ApplicationNOTE



Understanding iR Compensation

Introduction

Some of the most common technical questions we hear at Gamry have to do with iR compensation:

- Where does uncompensated iR come from?
- Do I need to use iR compensation with my experiment?
- How should I set up the iR compensation parameters?
- Why doesn't iR compensation work on my system?

In this Application Note we shall attempt to answer these questions and to give you a basic understanding of iR compensation.

Some initial background information describes the general issue of iR errors. The later information concentrates on the specifics of "current-interrupt" iR compensation as a means for measuring and correcting for iR error. Positive-feedback iR compensation is mentioned but not discussed in detail.

This Application Note presumes that you have a basic understanding of potentiostat operation. If not, please review our <u>primer on potentiostats</u>. Experienced potentiostat users should skip the primer and read on.

It would also help if you know some of the fundamentals of electrochemical impedance. We have a <u>primer on electrochemical impedance</u> on our website. Pay particular attention to how typical chemical processes are mapped into electrical circuit elements.

When Do I Need iR Compensation?

We give an approximate answer to this question here. A more complete answer requires the information discussed below and some information about the system that you are testing.

In general, iR compensation is needed when one or more of these is true:

- You are doing a quantitative test that yields a numerical result, such as a corrosion rate, equilibrium constant, or a rate constant.
- The solution in your cell is not very conductive.

- Your currents are fairly high.
- Your cell geometry is less than ideal.

Unfortunately, these criteria are subjective. For example, aqueous 0.5 *M* potassium chloride may be considered to be very conductive in an electrochemical analysis application and poorly conductive in a plating application.

One simple rule-of-thumb is often used:

Record some initial data-curves with and without iR compensation. If the shape of the curve changes significantly when iR compensation is applied, compensation is required.

IR compensation often adds additional noise to the data, so increased noise on the plot is not considered to be a significant change.

What Is the Origin of iR Error?

Let's look at a typical three-electrode electrochemical test cell. We have labeled some reference points in the cell to which we shall refer throughout the remainder of this Application Note.



Reference points:

- A: Counter electrode output at the potentiostat
- B: Metal surface of Counter electrode
- C: Electrolyte surface of Counter electrode
- D: Electrolyte at the tip of the Luggin capillary
- E: Electrolyte surface of the working electrode
- F: Metal surface of the Working electrode
- G: Electrolyte surface of the Reference electrode
- H: Reference electrode input at the potentiostat
- I: Working electrode output at the potentiostat

We can also think of the cell as a (simplified) network of electronic components something like this:



Almost any potentiostat will do a good job of controlling and measuring the voltage between points H and I.

Unfortunately, we really want to control and measure the voltage between points E and F. This is the potential difference across the electrochemical interface we are trying to study:

$$V_{\rm important} = V_{\rm f} - V_{\rm e}$$

For purposes of this discussion, point I is equivalent to point F. Point G is equivalent to H, except for a constant offset-voltage caused by the difference in potential between Working Electrode/Reference Electrode, also known as the open-circuit voltage, V_{oc} . And because there is no current flowing through the reference, the potential drop across R_{bridge} is 0 and point G is equivalent to D. So starting from:

$$V_{\text{measured}} = V_{\text{I}} - V_{\text{H}}$$

we can get to:

$$V_{\rm measured} = V_{\rm f} - V_{\rm d} + V_{\rm oc}$$

We are one step from finding the reaction potential, $V_{\rm f} - V_{\rm e}$. Point E is equivalent to point D except for that nuisance resistor, $R_{\rm u}$.

How is it related? By Ohm's law:

$$V_{\rm e} - V_{\rm d} = I_{\rm cell} \times R_{\rm u}$$

So substituting this in gives:

$$V_{\rm measured} = V_{\rm f} - V_{\rm e} + I_{\rm cell} \times R_{\rm u} + V_{\rm oc}$$

Rearranging the result gives:

$$V_{\text{important}} = V_{\text{f}} - V_{\text{e}}$$

= $V_{\text{measured}} - I_{\text{cell}} \times R_{\text{u}} - V_{\text{e}}$

We can measure the voltage, V_{measured} . We measure the current, I_{cell} . We can look up or measure the open circuit voltage, V_{oc} . But without knowing R_{u} , we can't find $V_{\text{important}}$!

