

Have you ever seen a phonon?

Anders Blom

Division of Solid State Theory
Department of Physics
Lund University
Sövegatan 14 A
S-223 62 Lund
Sweden

Anders.Blom@teorfys.lu.se

Introduction

Most basic courses in solid-state physics touch upon the concept of lattice vibrations, in order to explain sound and heat transport and various other important everyday phenomena occurring in solid materials. The students are exposed to some mathematical derivation of a so-called dispersion relation, and somewhere along the line the discussion mysteriously shifts focus from the original atomic vibrations, which are quite tangible, to something called a *phonon*. This object is attributed particle-like properties such as momentum and energy, and is therefore unfortunately perceived by many novices as yet another elementary particle (which it is *not*); if they pay close attention to the lecturer, they may even catch the fact that it is a spin-less boson.

This picture is nevertheless often sufficient to solve homework or examination problems, but it would certainly be desirable to bestow the students with a more correct picture of the concept of phonons. In particular, it is necessary to restore the connection to the actual lattice vibrations.

Since vibrations involve many atoms that influence each other through the interatomic elastic forces, one is dealing with a *collective phenomenon*, as opposed to the single-particle picture that pretty much dominates the description of most physical systems in undergraduate courses. Therefore, the proper understanding of phonons requires a higher level of abstraction than usually is assumed at this level.

At Lund University we have therefore designed a computer exercise devoted to phonons. The idea is that by combining numerical and graphical tools to visualise the temporal and spatial variations of the lattice vibrations, the students will obtain a more intuitive feeling for the relationship between the formulas and the physical reality that they attempt to describe. As an additional bonus, the numerical model can be extended beyond the traditional textbook treatment. Impurities have important consequences for the low-temperature electrical conductivity of metals, but even the simplest models, such as a single impurity in an infinite linear chain, cannot be solved analytically. This system can however easily be considered with our computer model.

Here we will summarize the experience gained during the development and teaching of this computer exercise, touching on topics such as computational details and how the computer work can be integrated with the lectures. In addition, the numerical model used and the results derived from it will also be presented, with emphasis on the pedagogical aspects.

Phonons—the hard way

Before we enter the realm of numerics and computers, let us first satisfy the impatient reader by providing an answer to the conundrum posed in the introduction, regarding the connection between phonons and lattice vibrations. The following simple picture is often offered to clarify this relationship.

If an atomic vibration is somehow initiated at one end of a solid body, the vibrations will spread throughout the rest of the crystal. One can thus say that we created some (probably many) phonons at the point of excitation, and that these travelled—in a particle-like manner, if we so wish—to other locations in the material, where we later detect them by noticing that the lattice is now vibrating there. The process just described may represent sound propagation or heat transport in the material.

Thus phonons appear to exhibit the same wave/particle duality that we have become used to accepting for quantum-mechanical objects. We must nevertheless still reject the notion of phonons as elementary particles like photons or electrons.

Since phonons are intimately linked to lattice vibrations, the existence of a phonon requires a medium, in fact a regular crystal lattice; sound does not propagate in vacuum. One should also carefully note that in the same way ocean waves do not move water around the globe, phonons do not give rise to any net transport of mass, since the atoms vibrate about their equilibrium positions—only energy is propagated. Therefore, phonons cannot be attributed any mass.

Unfortunately, at this point, the students are most likely completely confounded, contrary to the initial intention behind this ‘illuminating’ example, and the main question is still suspended in mid-air: So, what is a phonon, really?! To answer this, it is naturally utterly unpedagogical to attempt to introduce the language of many-body response theory, where a phonon is defined as a pole of a dressed one-particle boson propagator. Also the notice of collective modes and quasi-particles in the sense of Landau and Fermi liquid theory is beyond the reach of undergraduate students. However, but let us cling on to the keyword *collective*!

