

Superfluidity and Superconductivity

R. P. FEYNMAN

Norman Bridge Laboratory of Physics, California Institute of Technology, Pasadena, California

I AM sorry that Professor Landau was unable to come for two personal reasons. The first is that I have worked on this problem of helium upon which he has also done so much, and I would have liked to have spoken with him about it. The other reason is that I wouldn't have had to give this lecture.

Quantum mechanics was developed in 1926, and in the following decade it was rapidly applied to all kinds of phenomena with an enormous qualitative success. The theories of metals, other solids, liquids, chemistry, etc. came out very well. But as we continued to advance the frontiers of knowledge, we left behind two cities under siege which were completely surrounded by knowledge although they themselves remained isolated and unassailable. The first of the two very similar phenomena for which we still lacked a qualitative explanation is the superfluidity of helium. I may remind you that the discovery that helium flows without resistance through very thin tubes was made as far back as 1911. The second phenomenon is the superconductivity of metals. The fact that electricity flows through some metals without any resistance at low temperatures was discovered, I believe, in 1905. Of these two phenomena, I think we now understand qualitatively the superfluidity of helium, but we do not yet understand qualitatively the superconductivity. I propose to give a brief summary of views on liquid helium insofar as they may give some clue as to the kind of thing that is involved in superconductivity. It is clear, however, that the solution of the superfluidity problem does not give us a very *good* clue because we have solved the first problem but not the second.

First, I make some semiphilosophical remarks about the kind of problem facing us, and the kind of view that I take toward it. I do not want to discuss all of the aspects of superconductors and of superfluids. I want to discuss only the interesting qualitative features, i.e., the curious problem of how does it work more or less. In other words, we would like to make an analogy between the problem of, say, superconductivity and the problem of friction. After all, how is friction explained on the basis of the Schrödinger equation? No one has ever computed the coefficient of friction of two blocks of copper, but qualitatively we feel that

somehow when we rub the atoms together we can understand, more or less, how friction arises. But we do not understand, more or less, how superconductivity works and I would like to address my attention to this problem of understanding it *more or less*, not of understanding the details of a lot of special phenomena. In other words, I would like to concentrate here on the problem of interpretation from first principles. We would like to connect the Schrödinger equation directly to some experimental facts. There are many interesting things that happen when one works on a problem in this way. For example, if one makes some approximations, they cannot be justified by comparing them with experiments. It is necessary to justify the approximations directly in terms of arguments from the equations as a mathematical problem.

In the first stage of the development of such explanations there is an enormous amount of phenomenological explanation made which is of great value, and I don't mean to say anything against it. That is just not what I want to speak about. This phenomenological development is vital in combining various experiments together and is helpful in giving us hints as to what to try to explain. If we were any good, however, we wouldn't need the experiments and we wouldn't need the hints. We could simply calculate everything directly from the Schrödinger equation. So I will assume some familiarity with the wide range of experimental properties and will only make a very brief mention of the salient points.

First we will discuss the properties of liquid helium. As a function of temperature, the specific heat appears something like Fig. 1 experimentally. The details near the transition are unknown. There is a certain temperature at which there is a transition, the so-called λ point, and the properties of the liquid below the transition are very peculiar. First, helium remains a liquid all the way down to absolute zero, which we know is due to the zero point motion of the atoms. The reason for the existence of the transition is undoubtedly the same as for the transition in the corresponding Einstein-Bose gas neglecting interaction between the atoms. The Einstein-Bose transition for an ideal gas, however, has quite a different shape. It looks something like Fig. 2, and the problem of why the actual transition looks like

FIG. 1. Specific heat of liquid helium as a function of temperature.





FIG. 2. Einstein-Bose transition for an ideal gas.






FIG. 3. Coefficient of viscosity of liquid helium *versus* temperature as determined by rotating two concentric cylinders relative to one another.

Fig. 1 when there are forces involved is not yet completely solved. The detailed behavior near the transition, however, is a very delicate and difficult proposition. One is convinced that the transition is due to an Einstein-Bose condensation, but exactly how interatomic forces modify the Einstein-Bose specific heat is not yet thoroughly understood.

