the material melts, and the final product is a molten metal and a slag. Smelting is done so that the impurities are either carried off in the slag, or are burned off as a gas. In some types of furnace (blast furnaces, flash smelters) the roasting and smelting operations are combined.

Blast Furnaces:

The function of a blast furnace is to produce pig iron, not steel. A blast furnace is a shaft furnace in which a blast of preheated air is blown in through tuyeres at the bottom. A schematic of a blast




Figure 7: Flash Smelting. The heat produced by the roasting reaction is used to dry the incoming ore and bring it up to roasting temperature. If the reaction is sufficiently exothermic, the ore can even be melted.

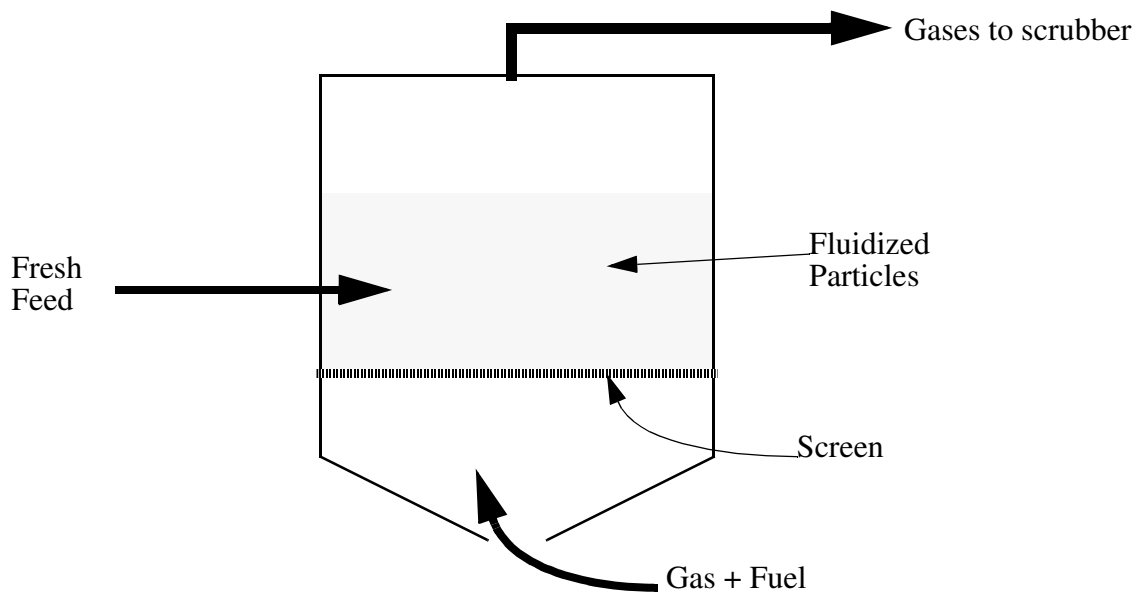


Figure 8: Fluidized Bed Reactor.

furnace used for making iron is shown in Figure 9. A description of the reactions that take place in the furnace is shown in Figure 10.

Components of the charge to a blast furnace:

- Iron Ore
- Coke
- Fluxes

Iron Ore:

Iron ore can be hematite (Fe_2O_3), magnetite (Fe_3O_4) or taconite (a colloquial term used for Minnesota ores which can refer to either hematite or magnetite). Iron ore is beneficiated to 65 - 72% Fe, and pelletized to form pellets 3/8 inch to 5/8 inch in diameter. In some operations, concentrated fines may also be sintered, that is, fused into porous lumps that are broken into one or two inch pieces. Pellets are durable and ship well, and sinter does not. Pelletizing plants are therefore often built near the mine and pellets are transported to the blast furnace by rail or ship, while sinter is usually produced at the steel mill.

Coke:

Coke serves three functions:

- Supply chemical reactants for reducing iron ore to metallic iron.
- Act as a source of carbon in the pig iron and eventually in the steel.
- Provide a source of heat in the blast furnace (fuel).

Fluxes:

When metal is smelted, the metal is separated from its impurities by melting, with the impurities forming a molten slag on top of the metal. Many of the impurities associated with iron ore are difficult to melt, and so they will not form a proper slag easily, which retards the smelting process. To make these impurities easier to melt, fluxes are added.

Limestone (CaCO_3) or dolomite ($(\text{Ca,Mg})\text{CO}_3$) are two typical fluxes used in blast furnaces. When a large amount of sulfur needs to be removed from the furnace charge, limestone is the preferred flux. Limestone is also a better flux to use if slag from the blast furnace is to be used as a raw material for cement manufacture. An important criterion for flux selection is availability and cost, and dolomite is often more readily available and less expensive than limestone.

Components of a blast furnace:

The components of a blast furnace are shown in Figure 9, and brief descriptions of the functions of each are as follows:

Stack. In the blast-furnace stack, there is a countercurrent flow of gas and solids. In the stack, hematite or magnetite is reduced to sponge iron by the action of carbon monoxide.

Bosh. This is the part of the furnace where the contents melt. As shown in Figure 10, several chemical reactions take place in the bosh. This is where the slag is formed.




Figure 9: The iron blast furnace. Heated air is pumped into the bustle pipe, and injected into the furnace through the tuyeres. The bell acts as an air-lock, so that the ore charge can be added to the furnace without letting all the gas escape.




Figure 10: Temperatures and reactions in the iron blast furnace

Notes:

1. Reactions 7, 8 - SiO_2 , Si, and CaS form the slag.
2. Reactions 9, 10 - some of the P and Mn go into pig iron, some goes into the slag.

Hearth. In the hearth, molten pig iron and molten slag segregate from each other and are tapped off separately. Slag is lower density than iron, and so it floats on the molten metal.

Blast Furnace Operation

The furnace is charged with iron ore lumps, pellets, and/or sinter; coke; and possibly extra flux. These are carried to the top of the furnace with skips or conveyors, and are tipped, or charged, into the furnace. Meanwhile, air preheated to 900°C is injected through the tuyeres, which are nozzles at the bottom of the furnace. The coke is partially burned by the injected air both to produce heat, and to generate carbon monoxide. Since coke is relatively expensive, some furnaces inject coal or oil along with the air as supplemental fuels to reduce coke usage. The carbon monoxide travels upward through the column, and removes oxygen from the iron ores on their way down, leaving metallic iron. By the time the charge reaches the base of the furnace, the heat generated there melts the iron. The resulting molten “hot metal” is tapped at regular intervals by opening the “tap hole” in the bottom of the furnace so that it can flow out. The fluxes combine with impurities in the coke and ore to form the slag, which floats on the iron and is removed through the “cinder notch”.

The hot metal from the furnace is collected in specially-constructed railway containers, called “torpedo cars”. The torpedo cars carry the molten iron to the steelmaking furnace.

Blast furnaces are operated continuously without shutdown for ten years or more. If the furnace were allowed to cool, thermal stresses can cause damage to the refractory bricks. Eventually, the refractory bricks in the furnace will wear away, and at that point the furnace is emptied and shut down so that it can be relined with new bricks. The period between shutdowns is referred to as a “campaign”.

Iron taken directly from the blast furnace contains about 4 - 4.5% carbon, as well as a number of other elements. This is referred to as “pig iron”, and if it is allowed to solidify it is brittle, difficult to work with, and has poor structural properties. The pig iron can be converted to steel by refining in the steelmaking process, which reduces the carbon content and removes other impurities, to make a stronger, tougher, and more generally useful product.

Steel Production

In the United States, we produce about 100 million tons of steel per year. Steel is a very valuable commodity.

Blast furnaces do not produce steel, they produce pig iron. The differences between pig iron and steel composition are shown in Table 1

Table 1: Composition differences between pig iron and steel

	Pig Iron	Steel
Carbon	3-4%	0.04-1.7%
Manganese	0.15-2.5%	0.3-0.9%
Silicon	0.5-4%	Trace
Sulfur	Up to 0.2%	0.02-0.04%
Phosphorus	0.025-2.5%	<0.04%

Steel is made in two different types of facilities:

- Integrated steel mills (blast furnace operations using iron ore as feed)
- Minimills (electric arc furnace operations using steel scrap as feed)

Integrated Steel Mills

In the integrated steel mills, molten pig iron is produced in a blast furnace and transported to a Basic Oxygen Furnace (BOF) or Basic Oxygen Process (BOP) vessel (Figure 11), or a related unit. In the oxygen steelmaking process, high purity oxygen is blown under pressure onto or over a bath containing hot metal as shown in Figure 11.

The basic steelmaking reactions with this type of furnace are:

- Carbon reacts with oxygen to form CO and CO₂, which escape as gases.
- Silicon, Manganese, and Phosphorus react with oxygen to form their oxides, which are removed in the slag.
- Sulfur reacts with calcium (from the flux added to the furnace) to form calcium sulfide, and is removed in the slag.

The progress of refining in a top-blown basic-oxygen furnace is shown in Figure 12. This is a considerable improvement over the old open-hearth furnace process, which took several hours to accomplish what the basic-oxygen furnace can do in a matter of minutes.

The molten steel is sampled at intervals to determine its carbon content and content of other elements. When the desired composition is reached, the vessel is rotated to pour off the molten steel.




Figure 11: Basic Oxygen Furnace cross section. The oxygen is blown into the furnace to burn off excess carbon, and to convert manganese, silicon, sulfur, and phosphorus into oxides that will segregate into the slag.





Figure 12: Progress of refining in a top-blown vessel

The properties of steel depend not only on its chemical composition, but also on the heat treatment. At high temperatures, iron and carbon in steel combine to form iron carbide (Fe_3C), which is commonly known as cementite



The forward reaction is endothermic, so that the formation of cementite is favored at high temperatures. When the steel is cooled slowly, the equilibrium shifts back and the cementite begins to decompose back into iron and small particles of graphite. Steel that contains graphite tends to be gray, while steel that contains cementite is lighter colored, harder, and more brittle.

Heat treatment of the steel controls the ratios of graphite to cementite, and also causes other crystallization effects that allow the mechanical properties of steel to be altered over a very wide range.

Secondary steelmaking:

In the integrated steel mills, an increasing proportion of the steels undergo what is called “secondary steelmaking” before they are cast. The objective of secondary steelmaking is to fine-tune the chemical composition of the steel, improve its homogeneity, and remove residual impurities. This is done by first tapping the molten steel from the furnace into a ladle, which is fitted with a lid to conserve heat. The ladle of steel can then be treated by any of several different processes, including argon stirring, addition of alloying elements, vacuum degassing, and powder injection. If necessary, an electric arc can be used to keep the molten steel at the proper temperature for casting.

Summary of Integrated Steelmaking

The first step in integrated steelmaking is mining the ore. The second step is preparing the raw materials. Ores must be beneficiated to increase the iron content to 65-72% Fe. If the ore is produced far away from the steel mills, then the concentrated ore is mixed with a binder and formed into pellets with diameters between 3/8 inch and 5/8 inch. These pellets are easy to handle and ship well for long distances. If the iron-bearing fines are produced very close to the steel mill, then they can be sintered, that is, fused into porous lumps that are broken into pieces 1 - 2 inches in diameter. Sinter does not ship well, and so a sintering plant is usually a part of the integrated steel mill, while pelletization plants are usually located at the mine.

Preparation of raw materials also includes limestone crushing to a convenient size, and production of coke from coal. Coke is made by heating coal in the absence of oxygen, which drives off the volatiles and leaves a hard, porous, high-carbon coke that is crushed to an appropriate size and charged to the furnace.

