

Some considerations on future developments in ferro-alloy furnaces

I.J. Barker

University of the Witwatersrand, Johannesburg, South Africa

Keywords: Submerged-arc furnace, plasma furnace, d.c. furnace, scale up, Søderberg electrodes, smart grid, demand side management, load shedding

Abstract - This paper argues that the scale up of furnaces and the supply of electricity are going to be two major issues that will affect the ferro-alloy industry in the near future, and that economic factors will drive this development.

The most common type of furnace for producing ferro-alloys at present is the submerged-arc furnace with three electrodes, fed from a three-phase a.c. electrical supply. The scale up of this technology has now reached a fundamental constraint, which is caused by the electrical reactance of the secondary circuit. If the economy of any further scale up is to be achieved in the future then a different technology will have to be used.

The supply of electrical power in future is likely to become a more complex issue than it is at present. Existing submerged-arc furnaces tend to run at fairly steady loads, but the ability to swing the load under demand-side management may offer advantages, and may allow a furnace to get cheaper power from the organisations that supply this power. This will have to be counterbalanced against the nuisance factors incurred by having a varying load in the operation of the furnace.

Various options are therefore discussed in this paper. The possibility of scale up of Søderberg electrodes is discussed, as well as the use of d.c. power and multiple electrodes. Some of the issues with load swinging are also examined. A particular scenario is also briefly presented to show that considerable further scale up of ferro-alloy furnaces is still possible.

INTRODUCTION

The basic technology used for producing the bulk of the ferro-alloys worldwide has not changed dramatically for several decades. There have been some notable developments such as d.c. plasma furnaces, pollution reduction technology, pre-reduction technology, control and information systems, and various progressive refinements of the technology, but the standard workhorse still remains the submerged-arc furnace with three electrodes arranged in a circle and fed from a three-phase a.c. supply.

Up to the late 1960s, submerged-arc furnaces tended to be somewhat smaller than they are today. Then around the early seventies, new installations with significantly larger capacities started to appear. The main motivation for this

would seem to have been the economy of scale. As the sizes increased, other problems started to appear that had not been evident with the earlier smaller furnaces, and this effectively constrained further scale up. However, the production from one of the larger of the present ferro-alloy furnaces is of the order of only 10 000 metric tons per month, while in comparison the output of a larger-sized blast furnace is of the order of 300 000 tons per month. Although these numbers are very approximate, they do reveal a dramatic difference. This would seem to suggest that the motivation for further scale up of ferro-alloy furnaces would still be there if the various constraints could be overcome.

A further consideration is that the future environment in which ferro-alloy furnaces will need to operate is likely to be different from the present. One factor that stands out is that the supply of electrical energy is likely to become more restrictive, but also more machine-intelligent. Within such constraints, the ability to swing load is likely to have economic advantages, but present-day submerged-arc furnaces are not easily able to do this.

The above issues all suggest that further developments are now needed. The purpose of this paper is therefore to table a number of the factors involved, and to make suggestions as to how some of these factors might be addressed.

It has been argued that sustainability issues also need to be addressed by the industry in future (see, for example, Dos Santos¹). The ability to swing load as described here will be a fundamental requirement for the industry to use more 'green' energy than it is able to accept at present. Other 'black swan' factors² may also influence future developments, but, by their very nature, they are difficult to predict in advance. Not only are the two issues discussed in this paper clearly evident now, but also there are practical ways of evolving the technology through both of them.

PRESENT CONSTRAINTS ON THE SCALE UP OF SUBMERGED-ARC FURNACES

It is perhaps not surprising that problems associated with scale up tend to be found more often on larger furnaces. Generally, these problems are associated with the operational side of the furnaces. Many basic factors are involved in such problems, but with regard to scale up, two stand out in particular, *viz.*:

1. Scale up of the electrical circuit

As this type of furnace is scaled up, the reactances increase while the resistances decrease with size. When the reactances become comparable to or larger than the resistances, the behaviour of the electrical circuit becomes problematic³. This leads to difficulties with the measurement of the electrical state of the furnace, and consequently with the control of the electrodes.

