Microsegregation and "Current Burst" Theory

Why Electrochemistry of Silicon May Help to Understand Wootz Steel

Advanced

Take a piece of single crystalline silicon. Contact it on the backside, put the sample in a beaker containing some suitable electrolyte; diluted hydrofluoric acid would be best. Put some noble metal wire in the beaker, without it touching the Si sample. Platinum (Pt) is best. Now apply some voltage to the system (a few volts will be enough) by hooking up the contacted Si sample to the positive pole of a battery or power supply, and the noble metal to the minus pole.

Just to be on the safe side: hydrofluoric acid (HF) would be your best choice if you plan to disfigure or to kill yourself in a particularly painful and messy way. It is an extremely dangerous chemical! *Never ever* get close to the stuff if you don't know *exactly* how to handle it!

OK, now that this is cleared up, let's go on. You will find that some current will flow from the silicon into the electrolyte and out again via the metal electrode, and that the silicon will *dissolve* at a rate that is essentially proportional to the current density (Ampères per cm^2 sample surface). Quite unexpectedly, it does so in rather peculiar ways. Let's look at the basic processes for starters:

- The electrical current forces a chemical reaction of the silicon with the water in the electrolyte that produces silicon dioxide (SiO_2). In other words: it forces the silicon to "rust". If you would do that that in water without the acid, the SiO_2 layer grows thicker and thicker. Since it is a good insulator the current decreases to zero after a thickness of a few nanometers has been reached. Not very exciting.
- However, hydrofluoric acid dissolves the silicon dioxide chemically - it doesn't need voltage or current to do that.
- So under certain conditions a *steady state* can be reached for a constant voltage, requiring that just as much SiO_2 is dissolved by the chemistry as is produced by the current. The oxide thickness then stays the same everywhere, current and voltage are constant, the silicon dissolves with a constant rate, and the *solid-liquid interface* moves with a constant speed.
- So it appears from the outside. However, external measurements give constant macroscopic values because they *average* over what is going on on a microscopic scale. Microscopically, the oxide thickness changes *locally* in a rhythmic, quasi-periodic manner. We call that a local **stochastic oscillator**.
- As long as all local stochastic oscillators do their thing completely independent of each other, i.e. having random phases, everything averages nicely to constant values. It's like a bank manager, measuring the level of money in the checking accounts of many thousand customers and finding the average rather constant. Nevertheless, my and your microscopic money flow oscillates or pulses rhythmically, if somewhat stochastically. It goes up when the pay check is deposited and then tends to go down in leaps and bounds, then it goes up again.
- On occasion, however, the local stochastic oscillators on a silicon electrode synchronize - pulsing now in harmony. And this we will notice on the outside: The voltage is still constant because the "constant voltage" or "potentiostatic" power supply sees to that. The current, however, now **oscillates** like crazy!
  If you are a sophisticated experimenter, you will now try a *galvanostat*, a power supply that keeps the current constant. Then the voltage will oscillate! This is what it looks like:

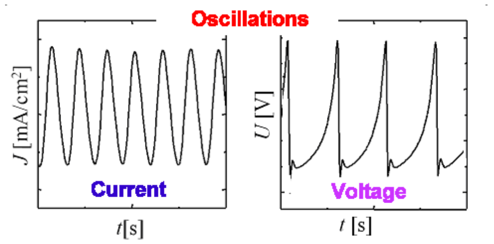

Measured current and voltage oscillations during silicon dissolution

- This is very unusual behavior! Almost any "Black box" with two wires coming out will not do that when hooked up to a constant voltage. It may do nothing, explode, make noises, or get hot but making the current oscillate it will not!
  Note that the word "oscillations" means usually that something changes rhythmically *in time*. Now look at this:

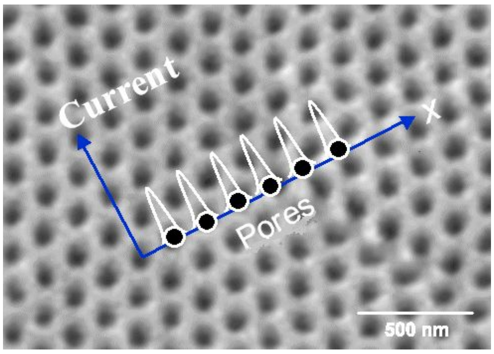

Very deep pores in a self-arranged periodic pattern

- The inset show the distribution of the current along some direction on the specimen surface. There is an obvious two-dimensional current oscillation *in space*, and since current dissolves the material, an array of deep pores (about 300 µm here) is formed.
  We are looking at self-organized **current oscillations in space**. Wherever current flows, the semiconductor dissolves, leaving an array of pores. In the picture it is indium phosphide (InP), but similar things happen in silicon and in many other semiconductors.
  In other words: The formerly planar solid-liquid interface became unstable and is now "cellular". Or is it? Let's look at some cross-sections of those pores:

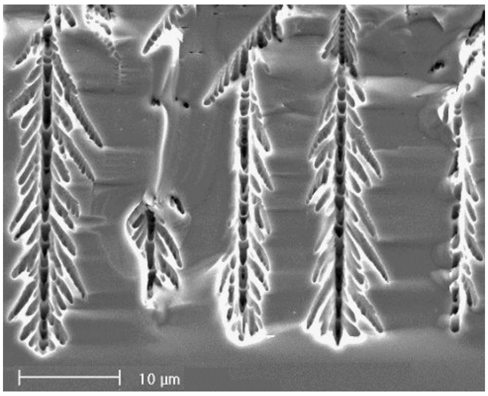

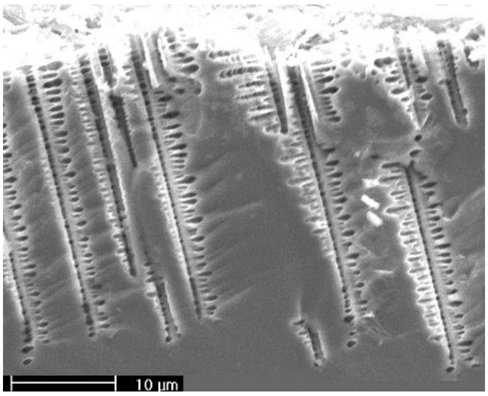

Dendritic pores in silicon

- Those are cross-sections for particularly well-expressed "dendritic" pores in silicon. In contrast, the interface going with the InP pores above would be cellular, indeed.

To be sure: dissolving silicon at room temperature in an electrochemical experiment is something very different from solidifying steel at very high temperatures. No doubt about that. However, there are obvious similarities. Indeed, on a more abstract level, both experiments have a lot in common. Let's see how that can be:

1. In both cases an *energy* current is driven across a solid-liquid interface. During casting it is a heat current driven by a temperature gradient, in the electrochemical experiment it is an electrical current, driven by a gradient in the electrical potential. The basic equations for these currents have the same mathematical structure, and one way the systems copes with this occurs in both examples: Increase the interface area to get the current *densities* down if you can't reduce the absolute current!
Now I give you the thought for your weekend. Can you guess what will happen if you reverse the process of solidification and start to *melt* a piece of iron? What would the solid-liquid interface look like? Would there be dendrites of liquid extending into the solid? Any other interface instabilities? Does anybody know? Or care?

- **2.** In both cases *particle* currents must flow towards, through, or away from the interface. For solidification this is described here, for the electrochemical experiment we need to bring the proper molecules to the interface so the necessary reactions can take place, and we need to get the reaction products away from the interface. In both cases diffusion and liquid flow are the means of transport, and once more it is the same math describing this.
- **3.** Energy and particle currents need to be balanced for achieving steady state. You can't produce more oxide than you can dissolve, and you can't solidify more (alloy) melt (with equilibrium composition) than given by the transport of the right atoms to and from the interface.
- **4.** Both systems are not in equilibrium. The larger the offset, the more drastic the actions to restore equilibrium. Both systems respond by changing the area of the solid-liquid interface and the interface velocity. We have, planar, cellular and dendritic crystal *growth* as well as *dissolution*.
- **5.** Both system contain a stochastic component. Certain features can only be described by *probabilities* and not with fully deterministic equations. This induces a "chaotic" component in both systems and thus also the possibility to create order out of chaos by *self-organization*.

To be sure, there are specific differences, too. Solidification generates heat, and that is an important quantity since it interferes heavily with the heat transport and the temperature gradients. Dissolution also generates heat but that is of minor importance. Dissolution, while not directly generating an electrical current, nevertheless interferes with the local potential gradient and the flow of the electrical current in a subtle way.
In a nutshell, the electrochemical system does have close analogues to solidification topics like interface instabilities, cellular growth, dendritic growth and whatever produces *quasi-periodic* structures *in space*. In contrast to casting, it also produces easily measurable *quasi-periodic* structures *in time* that are well understood.