The third step in an integrated steel mill is ironmaking. The prepared ore, coke, and limestone are charged into the top of a blast furnace in the correct ratio, while a blast of hot air is injected at the bottom of the furnace. The heated air burns the coke to carbon monoxide, which rises through the downward-moving ore to reduce the iron oxides to metallic iron. When the iron reaches the bottom of the furnace, the temperature is sufficient to melt the iron, producing liquid pig iron.

In the fourth step, the liquid pig iron is transferred to a steelmaking furnace, usually a basic-oxygen furnace (BOF) or an electric-arc furnace (EAF). In this step, the metal is refined by removing unwanted elements into a slag, and adjusting the concentrations of desirable elements. In particular, the carbon content is brought down to between 1.7% and 0.3%.

In the final step, the molten “raw” steel is poured into a ladle, where it may undergo secondary steelmaking processes before being poured into molds where it solidifies into ingots or slabs of steel that can weigh 60 - 100 tons. A schematic of the overall process is shown in Figure 13.

Electric Arc Furnace Minimills

Small steel mills, which use electric arc furnaces to melt scrap steel, are referred to as “minimills”. A schematic of a minimill flowsheet is included in Figure 13. Electric arc furnaces consist of a tiltable vessel with a removable lid. The lid contains electrodes which are lowered into the furnace and contact the scrap steel charge. An electric current is passed through the electrodes to form an arc. The heat generated by the arc rapidly melts the scrap.

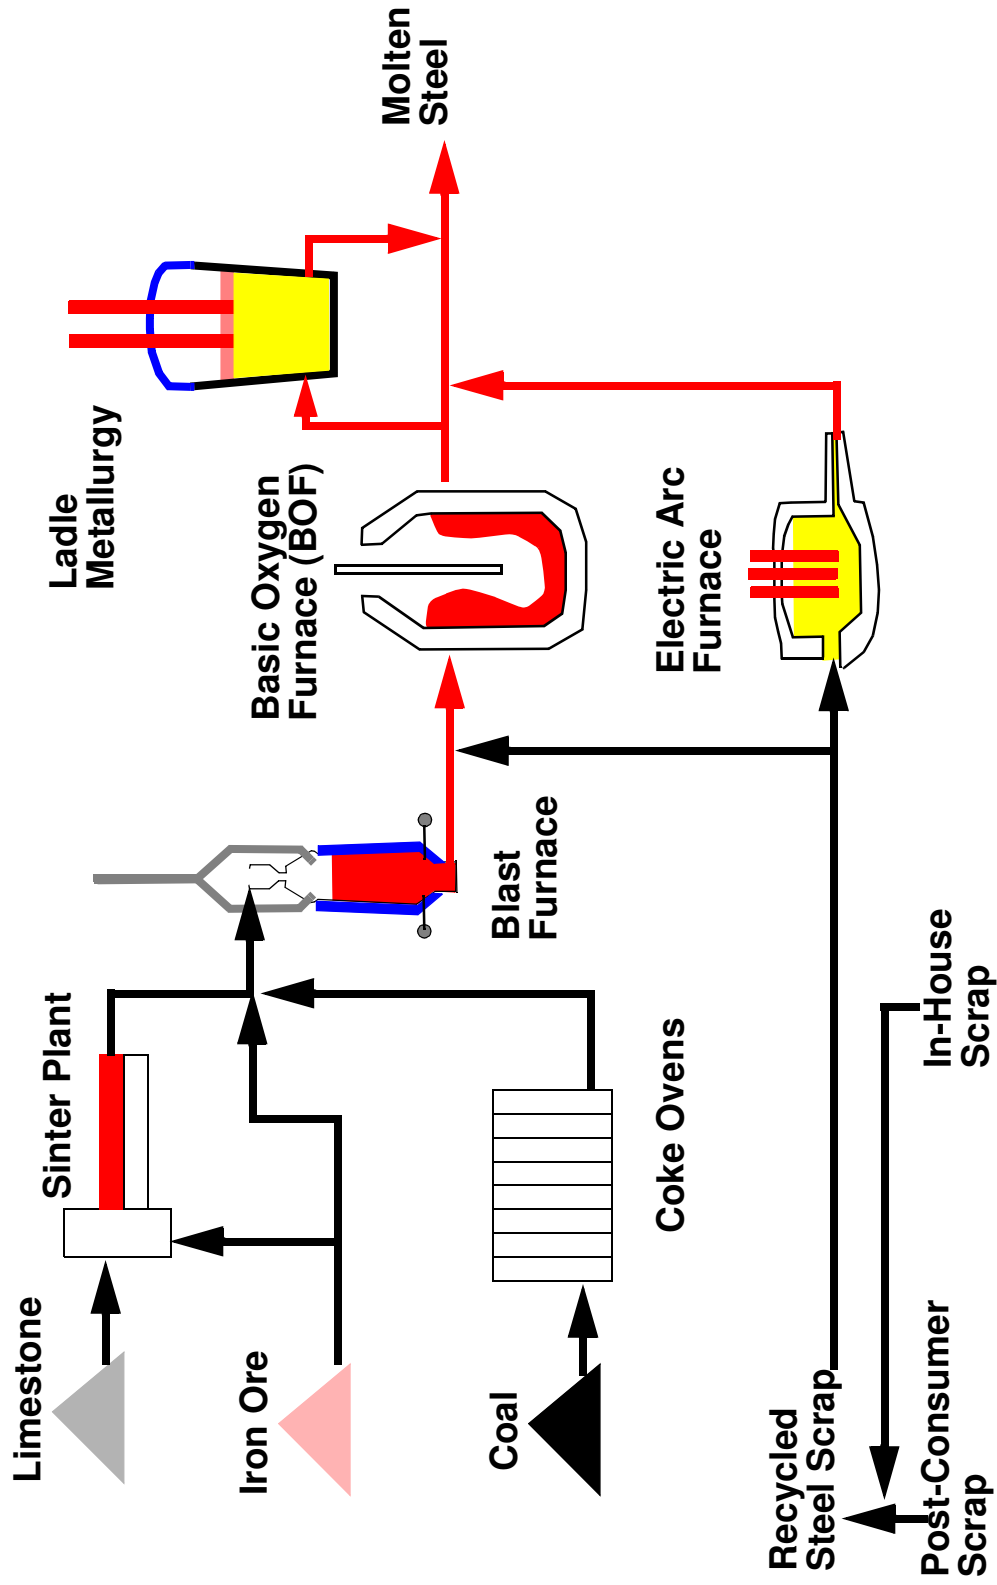


Figure 13: Steelmaking by both the integrated steel mill route, and by the electric-arc furnace minimill route

During the melting process, other metals are added to the steel to give it the desired chemical composition. As with the basic oxygen furnace, oxygen can be blown into the electric arc furnace to purify the steel. Fluxing agents such as limestone and fluorspar are added to combine with impurities and form a slag.

Once steel with satisfactory chemical composition is formed, the furnace is tilted to pour off the slag which floats on top of the molten steel. After the slag is removed, the furnace is tilted the other direction to pour the molten steel into a ladle. The steel is then transported either to secondary steelmaking, or cast into slabs or ingots.

A typical modern electric arc furnace melts 150 tons of steel in a charge, and a melting cycle takes approximately 90 minutes.

Primary Copper Production

The traditional method for producing metallic copper is to use high-temperature pyrometallurgical techniques to smelt copper sulfides (and, in areas like the UP, by smelting native metallic copper). In the last century, low-temperature hydrometallurgical techniques have been developed which have been very complimentary to the conventional smelting methods, as they allow the treatment of the low-grade copper oxide ores that often occur along with copper sulfide ores, as shown in Figure 14.

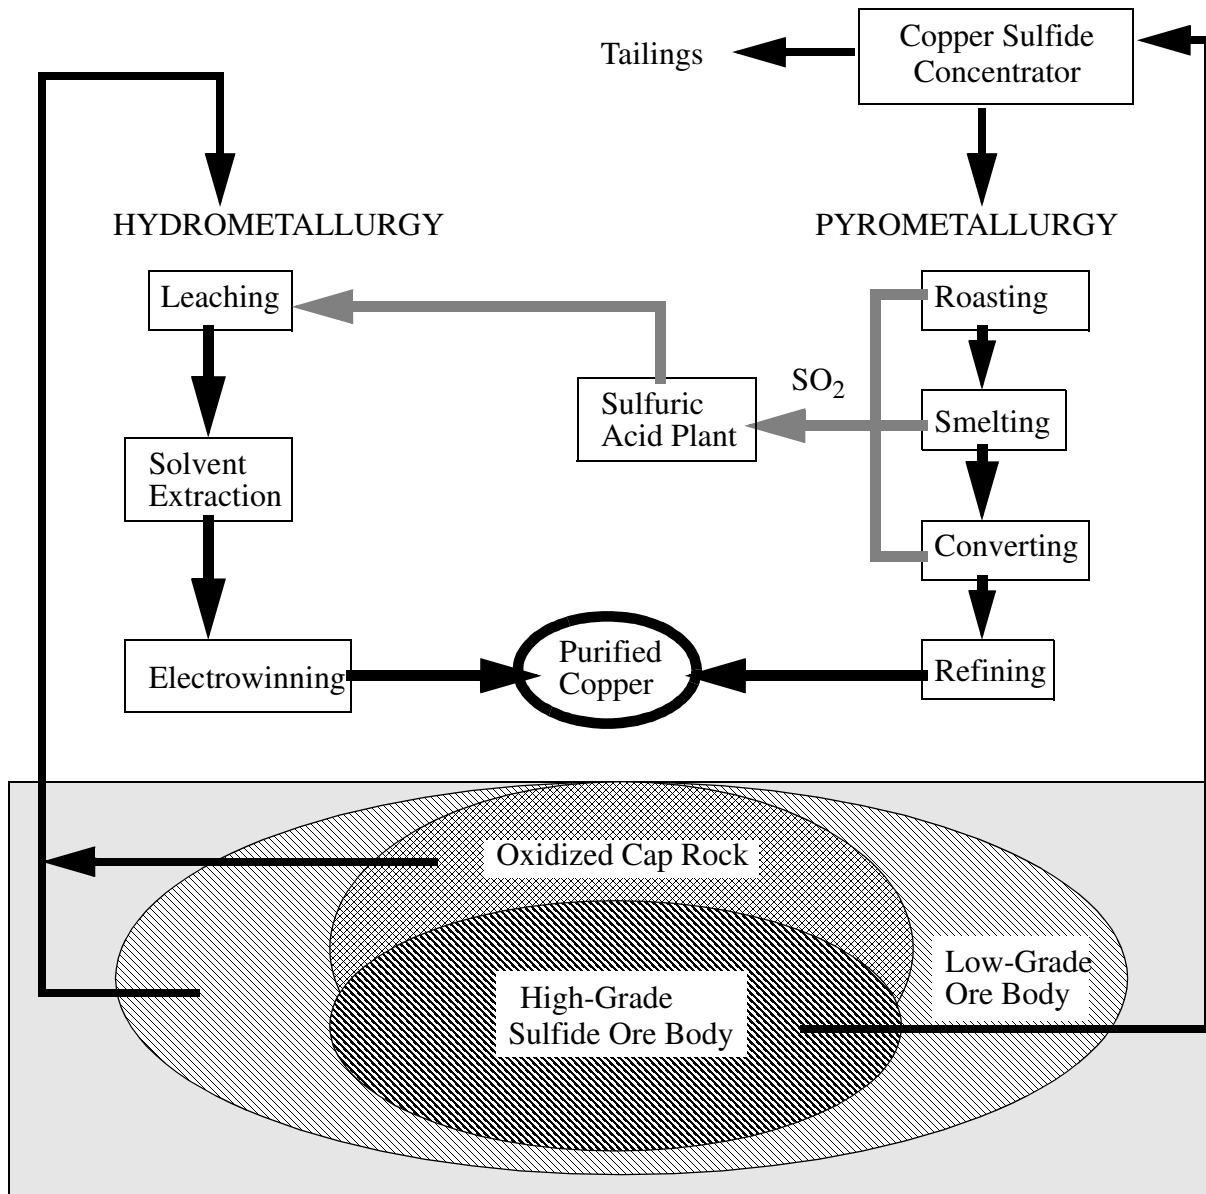


Figure 14: Processing of a typical copper oxide/copper sulfide ore body by a combination of pyrometallurgical and hydrometallurgical methods. The oxidized cap rock and the low-grade ore must be removed so that open-pit mining can reach the high-grade sulfides, and so they are readily available for hydrometallurgical processing.