2. Scale up of the electrodes

A given Söderberg electrode is limited in the current that it can carry. Scaling up to higher currents requires larger electrode diameters, but these larger electrodes would seem to be more prone to problems. These problems can probably be ascribed to two main mechanisms. Firstly, the mechanical strength of any structure gets weaker as the structure is scaled up in size. This is because the weight rises by the cube of the linear dimension while the strength rises by only the square of the dimension. Secondly, there is the issue of the skin effect associated with the a.c. current in a Söderberg electrode. This is an electromagnetic phenomenon in which the current density is higher nearer to the surface than within the bulk of a conductor. In the carbon part of a Söderberg electrode, the skin depth is of the order of 0.3 to 0.4 metres. Hence, in a small electrode the skin effect is not significant, but for diameters larger than about 1.2 to 1.5 metres the centre tends not to receive enough current to heat it at the same rate as the rest of the electrode.

To people directly involved with a larger furnace that is giving trouble, the above underlying factors usually do not reveal themselves as such but as a syndrome of endemic problems such as recurring broken electrodes including apparent tip breaks, unbalanced electrodes, low power factor, tapping difficulties, hot roofs, off-grade product, higher MWh/t, and so on.

The frequency of the a.c. electricity mains plays a part in both the above factors – the reactances are proportional to frequency, while the skin depth is inversely proportional to the square root of frequency. Hence, a lower frequency would benefit both. This also suggests that d.c. current might help to alleviate some of the problems. However, while d.c. current might obviate the skin effect, it plays no direct part in the mechanical strength of the Söderberg electrode. At this stage, we do not know whether it is the strength or the skin effect that is the main constraint on the scale up of the electrodes. This will only become apparent once large Söderberg electrodes have been tried in a d.c. application in practice.

THE SUPPLY OF ELECTRICAL ENERGY

Ferro-alloy furnaces use large amounts of electrical energy, and the future availability of such energy is becoming a concern. Firstly, throughout the world, electricity generation based on the combustion of carbonaceous fuels is becoming less favoured while widespread use of nuclear energy remains controversial, but there is a push towards renewable sources of energy for the future generation of electricity. Solar power and wind power are typical examples of such renewable resources. Secondly, in South Africa at present, the generating capacity is somewhat limited. The ferro-alloy industry has felt the effects of this in various ways. In particular there is an ongoing concern that there might be insufficient generating capacity to meet peak demand, particularly at times when the demand is high or when several generators have to be taken out of service simultaneously.

The purpose of this paper is not to discuss the energy problems facing the world, but it is necessary to briefly discuss the ways in which some of the options can be handled. For this purpose, Figure 1 shows a very simplified schematic of an electricity distribution system. This figure is intended to show how electrical power is generated by a number of suppliers and then fed into a common grid to be distributed to a number of users. At all times the production of electrical power must match the consumption – there can be no excess or shortfall. This is unlike, for example, a water reticulation system where there are intermediate reservoirs that serve as buffers between the supply and the demand. Of course, there are ways to store electrical energy when it is available in excess and to return it to the grid when it is needed. Pumped water storage schemes are a relatively common example of this. Such an arrangement is not shown explicitly in Figure 1, although it can be modelled as a combination of a user and a supplier. The main purpose of Figure 1 is simply to emphasize the balancing of power between suppliers and users at any given time.

