

8.2.3 It's a Long way to Nirvana

The Aluminum - Copper Experiment

Let's look at a classical example that uses our strategies in a particularly transparent way. Thanks go to [Michael Ashby, David Jones](#) and their grad students who most likely did the actual work.

The idea was to **harden Aluminum** (Al) by adding a small amount of copper (Cu). This is important for making superior **beer cans** (and the cans for less interesting beverages). You don't want to entrust your beer to any can, you want a can with properties that are "[just right](#)". Here is a short list:

1. **High** mechanical strength when filled with beer. You must be able to stack a lot of cans on top of each other; rough handling must not lead to leakage.
2. **Low** mechanical strength when empty. Empty cans should be easy to crush with one hand or with little effort, to reduce the volume for transport and easy re-melting.
3. **Extremely low** costs. That means the can is made from a minimum of aluminum alloy that is easy to shape.
4. **100 %** recyclability. No weird things must happen upon remelting and solidification, even if the material is a bit dirty.
5. No **bad influence** on the beer inside. A totally passive or chemically inert surface is necessary

That calls for an aluminum alloy with precisely tailored properties, in particular hardness. Actually, it calls for *two* aluminum alloys with *different* hardness: one for the bulk of the can, and a somewhat harder (and thicker) one for the top with the pick-open device.

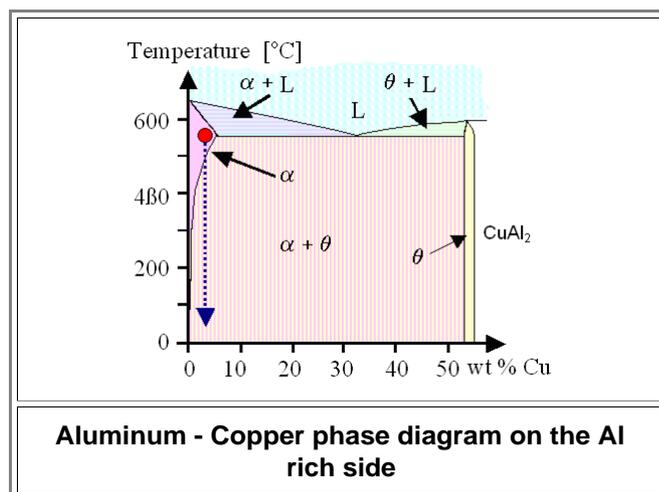
One of many possible ways to go is to use copper (Cu) as an alloying element. It must do the job already in small quantities since copper is expensive. In the "[Ashby and Jones](#)" an extensive systematic study is reported about what you can do with copper in aluminum. It serves as a paradigm, a model case for hardening metals by the mechanisms of

- [solid solution hardening](#),
- and [precipitation hardening](#).

I will guide you through this study because what is happening in the Al-Cu system is typical for most alloys—but far easier to understand than what happens in the Fe-C-some dirt system. Note that what follows is not the way the materials for real beer cans are made, it just provides basic and systematic knowledge for that venture.

First we alloy *very pure* aluminum (Al) with about 4 % of copper (Cu) and hold it at around 550 °C (932 °F) for a while. In other words: we make a rather ideal and expensive binary system.

Nothing helps, you now need to look at the [aluminum \(Al\) - copper \(Cu\) phase diagram](#). Use the link. Yes, it is horribly complex; a real nightmare. But if we just look at the *aluminum-rich* part, it is actually simple. Here goes:



In the Al-rich part, with just 4 % copper (Cu), at 550 °C (932 °F) we just have Cu *dissolved* in aluminum or the usual α -phase. That means that single Cu atoms are occupying the sites of Al atoms in its fcc lattice. The copper atoms thus are [extrinsic impurity atoms](#) of the substitutional kind. You just as well could call them *doping* atoms because they are intentionally introduced, but nobody does that. This term is more or less reserved for changing **electrical properties** by introducing impurities. We do call them "**alloying element**", though.

Our starting material at 550 °C thus is poly-crystalline aluminum with 4 % of copper (Cu) atoms fully dissolved, some dislocations and vacancies, and not much else.

