TechnicalNOTE

GAMRY INSTRUMENTS

Understanding iR Compensation

Introduction

Probably 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?

In this Tech Note we'll attempt to answer these questions, and leave you with a basic understanding of iR compensation.

Background

This application note presumes that you have a basic understanding of potentiostat operation. If not, please review our Primer on Potentiostats Tech Note (http://www.gamry.com/App_Notes/Potentiostat_Primer. htm). 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 Primer on Electrochemical Impedance on our website. Pay particular attention to how typical chemical processes are mapped into electrical circuit elements.

Where does iR come from?

Lets look at a typical 3 electrode electrochemical test cell. I've labeled some reference points in the cell that we'll refer to throughout the remainder of this application note.

Reference Points:

A: Counter electrode output at the potentiostat

B: Metal surface of Counter electrode
   electrode
   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
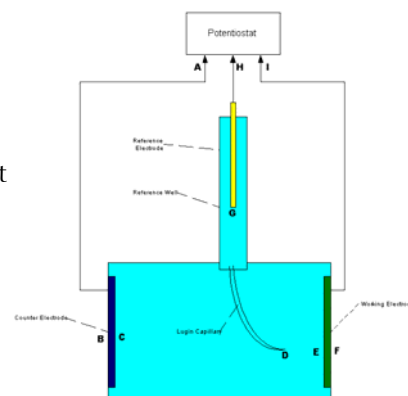
We can think of the cell as a (simplified) network of electronic components something like this:
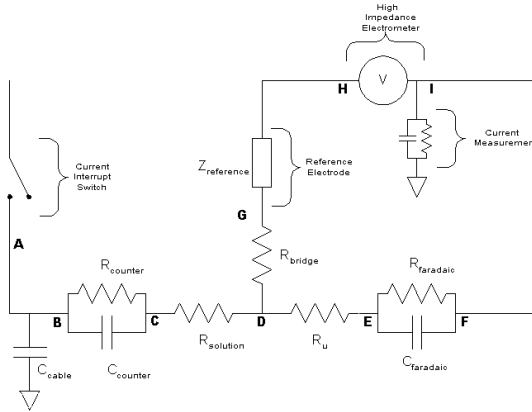
The potentiostat does a good job of controlling and measuring the voltage between Points H & I. Unfortunately, we really want to control and measure the voltage between Points E & F, across the electrochemical reaction we're trying to study.

$V_{important} = V_f – V_e$

For purposes of this discussion, Point I is equivalent to Point F. Point G is equivalent to H, except for a constant offset voltage due to the Working Electrode/Reference Electrode potential difference, also known as the open circuit voltage, V_{oc}. And since 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:

$V_{measured} = V_I – V_H$

we can get to:

$V_{measured} = V_f – V_d + V_{oc}$

We're one step from finding the reaction potential, V_f-V_e. Point E is equivalent to Point D except for that nuisance resistor, R_u. How is it related? By Ohm's law:

$V_e – V_d = I_{cell} * R_u$

So substituting this in gives:

$V_{measured} = V_f – V_e + I_{cell} * R_u + V_{oc}$

Rearranging the result gives:

$V_{important} = V_f – V_e = V_{measured} – I_{cell} * R_u – V_{oc}$

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 $V_{important}$!

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

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

What kinds of cells/systems do we need to worry about?

Most of the time, electrochemists can rig their experiements 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.
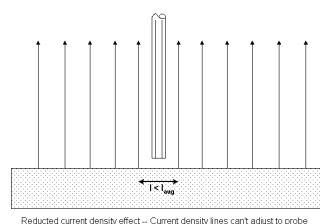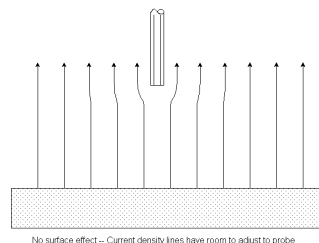
Conductivity up --> Ru down -> iR down

Another way is to minimize the distance between the tip of the Lugin 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 (Cf). 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 cell design solution often is not a solution. Many cells designs are limited mechanically. A more subtle problem is that putting the reference too close to the working electrode surface will change 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

Electroplaters are familiar with this type of effect and have a special term to measure it – Throwing Power.

So do you need to worry about iR? 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 measured?

The following circuit has some important clues for measuring iR and from it, Ru.