For in fact, a phonon is actually nothing but the vibrations themselves; when ‘the phonon travels somewhere’, this simply means that ‘the lattice vibrations spread out (in a certain direction)’. A phonon is a quantum of vibrational energy, in the same way a photon is a quantum of electromagnetic energy, and the dispersion relation describes the relationship between the energy of each quanta and the wave length (actually the wave vector) of the corresponding lattice vibrations. Like the photon, a phonon is a boson, since we may excite infinitely many quanta of the same energy without restrictions.

So, the important conclusion is that phonons do not really exist, but they are merely a useful way of describing a collective phenomenon. There are, in fact, many different collective excitations to which it is convenient to assign particle-like properties, and excitons (electron–hole pairs), bogolons (superconductivity), magnons (spin-waves), plasmons (charge oscillations), polarons, skyrmions, and many other -ons make up a whole quasi-particle zoo.

Visualising physics

The reason for including this admittedly somewhat overloaded discussion is certainly not to dazzle the reader with fancy words, nor to cause confusion and despair regarding the possibility to properly understand the concept of phonons. The point is, that sometimes words are not the best way to explain a complicated picture. This is in some sense obvious: it is clearly superior to show a painting than to attempt to describe its colours and scenery.

It is this lesson that teachers of physics (and other subjects) should take to their hearts: sometimes things which are so painfully complicated to explain by use of special terminology, formulae and reasoning, can be so immediately apparent when cast in a different, for instance graphical, shape. Notably, this does not mean—as some prefer to argue—that the formulas should be abandoned; they are always required to produce quantitative results and, importantly, to enable further development of so-far unknown relationships. The tools of teaching could

however be diversified to include the possibilities that computer animation and visualisation offer.

The use of computers as an aid in education has of course already been exploited in many ways: to plot complicated relationships, solve difficult equations, illustrate complex geometries, and display dynamic phenomena. Many of these topics are particularly applicable to the teaching of solid-state physics, which by construction involves all three spatial degrees of freedom and time.

For maximal pedagogical impact, computer visualisations should ideally be closely integrated with the lectures. Although many classrooms are still not equipped with the proper technical facilities, this is probably just a matter of time (and money). A greater hurdle to overcome may rather be the reluctance that many can feel towards introducing new elements in their teaching methodology. However, using ‘electronic slides’ instead of plastic transparencies, allowing for easy updating and sharing via e.g., dedicated course web pages, and displaying dynamic visualisations during the lectures, may soon become commonplace.

On the other hand, one of the great advantages of interactive simulations is the possibility for the students to fiddle with the software themselves, and get hands-on experience with the subject. A simple solution is to provide them with the relevant software, and trust that at least someone will go home and try it themselves.

An alternative, or complement, to these two methods is to create a *computer exercise*. This means that the hands-on part is done in groups with a tutor present, and the computer work is interlaced with a sort of lecture. This approach has so far worked excellently at Lund University, and we are particularly pleased to note that almost all students state on the course evaluation forms that the exercise was very useful, and helped them understand the concept of phonons much better. An important ingredient was also that a detailed manual was developed, which contained details on the theory of phonons, as well as clear step-by-step instructions on how to execute the computer work itself; in fact, if necessary one could carry out the exercise as a self-study.

Nevertheless, there are some complicating issues that could make the computer exercise approach less attractive in some situations. Laboratory work is traditionally supervised not by the lecturer but by graduate or senior students. The lecturing element of the kind of computer exercises we are discussing here demands skills that they may not possess, due to lack of pedagogical training. Many computer halls are not equipped to allow comfortable lecturing; the students have great big monitors in front of them and will easily be distracted by it. Also, two or three students must often share a computer.

It is also a bit more difficult for a new assistant to inherit a computer exercise from someone else, in particular if the software was developed locally. The assistants must not only manage to adapt the lecturing to their own liking and abilities, but also learn to maintain the code. After a few generations of assistants and the parallel evolution of

computer program versions, modifications (perhaps small, but crucial) may more than likely be required.