To discuss superfluidity, I would like to discuss the properties of helium near 0°K, far from the transition region. If one measures the viscosity of the liquid as a function of temperature by rotating two concentric cylinders relative to one another, the coefficient of viscosity behaves as in Fig. 3. Near the transition temperature nothing in particular happens.

On the other hand, if one measures the viscosity by making the liquid flow through a very narrow tube, the liquid flows through such a tube with an apparent viscosity above the transition temperature, but below the critical temperature it flows with apparently no viscosity. Incidentally, if one measures the resistant force to push the liquid through a tube at different velocities then above a certain critical velocity a resistance does arise but below the critical velocity the resistance is zero. A relatively thick tube has a lower critical velocity than a thinner tube as is shown approximately in Fig. 4. The explanation of these things at least in a qualitative way, has come essentially from Landau. He pointed out that we should imagine that the liquid at, say, absolute zero is a perfect fluid for some reason (or rather it is a perfect liquid when the flow is potential flow and has zero curl). As we heat the sample up, the energy goes into local excitations of some kind, for example, quantized sound waves called phonons and excitations of higher momentum and energy called rotons. If, as the temperature rises, the excitations are few and far between (a fact directly indicated by the very small specific heat), then they will be localized inside the liquid and bounce around within it. The local excitations behave very much like a gas, and the system behaves in many respects like a mixture of two fluids. Thus we have a basis for the two-fluid model, which was a phenomenological model.




FIG. 4. Coefficient of viscosity *versus* velocity in thin and thick tubes.

The energy of the excitations required as a function of their momenta to produce the specific heat and general properties seems to be a curve of the general nature of Fig. 5. The curve is linear for smaller momentum, meaning that at long wavelengths sound waves are excited. But the states of higher momentum excitations are of some other kind. At any finite temperature one excites states mainly in two regions; small p near where the curve is nearly straight, and near the bottom of the concave part of the curve. The lower states are called phonons and the higher states are called rotons.

The reason why the flow is perfect, from the point of view of Landau, can be seen in the following way. Put a little ball inside the liquid at absolute zero and try to make it move through the material, allowing the liquid to flow around it. How can the ball lose energy to the liquid? Only by exciting states inside the liquid, and if the aforementioned states are the only ones available, it is easy to show by the conservation of energy and momentum that a very high velocity is needed before we have the necessary energy at the given momentum transfer to produce such a state. The ball must have a velocity v high enough that the straight line $E = pv$ intersects the curve in Fig. 4. Actually, the




FIG. 5. Energy of excitations required to produce specific heat and general properties of liquid helium in the two fluid model as a function of their momentum.

resistance sets in at a very much lower velocity by a factor of about 100, and this fact has to be explained. We have to explain *not* the superfluidity, but why it is not quite so perfect a superfluid as it should be from this point of view.

The essential feature of the liquid is that there are very few states of low energy, a fact directly shown in the specific heat and one which also explains the superfluidity ultimately. So the question is, why are there so few states of very low energy? Why in this liquid, for example, can't we get a very low energy quantum state, say, by letting some portion of liquid slowly rotate? The portion of liquid could have a big moment of inertia and one unit of angular momentum would have a very low energy. Or alternatively, say, one of the particles inside the liquid decides to drift slowly through the liquid while the others remain stationary. The explanation of why there are no such low energy states in the liquid has been worked out and published in different places.¹⁻³ Time does not allow me to describe it here.

¹ R. P. Feynman, *Phys. Rev.* **91**, 1301 (1953).

² R. P. Feynman, *Phys. Rev.* **94**, 262 (1954).

³ R. P. Feynman, *Progress in Low Temperature Physics* (North Holland Publishing Company, Amsterdam), Vol. I, p. 17.