Randle's Cell Model
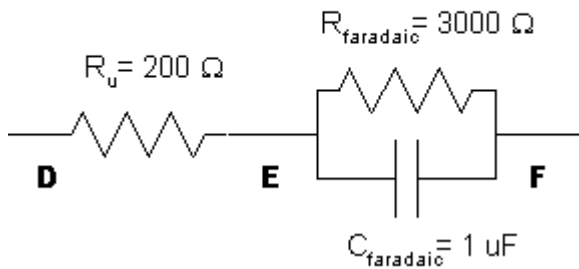
Look at Rfaradaic. It has a capacitor, Cfaradaic in parallel with it. Ru doesn't. This should suggest that an AC experiment can differentiate the two. High frequency signals pass right through Cf with no voltage drop while they are forced to drop Icell * Ru volts going through Ru just like low frequencies.

Indeed, if you take an AC Impedance Spectrum of the cell, the result often looks like the following:
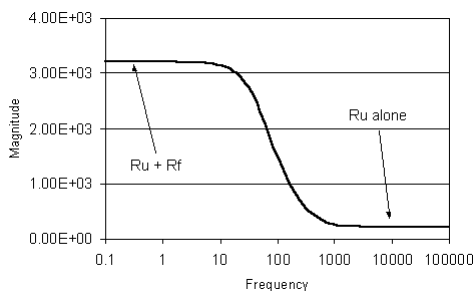
At low frequencies where Cf is effectively an open circuit, the measured impedance is the sum of Ru and Rf. At high frequencies where Cf is effectively a short circuit, the measured impedance is Ru.

So do you need to worry about Ru? Measure it. Then multiply it by the cell current. The resulting voltage from the diagram is the uncompensated iR, Ve-Vd. If iR is smaller than a few millivolts, don't worry about it.

For example. Suppose from the impedance spectrum, Ru = 100 Ohm. Suppose Icell = 10 uA.

iR = 100 * 10 * 10-6 = 1 mV.

For most electrochemical phenomina, that is a small error.

Another way to think about it is - if Rf >> Ru, don't worry about Ru.

Measuring Ru using a DC technique

Measuring Ru 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 running an IV curve.

Luckily, there is an equivalent DC method. Actually I shouldn't call it DC because it uses a rapidly changing signal to measure Ru. A better term is a "transient" technique. It is know as Current Interrupt.

Consider the simple Randel's cell model for an electrochemical reaction with solution resistance.
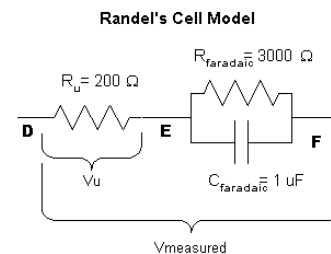
To make the current interrupt measurement, cell voltage (Vmeasured) is measured immediately before and immediately after the current has been interrupted. Ideally, the measured voltage of this cell would look like the graph:
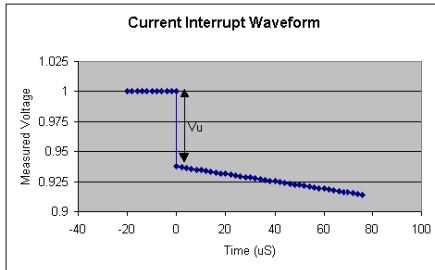
Suppose we are measuring 1.0 V while current is flowing. At time 0 we interrupt the current – very quickly! The voltage immediately drops by the amount across Ru. Then it starts dropping more slowly. Why? Because the faradaic capacitor is slowly starting to discharge. At short times, the capacitor can hold the voltage to Vmeasured – Vu which is our Vimportant. And this is what we're after.

As usual, things aren't so simple. Three problems interfere with this simple model; Sampling Speed, Output Capacitance, and 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 μsec (very fast). We can slow down the sampling considerably if we pretend the decaying curve is a straight line and back extrapolate it to the turn-off time.

Lets try that with the same Randles cell we've been using. Sample at 1 ms and 2 ms, back extrapolate to 0 ms when the switch was turned off. You get something like this:
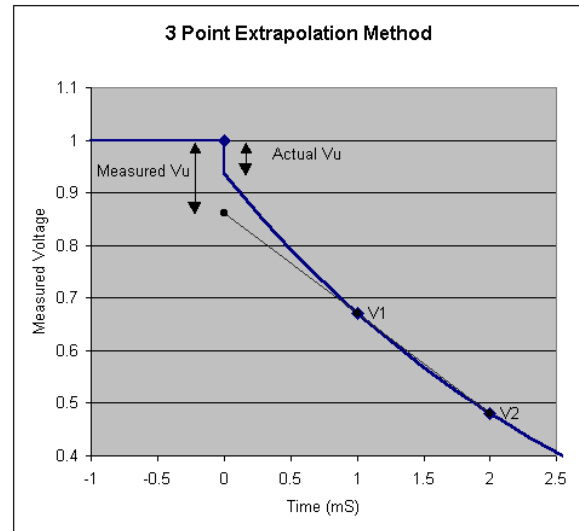
The estimated Vu is calculated by: Vu = V1 + (V1-V2)

In this case V1 = 0.671 V, V2 = 0.481 V, and by the straight line extrapolation, Vu is estimated at 0.862 V.

