Which strategy to use when sampling anodes for coring and analysis? Start with how the data will be used

Keith A. Sinclair[1], Barry A. Sadler[2]
1 Sinclair Associates, Inc., 12620 Ridgepath Lane, Knoxville, Tennessee 37922 USA
(keith@sinclairassociates.com)
2 Net Carbon Consulting, PO Box 286, Kangaroo Ground, Victoria, 3097, Australia
(barry.sadler@bigpond.com.au)

Keywords: Anode, core, sampling strategy

Abstract

Most Smelters in the world take core samples from anodes for a range of purposes, such as:

- Characterizing product quality going to the Potrooms
- Anode performance troubleshooting
- Analysis of raw material, Green Carbon (GC) or Carbon Baking Furnace (CBF) performance problems
- Benchmarking anode properties against other plants

A number of different sampling strategies are used to select anodes for core drilling and analysis of baked anode properties. However, in many cases, the sampling strategy is not properly designed to yield the information desired.

This paper will outline appropriate sampling strategies and also discuss the issues associated with sampling to be considered when designing a sampling strategy suitable for your intended purpose.

Introduction - why are baked anode properties measured?

In the design of systems for anode core sampling and testing, the most important and first question to be answered is: Why do we actually measure anode core properties? E.g. how do we intend to use the information to control and improve the process? This is particularly critical, not only because of the need to get value from the significant cost of sampling and testing, but because any misalignment between the core sampling strategies used and the intended use of the data can lead to incorrect decisions and actions regarding product or process issues.

Data on baked anode core properties (e.g. Baked Apparent Density, Air Permeability, Electrical Resistivity, Flexural Strength, etc…) are commonly used to:

1. Identify and troubleshoot problems in Carbon Baking Furnaces (CBF's), Green Carbon (GC), and with Raw Materials (RM's).
2. Monitor on-going process behavior and performance.
3. Characterize the properties of the baked anodes sent to the Potrooms – to give confidence that what we are sending them satisfies their specifications.
4. Assess the effectiveness of process improvement actions.
5. Benchmark anode quality between plants.

In order to generate appropriate data that will satisfy each of these uses, it is necessary to use suitable strategies to select the anodes for core sampling. This will now be explained further.

What sampling strategies are used to select anodes for coring?

A variety of sampling strategies have been observed by the Authors for selecting the anodes to be cored for testing:

a) No strategy – select next available anode or anodes in sequence.
b) Structured – select anodes from defined locations within each CBF section as it is unloaded. Based on prior testwork or technology supplier recommendations, these locations are often selected to represent the hottest, coldest, and "average" baking temperature locations in the section.
c) Structured – select anodes from defined layers (e.g. Top, middle, bottom) within each CBF section, as it is unloaded.
d) Purely Random – select anodes from each CBF section as it is unloaded but from random locations within the section. Random selection may also involve simply selecting random anodes off a conveyor.
e) Structured Random – set up a "random" sampling plan that over time selects anodes from the CBF's for coring at a rate representative of the weighted proportion of anodes produced in that "category" of location in the furnace (e.g. straight section, inside pit, middle layer, middle anode). This strategy is rarely used as it quite complex to follow, not having a simple "rhythm" to guide actions.

The rationale of keeping sampling simple can lead to the fundamental issue encountered in most smelters - the sampling strategy for capturing anode core property data is not designed for purposes of answering questions, but rather to make the cumbersome process of sampling "as simple as possible." The result is that data are used for purposes for which it is patently wrong to do so. This can lead to poor quality decision making with a focus on issues that are not significant, and significant issues being ignored. The most common mistake seen by the Authors is that people who monitor and analyse anode core test results automatically believe that they are looking at data that statistically characterises the anode population sent to the customer. This assertion is simple to test – ask Potrooms "what do you think the core test results sent to you by the Laboratory or Carbon Plant represent?" The most common response is, "The quality of the anodes the Carbon Plant makes and delivers to us - why else would they send the results to us?" Unfortunately, core test results are rarely based on a sampling strategy that will make this understanding true.