Finally, a general hazard with numerical models is that it is 'too easy' to obtain the results; you just click here and there, and the computer does all the work. The students may not be equipped (or motivated) to proceed beyond the computer model and into the analysis of the results, which is the crucial step in order to attain a proper understanding of what the results actually mean. This is an issue which must be taken into consideration in all numerical work, and involves topics such as going to extreme limits to detect errors in the model (or its implementation), analysis of symmetries, interpretation of unexpected results, and so on.

As for the implementation itself, a prudent first step in any development process is to clearly define the desired final outcome; in our case, this means specifying what we wish to teach the students. For visualisation problems, it is however also important to already at an early stage consider what environment will be used for the implementation. In particular, one often finds that in the type of applications we are discussing here, there is actually only a very small amount of physics involved; most time and energy is spent on the graphical presentation. The algorithms describing the physical model itself are often quite simple to implement, and the important issues are instead things like which parameters the user should be able to change and how, what information should be presented on the screen, and so on.

The ideal software environment for implementing physics visualisation applications should therefore possess both strong numerical capabilities and flexible graphical elements. It is obviously difficult to make specific recommendations in this area, since many factors, including personal preference and financial issues, are involved. Nevertheless, in our case, *Matlab* was deemed most suitable, for many reasons. Apart from fulfilling the general requirements mentioned above, *Matlab* is highly capable at handling matrix manipulations, which was an important part of the program, and it is available for a large variety of operating systems without any need to rewrite the code.

Our numerical phonon model

The particular system we decided to consider was the linear harmonic chain, treated within the usual adiabatic approximation. Of course, this problem has simple analytical solutions given in every solid-state physics textbook. However, it is precisely the simplicity of those solutions, such as the dispersion relation, that so deceptively mask the underlying connection to the equations of motion. Therefore, even within the very simplest phonon model, the students are able to familiarise themselves with the relationship between the formulae and the tangible physical vibrations by using the visualisation program.

The textbook ideal linear chain is an infinite system. This is however numerically inconvenient, and instead we employ a typical trick, namely to apply periodic boundary conditions. In reality, this means that we put our atoms on a ring, but the result is the same: the vibrations can propagate in both directions in space forever. In order to make it

possible to discuss the influence of the boundary conditions, there is also an option to use fixed ends (for which standing waves are unavoidable), but we will not include that case in the present discussion.

Monatomic linear chain

Once the boundary conditions are specified, Newton's equations for the harmonic one-dimensional chain of atoms are easily cast into a matrix eigenvalue problem. The eigenvalues of the matrix describe the dispersion relation, and the eigenvectors correspond to the relative atomic displacements at a given reference moment in time. The matrix itself depends on the mass of the atoms, the force constants and the boundary conditions. The mathematical details of the model are presented in a longer version of this article, available at <http://science.uniserve.edu.au/pubs/callab/vol12/blomfullpaper.pdf>

We shall first assume that all atoms and force constants are identical, and consider an ideal monatomic linear chain. In order to be able to compare our results with the analytic solutions, we furthermore assume that only nearest neighbours interact with each other.

One of the most important lessons to be learned from considering this textbook model within the computer program is to clarify the meaning of the various *modes*, a term used in many contexts with different meanings. Here, a phonon mode means a point on the dispersion curve, and each such point corresponds to a solution of the equations of motion. In the simplest case, with a one-dimensional linear chain of identical atoms, there is a single dispersion curve or *branch*; with two different atoms in the chain we get two branches, known as the optical and acoustic branches (see below), and in real three-dimensional crystals there are multiple longitudinal and transverse acoustic and optical branches.

Each point on any given phonon branch is described by a given value of the mode wave vector (or in one dimension the wave number), and the corresponding energy. There are infinitely many solutions in the ideal infinite case, and a discrete number equaling the number of atoms in the ring in the case of periodic boundary conditions. The dispersion relation thus tells us the size of the energy quanta that can be transported by a vibrational mode with a particular wave vector. These energy quanta are, as mentioned earlier, precisely the phonons! The derivative or slope of the dispersion relation furthermore tells us the speed—the group velocity, to be specific—at which these quanta propagate.