- I have claimed that striations are rather universal and result from an interface that moves in a kind of "stop-and-go" fashion. This is just a flashy expression for saying: "the interface velocity *oscillates in time*".
  I will now make a daring claim:

Oscillations of the oxide thickness during the electrochemical dissolution of silicon, and oscillations of interface velocities during solidification, are different expression of the same underlying principles!

There is a difference, however: For the electrochemical oscillations we have developed a detailed theory called "**current burst" (*CB*) theory**, whereas I'm not aware of a general "interface velocity oscillation" theory. We thus might profit from exploring possible analogies.

Article Link

Pattern formation

In what follows, I will first give the basics of the current burst theory, and then draw the parallels to interface velocity oscillations and striations. Details on oscillations and pattern formation during the dissolution of semiconductors in general and silicon in particular can be found in the links

Article Link

Oscillations

The Current Burst Model

The easiest way to understand the essentials of the CB model is to get a (p-type) silicon specimen covered with a perfectly uniform oxide (SiO_2) with a thickness of 50 nm or so. That is standard stuff in semiconductor technology, so the following experiment is easily done. Contact the backside of the sample, isolate backside and contact (cover it with wax, for example), put it into an electrolytic cell with some noble metal counter electrode and a fluoride bearing electrolyte, apply some voltage (about 5 V, **+** pole on the Si). and measure the current. It looks like that:

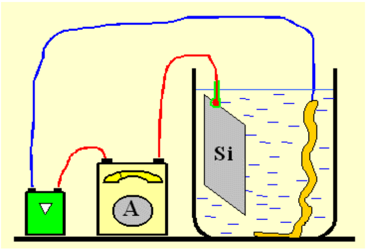

- In the beginning of the experiment the current will be zero. That is not surprising because silicon dioxide is an extremely good electrical insulator, so no current can flow through the 50 nm thick layer. However, the electrolyte will slowly dissolve the SiO_2 just as water dissolves salt. The SiO_2 layer gets thinner and thinner, and the *electrical field strength* in the layer thus gets larger and larger - and something is bound to happen before it reaches infinity!
- Before I get to that it is time to remember (or learn) that voltage ***U*** and current ***I*** are completely meaningless quantities from the viewpoint of materials! It is **electrical field strength** and **current density** that counts. I have already cleared that up with respect to the current density; it is simply the same thing with the field strength. For our case the field strength ***E*** in the oxide layer is given by $E = U/d_{ox}$; d_{ox} is the oxide thickness.
- In our experiment the oxide thickness decreases due to dissolution. The field strength then goes up, and for ***U*=5 V** reaches a value of about **10 MV/cm** as soon as the oxide thickness is down to **5 nm**. An electrical field strength of **10 MV/cm** is a kind of "magical" number because it is about the *highest* field strength that the very best insulators (like SiO_2) can bear. Go beyond that field strength and **electrical breakdown** or **electrical discharge** will inevitably occur in any material including vacuum.
  You have seen *and* noticed electrical discharge taking place. When the field strength in the air gets too high, electrical discharge occurs in a spectacular way: thunderbolts have scared you and me, no doubt about that.
  Then you have literally seen but probably not noticed electrical discharge taking place when you turn on the light from a fluorescent tube. You start a (continuous) electrical discharge in there because a little starter device produces a high voltage for a few second when you press the button, and that initiates the discharge.
  The two examples are perfect to elucidate two very special and very important aspects of *all* electrical discharge events - huge one like thunderbolts, or tiny ones on a nanometer scale in the SiO_2 layer :

  1. There is *no way* to predict exactly when and where a discharge will take place. There will be flashes of lightning between the two clouds about every 3 minutes but there is no way of calculating exactly when, and from where.

The best you can do is to work with *probabilities* and averages.

2. The field strength needed to keep a discharge burning is always considerably *lower* then the field strength needed to initiate it. This is well understood and text book stuff, just believe me on that.