At lower temperatures the crystal obviously needs to produce a few large CuAl_2 precipitates for achieving nirvana. The [chemical compound](#) CuAl_2 is also called the theta (θ) phase

- Next you [quench](#) to room temperature. You do that so fast that no precipitation occurs during the cooling. In the laboratory we can do that, effectively "freezing in" the high temperature structure. Then you temper by going up to a constant 150 °C (302 °F), i.e. you employ [strategy No. 1](#).
- When you start at the 150 °C, you now have **exactly** what you had at 550 °C: 4% copper atomically dissolved in an aluminium polycrystal. The only difference is that at 150 °C the crystal is unhappy about that. Very unhappy. It will now do whatever it can to remedy the situation and to establish the proper nirvana state, which is almost copper-free aluminum plus CuAl₂ precipitates. You also can call it a mixture of the α and θ phase; same thing.
- So all you do next is to let your specimen sit at the chosen temperature for a while (employing [strategy No. 3](#)). You picked 150 °C because at room temperature your aluminum cannot do much during your life time; everything is pretty much immobilized. If you temper at 150 °C, there is enough [thermal energy](#) to allow precipitation to happen, just rather s-l-o-w-l-y because the Cu atoms move rather sluggishly at that temperature.
- That is good because it gives us enough time to observe at our leisure what is going on.
- The important point is that the [driving force](#) for precipitation is large at this low temperature because the crystal can gain a lot of energy by getting rid of surplus copper atoms.
- A temperature of 150 °C combines an overwhelming need for making precipitates with some sluggishness in doing things. We're thus heeding the "[golden rule](#)" from before and control nucleation and kinetics. The crystal has no choice but to start a lot of very small precipitates all over the place. If there is a sufficient number of defects around, nucleation will now start at those defects, i.e. we have [heterogeneous nucleation](#). Otherwise nucleation will be [homogeneous](#) but the details don't matter for our purpose here. We just will produce a large number of nuclei soon after we reached the final 150 °C, and we do not need to care exactly how that happens and exactly how large that number will be.
- Before I go on let's reconsider the concept of the driving force; it is really important (you might want to go to the [bathroom](#) now if you remember the preceding chapter).
- The **driving force** for things to happen is just a measure of how much the crystal has to gain (in terms of energy) by doing something that gets it closer to nirvana. Consider steel just a few centigrade below the austenite-pearlite transition temperature. Yes, turning into pearlite would be better than remaining austenite, but just a *little bit*. You, the crystal, will go for it but not with real zest and vigor. The driving force for the transition is small. It may happen rather quickly anyway because everything is very mobile. Now consider that you quenched the steel very quickly to several 100 centigrades below the transition temperature. The crystal wants the same structure as in the case above but turning to pearlite at this temperature would be **far better** than being austenite, and transformation becomes supremely important. The driving force for the transition is large. The ominous "*driving force*" invoked a [few times](#) already thus simply measures the reduction of the distance to nirvana (the gain in energy in science speak) that the crystal can realize under the circumstances. Needless to say, we can calculate that and express it in a cold hard number. The message is clear:

Large driving forces induce desperate measures

- Crystals are like humans in this respect, with the exception that they are **always** reasonable. Humans, just like crystals, are less inclined to do the right thing under circumstances where the benefit is small. For example, stopping overspending and being thrifty becomes a far more urgent business when your debtors are closing in with swords drawn. As I'm writing this (July 31st 2011) bankruptcy for the USA is just two days away and politicians become frantic. Greece, very close to bankruptcy for about a year by now, is extremely agitated and induces desperate measures: frequent meetings of European heads of states, transfer of more than hundred billions of Euros (> 10¹⁰ €) that do not really exist, and severe violations of about every European treaty. Two years ago, the situation was not all that different but the **political** driving forces were much smaller because the rating agencies had not yet drawn their swords. However, in contrast to the forces driving crystals, we cannot calculate **political** driving forces precisely, because large amounts of irrationality and stupidity are involved that do not surrender to (mathematical) reason.
- Like always, there is a [catch](#) to producing a high density of small precipitates via large driving forces, meaning large degrees of sudden **undercooling** by quenching. While this will produce the nuclei for many very small precipitates, at low temperatures things move slowly. Those nuclei will take forever to grow to the size you want. Growth means that Cu atoms move by [diffusion](#) from the depth of the Al crystal to the nuclei.. If they diffuse slowly, growth will be sluggish.