The problem of the sampling strategy not providing data to answer the questions that are being asked is not limited to the issue of "characterizing the anode quality." It can also extend to issues such as benchmarking exercises. Done properly, these

types of comparative analyses require that the sampling strategy to be used (e.g. random sampling of a significant number of anodes) is clearly defined and then followed by each benchmarking partner. Only then can differences in anode qualities (as indicated in the data) be attributed to differences in the process capabilities at the benchmarking operations. If the sampling strategies applied at each operation are profoundly different, it is equally possible that the differences in the data could be a function of different sampling strategies (false indications in the data). The danger here, of course, is when these data are mistakenly used to justify process improvements as a result of the study.

Once we have an appropriate sampling strategy, there are further issues with anode coring, such as the actual location of coring in the anodes, that impact on the quality of core test results. These will be covered later.

What are the sources of variation in the process and core test data?

To be able to use anode core test results for any (or all) of the purposes outlined in 1 – 5 above, the different sources of variation in the data (i.e. in the process that generated the data) must be recognized. For example, three important sources of variation have been identified in the anode baking process. In designing the sampling strategy to meet the intended use of the information, we will need to consider all three of these:

- Random or common cause variations that are at work in the process all the time. These include normal variation in pit bake out temperatures. There will also be random variations in the raw materials used and in the green anode production process.
- Special cause variations resulting from changes or external influences that occur and are unpredictable in their timing, effect, magnitude or duration. As an example, this would include the effect on anode properties of an unplanned loss of furnace draft for an extended time, or in the case of raw materials, a significant change in properties from batch to batch.
- Structural variation associated with inherent differences in the process such as different baking conditions at various locations within the furnace: e.g. straight sections versus sections near the crossovers, or top layers versus bottom layers (See figure 1 below). If disregarded in the design of a sampling strategy, structural variation can appear as special causes (and more specifically, special causes that you can do little about), making the use of the data for the intended purpose more difficult and less insightful.

A discussion of anode sampling strategies.

There is a common thread within the reasons why we expend resources to core and test anodes - we are looking for signals that a significant change in anode properties may have occurred. There are fundamentally four ways in which we can improve our ability to identify such a signal, all of which assume we are using properly designed control charts to monitor the process and detect changes in process behaviour. In other words, when using properly designed control charts – we can increase the sensitivity of these charts to changes in the process by:

i. Increasing the frequency of sampling (this will reduce the time span over which you need to collect the number of groups of samples ("subgroups") to detect a shift in anode properties)
ii. Increasing the number of samples taken at each point in time (increase subgroup size).
iii. Using a sampling strategy to isolate the sources of variation and as a result, tighten up the control limits (Limits shown on control charts that are calculated from the data and represent the upper and lower limits of normal, or common cause variation, in the parameter being measured) making the data (as shown on the control charts) more sensitive to anode property changes and able to identify significant changes faster.
iv. Reducing the variation in the parameter by improving the process.

In this paper we will examine how we can apply option (iii) to our best advantage. If this is insufficient to give us the data we need to meet our purposes, then we would consider increasing sampling (frequency first, and then the number of samples taken in each subgroup). We always recommend that option (iv) is part of your efforts, on an ongoing basis.

Figure 1. Schematic view of a baking furnace (from Riedhammer)

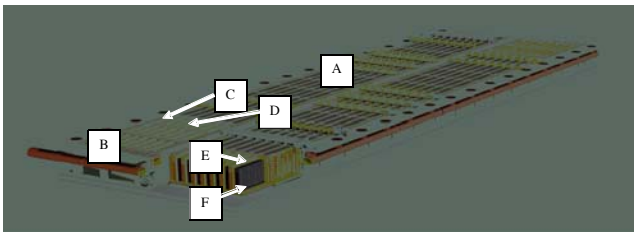

showing sources of structural variation: Straight sections (A) versus Crossover sections (B); Pit location – Outside (C) versus Inside (D); and Layer within a pit – Top (E) versus Bottom (F).