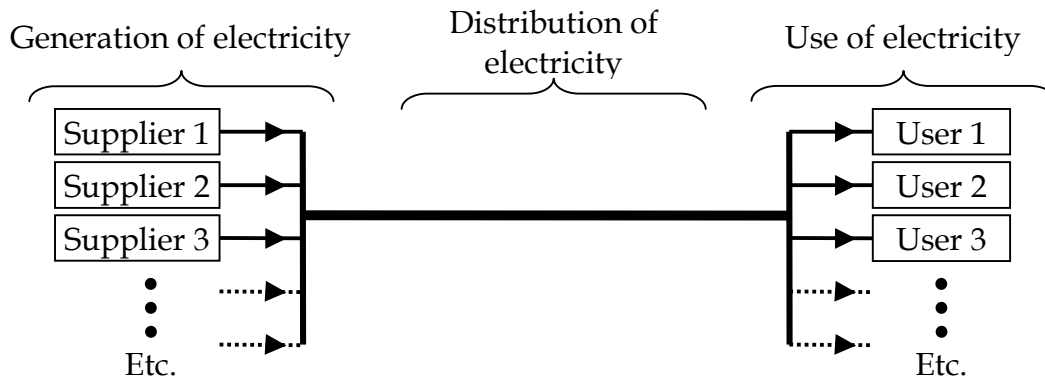


Figure 1: Simplified schematic representing generation and use of electricity

This balancing of the power requires that the output of the generators be adjusted continuously to match demand. The actual control of this over a number of suppliers is a fairly complicated issue. It also depends on the time scale involved, as the sub-second control has to be done with the equipment that is already running, whereas for hour-to-hour up to day-to-day control it is also possible to bring in or take out modules in the suppliers or the users.

The sub-second control is based in essence on the frequency of the a.c. as well as on the voltage. If, for example, a user starts to draw more power, then the rotating generators of the suppliers will start to slow down as much of this energy will come initially from the kinetic energy in their rotation. This decrease in frequency will be detected and the control system will then automatically increase the steam fed to the turbines to make these generators speed up again.

Often, some of the suppliers cannot easily have their output changed. In this case, other suppliers have to handle all the changes when there is a variation in demand. The biggest problem that can occur is when there is insufficient

reserve capacity on the generating side to increase the output any further when it is needed. This can happen when all the suppliers are generating at their respective maximum outputs.

Unfortunately, there is a problem with most large-scale renewable sources of energy in terms of their controllability. Consider a wind turbine for example. Not only can it only produce electricity when the wind is blowing sufficiently, but also it cannot increase its output if it is already generating at maximum. Such a turbine would therefore have to be run below its maximum output if it is to have the reserve capacity to handle fluctuating loads. Alternatively there would have to be some other generator that handles the fluctuations.

Another way to handle variations in demand is to shed some of the loads drawn by the users. This is the concept of demand-side management.

The Smart Grid

This has been a subject of much debate recently. In essence, the technology is now becoming available to pass information at fairly high speed between the various stakeholders on the grid, and this opens up the possibility of doing a number of things that were not previously possible. For example, 'intelligent' demand side management on ferro-alloy furnaces becomes a distinct possibility, whereas up to now it has been virtually out of the question.

The use of ferro-alloy furnaces for demand side management

Ferro-alloy furnaces generally have been regarded as reasonably steady base loads on the power grid, unlike many of the other loads on the grid that fluctuate with the time of day, the weather, and various other factors. It used to be argued that a steady base load like this is beneficial to the power producers as it lessens the ratio of the fluctuations to the total load.

However, if a large base load like a ferro-alloy furnace can be swung as needed, then it can also be used to absorb some of the fluctuations caused by other loads. Organisations that supply and distribute power should react favourably to such users, and it is not unreasonable to expect cheaper tariffs in return. The problem is that present-day submerged-arc furnaces are rather limited in their ability to vary their load in this way.

It is tempting to dismiss the concept of varying the load on a ferro-alloy furnace as an unnecessary nuisance, but eventually economics will decide - the plant that can tolerate fluctuations on demand will get cheaper power than the plant that has to run at a steady load. Also, power users that can offer a variable load on demand to their electricity supplier may get treated more favourably in return.

SOME POSSIBLE WAYS FOR THE FERRO-ALLOY INDUSTRY TO ADAPT TO CHANGES

Submerged-arc furnaces

Figure 2 shows the schematic layout of the existing common type of submerged-arc furnace, with three transformers feeding the three electrodes in the common knapsack arrangement.