In our model, the energy values of the dispersion relation are simply obtained from the eigenvalues of the matrix problem. The matrix solutions, however, do not *a priori* contain any information about the wave numbers. This is not surprising, since the wave numbers were never needed to set up the model. We do however know from Bloch's theorem that the wave number is a good quantum number in a perfectly periodic system, and thus somehow the eigensolutions to the matrix problem must contain information about the wave numbers. In principle we could extract them by just comparing the matrix eigenvalues to those predicted by the analytic theory, but our computer

model is able to consider more general problems, such as a chain with an impurity (see below) or different boundary conditions.

In both these cases, the periodicity is broken, Bloch's theorem no longer applies, and there are no analytic solutions to compare with. Of course, at the same time the wave numbers are no longer good quantum numbers, but they are still 'approximate' quantum numbers, and we would like to be able to at least estimate them. Therefore, a method was developed, in which we take the spatial Fourier transform of the eigenvectors, and we then associate each eigenmode with its dominant Fourier component. A detailed analysis shows that this method will exactly reproduce the correct wave numbers for the pure chain, and experience has shown that it gives good approximate quantum numbers in other cases as well.

Diatomic linear chain and impurity atoms

Obviously, the linear monatomic chain is an extreme oversimplification of a real solid. However, to make a more realistic, three-dimensional model including interactions and so on is a huge step, and requires extremely elaborate calculations. Instead, we try the next simplest thing, namely the diatomic linear chain, in the hope that some interesting physics will emerge.

Again, this problem has a simple analytic solution, but also in this case can the computer model be used to understand exactly those solutions, and in particular to work out the fundamental difference between the *optical* and *acoustic* branches.

The numerical model is the same as before, but we now replace every second atom with an atom of a different kind, although we still keep the force constant between all atoms equal. It is then not difficult to see (again we refer to the longer version of the article at <http://science.uniserve.edu.au/pubs/callab/vol12/blomfullpaper.pdf>) that the only modification needed in the matrix representation of the problem is to replace the constant atomic mass by a mass matrix. In fact, this approach allows us to consider any configuration of masses we like. Also, we can also easily include interactions between not only nearest neighbours, but between all atoms, with different force constants. The main difficulty will then rather be to interpret the results.

We can however make a particularly interesting modification, which is quite simple to analyse. If we replace a single atom, either in the monatomic or diatomic chain, with an atom with a deviating mass, we will be able to study the formation of a localised phonon mode. This will be discussed in more detail below.

The phonon computer program

As shown in Figures 1 and 2, the phonon computer code, developed in *Matlab*, displays the dispersion relation and the atomic displacements for the selected eigenmode frozen in time, and is also able to animate these displacements in time, in order to visualise the vibrations. The adjustable physical input parameters are the atomic masses (or the mass ratio for the two-atomic case), the force constant, the number of atoms in the chain, the boundary conditions, plus the possibility to introduce an impurity atom.

