


5.3.4 Bulky Things and Nucleation

Three-dimensional defects are next on our list. They are rather common and you simply cannot have decent steel without them. Since they are large in all three directions, they can be rather bulky and we might call them **bulk defects** on occasion.

Voids, small holes in the crystal, are also three-dimensional defects. They might come into being because a lot of vacancies cluster, or by some other mechanism. Voids are not important for us; I just mention them to be complete. The three-dimensional defects that are important to us, very much so, are **precipitates** or small particles of a different phase contained in a host crystal.

To get a first idea of what it's all about, we give the steel with about 1 % carbon as [shown before](#) a second look; the picture here is a reminder. The news is that there is no iron - 1 % carbon steel that looks like that at room temperature. Neither schematically nor in reality. It is important to appreciate that while the iron crystal is happy with 1 % *dissolved* carbon at *high* temperatures, it really doesn't want to have that much loose carbon around at *room* temperature. Disorder does not much good at low temperatures and it takes simply too much energy now to keep the carbon atoms dissolved.



Nirvana-wise, one atom, perhaps, would be fine in the area shown above at room temperature, but no more! Since the crystal can't get rid of the surplus carbon anymore than you can get rid of your [unruly offspring](#), it will round up the surplus carbon atoms and imprison them. Some will be locked up in [grain boundaries](#) but 1 % carbon is too much for the grain boundaries.

The crystal, searching for but not getting the best—nirvana—goes for the [second best](#) and invents something new:

**A new phase will form,
consisting of iron-carbide (Fe_3C)**

We call this new phase a carbon **precipitate**, and the process by which it forms **precipitation**. Note that this involves a lot of iron too. In an iron-carbide precipitate there are three iron atoms hooked up with one carbon atom. They are still iron atoms but no longer iron, just like the oxygen atoms in water (H_2O) are still oxygen atoms but no longer air. In principle the surplus carbon atoms could also precipitate as pure carbon (graphite in this case) and they actually do it: in cast iron but not in wrought iron or normal steel. Iron-carbide precipitates forms because the surplus carbon atoms, while wandering around at random, will get stuck at some tiny Fe_3C carbide [nuclei](#) that then grows into a large precipitate.

Of course, we now have a hen-egg or **nucleation problem**:

You need a precipitate to make a precipitate?

Carbon atoms only can get stuck somewhere if there is at least a small precipitate, or, to give it its proper name, a **nucleus** (singular of [nuclei](#)). So the question is: how does the *nucleus* come into being?

The first tentative answer is: The carbon atoms move around at random. On occasion two will accidentally run into each other and then stick together. Six of the surrounding iron atoms join the two carbon atoms, so that two Fe_3C molecules can be formed. The 8 atoms are somehow squeezed into the iron crystal. When another carbon atom drops by, the cluster grows to 3 Fe_3C molecules, and so on. If that sounds like it would be an unlikely and difficult process, that's because it is an unlikely and difficult process. It is not impossible, though, and nucleation in this way does happen on occasion. So take note:

**Nucleation is always
a very difficult enterprise.**

If nucleation rarely happens as described above, how does it happen? The answer is: nucleation is always much easier with some outside help. For nucleating a baby, almost all mothers need a bit of outside help too. For nucleating precipitates and other stuff, [grain boundaries](#) for example offer that outside help (they are "outside" the crystal grain, after all) and make good nucleation sites. Let's see how that applies to making an iron carbide (Fe_3C) precipitate:

Carbon atoms move around at random, and once in a while hit a grain boundary. Looking at all the pictures I provided in [chapter 5.3.1](#), you will appreciate that at a grain boundary there is just more space available for lonely carbon atoms that feel somewhat cramped-in and unloved between all the iron atoms. So the carbon atoms stays a while in the grain boundary, *far* longer than in the lattice, before it moves on again.

The chances of meeting a *second* carbon atom that happens to come by are thus much larger than in the lattice, where everybody always jumps around at high rates. More women / men come by and talk to you for a while when you hang out at the bar compared to jogging around at random in the park.

An additional advantage is that the iron atoms ~~at the bar~~ in the grain boundary are not so solidly bonded to other iron atoms as the iron atoms ~~at home~~ in the lattice, and thus more amenable to engage into some iron-carbon bonding, the very first step on the way to an iron carbide precipitate.