But there is a point that isn't discussed very often: the question of how is it that we do get some resistance even at a velocity lower than the velocity required to create an excitation directly. We must discuss the character of the wave functions representing the liquid in various states of flow. If the wave function for the ground state of the liquid is represented simply by Φ , then it is very easy to represent the same system in motion at a uniform velocity in a straight line. The wave function for the moving system is simply

$$\Phi \exp(i\mathbf{P} \cdot \mathbf{R}_0/\hbar) = \Phi \exp(im\sum_i \mathbf{V} \cdot \mathbf{R}_i/\hbar), \quad (1)$$

where $\mathbf{P} = Nm\mathbf{V}$ is the momentum of the whole liquid of N atoms moving with velocity \mathbf{V} , and $\mathbf{R}_0 = \sum_i \mathbf{R}_i/N$ is the position of the center of gravity of the atoms whose individual positions are \mathbf{R}_i . This wave function represents a whole system moving at a uniform velocity, but we would like to represent something more complicated. Suppose the liquid flows in some more elaborate fashion in which the velocity is not the same everywhere. If the velocity is a very slowly varying function of position $\mathbf{V}(\mathbf{R})$, and does not change appreciably in an atomic distance, we can try to represent the wave function by simply making the velocity of the previous formula dependent on position. The wave function now becomes

$$\Phi \exp(im\sum_i \mathbf{V}(\mathbf{R}_i) \cdot \mathbf{R}_i/\hbar)$$

which has the mathematical form

$$\Psi = \Phi \exp(i\sum_i S(\mathbf{R}_i)/\hbar)$$

where $S(\mathbf{R})$ is some function of position. We have said it is $m\mathbf{V}(\mathbf{R}) \cdot \mathbf{R}$ but this is not quite right. The momentum for a wave of slowly varying wavelength is the gradient of the phase, not the phase divided by the coordinate, so the correct relation is

$$\mathbf{V} = \nabla S/m.$$

This implies that the curl of the velocity is zero. We have thus only accomplished the description of flow which has zero vorticity by such a method, and the problem of how to describe the flow of the liquid with vorticity becomes very serious. There is no structure in the liquid with which we can tie down the atoms in any given region; they are all the same and identical.

I have already published some work on this problem and Onsager eight years previously⁴ also published the same type of discussion. So for variety I'll give a slightly different argument, the first one that I thought of. It is not a good one but it gives us a clue as to what's going to happen. I tried to *force* the liquid to have vorticity, to have some kind of nonpotential flow. I imagined that the liquid in region I of Fig. 6 is flowing to the right with the velocity v and that some kind of barrier or sheet of infinitesimal thickness separates region I from region II. In region II the velocity is

⁴L. Onsager, Nuovo cimento Suppl. 2, 6, 249 (1949).





FIG. 6. Sketch showing different fluid motions separated by an infinitesimal sheet.

zero. In region I the wave function should look something like Eq. (1), and the wave function in II is just Φ . If such a separated flow has been progressing for some time, what happens if all of a sudden I pull out the sheet? It's got to do something. In spite of the fact that the velocity according to formula (2) always has to have the form of a gradient of S , it cannot be a gradient in this case. I have to make a rule that the atoms in the upper section are in a different state than the atoms in the lower section. If I simply restrict the motion so there is no flow across the boundary I will get a possible kind of motion. There must be a layer here with an associated surface tension where the atoms in I can't wander into II or vice versa. Hence, there is a certain energy associated with relative flow with a sheet of surface tension in between.

Suppose that the velocity in I becomes smaller and smaller. Is it reasonable that, even with a very small drift in I, the surface energy per unit area is still the same no matter how small this velocity is? Why can't the atoms from I go into II? If they can't, then I need the surface tension, but if I can allow them to mix, I don't need that surface tension. They can't mix because the wave function must vary from one position to another. Suppose an atom moves along just above the interface of Fig. 6. The phase must change according to formula (1), from 0 to π and so forth as sketched in Fig. 7. At points where the phase is π , the wave function cannot be smoothly connected with the constant wave function in I, while at those points where the phase is zero, or 2π , etc., the wave function can be connected smoothly. The cosine wave and the constant 1 have places where they are equal. In those places I can erase the surface locally without any effect. What I imagine in the next approximation is that there are certain slots of surface tension with flow around the slots. What is the best length of slot? There is a certain energy of motion around the slot, and so on. If I continue the distortion further I probably lower the energy in every case until I come down to something like a string of holes instead of slots. The velocity distribution then is sketched in Fig. 8 and the wave function changes by 2π when it goes around one of the holes, as it went around the slot. But it is now con-

FIG. 7. The sheet of Fig. 6 is erased at those points where the wave functions can be joined smoothly.