And this question is at the heart of iR compensation. iR stands for $I_{cell} \times R_u$.

Luckily, it is possible to measure R_u and correct for it, if it is a problem.

What Kinds of Cells and Electrochemical Systems Do We Need to Worry About?

Most of the time, electrochemists can rig their experiments so that iR drop is not a problem. One simple way is to add a non-reactive salt, acid, or base (supporting electrolyte) which increases the electrolyte's conductivity.

Conductivity up $\rightarrow R_u$ down \rightarrow iR error down

Another way is to minimize the distance between the tip of the Luggin capillary and the working electrode. Just design the cell so that the distance is tiny.

Obviously, if it were that simple, we wouldn't be getting so many questions.

Adding a supporting electrolyte will affect the electrochemistry even if the ions don't directly become involved in the reaction. It will change the composition of the double layer (C_i). It may affect the solubility or structure of the reactants and products. It may change the structure of the surface atomic layers. In many cases, it may be important not to monkey with the electrolyte. For example, the corrosion chemist wants to study the corrosivity of "Gloop," not "Gloop" plus salt.

Likewise, the solution to cell design often is not a solution. Many cells designs are limited mechanically. A subtler problem is that putting the reference electrode

too close to the working electrode surface changes the current density at that surface, changing the very measurements we're trying to make.



No surface effect -- Current density lines have room to adjust to probe



Reducted current density effect -- Current density lines can't adjust to probe

So do you need to worry about iR error?

Yes, if your electrolyte isn't very conductive or your reference probe is far from the reacting surface. Often you can measure iR quite easily.

How Is iR Error Measured?

The following circuit has some important clues for measuring iR, and from it, R_u . It represents a common, simplified model of the electrical behavior of an electrochemical cell.



Model of Randles cell.

Look at R_{faradaic} . It has a capacitor, C_{faradaic} in parallel with it. R_u doesn't.

This should suggest that an AC experiment can differentiate between these two resistors. High-frequency signals pass right through $C_{\rm f}$ with no voltage drop, while they are forced to drop $I_{\rm cell} \times R_{\rm u}$ volts going through $R_{\rm u}$ just like low frequencies.

Indeed, if you record the Electrochemical Impedance Spectrum of this cell, its Bode representation looks like the following:



At low frequencies where C_f is effectively an open circuit, the measured impedance is the sum of R_u and R_f . At high frequencies where C_f is effectively a short circuit, the measured impedance is R_u .

So if you need to worry about R_u , you can measure it and multiply it by the cell current.

The resulting error voltage is the uncompensated iR, $V_{\rm e} - V_{\rm d}$. If iR is smaller than a few millivolts, don't worry about it.

For example: Suppose from the impedance spectrum, $R_u = 100 \Omega$. Suppose $I_{cell} = 10 \mu A$.

$$iR = 100 \times 10 \times 10^{-6}$$

= 1 mV.

For most electrochemical phenomena, that is a small error.

Another way to think about it is: If $R_f >> R_u$, don't worry about R_u .

Measuring R_u Using a DC Technique

Measuring R_u using AC impedance is good in theory, but sometimes we need a quicker, less expensive way to make the measurement. And very often, we want to make the measurement while we are doing something else, such as recording a current-versus-voltage curve. Luckily, there is an equivalent DC method. Actually a better term for it is a "transient" technique. It is known as Current Interrupt iR Compensation.

Again, consider the simple Randles cell model for an electrochemical reaction with solution resistance.

Gamry Instruments' potentiostats contain a circuit that can rapidly turn off (interrupt) current flow through the cell, wait a short time (10 μ s to 30 000 μ s), and then turn the current back on. Many other commercial instruments have the same capability.

To make the current interrupt measurement, cell voltage (V_{measured}) is measured immediately before and immediately after the current interruption. Ideally, the measured voltage of this cell would look like the graph:



Current Interrupt Waveform

Suppose we are measuring 1.0 V while current is flowing. At time 0 we interrupt the current – very quickly! The voltage immediately drops by the voltage drop across R_u . Then it starts dropping more slowly.