Sampling strategies for answering specific anode property questions

If the behavior of the baking process (and raw materials and green anode production) resulted in a population of anodes with homogenous properties, many of the issues associated with sampling and the usefulness of the data would be resolved. But, we do not have a homogenous population. This means that we must carefully design our data collection strategy (sampling strategy) based on the following:

- What questions are you trying to answer with the data?
- What sources of variation do you want to expose through the analysis of the data?
- What is the underlying structure of the data?

With these in mind, we can then set out a sampling plan to address these issues.

<u>To study the effect of Green Carbon and Raw Materials on Baked Anode properties:</u>

If the intent of coring anodes is to study the effect of changes in Raw Materials (RM) and the Green Carbon (GC) process on Baked Anode properties, the appropriate sampling strategy requires that we "block" (to the extent possible) the variation in anode core properties associated with the baking process By doing this we can isolate RM and GC effects. In this strategy we are NOT trying to address the question related to the effect of baking process changes on anode core properties, or to characterize the

distribution of anode properties sent to the Potrooms. If the sampling strategy is designed specifically to isolate the effects associated with GC/RM, then we must recognize that this actually prevents the identification of changes in anode properties associated with the baking process, i.e. questions regarding the effect of CBF changes cannot be answered.

This strategy requires that we restrict the sampling regimen to anodes only from a CBF section that has the demonstrated capability to produce anodes of acceptable properties (This should be determined using data) and also to set anode locations within this CBF section. This effectively blocks out the structural variation associated with differences between sections or location within the CBF sections.

The primary benefit of this approach is that the data (chart) becomes more sensitive to upstream changes (i.e. in RM or GC) because the structural variation associated with the baking process has been "blocked."

A weakness in this strategy is if you are only able to sample anodes from the selected section every 20 days (or so), then any changes in RM or GC must be of at least this duration for the effect to be detected. Changes that come and go within this time window will be almost impossible to detect so we will need to think about what RM/GC changes we are actually looking for when designing this sampling regimen. Should the effects we are looking for be characterized by short duration, then we need to consider adding another section(s) to the sampling regimen. Historical data should be examined to find two or more sections that behave similarly over time and that give similar average and standard deviation results for the properties in question. This approach is a more restrictive variation on sampling plans (b) and (c) mentioned previously.

This sampling strategy also requires a system for traceability of anodes from Green Carbon (including Raw Material details) through the Baking process, such as a date or anode number stamp. The Green Carbon data would be captured and affixed (in the database) to this anode designation. The baked anode properties that are to be traced back to Green Carbon are measured and linked to the Green Carbon data in the database.

<u>To study the effect of the baking furnaces on anode properties:</u>

If the goal is to study the effect of changes in the baking process on Baked Anode properties the approach is very similar to that used to isolate Green Carbon effects. The issue in studying baking effects is one of providing sufficient data to detect changes in a reasonable time frame at an economic sampling level, while at the same time blocking out any structural variation. A straightforward way to deal with this is to stratify the data based on the sources of structural variation. In this discussion, the structural sources of variation as shown in figure 1 are considered to be:

- Inherent differences in anode properties associated with bottom versus middle versus top anode layers within a pit,
- Inherent differences from one pit to another within a section depending on whether the pit is on the outside or in the middle of the section ("inside" pit), and
- Inherent differences from one section to another such as represented by straight versus crossover sections.

In this discussion we will not specifically address the issue of pit location variation (i.e. inside (C) versus outside pits (D)). If however, tests show that there are significant differences in anode properties associated with this source of structural variation, the issues to consider in designing a sampling strategy would be the same as the following.

There is no one right answer to the design of this sampling strategy. Just as in the case of sampling to detect GC/RM effects, the process of designing a sampling strategy starts with the questions we want to answer, the variation sources we wish to expose and the underlying structure of the data.