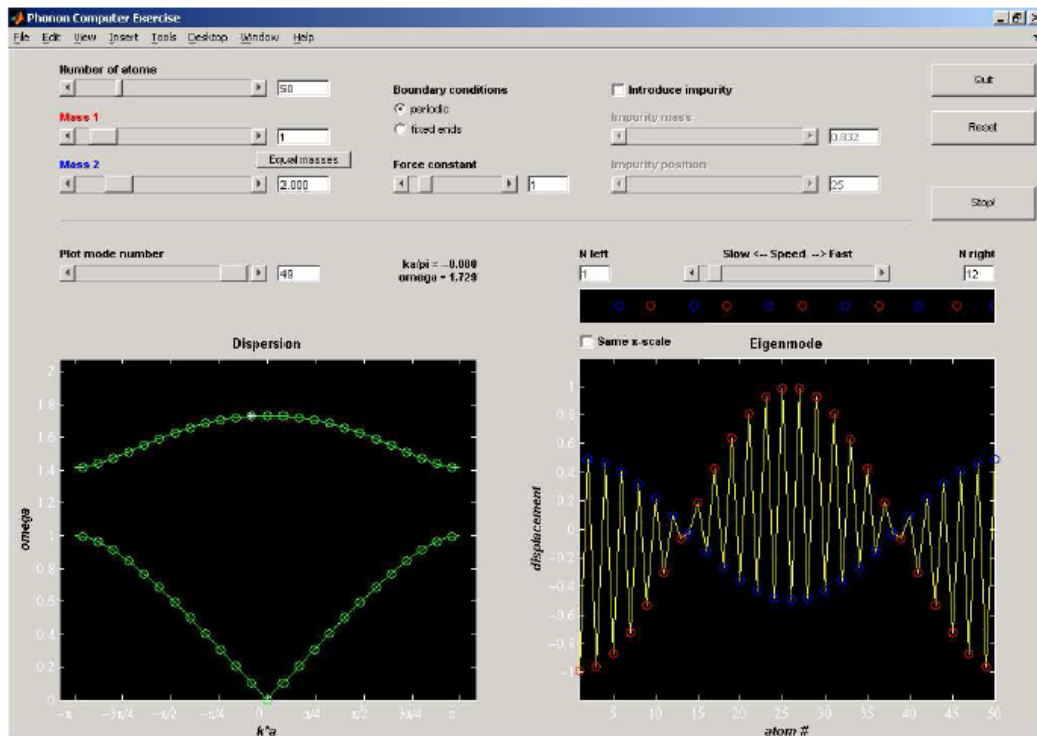


Figure 1. Screenshot of the phonon application running in *Matlab*. In the depicted situation a two-atom basis is chosen, and we see the dispersion of the acoustic and optical branches in the figure on the left. On the right is shown the relative atom displacements; note carefully that the vertical axis depicts the *longitudinal displacement*; we are considering a strictly one-dimensional chain without any transverse modes. The small window above the displacement plot is used to display the animated vibrations.