FIG. 8. Velocity distribution when the slots of Fig. 7 distorted into a string of holes.

tinuous everywhere except at the holes. The curl of the velocity is still zero, but S is a multiple-valued function of position (and is not defined in the holes). Thus the integral of the velocity around a circuit is $2\pi n\hbar$. In this way I was led to the idea that when liquid He flows, the vorticity is quantized and all the vorticity is along certain lines of one unit each. And when I was all finished I found a reference to a remark made by Onsager in a meeting in Italy, to a lecture on a different subject, in which he mentions the same solution!

What happens is this. If the object moves too quickly through the liquid vortex lines are formed. If the velocity is too slow there is not enough kinetic energy in the motion to create even one vortex line. For a very thin tube, for example, there is not enough kinetic energy to construct a vortex line. Before a vortex is made, there is no resistance and the whole system coasts together. With thin tubes and velocities that are not too high, we have not enough energy to create even the first line so the material maintains its superfluidity. There remains a theoretical problem to establish these lines firmly as they have not been established very well theoretically so far. The argument used is not very good so it is desirable to make a better theoretical argument about the quantized lines. In the meantime I would like to see an experiment to show that the lines of vortices really exist directly.⁵ The resistance, for example, would be irregular when it first starts if it comes in quantized units. How the vortex lines actually form from a surface has defied my analysis. This, with the omission of the part about why there are no excited states, is all that I have to say about helium.

Problems associated with this field still remain. There are no problems of understanding the phenomena qualitatively any more, but there are problems of quantitative understanding. For example, we need more details about the exact wave functions for the rotons and phonons and their properties. We would like to calculate the cross section for the collision of two rotons. The collisions determine, for example, the viscosity and absorption of second sound, as was pointed out by Khalatnikov and others. They use an arbitrary cross section to fit the experiment, but the question is how to calculate the cross section from first principles. A much more interesting problem probably is the behavior—statistical mechanically, or otherwise—very near the transition.

⁵ Experiments by Hall and Vinen on the attenuation of second sound in rotating helium give strong evidence for the existence of the lines. H. E. Hall and W. F. Vinen, Conf. de Phys. des basses temperatures, Paris (September, 1955), Paper 22.

Today we do not have a good description of what happens within a few tenths of a degree on either side of the transition. We have approximate views about the shape of the curve, but whether it has a vertical tangent or contains a jump in specific heat is not known.

Another problem is the problem of what is the ground state energy. Here we have a strongly interacting system of particles with known interaction and the problem is to find the energy, or other properties of the system in the ground state. It is a quantum mechanical problem and has never been solved, but it would be interesting to work that out from first principles. The energy of a roton has been worked out now. The curve of energy against momentum has been calculated experimentally with the data and also been computed theoretically. The energy of the roton theoretically is 11.5 deg and experimentally is 9.6 deg. To make the theoretical calculation we had to use some properties of the ground state, namely the correlation function of the ground state. To put everything on purely first principles, it would be necessary to calculate the correlation function of the ground state by solving the Schrödinger equation, and I don't know how to do that.

Another liquid of great interest is liquid He³. I do not believe that liquid He³ is a superfluid. I think that there are an enormous number of states reaching to absolute zero. I believe that the specific heat would be proportional to the temperature as it is for an ideal gas and that the forces would not modify it. I have no reason to believe all these things; I have thought a lot about He³, but I find it too difficult to analyze accurately from first principles. If there is a high density of states then He³ will not be a superfluid in the same sense as He⁴. I think that the same thing will happen in nuclear matter, so that superfluid is not a very good term to apply to the nucleus. It is a quantum fluid but I don't think it is a superfluid.

Another interesting problem is the following example from He³. The particles interact with strong forces in a Fermi system. The problem is to determine the temperature dependence of the viscosity of the system as T goes to zero. There are a large number of similar interesting problems in He³, and it would be fun to do them before the experiments, for the first time. I don't think anybody has ever computed anything in solid-state physics before the experimental result was out, so we have consistently predicted only what we have observed!