Pyrometallurgy of Copper

The pyrometallurgical extraction of copper from sulfide minerals traditionally consists of the following basic steps:

- Roasting of sulfide concentrate
- Matte smelting
- Converting
- Fire-refining



After fire-refining, the copper is about 99.5% pure, and is further treated by electrolytic refining to produce extremely high purity copper.

There are several different processes for carrying out these operations, and in some cases several steps are carried out simultaneously.



Reactions

The basic chemical reactions in copper pyrometallurgy are as follows. **Note: These reactions are grossly oversimplified, and are assuming a feed that contains no sulfides other than chalcopyrite.** In a real process, there are many other oxides, sulfides, and sulfates involved besides the ones shown, and many side reactions and alternate reaction paths occur which greatly complicate the chemistry.



(A) Roasting reactions - exothermic reactions that provide heat, and that oxidize excess iron so that it can be removed during the smelting process along with silicates that are present in the ore. About one-third of the sulfide in the ore is oxidized in this step, producing a mixture of copper and iron sulfides, sulfates, and oxides.




(B) Smelting reactions - any copper that was oxidized during roasting is re-reduced by part of the remaining iron sulfide so that the copper will not be lost in the smelter slag. The copper then forms a low-viscosity “matte” that melts and separates from the silicate slag.



(C) Converting reactions - after the matte is separated from the smelter slag, it is oxidized to produce sulfur dioxide, an iron oxide slag, and metallic copper. Silica is added to help form the iron oxide slag.





(D) Fire-refining reactions - the copper from the converter (“blister copper”) contains a considerable amount of dissolved oxygen and copper oxide. The oxygen is removed by adding a hydrocarbon as a reducing agent.



Roasting

Roasting consists of partially oxidizing the sulfide mineral concentrate with air, at between 500°C and 700°C. The main roasting reactions for chalcopyrite are:



It is important to note that only about one-third of the sulfides in the mineral concentrate are actually oxidized, with the rest remaining as sulfide minerals.

Objectives of roasting are to remove part of the sulfur, convert excess iron sulfides into oxides and sulfates that can be more easily removed during smelting, and to pre-heat the concentrate to reduce the amount of energy needed by the smelter.

The reactions that occur in smelting are all exothermic, and so roasting is an autogenous process that requires little or no additional fuel.

When roasting sulfides for smelting, the sulfide minerals are only partially oxidized. The objective is to convert the iron sulfides into iron oxide and iron sulfate, while keeping the bulk of the copper in the sulfide form.

During roasting, the gases produced are 5-15% SO₂, which is concentrated enough to be used for sulfuric acid production.

There are several different types of roaster in use. The most common types are:

- **Hearth Roaster** (Figure 15): This type of furnace is suitable for coarse particles, and is common in older facilities. The capacity is on the order of 100-200 tons/day.
- **Flash or Suspension Roaster** (Figure 16): These furnaces process very fine particles, and take advantage of the rapid reaction rate to provide autogenous heating. The capacity is approximately 3 to 4 times that of a similar-sized hearth roaster.
- **Fluidized Bed Roaster** (Figure 17): Gas is bubbled up through a bed of particles, with the particles large enough that they are not swept out of the furnace, but small enough that the gas can expand the bed so that it will behave as a fluid. These reactors provide excellent gas/solid contact, but tend to generate a lot of dust.

Smelting

The matte smelting process consists of melting the roasted concentrate to form two liquid phases: a sulfide “matte” which contains the copper, and an oxide slag which is insoluble in the matte, and contains iron oxides, silicates, and other impurities.

Smelting is carried out by melting the charge at about 1200°C, usually with a silica flux to make the slag more fluid. The silica, alumina, iron oxides, calcium oxides, and other minor oxides form




Figure 16: Flash Roasting. The heat produced by the roasting reaction is used to dry the incoming ore and bring it up to roasting temperature. If the reaction is sufficiently exothermic, the ore can even be melted.




Figure 15: Hearth Roaster. Typically, these have 6 to 12 hearths, and most of the roasting occurs as the particles drop from one hearth to the next.

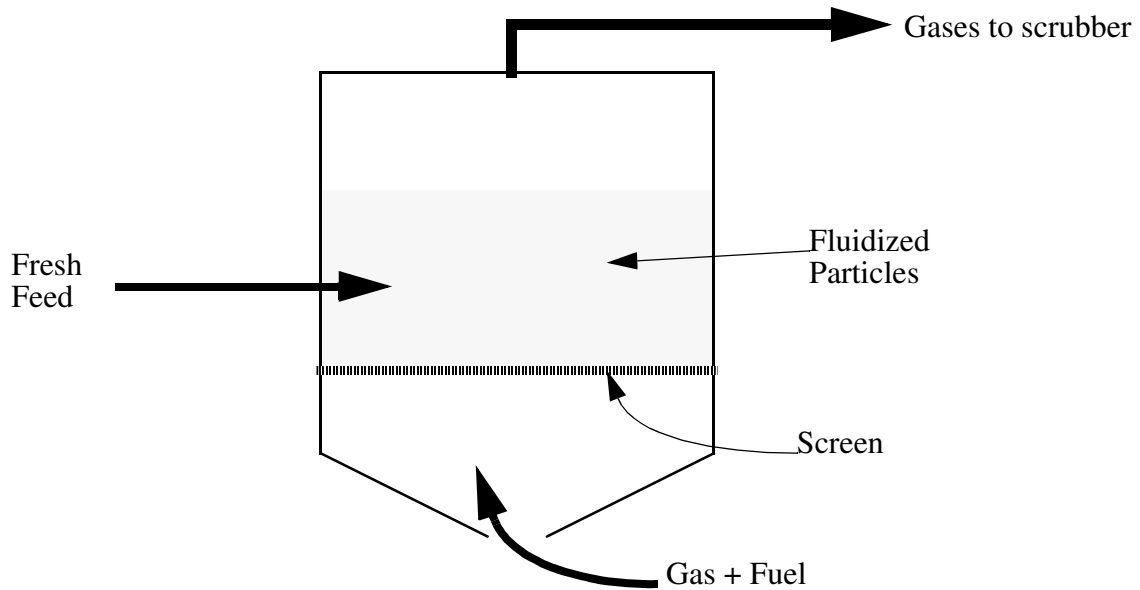




Figure 17: Fluidized Bed Reactor.

a molten slag, while the copper, sulfur, unoxidized iron, and precious metals form the matte. The slag is lighter than the matte, and so it floats to the surface to be tapped off and disposed of.

Relatively few chemical reactions are carried out during matte smelting, as its main purpose is simply to allow compounds to segregate into whichever phase they are most soluble in (slag or matte). The main important reaction is the conversion of copper oxides (that were formed during roasting) back into copper sulfide so that they will go into the matte phase:



In order for matte smelting to work, it is very important that the feed be only partially oxidized, and that enough sulfur remains in the charge for all of the copper to form copper sulfides. Matte smelting is carried out in a neutral or slightly reducing atmosphere to prevent overoxidation of the charge.

A typical matte consists of Cu_2S and FeS , and can be anywhere from 30% Cu to 80% Cu. At smelting temperatures, the viscosity of the matte is approximately 10 centipoise (cP), which is approximately 10 times more viscous than water.

Smelter slags typically have the following composition:

Fe (as FeO , Fe_3O_4)	30 - 40%
SiO_2 (from fluxes, or recycled slags)	35 - 40%
Al_2O_3	up to 10%
CaO	up to 10%

The slag must have the following properties:

- Immiscible with the matte phase
- Low solubility of Cu_2S in the slag
- Good fluidity, to minimize entrainment of droplets of copper-bearing material into the slag

In order to achieve these properties, the composition of the slag must be carefully controlled. It is particularly important that the viscosity be kept as low as possible. Slags are very viscous (500 - 2000 centipoise), which is many times higher than that of the matte phase, and high viscosity results in the slag entrapping more droplets of the matte.

Slag viscosity increases as it becomes more oxidized, because the iron becomes particles of solid magnetite which are not molten at copper smelting temperatures. It is therefore important to avoid overoxidizing the slag, so that the iron remains as liquid FeO .


Matte smelting can be carried out in a wide range of furnace types. Most commonly, it is done in reverberatory furnaces, although it can also be done in blast furnaces or electric furnaces. It can also be combined with roasting, as in the Outokumpu or INCO flash-smelting furnaces. In flash-smelting, the exothermic roasting reaction provides much of the energy needed to melt the matte and slag, and so the energy cost is reduced.

Converting

After the slag and matte are separated, the matte must be converted to metallic copper. This is done by selective oxidation of the matte with oxygen, to oxidize the sulfur but leave the copper as metal. This is a batch process, carried out in a horizontal cylindrical reactor called the Pierce-Smith Converter, which is shown in Figure 18. This converter can be rotated into different positions for charging with matte, blowing with air, skimming off slag, or dumping finished copper. This is a batch reactor, rather than a continuous process

Converting is carried out in two stages: an iron-removal stage and a copper-making stage.

Iron removal. For iron removal, silica is added as a flux to keep the slag molten, and air is blown into the converter to oxidize the iron sulfide, as shown in the following reaction:



The oxidized iron and silica form a slag that is insoluble in the matte. Since iron oxidizes more readily than copper, this reaction continues until the matte contains less than 1% Fe.

The converter is then rotated to skim off the iron-bearing slag, and the slag is disposed of.

Once the slag has been skimmed, there is room for additional matte in the converter, and it is re-charged and the iron is oxidized and skimmed from the fresh matte.

This is continued until the converter is filled with Cu_2S , with very little remaining iron.

Because of the viscosity of the converter slag, it always contains about 2-10% copper. The slag is therefore generally returned to the processing plant to be crushed, ground, and treated along with the incoming ore, to recover the copper.






Figure 18: Schematic of a Pierce-Smith copper converter. During the converting process, air is injected through the tuyeres, and off-gases are collected by a removable hood that covers the charging port. The tuyeres are equipped with pneumatic “punchers” to break through any slag or other solids that may solidify over the tuyere opening.

Copper making. Once the iron is removed, blowing is resumed. Since the iron is no longer present to consume the oxygen, the sulfur can be selectively oxidized to leave behind metallic copper, as shown in the following reaction:



Blowing is continued until the copper sulfide has been completely oxidized to metallic copper and sulfur dioxide. The resulting product is called “blister copper”, and is approximately 98.5% pure copper. The name is due to the fact that, if it is allowed to solidify, this copper will contain “blisters” due to evolution of SO_2 .


Fire Refining

Copper is fire refined by transferring it into a second converter, where it is first blown with air to completely oxidize any sulfur present. In the process of doing this, a significant amount of the copper also oxidizes, and a great deal of oxygen is dissolved in the molten copper as well. Traditionally, the oxygen was removed by a process known as “poling”. This consists of adding a “pole”, which is a green tree trunk (usually a softwood such as pine) to the copper. The pole is consumed, and the hydrocarbons it releases then reduce any copper oxide back to metallic copper. Poling was widely used in the past, but modern smelters typically inject gaseous hydrocarbons, which gives better control of the process.