Why?

The decrease in voltage is caused by the slow discharge of the faradaic capacitor through R_{faradaic} . This only becomes important at longer times. At short times, the capacitor can hold the voltage to $V_{\text{measured}} - V_{\text{u}}$ which is our $V_{\text{important}}$. And this is what we seek.

As usual, things aren't so simple. Three problems interfere with this simple model:

- Sampling speed
- Output capacitance
- Noise

Sampling Speed

One problem with the Idealized Current Interrupt waveform is the speed of sampling required. In the above figure, the sampling is 2 μ s (very fast). We can slow down the sampling considerably if we pretend the decaying curve is a straight line and extrapolate it back to the turn-off time.

Let's try that with the same Randles cell we've been using, so we sample at 1 ms and 2 ms, and extrapolate back to 0 ms when the switch was turned off. You get something like this:

3 Point Extrapolation Method



The estimated V_u is calculated by

 $V_{\rm u} = V_1 + (V_1 - V_2)$

In this case, $V_1 = 0.671$ V, $V_2 = 0.481$ V, and by the straight-line extrapolation, V_u is estimated at 0.862 V.

But wait: V_u is actually closer to 0.938 V. We were misled by picking too slow a time base. You can see the curvature in the turn-off trace in the picture. Of course it's curved: it's an exponential decay.

By the way, these are actual numbers from our original Randles cell model, calculated in Mathcad[®] and plotted in Excel[®].

So we've gone too far. Millisecond timing is too slow for iR measurements on this cell. Something faster (but not too fast) ought to work.

How do you tell?

Mathematically the decay time constant r is $R_{\text{faradaic}} \times C_{\text{faradaic}}$. In this cell,

$$\tau_{\text{faradaic}} = 3000 \ \Omega \times 1 \ \mu\text{F}$$

= 3 ms

If you know, roughly, these values, you can pick a short time τ to be *RC*/10. Or you can sneak up on the correct answer by decreasing t_{sample} until the number stabilizes.

But there's another problem lurking as t_{sample} gets shorter: Potentiostat output and cable capacitance.

Cable Capacitance

Think back to the original cell model:



That innocent-looking capacitor, C_{cable} , can cause a lot of problems. If you have a typical shielded cable, the value of C_{cable} can be 50 pF per foot (1.6 pF/cm). For a 5-foot (152 cm) cable that's 250 pF. Add another 100 pF or so for capacitance in the switch itself, on the circuit board, and in the driving amplifier.

We can use this circuit as a model:



The cable capacitance forms an *RC* section with R_u and $R_{solution}$. This means the voltage across R_u doesn't disappear infinitely fast.

For purposes of this discussion, we must assume the counter electrode's capacitance is big and acts as a short-circuit at these time scales. Fortunately, it's a reasonable assumption.

Suppose you decided to set iR sampling at 50 μ s and 100 μ s. These points are shown as the open diamonds on the following graph:

Current Interrupt Waveform with Cable Capacitance Effects



Using those two measurements, obviously the iR estimate is again highly inaccurate. You must wait until the cable-capacitance transient has died out before you make the measurement.

It may be helpful to see this on a log(time) scale so you can see the cell-cable and faradaic capacitors being discharged.



Current Interrupt Waveform with Cable Capacitance Effects

You need to find a time-range between the two limiting discharge curves.

The cable capacitance must be fully discharged, but the faradaic capacitance must still be in the approximately linear region. If your faradaic capacitance is not much, much larger than the cable capacitance, current-interrupt iR compensation *will not work*.

Noise

Noise can be a major problem when current-interrupt iR compensation is used on a real-world system.

Basically, current-interrupt iR compensation is a differential measurement. Remember the equation we use for estimating V_u :

$$V_{\rm u} = V_1 + (V_1 - V_2)$$

The differential term, $V_1 - V_2$, is very sensitive to noise.

In systems suitable for current-interrupt iR compensation, the difference between V_1 and V_2 is small, from a few millivolts to a few hundred millivolts.