- We are done. The two points above embody the basic principles of the CB model. All we need to do is to apply it to our Si electrode in the experiment and write a little (actually a rather sophisticated and long) program to do the number crunching. The essential points for doing this are:
  - **1.** Derive two equations for the probability $P_{start}(d_{Ox})$ that discharge starts or nucleates as as function of the oxide thickness d_{Ox}, and a second equation for the probability $P_{stop}(d_{Ox})$ that some "burning" discharge stops as as function of the oxide thickness.
    There is some leeway in choosing specific equations but no doubts about the general structure.
  - **2.** Divide the specimen surface into small pixels. Assign an oxide thickness value to any pixel as a starting value. You can give any pixel the same value or anything you fancy; it doesn't matter.
  - **3.** Define the etch rate, i.e. how much oxide is dissolved per time increment. This depends essentially on the composition of the electrolyte and its temperature; the necessary relations are well-known.
  - **4.** If some current flows through the pixel let the oxide grow with a rate given by the current, and in the shape of a (more or less spherical) "bump" that spreads into neighboring pixels. The strength of the current is given by known external parameters.
- Now start the calculations. Starting at some time t_0 the computer "looks" at every pixel and notes the present oxide thickness. Then it looks at the two basic options:

  1. If current is flowing through the pixel, i.e. if there is an electrical discharge through the oxide, it casts a die biased by the "turn-off" probability function. If "turn-off" comes up, it stops the current flow, if "keep going" comes up, it increases the oxide thickness of the bump by the proper amount.
  2. If no current is flowing through the pixel, it casts a die biased by the "turn-on" probability function. If "turn-on" comes up, it starts current flow and increases the oxide thickness, generating the first stage of an oxide bump. If "do nothing" comes up, it does nothing

  Next it reduces the oxide thickness in *every* pixel by the prescribed amount due to dissolution. Finally, it counts how many pixels are in an "on" state and adds up the local current to the total current. Then it computes the average oxide thickness and plenty of other parameters. Finally it plots desired parameters in a graphic representation of the electrode surface by assigning colors to numbers.
  - The system is now in a state that belongs to the time $t_1=t_0+\Delta t$. Δt is the time increment that goes with the changes in oxide thickness.
    The next few thousand steps in the calculations simply repeat the procedure above over and over again, keeping track of the parameters and plotting the result. This looks like this:

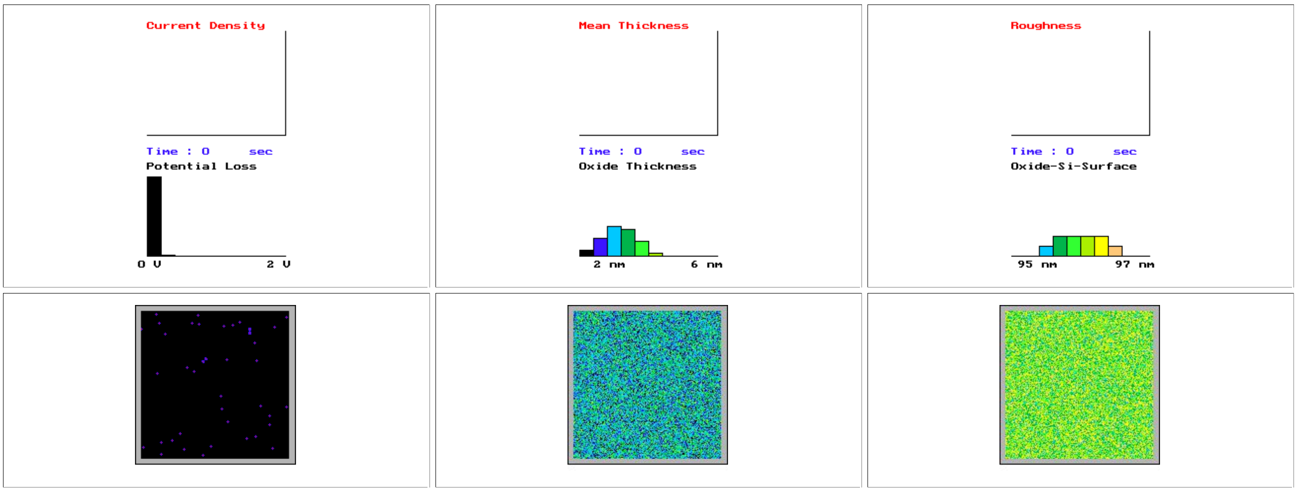

Black: no active CBs

Oxide thickness

Surface roughness

Calculating current oscillations in silicon

This is a *movie* and not an actual calculation (your computer would be far too slow for this). That's why things repeat