In the past, fire-refined copper was the finished product, and was cast into ingots for sale. Today, almost all copper is electrorefined, and so the fire-refined copper is cast into electrolytic anodes.

Combined Processes

Each of the steps described in coppermaking (roasting, smelting, and converting) are controlled oxidation reactions with air. This suggests that, with proper design and control, it should be possible to combine individual steps to simplify the process, and ultimately to carry out the entire process in a single continuous process. The overall reaction for the entire continuous process can be written as:



Since the late 1950s, processes have been under development that combine roasting and smelting into a single continuous process, and that combine continuous converting with smelting and roasting.

Several processes have been developed that accomplish this to varying degrees. These are the Outokumpu Flash-Smelting process, the Isasmelt process, the Noranda process, the Worera process, and the Mitsubishi process. All of these processes are generally similar to each other.

Flash Smelting

The Outokumpu flash smelting process combines a flash roast with a smelting bath, as shown in Figure 19. In this unit, sulfide concentrate is injected through a burner along with preheated air, and partially burned as it falls through a flash-roasting zone. The heat produced by roasting the




Figure 19: Outokumpu flash smelting process. Incoming ore concentrate is entrained in the incoming air, and partially combusted to simultaneously roast and melt the concentrate

sulfide helps to melt the concentrate, with supplemental fuel injected into the flash-roasting area to provide the remaining heat.

By the time the ore reaches the bottom of the smelter, it has been melted. It then segregates into a slag phase and a matte phase, which are tapped off separately.

This technology has been available for nearly 50 years, and is widely used. In addition to the flash smelting, Outokumpu and Kennecott have developed a converter system, known as “flash converting”, which works on similar principles. A combined flash smelting and flash converting system is shown in Figure 20.

Isasmelt Process

The Isasmelt process was developed at Mount Isa Mines, originally to smelt lead ores. The heart of the process is a top-blown vessel, as shown in Figure 21. Concentrate is pelletized along with supplemental fuel (coal), fluxes, and any other necessary additives. The pellets are then fed to the vessel, and air that has been enriched to about 45% oxygen is blown into the top using a lance. Roasting and fuel combustion occur in the molten bath, and the molten material is sent to a holding furnace where the slag and matte are separated.

Work is currently underway to adapt the Isasmelt furnace for use as a converter, so that the entire process could be run continuously.

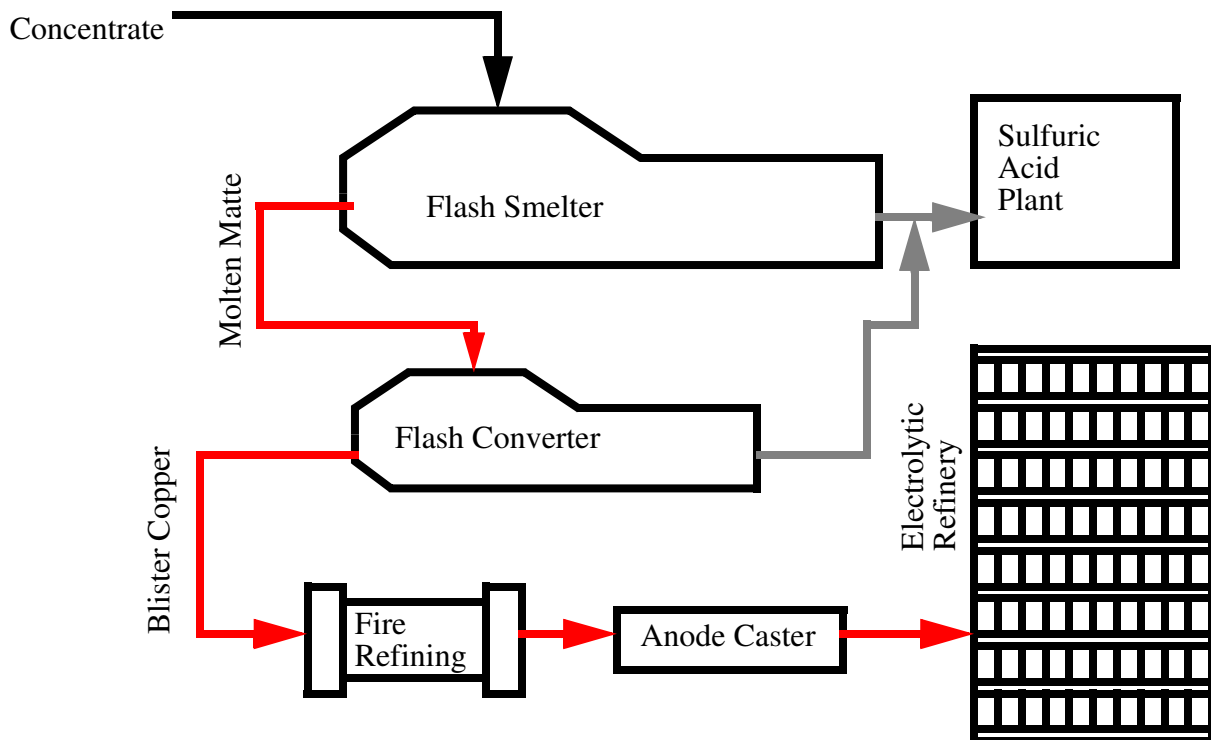


Figure 20: Integrated flash smelting/flash converting copper production facility.

Noranda Process

The Noranda process uses a cylindrical reactor similar to the Pierce-Smith converter, as shown in Figure 22. Pelletized feed is added along with coal fines at the feed-end burner, where it is roasted, melted, and forms into slag and matte phases. The slag is tapped off, while air is injected into the tuyeres to continuously produce metallic copper.

Worcra Process

The Worcra process was developed at the same time as the Noranda process, and is similar in most respects. The main difference is that it is a stationary furnace rather than a tiltable vessel, and the air is injected from the top through air lances, rather than through submerged tuyeres as is the case in the Noranda vessel. It also includes a settling basin at the slag discharge end, to remove droplets of matte that were entrained into the slag.

Mitsubishi process

The Mitsubishi process uses three furnaces in series rather than a single furnace. The first unit is essentially a flash-smelting unit, used to simultaneously roast and melt the ore into a slag and matte. The molten material is then sent to an electric furnace which keeps it molten using electric arc heating under a reducing atmosphere, to allow the slag and matte to separate completely. The matte is then run to a continuous converter to be converted into metal. The use of three furnaces




Figure 21: Isasmelt smelting vessel.

instead of one gives better control of the process than in either the Noranda or Worcra methods, and the design of the furnaces to provide continuous flow rather than batch flow gives a higher throughput than is the case with the traditional methods.




Figure 22: Schematic of the Noranda single-step reactor. Feed at the feed end first melts and segregates into slag and matte. The slag is then tapped off, while the matte is converted to copper.

Hydrometallurgy of Copper

The most common application of hydrometallurgy to copper ores is to leach the low-grade and oxidized rock that surrounded the higher-grade copper sulfide deposits. With increasingly stringent environmental regulations, it is becoming more difficult to get the permits for a copper smelter, and so there is interest in developing hydrometallurgical processes that can completely replace conventional smelters. At the moment, hydrometallurgy is not competitive with pyrometallurgy for producing metal from high-grade ores.

Ore Preparation

The ore will require some degree of comminution, so that the leaching solutions will be able to reach the copper oxides and sulfides that are to be dissolved, as shown in Figure 23. The amount of comminution is limited by the value of the ore. For example, a very low-grade copper ore might only be crushed to a two inch top size, because the value of the added copper that would be recovered by grinding finer is not enough to justify the cost.




Figure 23: In uncrushed ore, valuable mineral inclusions are completely surrounded by rock, and cannot be reached easily by leaching solutions. Crushing the ore exposes these inclusions to leaching.

When possible, it is helpful to pre-concentrate the ore, to reject any rock that contains very little copper. This reduces the amount of material that has to be leached, and so reduces the size of the plant necessary to deal with it.

Leaching

Copper leaching is the dissolution of the copper in a solvent, while leaving the gangue minerals behind as undissolved solids. Copper is normally leached using one of the following methods:

- Dump leaching
- Heap leaching

Dump Leaching: The overburden and tailings dumps for copper mining operations often contain enough copper that it would be worthwhile to recover some of it. Dump leaching is then used to recover as much of the copper as can be leached out without spending a lot of time and money on preparing the ore. This is done by trickling the leaching solution over a dump, and collecting the runoff solution to recover the copper that it dissolves, as shown in Figure 24. Dump leaching is quite slow, with periods of months or years needed before leaching is completed, and typically only about 60% of the copper in the dump is recovered.




Figure 24: Typical dump-leaching operation.

Heap Leaching: This is similar to dump leaching, but instead of simply dumping ore on a hillside, the ore is crushed approximately to the size of gravel (to improve leaching rates and recovery) and piled onto artificial waterproof pads, as shown in Figure 25. Once the ore has finished leaching (after approximately 6 months to a year), the leached gangue is removed from the pad for disposal, and replaced with fresh ore.




Figure 25: Typical heap-leaching operation. The pad is usually concrete, with a rubber or asphalt coating to make it waterproof.

Factors Affecting the Leaching Rate

The amount of time needed to dissolve the copper from the ore is influenced by the following factors:

- **Particle Size:** Larger particles take longer to leach completely than smaller particles. The reason for this can be seen from the shrinking-core model of leaching (Figure 26). When a particle first begins to leach, the leaching solution can diffuse quickly to the portion of the ore that contains metal, and then diffuse back out. As time passes, the core of unleached ore becomes smaller, and the shell of leached ore (which the solution must diffuse through to reach the core) becomes thicker. Since diffusion is a slow process, this means that the leaching of a particle goes more slowly as time goes by. Also, since larger particles will have a thicker shell of leached material, they will take much longer to leach completely than small particles will.





Figure 26: The Shrinking-Core model. As the shell of leached ore becomes thicker, it becomes more difficult for solvent to reach the unleached ore, and so the leaching rate slows down.

- **Reagent Concentration:** The speed of the dissolution can often be increased by adding more of the leaching reagents (acids, complexing agents, oxidizers, etc.). At very high reagent concentrations, the cost will become prohibitive.
- **Temperature:** Changing the temperature will increase the rate of metal dissolution, increase the rate of diffusion, and usually increase the total amount of metal that can be dissolved per liter of solvent. Increases in temperature are limited by the boiling point of the leaching solution.
- **Thoroughness of Mixing:** Improving the mixing between the ore and the solvent will reduce the number of “dead zones,” which leach slowly because they do not contact fresh solvent. By mixing the ore and the solvent thoroughly, the speed and completeness of leaching will increase. Eventually, the ore and solvent will be perfectly mixed, and further improvements will not be possible.


Types of Leaching Reactions

Depending on the nature of the ore, copper leaching can either be a non-oxidative process, or an oxidative process.



Non-oxidative leaching: In non-oxidative leaching, there is no change in oxidation state. One example is the dissolution of copper sulfate by water:



Another type of non-oxidative leaching reaction is the dissolution of alkaline materials, such as malachite, by an acid:






Oxidative leaching: Many ores are not soluble unless they are oxidized. For example, CuS (covellite) is made much more soluble if it is oxidized to copper sulfate:





In this case, it is the oxidation of sulfur to the more soluble SO_4^{-2} that makes the material soluble.