Now we turn to the unsolved problem of superconductivity and again talk from the point of view of what we will have to do from first principles to solve this problem. The first thing that is troublesome is that the metals that become superconducting are always the complicated ones near the middle of the periodic table. Superconductivity depends on the crystal structure, some metals undergo the transition in one crystal form

and not in another. Also, some alloys are superconductive even though neither of the constituents is by itself superconductive. The simplest metals, the alkalis and alkaline earth metals, are not superconductive. The temperature of the transition corresponds to an energy $kT=10^{-4}$ eV, while the electron energies are of the order 10 v or so. Hence, we have a very delicate problem here. It is not like the case of helium, where the energy of the transition is comparable to the thermal energies involved. In the helium case, when the thermal wavelengths, or the de-Broglie wavelength corresponding to a temperature of the transition, becomes equal to the spacing of the atoms then all the excitement occurs. On the other hand, in superconductivity the effect is very small and therefore quite subtle. I have to make some excuse, you see, for not having immediately solved the problem.

The specific heat of a superconductor looks something like Fig. 9. According to the theory of an ideal electron gas, the specific heat should be proportional to T . Experimentally, there appears to be a discontinuity in the specific heat at the transition, and the electronic specific heat at very low temperatures may be exponential. The curve is certainly not linear with T ; it is much, much less.




FIG. 9. Sketch of specific heat of a superconductor versus temperature.

It is possible to maintain the material in the normal state with a sufficiently large magnetic field. If one applies a magnetic field for different temperatures the stronger magnetic field for lower temperatures, the curve of the transition temperature against magnetic field looks like Fig. 10. Even at absolute zero, the superconductive state can be destroyed by an application of a sufficiently strong magnetic field, and the specific heat of the normal state behaves linearly right down to zero temperature.

The electromagnetic properties of the superconductor have been studied in great detail and summarized in a very simple way by London, at least for low frequencies. London deduced that the law stating that the current density is proportional to the vector potential correctly describes all of the properties. I have to omit all the details of the experiments and the arguments which led London to this law. It is necessary that the vector potential \mathbf{A} be in such a gauge that the normal component of the current density at the surface be zero everywhere. London also gave a kind of explanation of how this may have come about. The formula for the current density of the electrons in the quantum mechanics is

given by

$$\mathbf{J}(x) = -\frac{e}{m} \sum_{k=1}^N \int \Psi^* \left(\frac{\hbar}{i} \nabla_k - \frac{e}{c} \mathbf{A}(x_k) \right) \times \Psi \delta(x-x_k) d^N x,$$

where Ψ is the wave function of the electrons and the integration is carried out over all electron coordinates. Now suppose that for some reason a small electromagnetic field \mathbf{A} changes the wave function Ψ very little. For the wave function without the potential of course, the mean value of the current would be the expectation value of $\sum_k \delta(x-x_k) (e\hbar/mi) \nabla_k$, alone and would be zero. If the wave function does not change much, this term is still essentially zero. Then all one has is the vector potential at the point x averaged over all the electrons. In other words, one has $\mathbf{J}(x) = -ne^2/mc\mathbf{A}(x)$. London would then explain his law on the assumption that the wave function does not change much when a small magnetic field is applied.

Now what is it that we would like to explain about superconductivity and with what are we going to explain it? There are a whole set of properties such as the specific heat, the London equation, etc. to be explained in terms of the Schrödinger equation. The first question is, which property shall we choose first to try to explain? The next question is whether we need the full Schrödinger equation, or can we take an approximate Hamiltonian that is simpler to work with as a reasonable model of the actual situation.

I would like to maintain a philosophy about this problem which is a little different from usual: It does not make any difference what we explain, as long as we explain some property correctly from first principles. If we start honestly from first principles and make a deduction that such and such a property exists—some property that is different for superconductors than for normal conductors, of course,—then undoubtedly we have our hand on the tail of the tiger because we have got the mechanism of at least one of the properties. If we have it correct we have the clue to the other properties, so it isn't very important which property we explain. Therefore, in making this attempt, the first thing to do is to choose the easiest property to handle with the kind of mathematics that is involved in the Schrödinger equation. I want to summarize some thoughts on this question, although they do not represent a solution. They represent a statement of the problem and a little bit of a personal view.