But wait, Vu is actually closer to 0.938 V. We've been 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 - its an exponential decay. (By the way, these are actual numbers from our original cell model, calculated in MathCad and plotted in Excel.)

So I'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 is Rfaradaic * Cfaradaic. In this cell,

taufaradaic= 3000 Ohm * 1 μF = 3 ms

If you know, roughly, these values, you can pick a short time tau to be RC/10. Or you can sneak up on the correct answer by decreasing Tsample until the number stabilizes.

But there's another problem lurking as Tsample gets shorter, -- Potentiostat Output & Cable Capacitance.

Cable Capacitance

Think back to the original cell model:

That innocent looking capacitor, Ccable, can cause a lot of problems. If you have a typical shielded cable, the value of Ccable can be 50 pF per foot. For a 5 foot cable that's 250 pF and add another 100 pF or so for capacitance in the switch itself, on the circuit board, and in the driving amplifier.
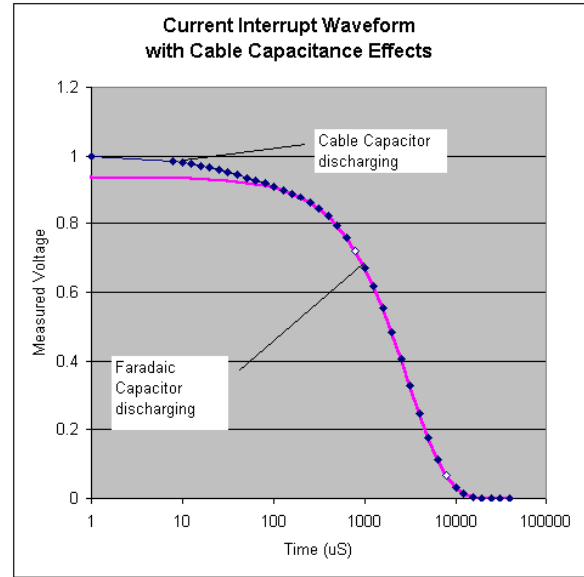
We can model this with this circuit:
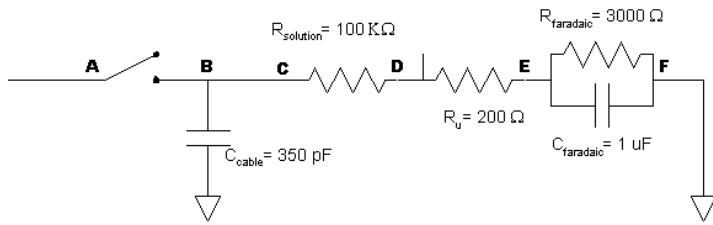
The cable capacitance forms an RC section with Ru and Rsolution. This means the voltage across Ru doesn't disappear infinitely fast.

For purposes of this discussion, we have to assume the counter electrode 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 μsec and 100 μsec. I've put in some sample values to calculate what this looks like on a scope. Using those two measurements, obviously the iR estimate is again bogus. You must wait until the cable capacitance transient has died out to make the measurement.

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

So the moral of the story is that you need to find a time range between the two discharge curves. The cable capacitance must be fully discharged but the faradaic capacitance must still be in the approximately linear region.
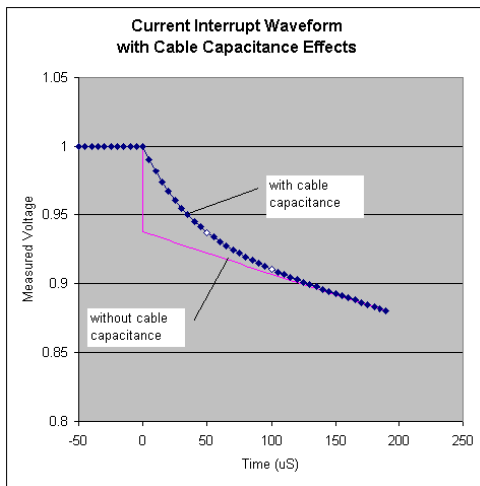
Noise

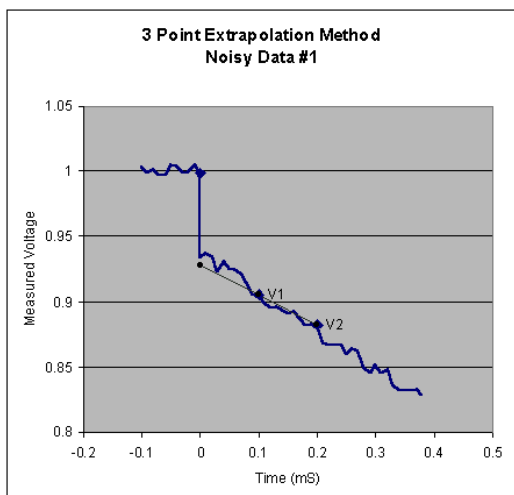Noise is a real problem for iR compensation.