The oxidative leaching reaction above uses oxygen as the oxidizing agent. Unfortunately, gaseous oxygen is not very soluble in water (at 20°C, water in equilibrium with air only dissolves about 8 parts per million oxygen). Because of this, it is difficult to dissolve enough oxygen in the leaching solution to give a rapid leaching rate, and violent agitation of the slurry is needed to bring air in contact with the solid particles that need to be oxidized. A solution to this is to use other oxidizing agents, that are more soluble in water than oxygen is. One such oxidizing agent is ferric iron (Fe^{+3}).



There are two drawbacks to using ferric iron as an oxidizing agent. The first is that it is only soluble in acidic solutions (pH less than 4). If the solution approaches neutral pH, then the ferric iron will rapidly precipitate out as iron hydroxide, and be useless. The second problem is that, unlike air, it is not free, and so to be economical it should be regenerated by oxidizing the ferrous iron (Fe^{+2}) with air to convert it back to ferric iron.

Bacterial leaching: The re-oxidation of ferrous iron to ferric iron would be too slow to be of practical use, except for the action of the bacteria *Thiobacillus ferrooxidans*. This bacteria is unusual, in that it gets its metabolic energy by oxidizing ferrous iron to ferric iron, rather than by eating organic material or by collecting sunlight like other organisms do. In nature, this bacteria and its relatives oxidize pyrite (FeS_2) to produce dissolved ferric sulfate and sulfuric acid, according to the reactions:



Both the ferric iron and the sulfuric acid are useful leaching reagents for copper sulfides, and pyrite is a common accessory mineral in copper sulfide mineral deposits. As a result, these bacteria can greatly reduce the cost of leaching operations by producing many of the necessary reagents directly from the ore. *Thiobacillus ferrooxidans* and related species are widespread, and applications such as heap and dump leaching (which are exposed to the open air) will rapidly pick up an appropriate bacterial culture without being inoculated.

Solution Purification

Leaching reactions are not perfectly selective, and so there will be many elements in solution after leaching, and not just copper. These other elements can cause problems in recovering the metal, and so they sometimes need to be removed from the solution. Also, the copper can be very dilute, and so a means is needed to concentrate it. This is generally done using ion exchange processes.

Ion-Exchange Groups

Ion-exchange separations are based on the electrostatic attraction of certain molecular groups for ions in solution. These groups can be either positively charged or negatively charged, as shown in Figure 27. The groups are attached to any of several types of larger molecular structures, so that they will not go into solution and be lost. The groups are commonly attached to either solid resins, or to organic liquids that are not soluble in water.




Figure 27: Mechanism of ion exchange in solvent extraction. Active groups such as carboxyl are attached to the hydrocarbon chain of a non-water-soluble oil. When the oil is contacted with water, the carboxyl exchanges the ion it currently holds (H⁺ in this case) for a copper ion. It can be made to release the attached copper by contacting it with a stripping solution that has a very high H⁺ concentration.

Solvent Extraction

In solvent extraction, the ion-exchange groups are attached to the molecules of an organic liquid that is insoluble in water. This makes it possible to make a continuous process, as shown in Figure 28.

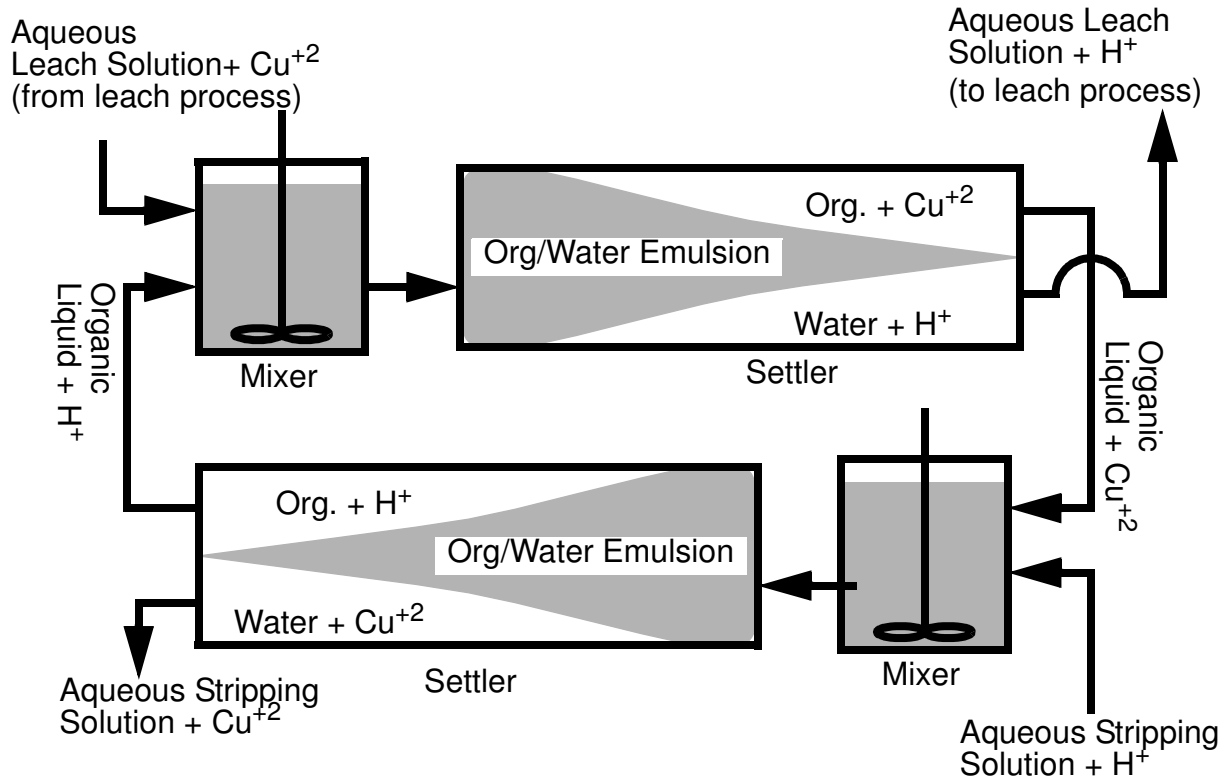



Figure 28: Solvent-Extraction process. The “mixers” ensure good mixing between the water solutions and the organic liquid so that ions can be freely passed back and forth. The “settlers” give the organic liquid the chance to separate from the water

Metal Recovery

Once the metal has been dissolved, the solids have been removed, and the solution has been purified, the metal must be recovered in a solid form. This can be done chemically, or electrochemically.

Cementation

Dissolved copper will plate out on an iron surface due to the following reaction:



So, if a solution with dissolved copper is run through a bed of shredded scrap iron, the iron will oxidize and dissolve, while reducing the copper ions and plating them out as solid copper.

For cementation to be efficient, the iron scrap must have a high surface area.

Electrowinning

Electrowinning is an electrochemical process for precipitating metals from solution, as shown in Figure 29. The anode is made out of a material that will not easily oxidize and dissolve, such as lead or titanium. The advantage of electrowinning is that once the metal is plated out of the solution, the metal-depleted solution can usually be recycled directly to the hydrometallurgy operation.




Figure 29: Basic electrowinning cell. Oxygen gas is produced at the anode, while metal is plated onto the cathode.

Electrorefining

Most copper is used for electrical applications, and it is therefore important that the copper be high-purity so that it will have good electrical conductivity. This high-purity copper is produced by electrorefining. In this process, impure copper (generally from the fire-refining stage in a pyrometallurgical operation) is cast into copper anodes. These are then placed into an electrochemical cell which dissolves metal from the anode, and redeposits high-purity copper on the cathode, as shown in Figure 30. Metals that are more electropositive than copper (such as zinc and iron) dissolve and remain in solution, while metals that are less electropositive than copper

(such as gold and silver) accumulate as an “anode sludge” that is collected and sold for further refining.




Figure 30: Basic electrorefining cell for producing high-purity copper.

Layout of a typical hydrometallurgical operation for copper

A copper dump-leaching operation would generally be laid out roughly as shown in Figure 31. Sulfuric acid would normally be produced by the associated copper smelter, and other reagents are recycled as much as possible both to minimize effluents and to reduce reagent costs.

Since organic liquid from the solvent extraction will cause problems if it gets into the electrowinning cells, a froth flotation stage is often included between solvent extraction and electrowinning to completely remove any remaining droplets of organic liquid.

For a simpler circuit, the solvent extraction/electrowinning stages could be replaced by a cementation step, which would produce a lower-purity copper product that would need to be refined further.

Primary Copper Production Byproducts

There are a number of other metals that are frequently associated with copper ores, and that are produced as byproducts of copper production. The most important byproducts are:

- Nickel (recovered during conversion of the matte during smelting)
- Arsenic (recovered along with nickel)
- Gold (recovered from anode sludge in electrorefining)
- Silver (recovered from anode sludge in electrorefining)
- Molybdenum (separated as molybdenite during sulfide concentration)

The value of these byproducts is often enough to make it profitable to process a copper ore that would not be economic based only on its copper content.