FIG. 10. Transition temperature versus magnetic field.

I decided it would be easiest to explain the specific heat rather than the electrical properties. In the case of the electrical properties we have to consider both the temperature and the magnetic field, which is more complicated than the temperature alone. But we do not have to explain the entire specific heat curve; we only have to explain any feature of it, like the existence of a transition, or that the specific heat near absolute zero is less than proportional to T . I chose the latter because being near absolute zero is a much simpler situation than being at any finite temperature. Thus the property we should study is this: why does a superconductor have a specific heat less than T ? Let us say that the specific heat is essentially zero. Thus the density of states above the ground state is very much less than that of the ideal gas. It is orders of magnitude of infinity less, to get a different power of T . The property to be explained for superconductors is that the ground state is separated from the higher excited states by a region where the density of states is low. The simplest thing would be that there are no states in between and in fact it may be that that is the case. The specific heat curve looks like it. Let us suppose that the density of states is very low just above the ground state. All we have to do is to explain why the excited states have a definite separation energy from the ground state, or that there are only a very few states near the ground state.

I have not got much further and this conclusion is well known to many others, but I would like to make some remarks on some work in that direction. I tried to state this property directly in terms of some property of the Hamiltonian. First why is the density of the excited states so high in the ideal gas, and how could it fail to be so high in a superconductor?

The ideal gas is easily represented in momentum space by drawing a Fermi sphere with a lattice of momentum that is very fine. The lattice spacing depends on the size of the container and the spacing, ϵ , is inversely proportional to the cube root of the volume of the container. If one calculates the excitation energy of the system at temperature T , one realizes that all the states within a range of kT of the Fermi surface are excited. Thus the number of states excited is kT/ϵ times the number of states around the sphere, which is the square of the radius times $4\pi/\epsilon^2$. All of these states are excited on the average by the energy kT , so this total excitation energy is proportional to the volume and to T^2 so the specific heat is proportional to T . Now where did this enormous specific heat come from? Essentially, it came from the fact that there is a state in any direction from the center of the Fermi sphere. It seemed to me likely that here is the feature that fails in the real situation. How can we define this failure? In order to define one of these states, and to distinguish one from the other, we have to be able to locate points on the sphere. That is to say if we expand our functions in terms of angular variables we must be

representing states of enormous angular momentum of very low energy.

Consider a big box, the metal. If one excites a particle, that is, takes it out of one state and puts it somewhere else, then this excitation has a tremendous angular momentum around the center of the box, because it has finite momentum (the Fermi momentum) and is a long distance from the center. It takes practically no energy to create this large angular momentum, according to the ideal gas picture. One hardly has to lift it out of the shell. Thus just above absolute zero, with an infinitesimal energy, one can create very high angular momentum. If the distance required to get superconductivity in small grains corresponds to 10^4 atoms or so, then this means that in a normal conductor in order to create $10^{5\frac{1}{2}}$ units of angular momenta requires infinitesimal energy.

Therefore, in the superconductor we probably have the following property. The difference between the energy of the lowest state with large angular momentum and the ground state, must approach a finite number as the volume of the box goes to infinity. This is the property that the superconductor has that I now look for. But the angular momentum of the electrons does not commute with the Hamiltonian. Let L be the operator representing the angular momentum of all the electrons, the z component of which is $\sum_i (x_i p_{iy} - y_i p_{ix})$. Take E_0 to be the ground-state energy. If f is any normalized wave function, for which the expected value of L is greater than 10^5 (the exact value is unimportant), then

$$\int f^*(H - E_0)f dv > \eta,$$

for

$$\int f^* L f dv > 10^5,$$

where η is finite as the volume goes to infinity. This is not true of an ideal gas, but I believe it is true of the Hamiltonian of the superconductor.