Suppose there is a positive noise contribution to V_1 and a negative noise contribution to V_2 . The average noise is 0 but the error in V_u is twice as big!

You might say "So what? Just turn on some filtering. The noise will go away." But we are trying to measure a rapid $(10-1000 \ \mu s)$ phenomenon. We cannot insert a 5 Hz filter, or the whole transient gets distorted—if not disappears altogether.

All of these things can help to reduce the noise:

- Use a Faraday cage to keep extraneous noise out of the measurement.
- Use signal-averaging to make the noise terms average out while keeping the real values.
- If your noise source's frequency is known, use a synchronous sampling method so that all the noise causes errors in the same direction.
- Finally, if the noise is still too large, don't use the extrapolation method. Just settle on an average, such as:

$$V_{\rm u} = (V_1 + V_2)/2$$

The noise situation is worse when you're trying to measure low currents. Under these conditions, when the current-interrupt switch is opened, the Reference electrode and Working electrode pick up more noise than under higher-current conditions.

How Is It Corrected?

So far we have only discussed measuring iR. If we know the value of V_u we can subtract off the V_{measured} value and get $V_{\text{important}}$. This is known as Post-Processing Correction.

One problem with post-processing correction is that we cannot predict the true potential prior to applying it. This is particularly problematic in experiments that sweep the potential. In such experiments, the scan rate across potential is not constant and the scan limits may be wildly inaccurate.

We would prefer that the potentiostat measure and correct for V_u continuously. After all, when you want to apply 1 V to a cell, you intend that $V_{important} = 1$ V, not $V_{measured} = 1$ V.

The situation is much simpler when you use a galvanostat. A galvanostat does not have to apply a correction for V_u , for its job is to control current, not voltage. It still has to measure V_u .

The simplest, although not the most effective, way to automatically correct for iR using Current-Interrupt is to have the potentiostat add its best guess for V_u onto the applied signal. This can be stated as a formula where the indexes in brackets represent measurement points:

$$V_{\text{applied}}[i] = V_{\text{requested}}[i] + V_{\text{u}}[i-1]$$

Initially, there is no correction. The measured error from the first data point is then added to the applied potential of the second data point. This process continues, with the correction becoming more accurate as data are accumulated.

Note that this is a dynamic correction. R_u can change during an experiment, and the system automatically compensates for that change.

Control-loop Algorithm

Direct feedback of the error voltage into the next data point is not the most sophisticated way to apply the correction. A better approach to understanding the feedback mechanism is to treat iR correction as a control loop.

The control-loop algorithm treats the potentiostat as a loop within a loop. The inner loop is the potentiostat itself, which measures V_{measured} and controls it in a feedback mechanism. This loop is made purely from analog electronics as shown:



(We have eliminated a few components that are irrelevant to this part of the discussion.)

The potentiostat is a control loop. It measures V_{measured} and compares it to V_{applied} , making corrections to the counter electrode voltage until the difference between the two is 0. All of this happens continuously.

The iR correction also happens in a control loop outside of the potentiostat loop. It is shown in this figure:



The outer loop does a very similar job to the inner loop but it is implemented digitally in a computer. Its job is to see that $V_{important} = V_{applied}$, and now the job of the inner loop is to see that $V_{measured} = V_{actual}$. V_{actual} is a new value that comes from the outer loop.

We also show an interesting block that produces V_{actual} . It is known as a gain block or controller block, depending on to whom you talk. Its output is given by the expression:

$$V_{\text{actual}} \ = \ K_{p} \ \cdot \ error \ + \ K_{d} \ \cdot \ \frac{d \ error}{dT} \ + \ K_{i} \ \cdot \ \int error \ \cdot \ dT$$

This is known to control engineers as a PID loop. Each of the gains is individually controllable. By adjusting the gains, we can get iR compensation to perform better than the simple feedback algorithm. In practice we tend to use the integral control for iR compensation.