Figure 2: Schematic layout of a conventional submerged-arc furnace fed from a 3-phase a.c. supply, with the common knapsack connection to the electrodes

The first obstacle to scaling up is the supply of 3-phase power from the a.c. mains direct to the furnace electrodes, in the manner shown in Figure 2. This is because the inductances in the furnace circuit at the frequency of the a.c. mains are simply too large. This obstacle is a fundamental problem with the basic circuit, so some alternatives are going to have to be found if the industry is to get around this constraint.

A related problem with this arrangement is that all three electrodes are part of the same circuit, so that when there is a problem on one electrode, the other two electrodes are likely to be affected as well. This interaction effect³ between electrodes gets worse with decreasing power factor, and hence is also a constraint that inhibits scale-up.

One fairly obvious possibility is the use of d.c. on existing submerged-arc furnaces in place of a.c. power. Figure 3 shows a suggested arrangement for doing this, with each electrode fed from its own d.c. supply. Note that the three boxes shown in Figure 3 are now d.c. power supplies, not simply a.c. transformers as shown in Figure 2. The advantage of this configuration is that each electrode can then be handled on its own, as each electrode and its power supply is a separate circuit.

One way to test this arrangement might be to retrofit this to an existing a.c. furnace, *i.e.* to use the infrastructure of an existing a.c. furnace, but in particular replace the hearth by one with anode connections, and replace the transformers with d.c. power supplies.




Figure 3: A rearrangement of the furnace in Figure 2, with each electrode now fed from its own separate d.c. supply.

The d.c. power supplies on furnaces like this should be fairly amenable to fast demand-side management should the need arise in future. This could be an attractive feature to organisations that are involved in the supply of power.

Plasma furnaces

D.c. plasma furnaces (also known as d.c. arc furnaces) offer another approach, and also they are now proven in commercial existence. Being inherently d.c., they avoid any problems with a.c. reactance. However, the metallurgy is somewhat different, so a direct comparison with a submerged-arc furnace is difficult. A 60 MW furnace of this type has recently been commissioned at Middelburg^{4,5}.

In comparison with submerged-arc furnaces of similar power, d.c. arc furnaces tend to run at higher voltages and lower currents. Because of these lower currents, they have smaller electrodes. They also normally use graphite electrodes, but already the diameters of these graphite electrodes on the larger furnaces are close to their upper limit. However, further scale-up would still be possible using multiple electrodes, and arrangements like that shown in Figure 3 might also work for d.c. arc furnaces.

The d.c. power supplies on plasma furnaces should also be fairly amenable to fast demand-side management.

Constraints on electrodes

Søderberg electrodes are reputedly cheaper than pre-baked or graphite-type electrodes, but they tend to be more complicated to use. They are the normal type of electrode used in existing a.c. submerged-arc furnaces, but potentially they could also be used with d.c. in submerged-arc furnaces, and possibly also in d.c. plasma furnaces. These other applications need further development as they are still largely untried.

This use of d.c. current will overcome the skin effect, and so might allow Söderberg electrodes in submerged-arc furnaces to be scaled up beyond the present limit of about 2 metres diameter. Larger electrodes will allow higher current and higher power per electrode, and so are one avenue towards further scale up of furnaces.

During a furnace shutdown, after the current through a Söderberg electrode is turned off, the electrode cools down slowly with time. Later, when the current is restored, the electrode proceeds to heat up again. The contraction and expansion that results from this thermal cycling during furnace shutdowns tends to stress an electrode mechanically and can lead to breakages. Computer modelling⁶ has made it possible to study the time evolution of this heating and cooling and the resulting mechanical stresses. This has shown, for example, that with electrodes of about 1.5 metres in diameter, the time constant of this cooling and heating in the bulk of the electrode is of the order of about 4 hours, although some small regions inside the electrode near to the lower ends of the contact shoes do tend to heat and cool faster than this. It is perhaps not surprising that common wisdom in the industry advises that furnace shutdowns longer than about 4 hours have to be followed by baking in of the electrodes while for shorter shutdowns one can often get away without baking in.