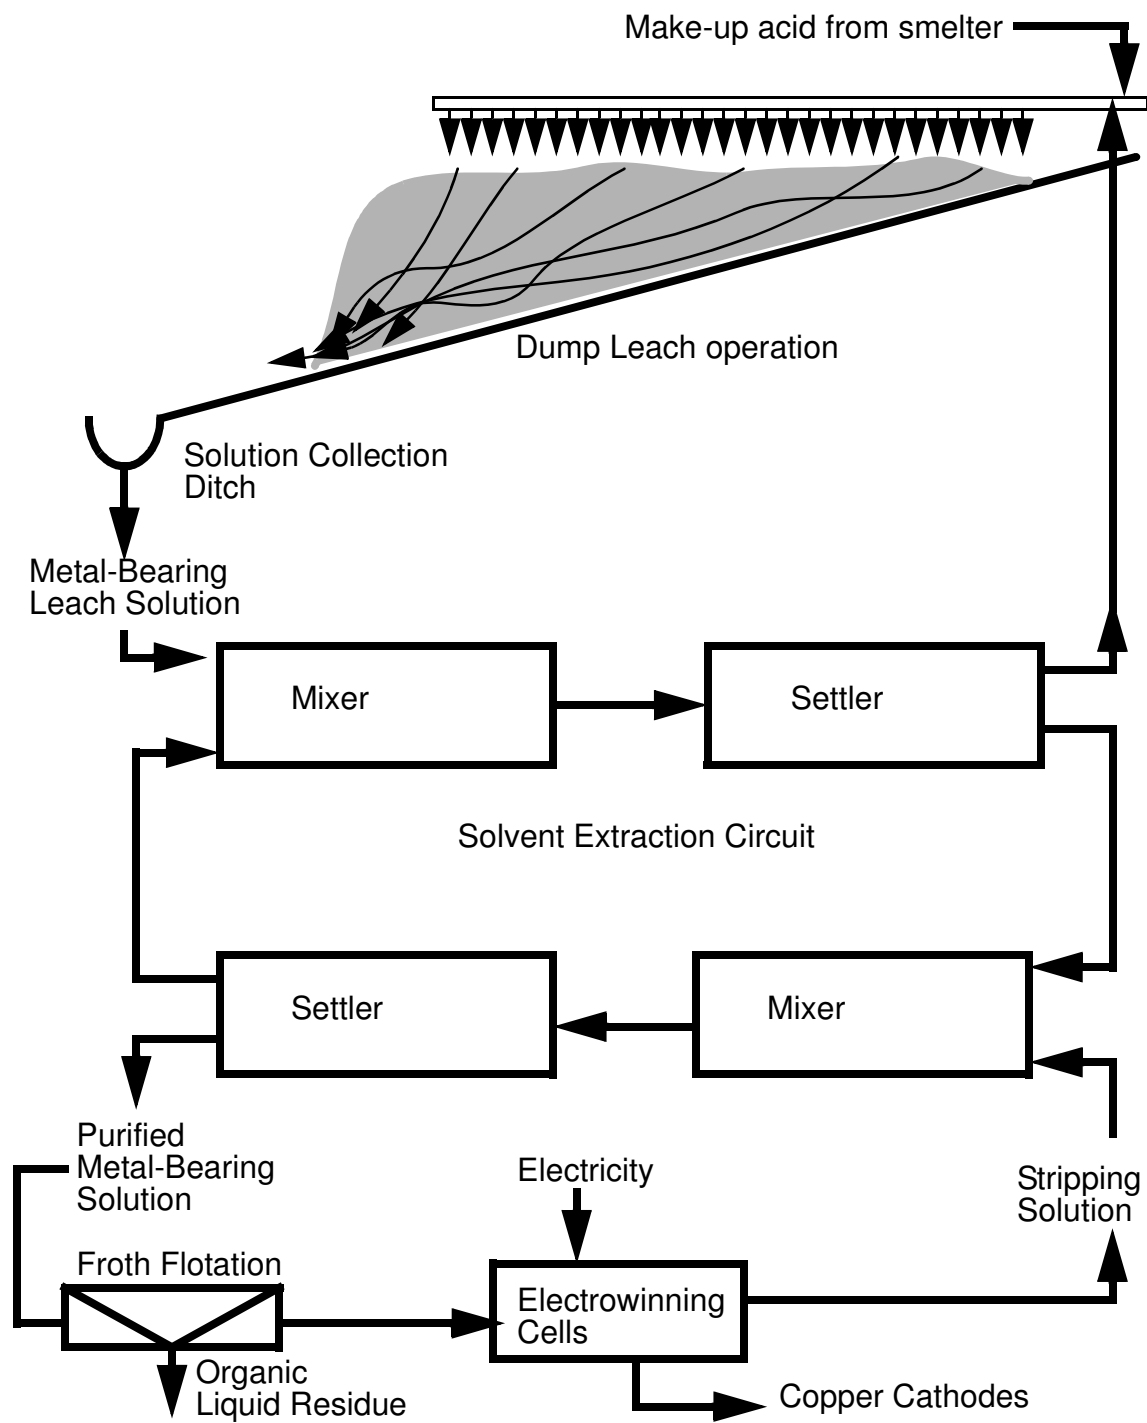


Figure 31: Typical layout of a hydrometallurgical copper extraction process.

Primary Aluminum Production

The primary ore of aluminum is bauxite, a mixture of hydrated aluminum oxides:

- Gibbsite - $\text{Al}(\text{OH})_3$ (most extractable form)
- Boehmite - $\gamma\text{AlO}\cdot\text{OH}$ (less extractable than Gibbsite)
- Diaspore - $\alpha\text{AlO}\cdot\text{OH}$ (difficult to extract)

It is formed by weathering of aluminum-bearing rocks by rainwater, and so bauxite deposits tend to be found in areas that are now, or were in the past, tropical high-rainfall areas.

Aluminum is contained in many minerals, but bauxite is the preferred ore because it has the highest aluminum oxide content and is therefore the cheapest to process.



Although aluminum is an extremely common element on earth, it was not practical to produce aluminum metal at a reasonable cost until two breakthroughs had been made: a method for producing purified aluminum oxide from bauxite (the Bayer process), and a method for converting aluminum oxide to metallic aluminum (the Hall-Heroult process).

Reactions

Bayer Process

The Bayer process consists of three steps:


1. Dissolving the aluminum hydroxides from bauxite with hot (250°C) sodium hydroxide solution




As an undesirable side reaction, Kaolinite clay dissolves, reacts with alumina and caustic, and precipitates as “red mud”:



2. Crystallizing purified aluminum hydroxide by addition of seed crystals:




3. Calcining aluminum hydroxide to aluminum oxide:



Hall-Heroult Process

This process consists of dissolving aluminum hydroxide in molten cryolite (Na_3AlF_6) and electrolyzing the molten salt at 1.5 - 1.7 volts with a consumable carbon anode, to produce molten metallic aluminum and carbon dioxide:





Alumina production using the Bayer process

The Bayer process is the primary method for producing purified alumina (Al_2O_3) from bauxite. It is a hydrometallurgical process, using concentrated alkali solution to dissolve hydrated aluminum oxide from bauxite in an autoclave at elevated temperatures, followed by solid-liquid separation and precipitation and calcination of purified alumina. Processing before the Bayer process is minimal, usually being little more than a rotary breaker to remove coarse rock from the bauxite.

Dissolution

The first step is to selectively dissolve the hydrated aluminum oxides from the bauxite, while leaving behind the bulk of the silicate minerals and other impurities:




In most bauxites, Gibbsite and Boehmite are the main aluminum-bearing phases. Gibbsite dissolves most readily in the alkali solution, and Boehmite dissolves with some difficulty. Diaspore, the third form of hydrated aluminum oxide, is difficult to dissolve and is generally not recovered in the Bayer process. When Diaspore is an important phase in a bauxite, it can only be dissolved if lime is added to the digester, and the digestion temperature is increased to about 300°C .

The conditions for efficient dissolution of the aluminum oxides are as follows:

- Caustic Soda (Na_2CO_3 equivalent): 150-255 grams/liter (Gibbsite), 205-445 grams/liter (Boehmite)
- Temperature: $100 - 150^\circ\text{C}$ (Gibbsite), $200 - 315^\circ\text{C}$ (Boehmite)
- Autoclave Pressure: 56-70 psi (Gibbsite), 280-1500 psi (Boehmite)
- Feed Slurry Solids Content: 50% wt. Bauxite
- Digestion Time: 30 minutes - 4 hours

Note that the temperatures are well above the boiling point of water at atmospheric pressure. An autoclave is therefore needed for the digester to raise the boiling point high enough.

Alkali is normally purchased as caustic soda (Na_2CO_3), but it needs to be converted to NaOH in order to work properly. The conversion is done by causticizing with lime ($\text{Ca}(\text{OH})_2$):



This has the added benefit that any phosphate dissolved from the bauxite will be precipitated by the lime as $\text{Ca}_5(\text{PO}_4)_3\text{OH}$.

If Kaolinite or other clays are present, they also dissolve in the alkali solution. For Kaolinite, the dissolution reaction is:



The precipitate from this reaction is a material referred to as “red mud”, which is a fine material that is difficult to remove from the solution. Note that the formation of red mud precipitate consumes dissolved aluminum, and so the red mud represents an aluminum loss.

Flash Cooling

The temperature of the slurry is reduced by progressively reducing its pressure in a series of tanks, allowing a portion of the liquid phase to flash to steam in each tank.

The steam goes to heat exchangers, where it preheats leaching solution returning to the digester. This recycling of heat is critical for keeping the fuel costs for the process as low as possible.

Once the slurry temperature is reduced to about 100°C, it does not need to be kept in pressure vessels any longer. It is then diluted to about 3% solids for solid/liquid separation.

Solid-Liquid Separation

Once the bauxite has been digested, the slurry contains one liquid phase and two solid components:

- Caustic liquid solution with the dissolved aluminum
- Undissolved coarse material (sand)
- Precipitated fines (red mud)

The solids must be separated from the liquid phase so that they will not contaminate the alumina product when it is precipitated from solution.

The coarse sands, which are mainly undissolved silicate grains from the bauxite, settle from suspension very rapidly and are easily removed by gravity classifiers or hydrocyclones. They are washed in rake or spiral classifiers to remove dissolved alumina from them, and then disposed of.

The fine red mud has a very slow settling rate, and is removed by large-diameter thickeners after adding a flocculant to increase the red mud settling rate.

The liquid overflow from the thickener still contains about 50-100 mg of solids per liter. The remaining solid particles are removed from this stream by filter presses.

The red mud underflow from the thickener is 15-25% solids, and contains a great deal of dissolved aluminum that needs to be recovered. This is done using a counter-current decantation (CCD) mud-washing circuit. This is a series of thickeners can be used to remove the liquid from the solids with a minimum of dilution by washwater (Figure 32). The effect of the arrangement shown is that the solids travel down the series of thickeners in one direction, while the wash solution travels up the series in the opposite direction. In each tank, the solids are being washed by a solution with a lower metal concentration than in the previous tank, until it reaches the final tank where the wash solution contains no dissolved metal at all. The final solids product is often either centrifuged or filtered to remove the last bit of liquid with dissolved metals.

The red mud is a disposal problem for the aluminum industry. No practical uses have been found for the material yet, and so it must be pumped into impoundment areas.

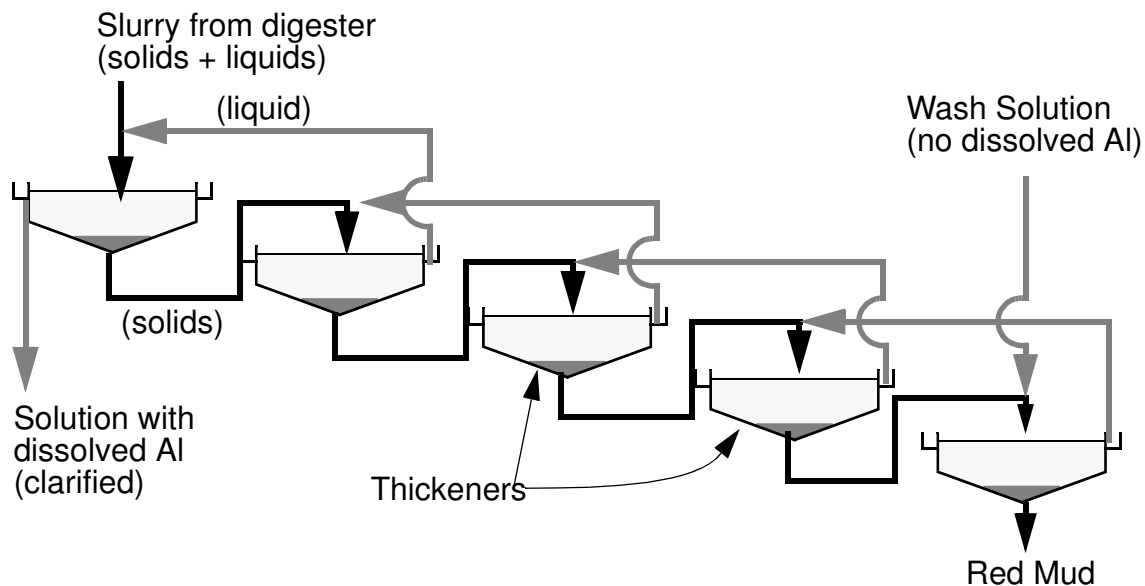


Figure 32: Series of thickeners, being used for continuous countercurrent decantation.

Precipitation

Once the liquid solution has been clarified and cooled, it is a supersaturated solution containing about 100-175 grams of dissolved alumina per liter. Since $\text{Al}(\text{OH})_3$ does not crystallize easily on its own, it is precipitated by adding seed crystals in long-residence-time agitated precipitator tanks. Crystallization takes many hours, and so the tanks must be large (24-36 feet in diameter, 60-80 feet tall).

The seed crystals of $\text{Al}(\text{OH})_3$ are typically added at a rate of 40-300 grams/liter, and are 50-95% passing 74 μm . The pregnant liquor contains 100-175 grams of alumina per liter, and about 45-75 grams per liter are precipitated out on the seed crystals.