This first pre-nuclei (it's not an iron carbide *crystal* yet) is now immobile. All it needs to do is to wait until more carbon atoms come by, to grow into a nuclei proper (tiny crystal), and finally into a precipitate with a diameter considerable larger than the unit cell of the Fe_3C crystal. [This link](#) gives a first glance at what an iron carbide or Fe_3C crystal looks like. Just get used to the fact that iron carbide or *cementite*, as we will call it, is an integral part of steel.

Needless to say, we can calculate all of that, using equations that contain the energy of interstitial carbon in a grain boundary and in the lattice, the probabilities for jumps in the lattice and at the grain boundary, and so on.

Of course, the meeting of carbon atoms and so on inside the grain, as described above, is not forbidden. We even have a name for it, we call it like **homogeneous nucleation** in contrast to the **heterogeneous nucleation** that relies on some **defects as nucleation** sites.

Homogeneous nucleation as described could and would happen, it's just unlikely. Two or a few more carbon atoms meeting accidentally will easily stick together but just as easily break apart again. If they meet at a grain boundary or some other suitable defect, they just like it there better and hang around until a third and fourth one comes along. There is simply a *competition* between homogeneous and heterogeneous nucleation. Heterogeneous nucleation at defects *always* wins, provided that there are enough suitable defects like grain boundaries around.

If you remember what I claimed way back about the [silicon crystal](#) for micro electronics (or even looked up the "[single crystal silicon](#)" module) you have a question now: What happens nucleation-wise in the rather perfect single silicon crystals? There are no grain boundaries or other big defect that could help nucleating whatever needs to nucleate. Tricky! Look up [this module](#) for an answer.

But iron and steel are *always* polycrystals, full of defects. If we talk about nucleation of something in iron and steel, it is sufficient to remember:

Precipitates need to nucleate somehow.
In iron /steel they *always* nucleate at grain boundaries or other defects
and rarely "just so"

I can't emphasize it enough, so once more: We need some process of "nucleation" to get the things that *want* to happen to *really* happen. Nucleation of carbide precipitates is one of the very important but also quite tricky parts when a sword blade is forged. It will typically happen at grain boundaries or other defects.

This is nothing special. "Seeding" clouds with silver iodine (AgI) so it rains, is working with nucleation in the same way. Before a water drop can grow to a size large enough to fall down (= rain), water vapor must condense on a nucleus, a very small water droplet, and those nuclei are very difficult to form. Some "defects", small silver iodine particles in this case, might be helpful to get nucleation going. By the way, all of that is difficult but there is nothing mysterious to it. I can calculate *exactly* what is going on. The science *super* module has a lot to say about nucleation, make sure to check it out. There is no problem in understanding nucleation.

The only problem is that the halfway decent education, that I [credit you with](#), does not (yet) include the basics of everyday life.


Now let's assume for the time being that the nucleation of carbon precipitates in our sword blade was somehow taken care off. The crystal can now precipitate as much carbon as it likes to. In other words, the crystal can now form an iron carbide crystal—something as different from iron or carbon as sodium chloride (common salt) is from sodium or chlorine—that is completely contained inside the iron crystal lattice.

How will it proceed?

[Science](#)
[Super](#)
[Link](#)
Nucleation

- Locking up the surplus carbon or, using our new and fancy word, *precipitating* it, can be done in many ways after the nucleation problem was solved.

For example, you could produce a lot of small precipitates or just one big one, as shown schematically below.



- ▶ Now a word of warning is in order ([once more](#)). Looking at the figure above or others like these, you should realize that there is no such thing as a fully correct drawing of precipitates in crystals.