In some parts of the world, ferro-alloy plants shut their furnaces off for a short period during the peak demand time each evening. This period is usually about 2 hours, which ties up as being safely less than the 4 hours discussed above for cooling and heating of the electrodes.

Fundamental energy-transfer theory tells us that the time constant for cooling and heating of a long cylinder should be proportional to the square of the diameter. Hence, if the time constant for an electrode with a diameter of 1.5 metres is about 4 hours, then for an equivalent electrode of, for example, 3.0 metres diameter it should be about 4 times longer, *i.e.* 16 hours. (This would only apply to these electrodes carrying d.c. current, not a.c.) This suggests that larger electrodes, if they are viable, may be able to tolerate significantly longer shutdowns. Hence for load shedding in demand-side management, it may be possible for plants to offer their electricity suppliers longer shutdowns if they have larger electrodes together with d.c. currents.

Swinging the load

Most furnaces at present run close to full power for most of the time, but if a ferro-alloy furnace is going to be used for demand-side management, then obviously the load is going to have to be varied.

Possibly the biggest problem in implementing such a concept will be to get acceptance from the plant management team to allow some external factor to vary the furnace load seemingly at random!

Consider the issue of selecting an electrical operating point when a submerged-arc furnace has to be run at low load for an extended period. One option is to adhere to a constant resistance so that power scales down by simply the square of the electrode current. Another option is to keep to a constant Westly C_3 factor⁷, which effectively increases the resistance at lower load. A third option is to decrease the resistance at lower load so that the electrode currents stay relatively higher in order to keep the electrodes baked. Yet another option is to switch the furnace off altogether, but this can be done only for a relatively short period. It is perhaps surprising that this difficulty in choosing an operating point for a reduced load is not a new problem – Westly himself was already well familiar with this issue over 37 years ago⁸ – and yet it still remains an area of uncertainty.

Swinging the load on a submerged-arc furnace is actually a more complex issue than just choosing an operating point as discussed above. When a furnace's load decreases, the metallurgy of the process is going to be affected. But this is not all; in a conventional a.c. furnace, changing the resistance set point will alter the power factor and so affect the electrical measurability and controllability of the furnace as well. Indeed it will probably be easier to swing load on d.c. furnaces than on a.c. furnaces because d.c. decouples the electrodes and so avoids this measurability and controllability issue. Also the baking of the electrodes is going to be affected as well as the slipping and erosion of the electrodes. This will make tracking of the electrode length difficult, which will aggravate the electrical measurability and controllability problems. Going from one load level to another may be done as a step but the responses will have transients with a wide range of time constants. Furthermore, other parts of the plant like the gas cleaning may work differently at lower loads, and the management of the tapping floor and the raw materials would have to be much more flexible.

From this it is clear that any implementation of load swinging is going to have to be done gradually and in an evolutionary manner, as much still remains to be learnt about the consequences. It will not be something that can be implemented in a large step-like change. Also the entire plant should be designed to allow the process to run at differing loads.

This evolution of load swinging would benefit if improved online information systems could be made available to assist with the more complex operational requirements. This would include facilities such as online computer-based simulation models that enable the operations team to 'see' such things as the baking of Söderberg electrodes and the tracking of electrode length. This would require more rugged and reliable instruments for measuring plant variables like the slipping of the electrodes. The information system would also tie in to the Smart Grid network, and include various predict-ahead planning facilities. Such information systems will need time for development and refinement.

A LARGER FURNACE

The furnace design outlined in this section is intended as a hypothetical exercise into what might be the next ceiling on the scale up of ferro-alloy furnaces. It is one of several possible scenarios, most of which are still relatively unproven at present.