Perhaps the most rigorous outcome of this approach involves a sampling strategy that would prevent differences between layers from being included in (or "getting inside") the sampling subgroup as well as blocking the effects of RM/GC changes (so you must look inside RM batches and inside GC operating windows). For this strategy to be effective it requires careful planning and forward thinking when loading the anodes into the sections to be sampled. Anodes to be cored must be identified early in the process - coming from the same RM batch and GC run. They are coded and loaded into set locations within the furnace (that cover the CBF structural variation) and then stratified to block out the CBF structural variation.

For example, a subgroup of anodes from a specific RM batch and GC run might be made up of multiple anodes taken from a single pit within only the middle or the top or the bottom anode layers and from only set anode positions. By doing this, we improve the sensitivity of the control chart to picking up shifts in the mean of a given anode property that may be associated with changes occurring in the CBF process.

<u>To characterise the population of anodes delivered to Potrooms:</u>

If the goal is to characterize the properties of the population of anodes delivered to Potrooms, the sampling strategy is quite different. The preceding sampling strategies have specific applications when investigating aspects of the anode production process, however, as discussed previously, the most common answer given to the question "why do we core anodes and test the cores?" is so that we "know the quality of the anodes we have produced and sent to Potrooms." While this is undoubtedly important, we must consider how these data are to be used before heading down this path. The ability of a typical anode sampling, coring and testing program to deliver on this objective is limited by the cycle time to go from unloading a baked anode to be cored from the CBF to getting the core test results back from the Laboratory. It is not unusual for this cycle time to be 2 weeks or more. This means that often the anodes are in the cells and half consumed before the core test results are received. Even in plants with a short cycle time for core testing (e.g. 1 week), the anodes are still usually in the cells before the results are obtained – too late to take any meaningful action in Potrooms based on the test results. This casts some doubt on if it is possible to quickly associate changes in anode properties with changes in cell performance. Further, without traceability linking anodes in the cells, back through the Rodding Room/CBF and further to GC or RM (i.e. without a comprehensive anode tracking system), then even if changes in cell performance could be attributed to changes in anode properties, the investigation cannot be easily taken further.

In most cases, this leaves us with the only appropriate use of data from anode core testing in a way that looks at the overall anode population to be largely historical – "was there a longer-term trend or change in anode quality that could have contributed to the change in cell performance (that we have already seen)?"

So what might a sampling strategy designed for this purpose involve?

A. Start with an analysis of the population by sampling all areas of the CBF's in a manner that ensures you capture and link each anode sample to a specific location in the CBF's.

B. Stratify the data by "structural" variation source, such as straight sections versus crossover sections. Examine the data, test for differences in the means and variances. If the "structural" differences are significant and important, consider this as two different sub-populations within the whole.

C. Repeat this examination for all relevant sources of structural variation (pit to pit, layer to layer, etc.)

D. With the different "sub-populations" identified, determine the percentage of the total population associated with each.

E. Establish a random sampling plan within each sub-population at a frequency that reflects the relative contribution of anodes from this group to the whole.

For example, if test work shows that the properties of the anodes baked in the first section coming out of the crossover are significantly different from the total population, then (in a 36 section furnace), 2/36 (i.e. about 6%) of all of the anodes cored should be randomly sampled from these sections.

This means we have a structured sampling plan (across different sub-populations) that is also random (within each of these sub-populations), and again structured – in terms of the number of samples taken from each sub-population.

Clearly, this is a rigorous approach that requires a commitment to sample enough anodes to be able to characterise not only the sub-populations, but also to capture enough data to make a reasonable inference regarding the total population. The experience of the Authors suggests that this will often require an increase in sampling from the current levels observed at most smelters.

If done properly, with this strategy there is also the possibility that by stratifying the data as outlined previously; the sub-populations may also be characterized. This will effectively enable you to make statements not only about "the properties of the anodes sent to Potrooms" but also to watch for changes associated with RM/GC or CBF.

This approach provides significant advantages in sensitivity to changes in the process as well as providing the best approach to characterizing the population – when contrasted with the purely random approach to sampling the full population.