- What can you do about this? Well, here we don't want to do anything about it because we like whatever happens to happen slowly. If we would want many nuclei but not-so-sluggish growth, we employ some smarter **temperature profiling**: Quench to a low temperature that produces a large driving force and thus produces a lot of nuclei. After enough time has elapsed to form many tiny nuclei, you go up to a specified higher temperature. Growth now is brisk but some of the smaller precipitates will dissolve and some others may get too large because [Ostwald ripening](#) cannot be avoided.

If needs be, repeat the process, perhaps with some in-between steps.

That's exactly what we do with [silicon](#) if we want to grow many silicon dioxide precipitates, which we may like to have in certain parts of the crystal.

- Does this sound a bit like working with wootz steel? Yes, indeed - all reports agree that the smith was employing many temperature cycles.

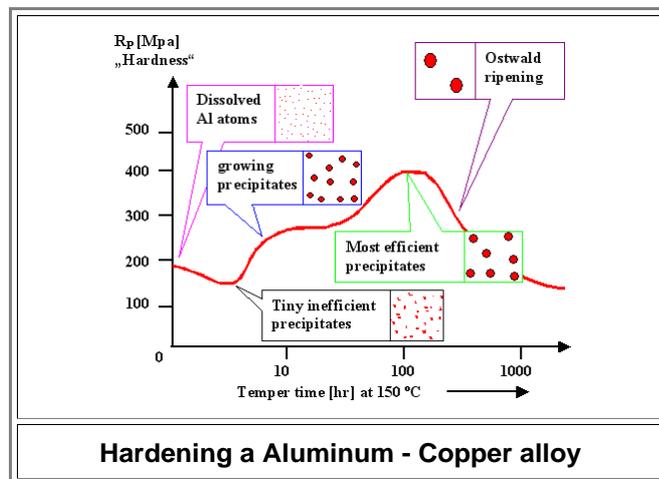
- Except, he did it in order to obtain just a **few large** precipitates, the exact opposite of what I described just above. He needs to do this so he can produce the famous "water" structure in this high-carbon or hypereutectoid steel. We'll come to that later in more detail.

Here you just can take it as hint that the golden rule is rather powerful, indeed, provided you know what you are doing. If you don't know what you are doing, you need to have a recipe, based on centuries of experience and untold failures, that works at least occasionally. "Untold failures" means exactly that: we have never heard, read or found the results of forging attempts that **failed** to produce a superior wootz blade!

- Back to our Aluminum-copper alloy. After **tempering** for some specified time at 150 °C (302 °F), you stop whatever is going on by now quenching to room temperature (strategy 1), where not much happens anymore. Then you test your tempered Al(4%Cu) alloy for tensile strength or hardness. That gives you just **one** data point for the curve in the figure below. Next you analyze your specimen by everything you have (metallography, electron microscopy, ...) to ascertain exactly what kind of structure was produced.

- Then you repeat the experiment with a **new** sample, waiting a little **longer** at 150 °C (302 °F) before you do the tensile test and the structure analysis.
- Then you repeat the experiment....
- Then you ...

- When you're done, you get the curve shown below.



- Then you repeat the whole **set** of experiments but with specimen where only **3 % Cu** was used for alloying.

- This gives another curve for this particular composition.
- Then you repeat the whole **set** of experiments with **5 % Cu**, and so on.
- When you're done, you repeat **everything** once more but now you hold the samples at **170 °C**, or **140 °C**.
- And so on at (almost) infinitum.

- A high and mighty director of a large research establishment once put it like that: "After doing this kind of work for a while, you start to realize that this not a suitable occupation for a self-respecting human being. That's when you realize that you need **graduate students**".