The proposed design is a 6-electrode furnace with the electrodes arranged in a circle, as sketched in Figure 4. Each electrode would be fed by its own power supply module as shown, using d.c. power in much the same way as the arrangement in Figure 3. Let us consider a submerged-arc furnace here, although the same layout and supply could probably also be used with a plasma furnace.

Let us also consider using Söderberg electrodes. Let us take a nominal diameter of 3.0 metres for these electrodes, as this is only about 50% larger than the largest electrodes that are presently in use on actual submerged-arc furnaces. Time might prove that even larger electrodes than this will work, but it might also show that even just this diameter is simply not attainable.




Figure 4: Layout of the proposed larger furnace

Table 1 summarises a comparison between a typical submerged-arc furnace at present and a furnace of the design proposed here.

Table I: A comparison between a typical furnace at present and the proposed larger furnace

	Typical existing 3-electrode a.c. furnace	Large 6-electrode d.c. furnace considered here
Electrode diameter, metres	1.6	3.0
Maximum electrode current, kA *	100	257
Burden resistance per electrode, mΩ *	1.0	0.533
Total furnace power, MW	30	211

*Note: According to the Andreae⁹ and Kelly¹⁰ formulae, the electrode current should be scaled by $d^{1.5}$ and the resistance by $1/d$, where d is the diameter of the electrode. We use these

formulae here without justification except to note that, while the predictions are probably not exact, any errors incurred are likely to be relatively small as the extent of scale up here is not particularly large.

Table 1 reveals a significant increase in power level with this proposed design, which suggests that the scale up of ferro-alloy furnaces to well beyond the present production levels should be a goal that the industry can reach.

CONCLUSIONS

1. The economy of scale creates an inexorable push towards larger furnaces in the ferro-alloy industry. However, the a.c. type of submerged-arc furnace as is commonly used at present has certain fundamental limits, and so conversions to other forms of furnace are likely. Once this 'a.c. barrier' has been overcome, scale up will probably involve larger electrodes and more electrodes. This could be applied to either plasma furnaces or submerged-arc furnaces. The potential for the scaling up of Söderberg electrodes will only be known once this a.c. barrier has been overcome and larger furnaces have been tried in practice.
2. The supply of electrical power to ferro-alloy furnaces is likely to become a more complex issue in the future because of constraints and other factors external to the industry, and because of the evolution of information technology. The ability to swing the furnace load is likely to bring significant economic advantages, and so might become common practice in spite of the nuisance factors that this will incur with the operation of the plant. This will also require more computer-based support for the supervision and control of the plant, as well as a tie-in to the so-called 'Smart Grid'. It would seem that larger furnaces might also be better than present furnaces at handling power dips and outages for a longer time, and d.c. arrangements could offer significant advantages when it comes to easy and fast swinging of the load.

REFERENCES

1. M. Dos Santos, Meeting the challenge of sustainability through technology development and integration in ferroalloy submerged arc furnace plant design, *Proc. 12th Internat. Ferroalloys Conf. / Sustainable Future* (Infacon 12, 6-9 June 2010, Helsinki, Finland), Outotec Oyj, Finland, 2010, pp. 71-80.
<http://www.pyrometallurgy.co.za/InfaconXII/071-DosSantos.pdf>
2. N. N. Taleb, *The Black Swan*, 2nd edition, Nassim Nicholas Taleb, 2007, 2010 (Penguin Books, London, England, 2010).
3. I. J. Barker and A. B. Stewart, Inductive reactance and the operation of large submerged-arc furnaces, *J. S. Afr. Inst. Min. Metall.*, Vol. 80, No. 3, Mar 1980, pp. 123-128.
<http://www.saimm.co.za/Journal/v080n03p123.pdf>
4. D. Sager, D. Grant, R. Stadler, and T. Schreiter, Low cost ferroalloy extraction in DC-arc furnace at Middelburg Ferrochrome, *Proc. 12th Internat. Ferroalloys Conf. / Sustainable Future* (Infacon 12, 6-9 June 2010, Helsinki, Finland), Outotec Oyj, Finland, 2010, pp. 803-814.
Reprinted in *J. S. Afr. Inst. Min. Metall.*, Vol. 110, No. 12, Dec. 2010, pp. 717-724.
<http://www.saimm.co.za/Journal/v110n12p717.pdf>