At Gamry Instruments, we commonly use a control-loop algorithm in our DC software. It is set using some lines from our experiment control language, Explain[™], that look like this:

if (IRToggle)

```
Pstat.SetIruptMode (IruptClfg,
RuExtrap,IruptTime, POTEN.Eoc (),
0.8)
Pstat.SetVchFilter (100000.0)
else
```

Pstat.SetIruptMode (IruptOff)
Pstat.SetVchFilter (5.0)

The value, 0.8, is the control-loop gain, K_i . K_p is fixed at 1.0, and K_d is 0. Note that the voltage channel filter (VchFilter) is set to pass 100 kHz signals. If we're not using iR compensation, the filter is set to knock out as much environmental noise as possible.

One can modify the parameters including the controlloop mode, V_u calculation, current interrupt timing, and gain to suit the dynamics of the reaction and cell under test.

Advantages of Current-Interrupt iR Compensation

Current-interrupt iR compensation offers some advantages compared to other iR compensation methods. These include:

- No prior knowledge of *R*_u is required
- *R*_u can change during an experiment without creating errors in compensation
- Compensation is independent of the current range used to measure the current, so it works in autoranged experiments.
- Scan parameters such as ramp limits and scan rates are corrected automatically.

Practical Limitations of Current-Interrupt iR Compensation

Current-interrupt iR compensation works well on some electrochemical systems, but fails to work properly when applied to other systems. The failures are generally explainable in terms of the practical limitations of the technique.

The limitations include:

- Need for a large faradiac capacitance
- Time-per-point limitations
- R_{faradiac} should be larger than R_{u}
- Value of R_u must be less than a limit

Each of these is explained below.

Need for a large Faradaic capacitance

As discussed above, C_{faradaic} maintains a "DC" potential when the current is interrupted. If the faradaic capacitor is missing or too small, current-interrupt usually drives the system to a large potential and current. The most

obvious symptom of this problem is measured current many times higher than the expected cell current. Overload indications may also be present.

Current-interrupt using Gamry Instruments' potentiostats works best with a faradaic capacitance of greater than 20 μ F. For a "bare metal" electrode you can estimate the capacitance as 20 μ F/cm², so the electrode area must be 1 cm² or greater. We do not recommend that you use current-interrupt iR compensation if the electrode is covered with any sort of insulating coating.

This requirement generally limits current- interrupt iR compensation to corrosion testing and battery and fuel cell research. Current-interrupt does not work well with the electrode sizes commonly used in physical electrochemistry cells.

Time-per-point limitations

Current-interrupt iR compensation assumes that you apply a DC potential and current most of the time. The interrupt time should be much smaller than the time required to measure each data point in a data curve.

By default, Gamry Instruments' software automatically selects a total current-interrupt time and sampling time for the interrupt. Theses time are adjusted whenever the current range changes, with longer interrupt times and slower sampling on more sensitive current ranges. The range of interrupt times is generally 10 μ s to 64 ms.

Gamry Instruments recommends that you only apply current-interrupt iR compensation when the time per data point is 1 s or greater.

If you are sweeping the potential, this limits the sweep rate to 5 mV/s or less.

$R_{faradaic}$ should be larger than R_u

There are also limitations on the ratio between R_{faradaic} and R_{u} . Because the same current flows through both resistors, this is also a limitation on the ratio between the voltage across the electrochemical interface and the error voltage.

Gamry Instruments' experience with current-interrupt on model systems shows that R_u should not be more than ten to twenty times R_{faradaic} . If R_u is bigger than this limit, the electronic measurement of the iR error starts to have significant errors.

Usually a more severe limitation on this ratio occurs in the cell. Most electrochemical cells have a non-uniform current distribution across the electrode surface. Some portions of the working electrode see more current than others. Under these conditions, the simple Randles cell model doesn't apply. The interface cannot be described using a single potential.

Unless your cell has a geometry designed for uniform current distribution, we believe you should keep R_u one tenth of R_{faradaic} or smaller. If the ratio is greater than 1/10, we believe that any quantitative results obtained on your system will be in error. Note that this is only a "gut feel" approximation. We cannot guarantee this approximation will apply to your electrochemical system.

Value of R_u not too large

There is also a limit on the value of R_{u} , independent of the value of R_{faradaic} . Our experience has shown that errors occur when R_u exceeds some upper limit. This limit is about 10 k Ω for most Gamry Instruments systems.