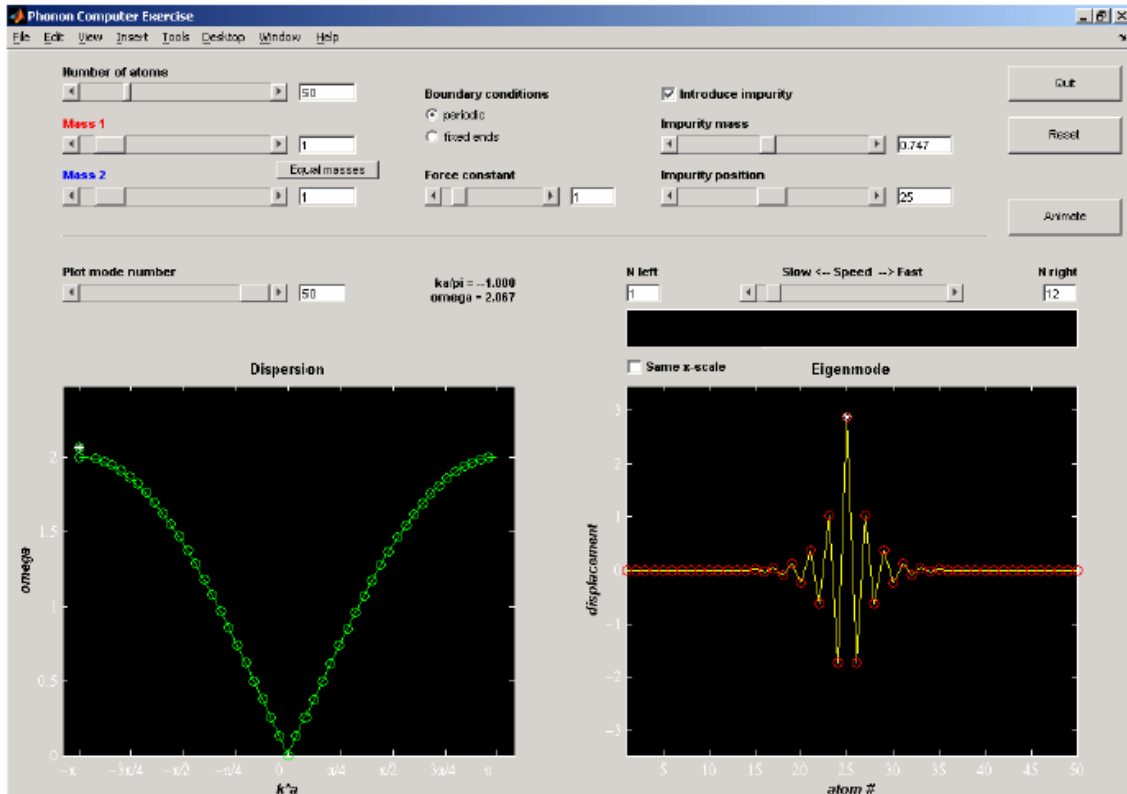


Figure 2. Another screenshot, this time with an impurity present. The selected eigenmode is the localised mode, as is clearly seen in the displacement plot.

Ideal linear atomic chain

After becoming acquainted with the dispersion relation and the various phonon modes, and understanding the operating principles of the numerical model, the first real task for the students is to attempt to visualise the phonon itself. This is a bit tricky since, as we have pointed out, the phonons are actually nothing but the lattice vibrations, and thus in some sense trivial to envision. This triviality is however still deceptive, since for some reason most of us picture *vibrations* as ‘static’ standing waves; the atoms may

vibrate, but they do not move anywhere, so how can the phonon propagate in space? Therefore, when a mode is selected, the program displays not only the atomic displacements frozen in time (mathematically, this corresponds to the mode eigenvector), but also animates how the vibrations evolve in time. One should then be able to make out how contractions and elongations propagate along the chain (cf., Figure 3) i.e., see the phonon move!




Figure 3. Can you see the phonon move? Each horizontal line depicts a snapshot of the vibrating atoms (represented by the black balls), with time increasing downwards. By focusing on a contraction (such as the first 5 or 6 atoms) and following it as time goes on, one sees that the contraction moves towards the right. This corresponds to the phonon propagating in the positive direction! Note that the amplitude of the atomic vibrations is unrealistically exaggerated for clarity.

The second important point of the computer exercise is for the students to study the diatomic chain, and to understand where the names of the acoustic and optical branches derive from by observing how the atoms vibrate in modes belonging to each respective branch. They should be able to notice how the two atoms in the basis are vibrating in phase in the acoustic branch, while in the optical branch they are out of phase.

Acoustic phonons can be excited with practically any small amount of energy, and in these modes the phonon wave lengths are macroscopically long. These features nicely demonstrate how acoustic phonons are mainly responsible for sound and heat propagation in solids. Moreover, the very near linearity of the dispersion relation for small wave vectors in the acoustic branch ensures a common group velocity for all frequencies. In contrast, the dispersion of the optical branch is flat near the zone center, and thus optical phonons move rather slowly. The optical phonon energies are comparable to the thermal energy at room temperature, and to the energy of photons in the infrared/visible range. The out-of-phase vibrations imply that dipole fields, including photons, can excite optical phonons.

Impurities and localised phonon modes

Apart from the pedagogical merits, which by no means should be underestimated, it is honestly speaking rather pointless to consider a numerical approach to such a simple model as the ideal linear chain, since we already know the analytical solutions. Things immediately get much more interesting when we introduce an *impurity* in the chain.

Let us consider the most simplistic model for conductance of electricity. By applying a voltage to the ends of a metal bar, we create a constant electric field inside the solid. Hence the electrons are subjected to a constant force, and will thus be perpetually accelerated. However, experiments show that the electrons rather travel at a constant, relatively low, speed.