- You appreciate now why we need to have all those students. They enjoy (we believe) doing this kind of work for their Master or PhD theses.

The figure above shows just one of the many curves produced by scores of graduate students. It shows what you get for 4 % Al in Cu "tempered" at 150 °C (302 °F) for quite a while. Note that 1000 hours equal 6 weeks. What was measured is the yield strength or [hardness](#) as we know by now.

The Aluminum - Copper Experiment Explained

What we did was **hardening our Al - Cu alloy by ageing**. What happened is very clear (and supported by much more data). It illustrates quite nicely what hardening any metal is all about. You know, of course by now that it is all about the obstruction of [dislocation movement](#) by obstacles.

Now let's look at the data.

For the first few hours of tempering the tensile strength (or hardness) goes **down** because extremely tiny precipitates or nuclei form. That reduces the number of obstacles a lot. For example, if you form agglomerates that of just three Cu atoms, you reduce the number of obstacles threefold. In a volume that originally contained 90 Cu atoms for example, you now have 30 three-atom agglomerates. Each tiny precipitate or nuclei might obstruct dislocation movement a bit more than a single Cu atom but since you don't have that many obstructions any more, the total effect is **less** obstruction to the [onslaught of the dislocations](#). Look at the [analogy](#) again.

Eventually - after about 7 hours of annealing - the precipitates grew large enough to **seriously** obstruct dislocation movement and the "hardness" goes up. "Seriously" means that for example a 100 Cu atom precipitate obstructs dislocation movement more than 100 single Cu atoms.

The hardness reaches a **maximum** after about 100 hours; obviously at this point in time we have the optimal number-size relationship for the precipitates. Around this time we may also have used up all the available copper atoms.

If we keep at it, the hardness will go down again. That happens because now **Ostwald ripening** takes over. All copper atoms originally dissolved in the aluminum are used up and imprisoned in precipitates. Further growth of precipitates can **only** happen by the bigger ones cannibalizing the smaller ones, and that is exactly what happens. We get fewer and fewer but larger and larger precipitates. Few and very large precipitates do not bother dislocations all that much, hardness goes down.

Of course, we can learn far more from an analysis like that. But that involves cumbersome details and an equation or two. The link provides for that. Here we just note that the [battle line scenario](#) given before is rather realistic and now we even know how to deploy the troops we have in the best possible way - in a simple little war and without regard to money.

Iron is a far trickier opponent than aluminum when it comes to hardening, and it fights dirty. We also need to get to terms with it on a limited money and time budget, so we aren't done yet.

But let's continue with looking at the results of the beautiful Al - Cu experiment and ponder some of the topics it touches upon.

All in all, the maximum hardness we could produce is almost **three times** larger than the minimum one, and that's quite an achievement. Moreover, we know **exactly** what is going on, **and** we can control it to a large degree.

If Ostwald ripening seems still strange, here is a comforting thought. You can watch it happen while drinking **beer**. Right after pouring the beer consists of a liquid and foam (the head).

The head consists of many tiny bubbles formed by a thin membrane of the liquid and the gas enclosed. The beer is very unhappy with this. Its nirvana state would be what we call "flat" - no bubbles! Why? Because foam means liquid with a lot of surface area and thus unnecessary energy.

So Ostwald ripening commences. Some bubbles collapse, giving the gas they contain to others that are now bigger. And so on. Big bubbles come into being by eating up small ones (provided you let it happen before you guzzle them up). The reason for that is the same as for precipitates: the driving force is the possible gain in [surface energy](#) by reducing the surface area as much as possible.

Isn't that great? Finally you have a good reason to drink a lot of beer because you are actually doing science, studying Ostwald ripening.

Since it's a complicated process when you look at the details, I recommend lots of studying. You might augment your studies with some [Swiss cheese](#), realizing that the holes in there are precipitates of some gas that actually must have started with lots of little ones but, thanks to the second law and Ostwald ripening, developed into the few big ones you are familiar with.