Basically, iR comp is a differential measurement. Remember the equation we use for estimating Vu –

Vu = V1 + (V1-V2)

The differential term, V1-V2, is very sensitive to noise. If we're doing this right, the difference between V1 and V2 is very small, a few tens of millivolts. Suppose there is a positive noise contribution to V1 and a negative noise contribution to V2. The average noise is 0 but the error in Vu is twice as big!

3 Point Extrapolation Method
Noisy Data #1

"So what", you say, "Just turn on some filtering. The noise will go away". But we're trying to measure a rapid (10-100 μsec) phenomenon. We can't put on a 5 Hz filter or the whole transient will be distorted, if not disappear altogether.

What do you do. Several things, actually. 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. We use all these methods.

Finally, if the noise is still too large, don't use the extrapolation method. Just be settle an average, such as:

$$Vu = (V1 + V2) / 2$$

By the way, the noise situation is made worse when you're trying to measure low currents. When the current interrupt switch is open, the Reference electrode and Working electrode are more susceptible to noise pickup than when the switch is closed.

How is it corrected?

So far we've only talked about measuring iR. If we know the value of Vu we can subtract off the Vmeasured value and get Vimportant. This is known as Post Processing Correction.

A better approach would be to have the potentiostat measure and correct for Vu continuously. After all when you want to potentiostat 1 V it is with the intention that Vimportant = 1 V, not Vmeasured = 1V.

The situation is much simpler when you use a Galvanostat. A galvanostat doesn't have to apply a

correction for Vu since its job is to control current, not voltage. It still has to measure Vu.

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 Vu on to the applied signal. This can be stated as a formula where the numbers in brackets represent measurement points:

Vapplied[i] = Vrequested[i] + Vu[i-1]

You should examine both the Vrequested[i] and Vimportant[i]. The difference can be quite significant.

A better approach to understanding the feedback mechanism is to treat iR Correction as a control loop.

Control loop algorithm

The control loop algorithm treats the potentiostat as a loop within a loop. The inner loop is the potentiostat itself which measures Vmeasured and controlis it in a feedback mechanism. This loop is made purely from analog electronics as shown:
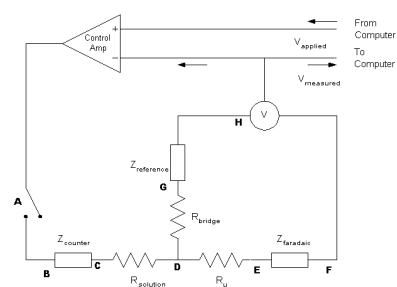
Again, we've eliminated a few components that are irrelevant to this part of the discussion.

The potentiostat is a control loop. It measures Vmeasured and compares it Vapplied, making corrections to the counter electrode voltage until the difference between the two is 0. All of this happens continuously – If it didn't you'd be very mad at me because it is the main job of a potentiostat.

The iR correction also happens in a control loop outside of the potentiostat loop. It is shown in this figure:
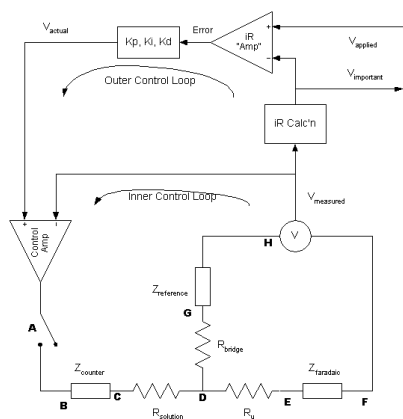
The outer loop does a very similar job to the inner loop but it is implemented digitally in a computer. Its job is to see that Vimportant = Vapplied, and now the job of the inner loop is to see that Vmeasured = Vactual. Vactual is a new value that comes from the outer loop. I've also shown an interesting block that produces Vactual. It is known as a gain block or controller block, depending who you talk to. Its output is given by the expression:

$$V_{actual} = K_p \cdot error + K_d \cdot \frac{d\,error}{dT} + K_i \cdot \int error \cdot dT$$

This is known to control engieers as a PID loop. Each of the gains is individually controlable. 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, we commonly use a the 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, EuExtrap, 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, Ki. Kp is fixed at 1.0, and Kd 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 control loop mode, Vu calculation, current interupt timing and gain to suit the dynamics of the reaction and cell under test.