- The current oscillates! Everything oscillates *stochastically*: things repeat but never exactly as before.
  Now why is that? There are *two* basic reasons for macroscopic oscillations to occur:

- **1.** We must have *local stochastic oscillators* or things that are on or off, growing fast or slow, awake or sleeping, very active or less active - whatever. The only requirement is that there is a kind of recognizable rhythm. In other words: We can define an *average* time for the duration of one cycle. For example, assigning an average time of 24 hours to your waking/sleeping cycle makes sense, but we cannot tell how long exactly your next cycle will be.
  If the local oscillators do their thing without any regard to what their neighbors do, there will be no macroscopic oscillations. At any time there are - on average - equal numbers in the "on" or "off" (or sleeping / awake) state.
- **2.** We must have some *synchronization* between the local oscillators. This can happen by
  - Interactions between neighbors.
  - Rhythmic disturbance from the outside that scale with the frequency of the local oscillators.

  This is easy to see if you look at your own quasi-periodic behavior. The change over from day to night has a well-defined frequency that matches your own disposition of having waking-sleeping intervals lasting about 24 hours. It induces most of us to sleep sometime during the night and be awake mostly during the day. There are more sleeping people at night, and their number now oscillates stochastically because of an outside influence.
  With respect to your eating pattern, it is the interaction with your neighbors that tends to produce an oscillatory behavior with respect to the number of people being in an eating (or drinking) mode. There just aren't that many shapely secretaries in the cafeteria if you lunch at 3:30, and that's why *you* tend to lunch between 12:00 and 1:00. There is your interaction! *My* interaction is entirely different but the effect is the same. As long as there is *any* interaction, chances are that some synchronization will occur.

The first ingredient is "automatically" present in the CB model, and so is the second one! There is a *next-neighbor* interaction in the model, and that produces synchronization. here is how:

- Let's look *schematically* at a tiny part of the system in cross-section. It might look like this:

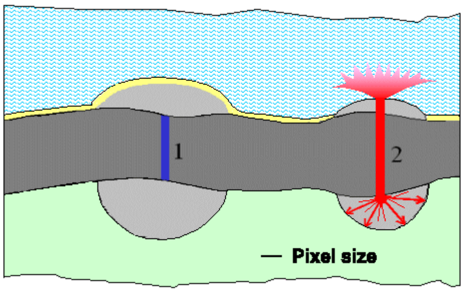

Time=t_0. One CB (1) just stopped and is now dormant, another one (2) is still active.

- Current flow is indicated in red. Oxide dissolving during the time interval considered is shown in yellow (but not on the growing oxide bump for clarity). Note that the oxide bumps extend into the bulk silicon *and* into the electrolyte since oxidation increases the volume by a factor of two.
- If we now move on a (few) time increments, CB 1 will become dormant since its oxide bump had almost reached a thickness were switching off will become very likely. In addition, a new CB (3) will have nucleated. This happens with high probability somewhere between the two old ones because there the oxide is thinnest.

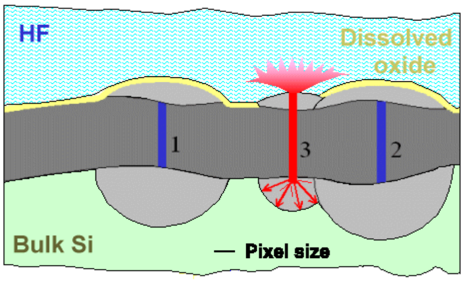

Time=t_0 + Δt. Both previous CBs are now dormant but a third one has nucleated and is actively growing an oxide bump

- Moving on a few more time increments, we reach the stage where CB 3 is just about to become dormant since its oxide bump has reached the thickness were switching off will become very likely; this is shown below

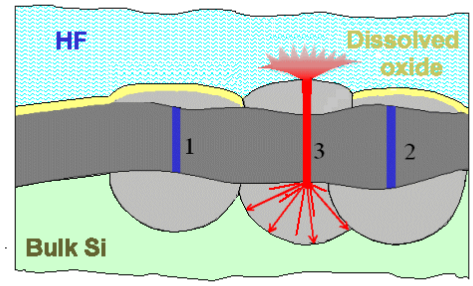

Time=t_0 + 2Δt. CB 3 about to extinguish. It has produced less oxide than the others and has a *shorter* life-time