When the slurry exits the precipitators, the crystals that have reached the target size are removed by a classifier, washed, and filtered. Crystals that are smaller than the target size are concentrated by a thickener, and recirculated back to the precipitators to act as seed crystals. The spent liquor is reheated, recausticized, and sent back to the digester.

Calcination

The product from the precipitation process is wet crystals of $\text{Al}(\text{OH})_3$. This must be dried and converted to Al_2O_3 before it can be used as feed to the Hall-Heroult process. This conversion is done by heating to approximately 1300 - 1500°C (fine, “floury” alumina requires the higher temperatures)

Calcination can be carried out using rotary kilns, or fluid-bed calciners. The fluid-bed units are reported to use 1500 BTU per lb of alumina produced, compared to 2200 BTU per lb for rotary kilns.

Bayer Process Flowsheet


A representative flowsheet for a Bayer process plant is shown in Figure 33.

Hall-Heroult Process

Many reactive metals, such as magnesium and sodium, can be produced by electrolyzing a molten chloride salt of the metal. Unfortunately, aluminum chloride can not be electrolyzed easily because it sublimates rather than melting. Even if the pressure is increased enough to allow it to have a molten phase, molten aluminum chloride is an electrical insulator and so cannot be used as an electrolyte. Because of this, aluminum was produced by an expensive sodium reduction process until a suitable electrolyte was found.

In 1886, Hall (in the U. S.) and Heroult (in France) independently developed processes using cryolite (Na_3AlF_6) as a molten salt electrolyte for producing aluminum. Cryolite is electrically conductive, and dissolves alumina.

The electrolysis cell used is shown in Figure 34. Both the anodes and the cathode are made of carbon. The anodes are gradually consumed by the oxygen that migrates to the anodes, and so the overall electrolysis reaction is:



The theoretical voltage for this reaction to occur is 1.15 volts, but due to anode overvoltages the potential in actual practice is 1.5-1.7 volts. In order to overcome the electrical resistance of the electrodes, conductors, and containers, the typical operating voltage is 4 -5 volts.

Operating conditions for an aluminum electrolysis cell are:

- Temperature: 935-975°C
- Alumina content of electrolyte: 2-6%
- Cell voltage: 4-5 volts
- Faraday efficiency: 85-90%

Additives to cryolite bath:

- AlF_3 (improves Faraday efficiency)
- CaF_2 , LiF , MgF_2 (reduce freezing point of electrolyte)

Raw material and power usage per ton of Al produced:

- Al_2O_3 : 1.90-1.95
- Electrolyte: 0.04-0.06

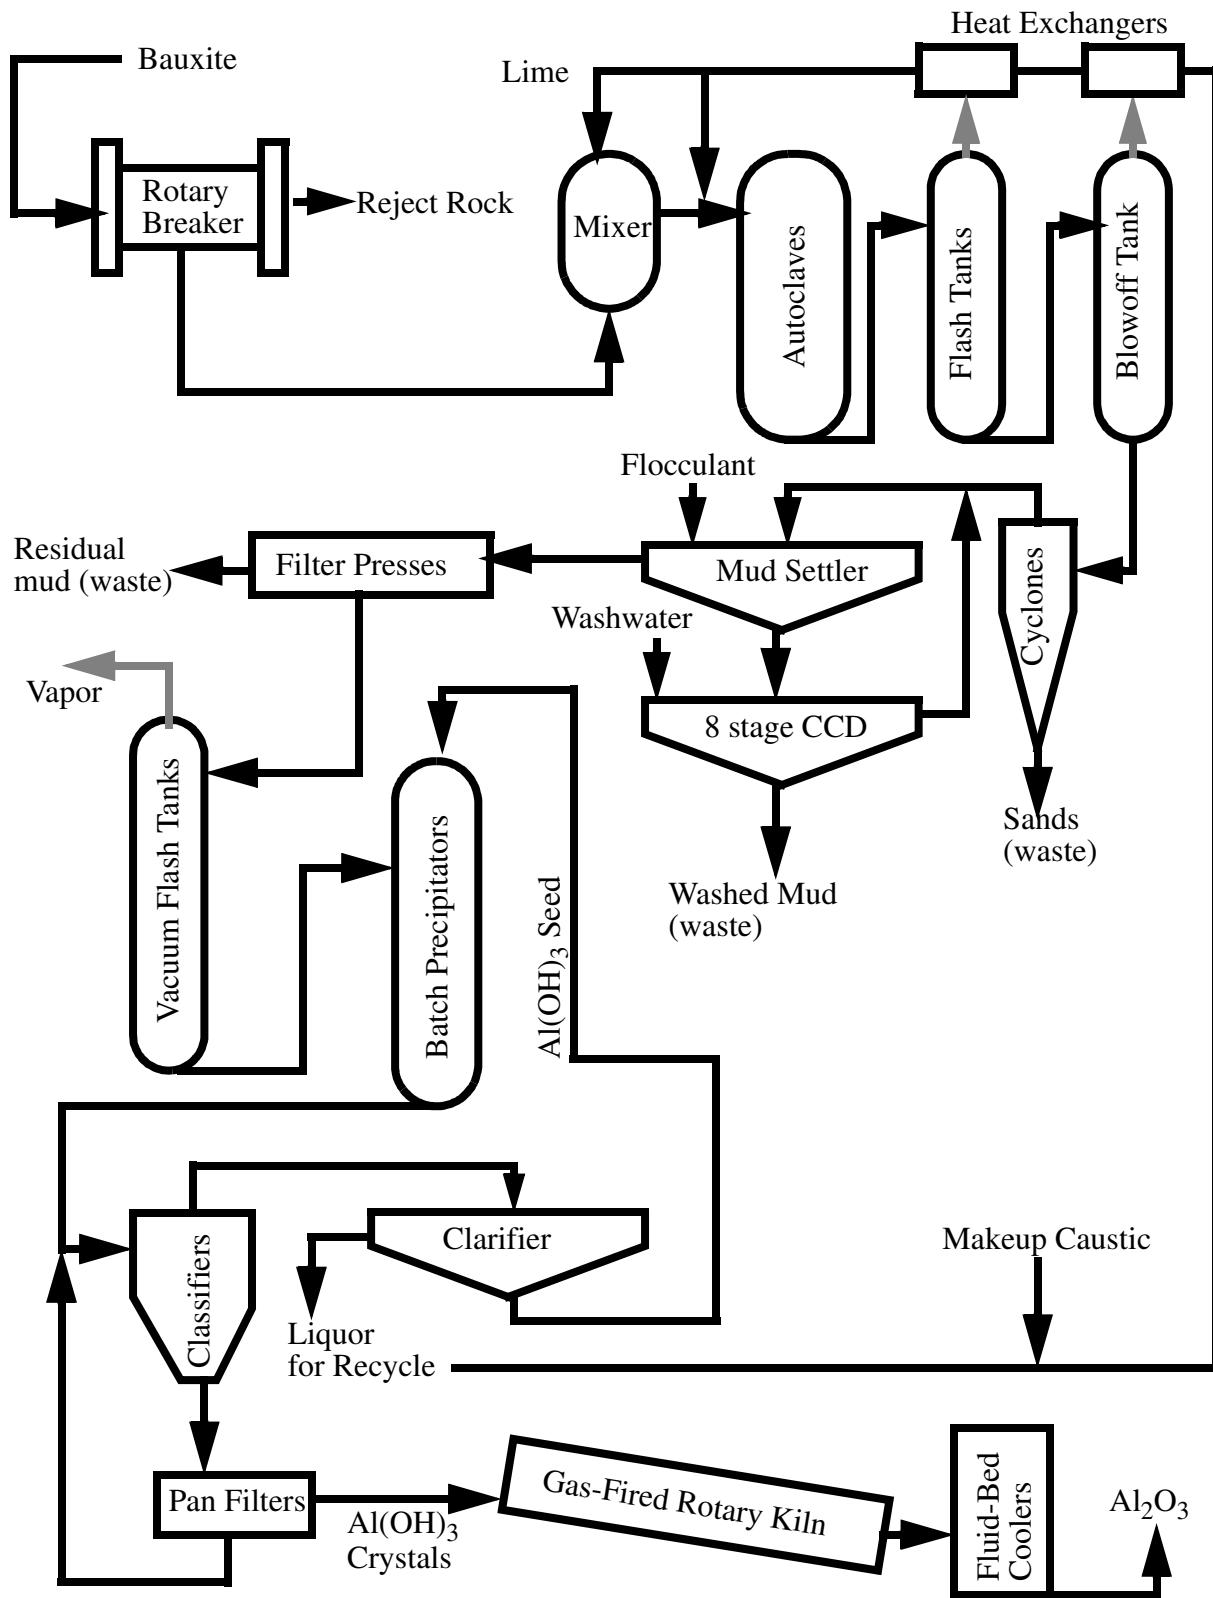


Figure 33: Flowsheet for a Bayer process plant, based on the Kaiser Aluminum plant in Gramercy, Louisiana.

- Anode carbon: 0.43-0.50
- Power, Kw-hr: 13,000-16,000

Wastes produced in aluminum electrolysis:

- HF, CF₄, and other fluorocarbons (from anode reactions)
- CO₂
- “Salt cake” (spent electrolyte, metal oxides, and entrained metallic aluminum droplets)

Note the tremendous power consumption of aluminum production. The electrolysis must supply three electrons for every atom of metallic aluminum, and so very high currents are needed to produce aluminum at a reasonable rate. Because of this, electric power is the single largest cost in aluminum production, and so aluminum smelters are typically located in areas where electric power is inexpensive, generally near major hydropower sites.

The high power consumption of primary aluminum production also encourages aluminum recycling, which requires much less electric power.





Figure 34: Electrolysis cell for production of metallic aluminum

Other Aluminum Production Processes

Because of the high energy costs of the Hall-Heroult process, there has been considerable interest in finding alternative aluminum production routes that are more energy-efficient. A few of the more promising approaches are as follows:

Carbothermal Reduction

The reaction of aluminum oxide with carbon




can theoretically occur at temperatures greater than about 2000°C. This would provide a relatively inexpensive method for producing aluminum, because carbon is a cheap source of reducing power. However, there are several side reactions that form oxycarbides, aluminum carbide, and volatile monovalent aluminum compounds, and these side reactions make the overall carbothermic reduction much harder to operate and control.

An example of a carbothermic process is one patented by Alcoa in 1976. This is a two-step process, using two baths of molten aluminum oxide. The first step is:




The Al_4C_3 produced is then sent to the second bath, where the reaction is:



to produce metallic aluminum

Toth Process

This process is based on the reduction of AlCl_3 by manganese, and consists of four steps. The first step is carbochlorination of alumina, where the alumina is provided as kaolinite clay:




The aluminum chloride is then reacted with manganese to make metallic aluminum:



The chlorine is then regenerated by reacting the manganese with oxygen:



and the manganese oxide is reduced with carbon to regenerate metallic manganese:



While this uses only about 5% as much electricity as the Hall-Heroult process, it consumes as much carbon as the carbothermal process. It is unlikely to ever be used industrially, because of the great difficulty in regenerating the manganese.

Alcan Process

This is based on the production of monovalent aluminum as AlCl, which is highly volatile. The Alcan process first melts the aluminum-bearing feed in an electric furnace, and reduces it with carbon to produce a liquid alloy of iron, silicon, and aluminum. This alloy is then reacted with AlCl₃ at about 1300°C, which produces AlCl:




The AlCl vapor is then condensed by a shower of liquid aluminum, and the AlCl disproportionates back into aluminum metal and AlCl₃:



This process does not require purified alumina, and so it can use materials such as raw bauxite or clay as feedstock, skipping the need for the Bayer process. Unfortunately, other volatile halides are removed from the alloy along with the AlCl, resulting in an impure product, and the chlorides produce a very corrosive environment, and so the overall process was less economical than the Hall-Heroult process.