Positive-Feedback iR Compensation

Current-interrupt iR compensation is only useful when you are measuring slow phenomena, such as corrosion reactions or characteristics of energy-storage devices. It cannot be used when very fast measurements are needed. One example of a fast experiment is 1000 V/s cyclic voltammetry used to measure chemical kinetics.

Fortunately, there is an iR-compensation technique that works on fast systems. It is called positive-feedback iR compensation.

This technique can be thought of as an additional analog feedback path in the potentiostat. All useful potentiostats measure the cell current. When positive feedback is enabled in a potentiostat, a fraction of the current signal is feedback as an additional voltage input.

The figure below is a highly simplified schematic diagram of a potentiostat.



Simplified Potentiostat with PFIR D/A Converter

In the lower-right portion of this diagram, current is measured by voltage drop across Rm. This drop is amplified to produce a voltage signal called Isig. In this diagram, Isig is 3 V at full-scale current.

In the upper right portion of the diagram we show a block labeled PFIR DAC. This is the Positive-Feedback iR digital-to-analog converter. Its output is a voltage that is a known fraction of Isig. When positive feedback iR is enabled, this voltage is applied to the control amplifier as an additional voltage input.

Note: In this diagram and the discussion, the PFIR DAC output is not scaled prior used as a source of feedback. This may not be true for all PFIR implementations, including some of Gamry Instruments' implementations.

Some simple math leads to:

$$lsig = 3.0 \times I_{cell}/IFS$$

$$= I_{\text{cell}} \times R_{\text{e}}$$

where R_{e} is the equivalent current-measurement resistor given by:

$$R_{\rm e} = 3.0 \text{ V/IFS}$$

 $R_{\rm e}$ tells us the maximum value of $R_{\rm u}$ we can compensate on any current range.

In positive-feedback iR compensation, you need to know R_u before you can apply the correction. Gamry Instruments' PHE software includes a convenient method for measuring R_u prior to running an experiment. This method uses a simplified form of EIS and assumes that the cell impedance at high frequencies is identical to R_u .

Once you enter a value for R_u and select positivefeedback iR correction, Gamry Insturments' software sets the PFIR DAC to generate a fraction of Isig identical to the ratio of R_u to R_e . Given this setting, the voltage feedback is:

PFIR out =
$$R_u/R_e \times \text{Isig}$$

$$= R_{\rm u}/R_{\rm e} \times I_{\rm cell} \times R_{\rm e}$$

$= R_{\rm u} \times I_{\rm cell}$

We are increasing the potential applied to the cell by the voltage across R_u . This voltage is created by the flow of the cell current.

The range of R_u values that can be corrected is 0 Ω up to R_e . The resolution in the value is controlled by the PFIR DAC resolution. With a 14-bit DAC (D/A converter), the resolution is R_e /16 384.

As an example, let's look as the 3 mA current range. On this range, $R_{\rm e}$ is 1000 Ω . The positive-feedback correction with a 14-bit DAC has a resolution of 1000/16 834 or about 0.061 Ω per bit.

Advantages of Positive Feedback iR Compensation

Positive-feedback iR compensation offers some advantages compared to other iR compensation methods. These include:

- Usable in very fast experiments
- Scan parameters such as ramp limits and scan rates are corrected.

Practical Limitations of Positive-Feedback iR Compensation

There are disadvantages in positive-feedback iR compensation compared to other compensation methods:

- Prior knowledge of the value of R_u is required
- Errors occur if *R*_u changes value during the experiment
- The current range must not change during the experiment
- Positive feedback can lead to potentiostat oscillation.

Application Note Rev. 3.0 1/16/2018 © Copyright 1990–2018 Gamry Instruments, Inc.

Explain is a trademark of Gamry Instruments, Inc. Mathcad is a registered trademark of PTC Inc. Excel is a registered trademark of Microsoft Corporation.



734 Louis Drive • Warminster PA 18974 • Tel. 215 682 9330 Fax 215 682 9331 • www.gamry.com •info@gamry.com