- Did you notice the next-neighbor interaction that now will start to synchronize all those current bursts? No? Well - it is subtle but powerful:
  - Current burst No. 3 does not have to grow as much oxide as its neighbors before it turns off since it "employs" the still present oxide of its neighbors to get out of work earlier than it would have if left to its own devices. Its switching off time thus is *correlated* to those of its neighbors, it is not what it would have been without the crowd. This is not yet synchronization but the beginning of it. If CBs are dense enough (i.e. the total current is high enough to space them at sufficiently high densities), interactions with the neighbors will rather quickly produce some degree of synchronization.
  - To be sure, there are more mechanisms of interactions than the one given above and they are contained in the full CB model. Modelling by a "Monte Carlo" computer program has the big advantage that all these effects come out automatically if one just describes the basic processes correctly.
    The problem with "Monte Carlo" programs is that they only can model a tiny part of the electrode surface since the number of pixels has to be small (a few millions) if you don't want to wait a long time for results. Moreover, the model might produce all kinds of *unexpected* results if you play with it, changing a parameter here and there.
- For example, in some computer runs the system produced structures like (shallow) pores! This was not expected but found to be another interaction feature that emerges under certain conditions. In a grand generalization the following statements result.
  - There are two basic kinds of interaction between CBs:

    1. Interaction in *space*, meaning that the probability for the nucleation / stopping of a CB in a given pixel depends to some extent on what is going on in its neighborhood.
    2. Interaction in *time*, meaning that the probability for the nucleation of a CB in a given pixel depends to some extent on what was going on in that pixel in the past.

    The two basic kinds of interaction produce two basic kinds of synchronization:

    1. Interaction in *space* induces synchronization in *time*. CBs, on average, do the same thing at a given time. This produces the current oscillations in time.
    2. Interaction in *time* induces synchronization in *space*. CBs, on average, tend to be in certain places and not in between. This produces the current oscillations in space that express themselves in ordered pore arrays.
  - There is something aesthetically pleasing about that. If you can't feel that, send me an e-mail and climb back up your tree. I'll send you a banana.
    The full story is given in this reference.
- It only remains to look at synchronization induced from the outside. This is easily implemented into the model and works exactly as expected; consult the references for details. Especially interesting is the observation that rhythmic disturbances from the outside can induce system to develop oscillations that would not yet oscillate by themselves.
  - I only want to show one picture that demonstrates the effects of outside interference rather spectacularly:

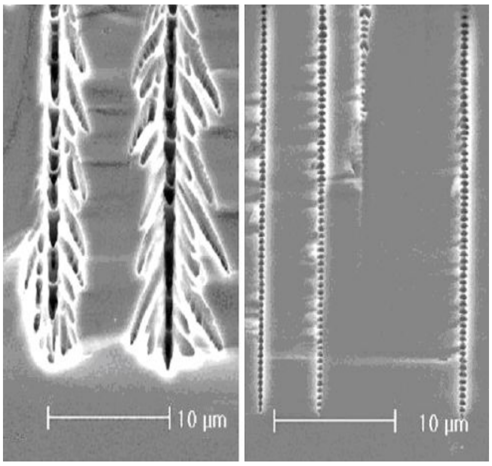

Pores in silicon; all conditions identical except the current modulation

Constant current or modulation with "wrong" frequencies

Current modulated with 33 mHz

- All one needs to know is that the CB model applies to the tip of those pores, where the dissolution takes place. CB's are not synchronized in the left picture, and every now and then side pores are generated.
  In the right picture, synchronization was achieved by modulating the current with just the right frequency. Doing the experiment with frequencies different by more than about 20 % has no effect whatsoever. We might call that a *stochastic resonance*.

Does a "Dendrite Burst" Model Explain Striations?

I wrote a whole module arguing that striations result if there is an "oscillating" solid liquid interface. It is not so much the position in space that needs to oscillate but the *growth speed* given by the average thickness of the material that crystallizes per time unit.
Now let's compare essentials of the two processes: crystallization and electrochemical dissolution.