From this we conclude that the electrons must be held back, braked, by some process. It is natural to make an analogy with *diffusion*: the electrons bounce back and forth, but with a net drift velocity towards the positive terminal. So the electrons collide with something (in proper terminology they are *scattered*), change direction, are accelerated by the electric field, scatter again, and so on. A careful analysis shows that the main scattering objects are impurities and phonons, at least if we broaden our mind and define an 'impurity' as any deviation from a perfectly periodic crystal lattice. In this sense, our impurity may represent a lattice dislocation, a vacancy, a different isotope of the same element, or most obviously an atom of a different kind. Ironically, semiconductors, which form the backbone of the entire modern electronics industry, actually do not conduct electricity properly *unless* impurities are introduced (a process known as *doping*, for rather obvious reasons), but we will not go into that here, except to say that naturally a balance must be found between the desired positive effects of doping, and the detrimental influence of the impurity scattering.

Obviously phonons also break the symmetry, as the ions are vibrating instead of sitting still in a perfect lattice, and phonons can thus in fact also be considered as 'impurities' under the broad definition given above. Nevertheless, one generally chooses to discuss impurity and phonon scattering separately, in particular since the temperature-dependence of their respective contributions to the resistivity are very different, with the dominant role being played by impurities at low temperatures and by phonons at high.

In our numerical model, we choose to represent the impurity as an atom with a different mass. With this impurity present in the chain, it is immediately obvious that the perfect periodicity is broken, Bloch's theorem does no longer apply, and the wave number is not a good quantum number any more. On the other hand, if there is only a single impurity among a large number of identical atoms, the influence should anyway be small. These predictions can now be verified by studying the results from the numerical calculation.

First of all we find that all phonon modes except one are more or less unaffected by the presence of the impurity (see Figure 4). There are some small shifts of the circles compared to the solid line, which corresponds to the case without any impurity, but they are hardly noticeable. There is however a single mode which behaves completely differently. Its energy is significantly higher, and does not fall on the dispersion curve at all. Moreover, its eigenvector reveals that only the impurity and its closest neighbors are vibrating, whereas the other atoms are almost static, as illustrated in Figure 5. Thus this mode is *localised*, as opposed to all the other extended or propagating modes, for which the wave number is still an approximately good quantum number.

The situation shown in Figure 4 corresponds to the impurity being *lighter* than the other atoms in the chain. For an impurity heavier than its companions, one again finds a single mode that behaves differently, but in this case there is a damping of the oscillation amplitude in a region around the impurity, which of course also corresponds to localization. In all other modes the heavier atom 'pinches off' the eigenvectors by having much smaller oscillator amplitude than its neighbors, without otherwise affecting the other atoms noticeably. Also, it is only for the lighter impurity that we get a split-off frequency, since a lighter atom, unlike a heavier one, can be made to oscillate faster than its neighbors. One could imagine a lower frequency in the case of a heavier impurity, but there is no possibility to drop below the dispersion curve, as this energy region is already 'occupied' by other modes.

We can now, as a finishing touch, give a simple explanation of why impurities are so effective in scattering electrons. First of all, impurities create localised modes not only for phonons, but also for electrons, since Bloch's theorem is an equally important ingredient when analysing the properties of electronic modes. Naturally, localised electron modes do not contribute to the electric conductance, and are thus also a source of resistivity.