An example of this sampling strategy is shown in figure 2:

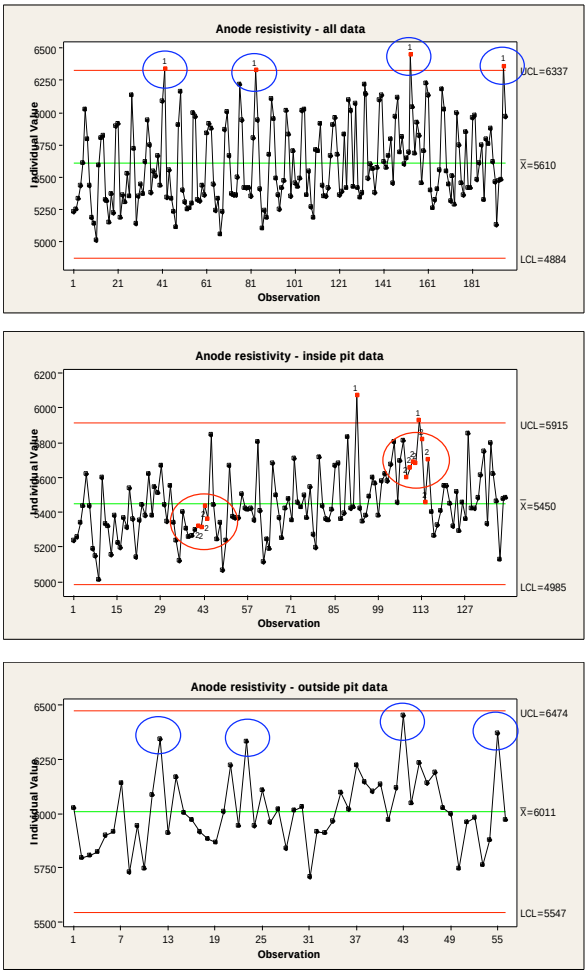

Figure 2. In these graphs we see data presented from a structured sampling plan for Resistivity (ELR) - all anodes taken from straight sections within the furnace, with sub-populations of inside and outside pits. Layer to layer variation is ignored for simplicity in this example. Random sampling is applied within each of the pits. Pits are sampled at the proportionate frequency; i.e. the furnace has 7 pits per section so every 6th and 7th core is from an outside pit. This outside pit data is stratified to reflect the view that it is a different sub-population of anodes. In the top "all data" graph, several special cause signals appear (circled), however, upon stratifying the data by inside and outside pits, we observe that these 4 points are all associated with outside pits, and therefore, more likely are false signals of special cause variation on the "all data" chart. While they are most likely just the result of structural variation, these points still appear somewhat "different" even on the "outside pit only" chart (bottom), suggesting that there may still be evidence of an external factor influencing ELR.

When we properly stratify the data to reflect the behaviour of a single population (e.g. inside pits only as shown in the middle graph), we observe that the data and the chart is clearly more sensitive to RM/GC effects. This reflects in our ability to detect a low run around point 43 as well as a high run around point 110 that were not evident in the "all data" chart. (These special causes were later found to be associated with changes in raw materials).

Additional considerations in anode core sampling.

Once the sampling strategy to be used has been agreed, and the limitations the strategy imposes on what decisions can be made with the core test data are understood by all users of the data, there are still ways that the value of anode core data can be reduced. These include:

- Not following your sampling strategy, ALL OF THE TIME.
- Not core drilling anodes in a way that captures all properties of interest.
- Not fully testing anode cores for properties that are important now.

Compliance to the sampling strategy: An additional problem arises when sampling procedures define a structured sampling strategy that is followed most of the time but not always (e.g. not on night shift – this is quite common!). This means that the data is a mixture of sampling strategies and has little value in helping with decision making for troubleshooting or process improvement - you cannot be sure if shifts in the data are due to real changes, or just deviations from the defined sampling strategy. To be confident of anode core test data, the agreed sampling strategy must be followed ALL OF THE TIME.