5. F. P. Greyling, W. Greyling, and F. I. de Waal, Developments in the design and construction of DC arc smelting furnaces, *Proc. 12th Internat. Ferroalloys Conf. / Sustainable Future* (Infacon 12, 6-9 June 2010, Helsinki, Finland), Outotec Oyj, Finland, 2010, pp. 815-823.
Reprinted in *J. S. Afr. Inst. Min. Metall.*, Vol. 110, No. 12, Dec. 2010, pp. 711-716.
<http://www.saimm.co.za/Journal/v110n12p711.pdf>
 6. M. S. Rennie, P. J. Brereton-Stiles, and N. Singh, Electrical energy management in the ferroalloy business, *Proc. 10th Internat. Ferroalloys Congr.* (Infacon X, 1-4 February 2004, Cape Town, South Africa), SAIMM, 2004, pp. 726-734.
<http://www.pyrometallurgy.co.za/InfaconX/076.pdf>
 7. J. Westly, Critical parameters in design and operation of the submerged-arc furnace, *Proc. 33rd Electric Furnace Conf.*, Iron and Steel Soc., Warrendale, USA, 1975, pp. 47-53.
 8. J. Westly, Resistance and heat distribution in a submerged-arc furnace, *Proc. 1st Internat. Ferroalloys Congr.* (Infacon), S. Afr. Inst. Min. Metall., Johannesburg, South Africa, 1974, pp. 121-127.
<http://www.pyrometallurgy.co.za/InfaconI/121-Westly.pdf>
 9. F. V. Andreae, Design and control of ferroalloy furnaces, *Trans. American Inst. Elec. Engrs.*, Vol. 69, 1950, pp 557-562.
 10. W. M. Kelly, Design and construction of the submerged arc furnace, *Carbon and Graphite News* (National Carbon Company/Union Carbide), Vol. 5, No. 1, April/May 1958.
-



Ian Barker

Visiting Professor, University of the Witwatersrand

Qualifications:

- B.Sc.(Eng.) — in Chemical Engineering, cum laude
- M.Sc.(Eng.) — in Metallurgy, studying froth behaviour in flotation cells
- Ph.D. — in Control Engineering , on the application of Kalman Filters

Career:

1. **Joined the National Institute of Metallurgy (NIM) 1975.**
Initially in the Process Development Division then in the Instruments Division, which later became the Measurement and Control Division. (*NIM was renamed the Council for Mineral Technology in 1981/82, and this was subsequently abbreviated to Mintek.*)
2. **Retired from Mintek 2008.**
3. **Joined the School of Chemical and Metallurgical Engineering, University of the Witwatersrand, 2008, as Visiting Professor.**

Technical/professional interests:

- (i) Mainly process control.
- (ii) Analysis of specific control problems on chemical and metallurgical processes and on electrical systems. Modelling, etc.
- (iii) Analysis of the control of power reticulation systems in the face of a variable supply from renewable technologies.
- (iv) Also inventing, developing, and commercialising new technology.
- (v) General interest in unusual phenomena.
- (vi) The interface between electrical and process technologies, including both electrochemistry and furnaces with their power supplies.

Some career highlights:

- (i) Inventor of the Minstral controller,
- (ii) Designed the power supplies for plasma furnaces,
- (iii) Recipient of a gold medal for work on furnace control,
- (iv) Co-inventor of the Blobulator (a type of granulator for making large lumps),
- (v) Inventor of the Arc Monitor or 'Arcmon'.

Hobbies:

Home workshop (electronics - both light current and heavy, metalworking, & woodworking).

Mountain biking.

Music.