Aluminum Chloride Electrolysis

There has long been interest in direct electrolysis of aluminum chloride to produce metal, according to the reaction:



The attractive features of this reaction are:

- Considerably lower working temperature than cryolite electrolysis.
- Higher current densities can be achieved before the “anode effect” halts production.
- Consumable carbon electrodes are not needed, which results in lower operating costs and less pollution.

The drawbacks are:

- Pure aluminum chloride sublimates at atmospheric pressure, and so it either needs to be pressurized to produce a liquid phase, or it needs to be dissolved in another salt, such as NaCl or KCl.
- Pure aluminum chloride is an electrical insulator, and so it needs to be dissolved in a conductive salt (again, NaCl or KCl can be used).
- Aluminum chloride is produced from pure alumina, which is made by the Bayer process. The cost of converting the alumina to aluminum chloride is a considerable expense.

A process of this type was developed by Alcoa, and was considered for industrial use, although it apparently has not yet been used on larger than a pilot scale. The process layout was as shown in Figure 35.

The Alcoa process required high-purity electrolyte from which all oxygen had been excluded, which is why the two filtration stages and the inert-gas-purged evaporator were needed. A special electrolytic cell was used that was lined with oxysilicon nitride to prevent corrosion. The

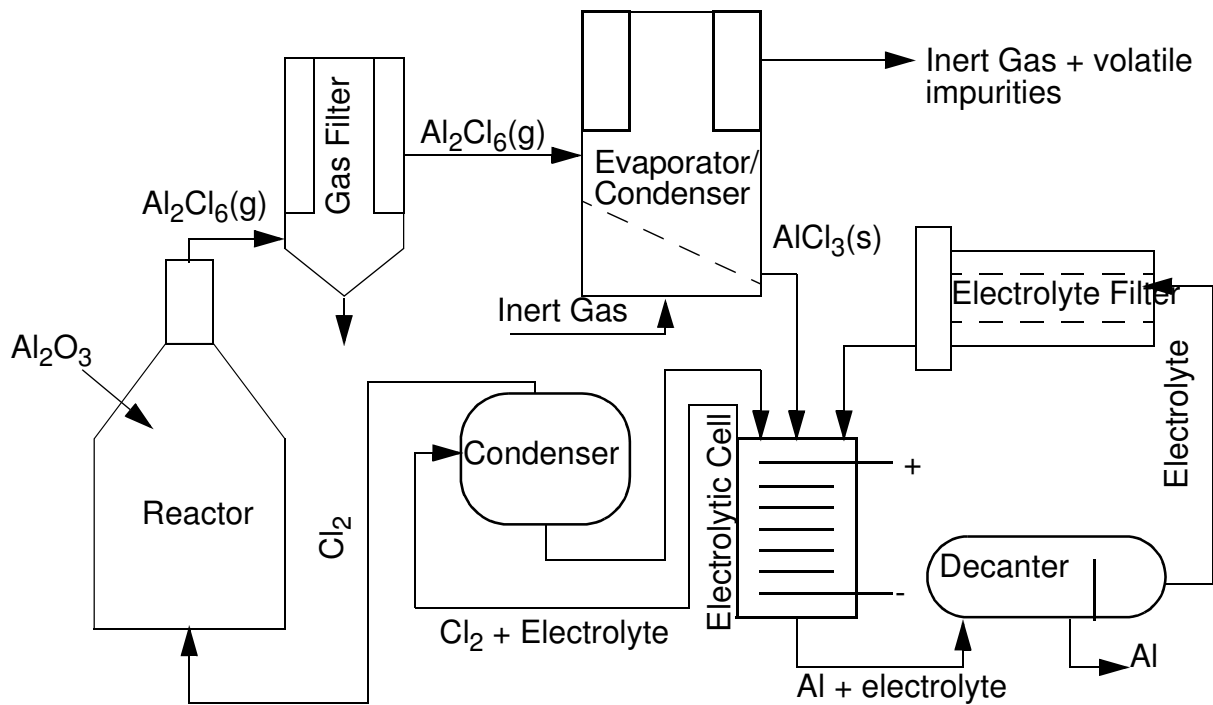


Figure 35: Flowsheet of the Alcoa aluminum chloride electrolysis process. “Electrolyte” signifies a mixture of alkaline chlorides, usually NaCl , LiCl , MgCl_2 , KCl , AlCl_3 , and CaCl_2 .

electrodes were graphite. Since chlorine was being produced at the anode rather than oxygen, the graphite anode was not consumed. A typical composition for the electrolyte was 5% AlCl_3 , 53% NaCl , 40% LiCl , 0.5% MgCl_2 , 0.5% KCl , and 1% CaCl_2 .

Titanium Production

Titanium is quite expensive for a structural metal, at \$3 - \$4 per lb for titanium sponge. Its cost is not due to scarcity (it is quite common, being the 9th most common element in the Earth's crust), it is expensive because of the difficulty in reducing titanium oxides to metal. Since the primary use of the metal is for aerospace, the demand fluctuates greatly depending on aircraft orders in a given year. U. S. consumption has varied between 18 million tons/year and 28 million tons/year since 1998.

Titanium is in high demand for aircraft because it has a better strength-to-weight ratio than aluminum or steel. It is also desirable for marine applications because it is almost completely corrosion-resistant due to the extreme chemical stability of the thin oxide layer that it forms on its surface.

Metal production only accounts for about 3% of all the titanium minerals mined. A much greater quantity of titanium is used for producing titanium dioxide pigment, with over 1.3 **billion** tons of purified TiO₂ produced in the U. S. annually. It is in great demand as a pigment because in purified form it is white, and its high refractive index gives it excellent hiding power. It is also non-toxic, and has completely replaced the toxic lead-based pigments that were used previously in almost all paint.

Sources

Titanium can be extracted economically from:

Rutile (TiO₂)

Ilmenite (FeTiO₃)


Leucosine (rutile contaminated with iron oxides, formed by alteration of ilmenite)

Both rutile and ilmenite are hard, dense minerals that resist weathering, and as a result they tend to concentrate in placer sand deposits where they are mined by dredging, and separated from other valuable minerals by a combination of density, magnetic, and electrostatic separations.

Purification and Metal Production

A typical flowsheet for producing titanium from rutile is shown in Figure 36 . The basic process consists of:

1. Roast the impure titanium dioxide with carbon in the presence of chlorine:



If the titanium dioxide is reacted with carbon alone, it will produce highly-refractory TiC, which is extremely difficult to process further. The chlorine is needed to convert the titanium to the highly-volatile chloride, which vaporizes readily.

2. Distill the titanium tetrachloride (TiCl₄) to remove chlorides of iron and other metals. Titanium tetrachloride is liquid at room temperature, and takes relatively little energy to distill.

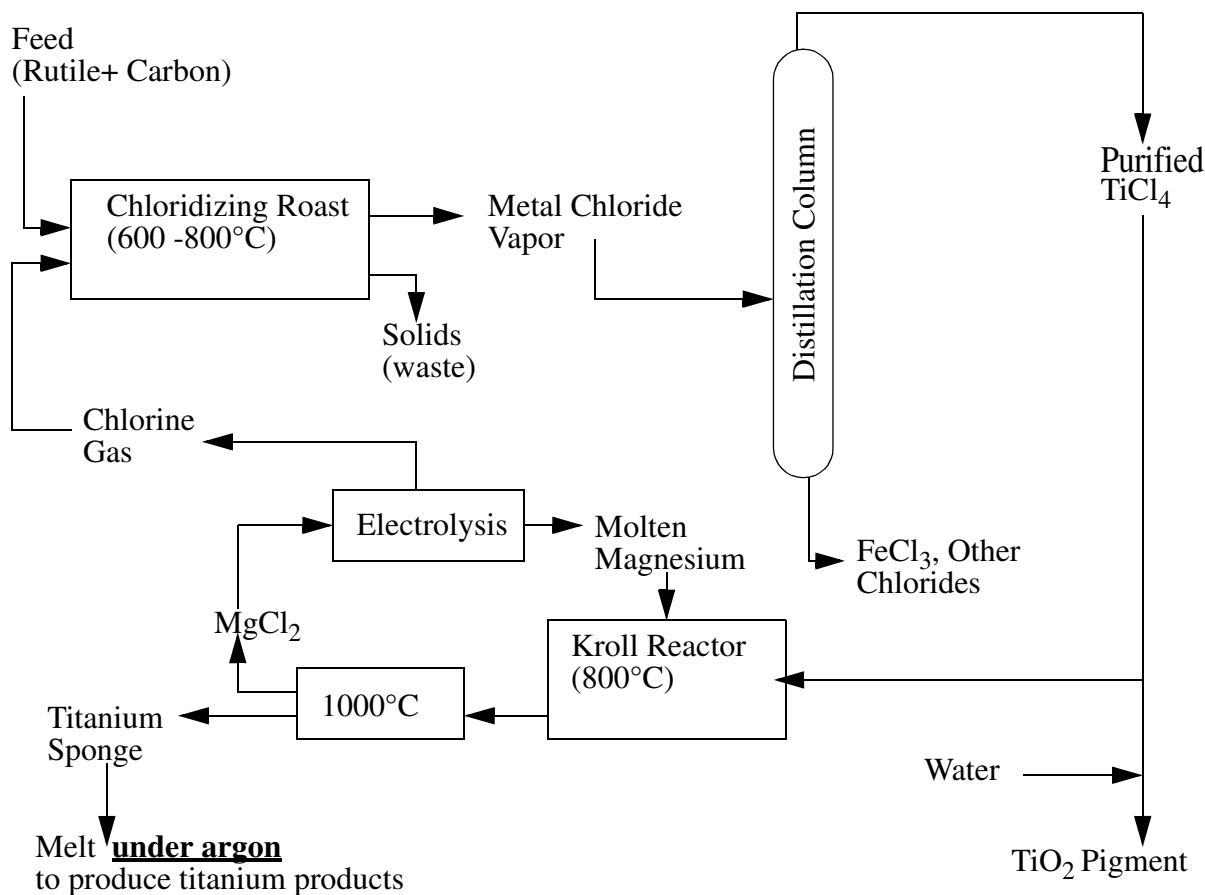




Figure 36: Flowsheet for producing titanium sponge from rutile. To produce titanium dioxide pigment, the purified TiCl_4 can be reacted with water to produce a white, finely-divided, iron-free product.

3a. If pigment is the desired product, react the liquid titanium tetrachloride with water to produce high-purity fine titanium dioxide:



3b. To produce metal, the titanium tetrachloride is reacted with molten magnesium metal, in the Kroll process:



This results in a porous sponge of titanium metal, filled with magnesium chloride.

4. The product from the Kroll reaction is heated to melt the magnesium chloride, and separate it from the titanium sponge. The molten magnesium chloride is then electrolyzed to produce magnesium metal and chlorine gas, which are recycled back into the process:



5. The titanium sponge is then melted under an inert atmosphere to form into ingots for further fabrication.

Critical Considerations

There are a number of important features that contribute to the high cost of titanium production:

- A. The chloridizing roast gases are extremely corrosive and toxic, and must be completely contained. Maintaining gas-tightness in high-temperature reactors is difficult and expensive.
- B. Titanium tetrachloride is highly reactive with water and oxygen, and so all air and moisture must be rigidly excluded from all process stages.
- C. Molten titanium metal will not only react with oxygen, it will **burn in nitrogen**, so all melting must be carried out under an inert atmosphere (usually argon).