- The amount of material crystallizing per time unit is essentially given by the rate of heat removal or the magnitude of the *energy current* flowing through the interface. It is just the opposite of the electrochemical experiment, where the energy current (in the form of electrical current times voltage) determines the amount of material dissolving per time unit. In both cases the currents can be increased by increasing either the temperature gradient or the electrical potential gradient right at the interface.
  How large these gradients will be depends on how the various thermal or electrical resistances are switched in series between the extremes of temperature or electrical potential. In other words: it is the *total* thermal (electrical) resistance between the hottest (highest potential) point and the coldest (lowest potential) point. The interface resistance is just a part of this total resistance. In either case the gradient increases if the difference between the extremes increases. The difference between the extremes in the electrical potential we call "**voltage**", by the way.
- In both cases there is a intermediate layer between the solid and the liquid. In the electrochemical experiment it is the oxide layer. The nature of the intermediate layer is not quite so obvious for casting or crystal growth, but there are several contenders:
    - The supercooled layer.
    - The layer containing all the dendrites, with a thickness measured from the tips of the dendrites to the completely solidified stuff the dendrites stick out off. One would guess that this layer scales with the thickness of the supercooled layer.
    - The layer enriched with rejected impurities and alloy atoms.
    - The boundary layer that always exists if a liquid flows across a solid.

  Of course, these layers are related and mostly you will have a "mixture". But whichever way you look at it, there is some layer with some defined average thickness that has properties that are distinctly different from that of the "regular" liquid or the solid.
- Dissolution and crystallization rates and morphologies depend heavily on the properties of these layers. In both cases the layers are subject to processes that *remove* them (oxide dissolution or diffusion / liquid flow / solidification between dendrites, ...) and processes that *produce* them (oxide formation or non-equilibrium solidification)

- In the electrochemical dissolution process, electrical breakdown supplies the mechanism for local stochastic oscillators, and interaction between the oxide bumps or the local intermediate layer non-uniformities, produce the synchronization.
  If we could identify local stochastic oscillators and next neighbor interactions for solidification, the analogy would be complete and we could expect that self-organized synchronized interface velocity oscillations will occur under certain conditions.
  - Well, there is at least one mechanism easy to conceive. All we need to assume is that the nucleation of a dendrite is a stochastic event, and that dendrites, after being nucleated, grow faster perpendicular to the interface for a while than the interface moves on average. What that implies is schematically pictured below.

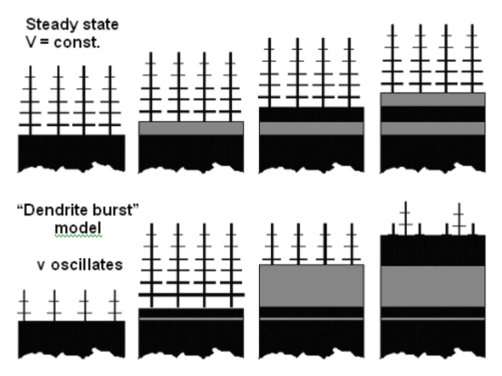

Comparison of steady-state growth and a "dendrite burst" model

  - The upper picture shows the usual steady-state model. Everything always looks the same as the interface moves upwards. The alternating black / grey layers symbolize the amount of solidified material after identical time steps.
  - The lower graphics illustrate a "dendrite burst" model. After nucleation, the dendrites grow faster than the rest, than the situation reverses. On average, exactly the same amount crystallizes as in the steady state model.
- I will rest my case at this point. Note that I did not say that solidification proceeds always in a "dendrite burst model". Even if it does sometimes, it will not do it for all conditions. It may never do it because I might be simply wrong.
  - Be that as it may, **I** strongly *believe* (but don't know for sure) that there are some mechanisms, roughly following the general arguments given here, that can and will produce interface velocity oscillations under some solidification conditions.

---

1) E. Foca, J. Carstensen, and H. Föll, "Modelling electrochemical current and potential oscillations at the Si electrode", J. Electroanal. Chem. 603, 175 (2007)

- **Link Hub:** **Segregation Science**
- Special Modules in Segregations Science:
  - **1. Basics of Segregation**
  - **2. Constitutional Supercooling and Interface Stability**
  - **3. Supercooling and Microstructure**
  - **4. Segregation at High and Ambient Temperatures**
  - **5. Striations**
  - **Segregation in Silicon**
  - **Microsegregation and "Current Burst" theory** This modul