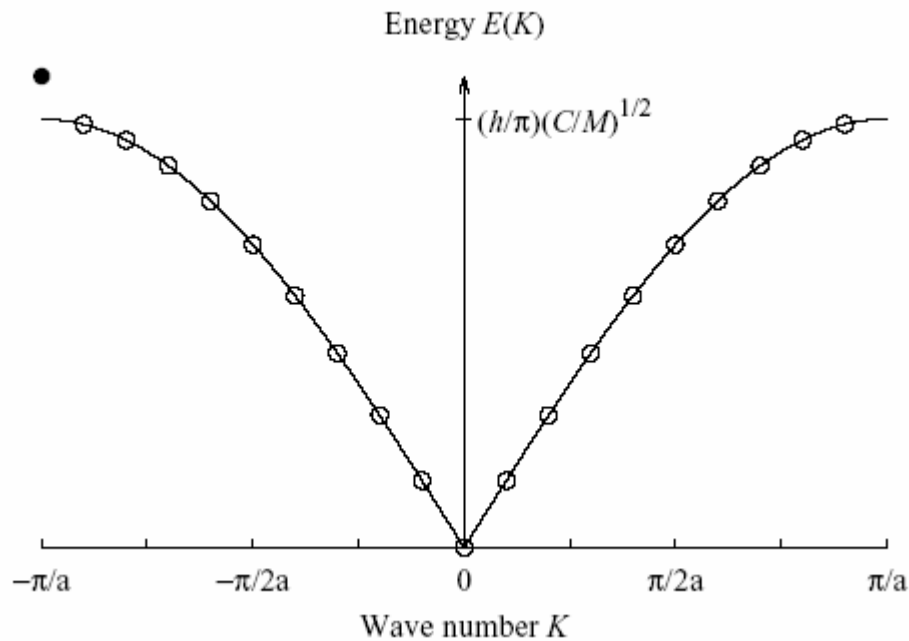


Figure 4. Dispersion relation for the linear monatomic chain with a single impurity. The solid curve applies to the ideal, infinite chain, whereas the circles represent the eigenvalues obtained from the matrix diagonalization with 20 atoms and periodic boundary conditions. The localised mode (solid circle) is easily distinguished, while all other modes are more or less unaffected. In the displayed case the impurity mass is less than the mass of the other atoms, so we get a split-off mode with higher energy. Note that the localised mode cannot really be associated with a specific value of the wave number, although we plot it in this way. Instead it is composed of a whole range of wave vectors, as discussed in the text.




Figure 5. The solid line with circles represents a typical eigenvector of a localised mode, whereas the broken line with squares corresponds to an eigenvector of an acoustic mode with a rather small wave vector. Each point represents an atom in the chain, with the impurity marked as the solid circle.

As shown in Figure 5, in the localised mode the oscillation amplitude decays in space as we move away from the impurity. By using Heisenberg's uncertainty relation, one may immediately show that in order to create a *localised wave packet* like this, it is necessary to involve many different wave vectors. In fact, by Fourier transforming the localised mode eigenvector, we can show that as this mode becomes more and more localised (as the impurity mass becomes smaller and smaller), it will contain an

increasingly broader range of wave vectors, ultimately filling the entire first Brillouin zone. When the impurity mass is very close to the mass of the other atoms, the localised mode contains wave vectors close to the zone boundary, and converges into the dispersion curve at precisely the zone edge as the masses become equal. However, as long as the impurity mass is different from the others, the localised mode cannot be attributed a sharp

CAL-laborate, November 2004

value of the wave vector (although, somewhat confusingly, it is plotted like that in Figure 4).

The laws of energy and momentum conservation, which always must be fulfilled in a scattering event, strongly restrict the available phonon modes against which an electron with a given wave vector can scatter. But, as we have just described, a localised mode, be it electronic or phononic, incorporates many different wave vectors, and can hence scatter electrons with a wide range of wave vectors, making it a particularly efficient source of scattering.

In this way, the concept of localised modes created by impurities, and studied in detail within our simple computer model for phonons, helps us understand the origin of the finite electric conductivity of metals at low temperatures. By the way, this has nothing to do with superconductivity, where in fact it is precisely phonons that mediate the attractive force between electrons that leads to the sudden disappearance of all electric resistivity below a certain temperature, also with impurities present. But that is a whole other story.

To obtain a copy of the program code and the manual for the computer exercise contact the author at Anders.Blom@teorfys.lu.se.