Back to beer cans. (Hey, another scientific reason to drink beer!). With all the work that produced the figure above, we now know **one** temperature profile that forces Aluminum plus 4 % of copper to do what we want - within the basic limits of the system.

It is a simple profile: Cool rapidly to room temperatures and then go up to the desired temperature. Keep that temperature constant for as long as you fancy, then drop quickly to room temperature.

Simple - but there is a catch: it takes too long for the best result - about **100 hours!** Far too expensive. So, maybe, we should:

- Raise the temperature a bit, from 150 °C (302 °F) to 170 °C (338 °F)?
- Start annealing at 150 °C but then go slowly up to 200 °C (392 °F)?
- Do it, perhaps, the other way around.
- Use 5 % or 3 % or ?% copper.

[Advanced
Module](#)

**The Al-Cu
system**

[Fun
Module](#)
**Beer bubble
science**

- Use some other atom instead of copper.
- *Your suggestion.*

● You can spend any amount of time tinkering around with optimizing beer cans, boring your female graduate students to tears. Nevertheless, you can never be sure that you found the best way to make your beer cans. So let's have another big epiphany:

**First law of Practical Science:
There is nothing more practical
than a good theory**

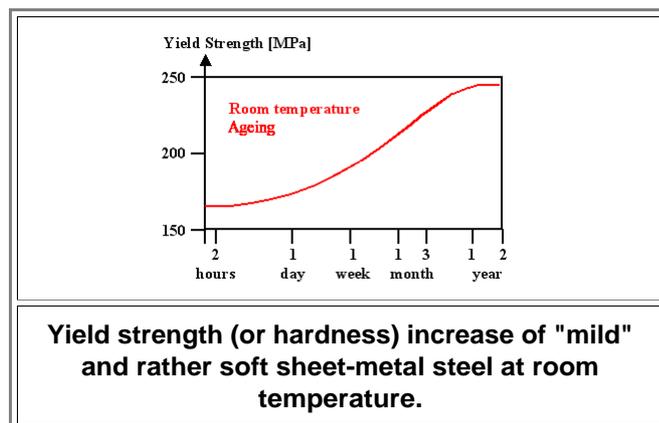
▶ If you have a (mathematical, of course) theory about hardening aluminum (and anything else) you could just *calculate* all of the above. It not only saves a lot of time and money, you also can be sure that you arrive at the best possible process - not to mention that your graduate students now can try their ken at something more challenging, like making superior steel for swords. And you have more time to study Ostwald ripening with beer, so everybody profits.

● You have a *why* question? Why, for demonstrating how to outwit the second law, did I pick an example from the boring world of aluminum and not from the exciting world of steel ?

Well, I invoked the [principle of supreme laziness](#).

▶ The iron -carbon system that we call (carbon) steel is far more complicated than the Al - Cu system. Looking at a comparable experiment for steel would have needed far more writing, which is hard work that keeps me from studying Ostwald ripening.

● OK. Just to keep you in good humor, here is an actual and interesting steel example:



▶ We have soft, low-carbon sheet-metal steel that shows the effect of **ageing at room temperature**. Its hardness changes ever so slowly at room temperature but after a year has passed it is substantially increased.

● You might think that this is good. No. it isn't. Hard is not always better than soft; hard steel sheets can be quite irritating if you want to shape the stuff easily. The Aluminum foil you use to wrap your food with shouldn't be too hard, for example.

● What has happened? First of all, you still have some excess carbon in this iron at room temperature. Not all of the stuff has precipitated into cementite, probably because it was cooled too quickly. That's easy to do with thin sheets: just letting it "sit" after hot-rolling might be sufficient.

At room temperature carbon interstitials and some other interstitial like nitrogen (N) are still a little bit mobile; look at the [numbers](#) I provided a long time ago. The carbon atoms move ever so slowly (roughly a few μm per year) until they get caught at some precipitate or in the core region of a dislocation, where they come to rest.

The net effect is that they make it more difficult for the dislocations to move when the need arrives, and that means the material is now [harder](#).

▶ Now let's get back to the topic we're actually looking at. You probably forgot it by now, so let's dedicate a whole paragraph to it.