Anode core drilling practises: In the same way as there are currently a number of different ways anodes are selected for coring, there are different ways that cores are drilled from anodes. These vary from only coring the first 30cm from the top or bottom of the anode, to taking cores right through the anode ("through coring"). Which of these approach is used has a very strong influence on the properties of the anode core (and hence the test results) and on the ability to detect anode defects. It is known that anode properties such as Density, Flexural Strength, and Electrical Resistivity normally vary within each anode in a generally predictable way top to bottom, and that "macro" properties such as significant internal cracking (a particularly insidious and damaging anode defect) also occur more frequently in some locations within anodes than in others. It would seem to make sense to core drill anodes in a way that captured the full range of anode properties that we are interested in. This means taking through cores. An additional advantage of taking through cores is that they are long enough to allow multiple measurements of key anode properties giving information on the vertical profile of anode properties. An example of how only coring the top 30cm of an anode can miss important anode defects such as anode cracking is shown in figure 2.

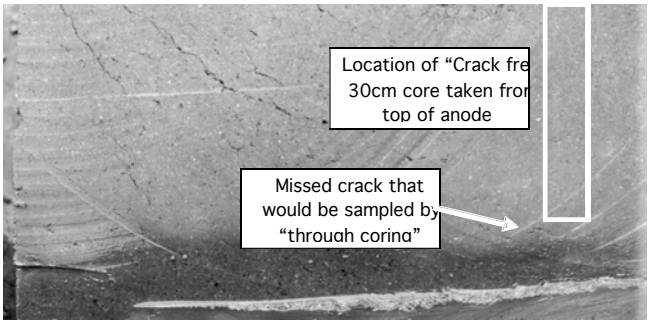

Figure 2. Section of an anode cut vertically with a concrete saw showing significant internal cracking that would be missed by just sampling the top 30cm of the anode. The test results from this core will not give results representative of the anode and a major anode problem will go undetected.

Relevant anode core testing: Just to confuse the current situation for anode core testing even further, there are different approaches used to determine which tests to conduct on the cores sampled from anodes. Some plants have a standard set of tests that are conducted on every core. Other plants have different sets of tests that are used for a proportion of the cores taken. These different sets of tests usually focus on anode characteristics such as physical properties, thermal shock characteristics, or anode reactivity. To ensure that anode core test results provide maximum value, it is important that the core tests undertaken adequately cover anode properties of current interest. For example, at present it is not unusual to find Laboratories still undertaking core tests that largely focus on properties related to anode thermal shock, despite the particular smelter using anode slots and not having experienced thermal shock problems since they were introduced. These anode core testing protocols were established before the use of slots and have not been changed to reflect the shift in anode quality issues. It is recommended that the tests undertaken on cores track all relevant anode properties, but focus on current concerns. In many cases, this means significantly reducing the amount of testing done of thermal shock related properties and reallocated the resources used for this toward the testing of anode density and reactivity properties, issues that are becoming increasingly important with changes in coke quality.

Conclusions

The issues surrounding the appropriate use of data generated from core samples are important in both process control and the assessment of process improvements. There is no single "correct" sampling strategy. Rather, there is a correct approach to developing the strategy that is right for each site. This sampling strategy must be designed to:

1. Answer your specific questions regarding the process:
   - Do you want to learn about changes in Raw Materials and/or Green Carbon performance?
   - Do you want to learn about changes in Carbon Baking Furnace performance?
   - Are you trying to characterize the "quality of the anodes sent to the Potrooms?"
2. What sources of variation are you trying to expose through the data?
3. And finally, what is the underlying structure of the data?

With these questions answered at the start, you can now design a proper sampling strategy making appropriate decisions regarding "your ability to learn from the data" and the "cost of data capture." Regardless of the final design of the sampling strategy, the value of the data can still be significantly diminished by a lack of discipline in sampling execution or through variation in coring methods. All these factors can impact the "value we get from the data we collect."

A structured + random sampling strategy is recommended to provide core testing results that represent the quality of the anode population sent to potrooms; these data can be used to answer in a historical sense, the question - "were there any longer-term trends or changes in anode quality that could have contributed to observed changes in cell performance?"