- First of all, the best I (and everybody else) can do most of the time is a two-dimensional cut through a three-dimensional structure. Ask yourself: is the big precipitate shown above a cut through a roughly spherical precipitate or a cross section through a long needle? There is no way of telling just from the picture. Worse, drawings that are both to scale *and* correct in the sense of showing the proper lattice of iron, the proper size of the atoms, the proper structure of the precipitate, the defects inside that provided for nucleation, and so on and so forth, would be so complicated that I simply refuse to do them. They also would not show much because the first layer of atoms would obscure everything behind. So take all those figures with a grain of salt. It doesn't matter much, however, that the figures cannot show all there is to precipitates. If you don't get the basic ideas I'm driving at with all the imperfect drawings (and I don't assume for a moment that you do not), changes are rather slim that you would get them with better drawings. Moreover, in real steel, as contrasted to the figures above, you might have to compare one precipitate containing a couple billion atoms (still only a few 100 nm in diameter), or 1000 precipitates with just a couple of million atoms each. Hard to draw.

- ▶ Whatever. The important thing to note at this point is that the *total* area of all phase boundaries enclosing the precipitates, is *always* much larger for many small precipitates compared to a few big ones that contain the same amount of precipitated stuff.

You know that. You get more potato peel from 20 small potatoes than from one big one.

- To make things a bit more complicated so we don't get bored: in three dimensions you could make the precipitates round and ball shaped ("**spheroidal**"), looking more like needles ("**acicular**"), plate-like as a coin, or just irregular like a ginger root. They all would look the same in the drawings above. So how do you, the crystal, like you precipitates? *Chunky* (a few big ones) or *smooth* (many little ones)? Spiky (shaped a like a sea urchin) or round like a ball? Donut shaped, perhaps?

- ▶ The long and short of the precipitation business is that a crystal has many choices for locking up surplus impurity atoms in precipitates. The way the crystal actually does it will decide, for example, if your sword blade is useless brittle steel or superior wootz steel.

- So what will *the crystal* do? Even more important: What can you do? How can you induce the crystal to do *your* bidding (you're the smith, after all—[remember](#)?). And what would your bidding be? Do *you* go for just a few big precipitates or lots of small ones? Round ones ("spherical") or longish ones ("acicular"), and so on? You can't give me an answer? So go on for a few more chapters.

- In contrast to you, the crystal knows exactly what it wants. It lives faithfully by the [second law](#) and thus knows what would be best for it.

For the best of all worlds or *perfect* nirvana at low temperatures like room temperature, it would get rid almost completely of the surplus carbon, just keeping the precise (extremely small) concentration that the second law demands for the particular occasion.

If that is not possible, it will go for the *second best* state of being and that is to have few and then necessarily big precipitates with spherical shape. If you don't see why, it will come to you later.

▶ The problem with the crystal's choice is that *we* don't like it. *Not at all!* You and I usually prefer a lot of small precipitates; just as we [like small grains better](#) than bigger ones.

- What we take from this is that a lot of what we do when we forge a sword is to keep the crystal from making large precipitates.

Except if we make a wootz steel sword.

- Then we like our iron carbide precipitates to be rather large, and in the "right" places; I'll come to that.

▶ Precipitates always have a [phase boundary](#) around them. They cannot avoid this any more than you can avoid having a skin around yourself if you like to stay "inside".

- The interface between the precipitate and the crystal is by definition a *phase boundary*, and phase boundaries are high-energy defects, just as grain boundaries and surfaces. We have an unavoidable *combination* of two kinds of defects here. Three-dimensional defects are *always* surrounded by two-dimensional defects. That is the reason why nucleation is always difficult as we shall see presently. [Go ahead](#), if you like; or jump to the [science module](#).

Phase boundaries come with unwanted energy. We have already ascertained that the crystal hates them just as much as it hates grain boundaries. It doesn't want them at all. However, while it could get rid of grain boundaries in principle if not in practice, it *cannot* get rid of phase boundaries as long as it wants precipitates.

We will learn presently that in certain conditions the crystal *needs* to have precipitates for achieving nirvana. But you can't have precipitates and not have phase boundaries just as you cannot have your cake and eat it. Compromises are called for.

- As a last and now rather trivial point, we realize that when a precipitate grows, its phase boundary has to move outward and its total area increases. Here we are [again](#): movement of our two-dimensional defects occurs quite often and it needs atoms to be mobile for that.

▶ Now the time has come that we can dare to give the last defects on our list a thorough look: one-dimensional defects or *dislocations*. These defects, while crucially important for sword making (and the whole [metal-bending industry](#) plus semiconductor industry) are also the most difficult ones to come to term with—so I saved them till the end.