I don't know if I have made any progress over the simple remark that the density of states above the ground state is very low. That is as far as I was able to get with the specific heat problem. If one assumes the truth of the previous equation one cannot necessarily explain all the properties of superconductors. But if one can find the *correct* explanation of one of the unique properties of a superconductor, one can then see the mechanism for all the other properties. Actually, we can almost see how the electrical properties will come out of such a procedure. If one considers a cylindrical piece of metal with a field along the axis of the cylinder, then the Hamiltonian of the system with the field is related to the Hamiltonian without the field by $H' = H e/c - B \cdot L$. That is, the $A \cdot p$ term is equal to $B \cdot L$ for such a cylinder. Let ϕ be the ground-state wave function. In the ground state the expected value of H is

the lowest energy, E_0 . $(\phi, H\phi) = E_0$. The expected value of the angular momentum is zero. For any other state, which has a reasonable amount of L , we get

$$(f, Hf) = E_0 + \eta - 10^5 (\hbar e/c) B.$$

If B is not too big, this energy is larger than E_0 , so with a small enough field, we would get the lowest state, which would be the ground state not the new state, and therefore the wave function would be the same as it was before. States with smaller angular momenta do not affect the argument since B has little effect on these. Thus, London's explanation of superconductivity would result. This argument is not very good. It must be more complete, since one has to show that if we took other states where L is larger than 10^5 , the energy rises fast enough for a sufficiently large L . When we do finally understand the reason for this gap, we will understand how the energies will vary with the angular momentum of the states that are created.

Now the next part of the problem is with what do we make this explanation? That is, what approximate Hamiltonian should we use? The question is, have we included enough junk in the model to include the effect of superconductivity? The first and simplest model is the model of the ideal gas of electrons. There we know we do not get superconductivity, because we do not get the right specific heat. The next thing to be added that seems to be very important is Coulomb interaction. Sometimes it has been thought that a Coulomb interaction would make some kind of lattice of electrons which would move through the material. However, everyone argues, and I think correctly, that this cannot produce the small effect, because it is too large. That is, the Coulomb interaction energies in a material are of the order of a volt or so per electron, while the energies we are looking for amount to 10^{-8} v per electron. We have all kinds of other effects, such as spin-orbit couplings, spin-spin couplings, magnetic interactions, and so on. Because we have factors v^2/c^2 in these cases, we get the right order of magnitude.

But there is one little piece of experimental information of vital importance, which shows that we will not get superconductivity even if we include all these effects. That fact is that the transition temperature depends upon the ionic mass. Or better, the energy difference between the superconductive state and the normal state, depends upon the mass of the ions. If these v^2/c^2 effects were the cause of superconductivity, it would not depend upon the mass of the ions.

Could it be that the superconducting state depends upon the mass, not because the phenomenon which produces the energy depends upon the mass, but because when you heat the material up, the shaking of the lattice breaks the state down at a temperature that depends upon the mass? I think this approach is wrong. Since the difference in the energy even at absolute zero depends on the mass, I think that the

energy must involve dynamically the motion of the lattice. So we have to include lattice interactions. Now the hope is that if we include lattice interactions, we will not have to include these other interactions.

One of the most popular models of a Hamiltonian to explain superconductivity is the one used so much by Frohlich and Bardeen. Their model consists of a jelly of positive charge, an elastic continuum which can vibrate, in which the band structure of the electrons is neglected. In the jelly of positive charge, one considers a gas of electrons of the same density as the charge density of the positive charge. The electrons do not interact with each other, but they all interact with the vibrating jelly. It is the simplest Hamiltonian which has been proposed that would seem to be able to represent superconductivity.

A lot of people have tried this model. However, they are not absolutely satisfied. Some still worry whether the band structure has something vital to do with the phenomenon. Perhaps superconductivity does not occur in the case of sodium etc., because the band is so very simple and it occurs only for the complicated metals, because their band system is complicated. There may be small regions in momentum space, for instance, where the electrons behave as positively charged particles, that is, places where the conductivity is by holes and other regions where they behave normally. There is some indication that this is the case because it has been noticed that the Hall effect is very small when the material has a tendency to be superconductive. The Hall effect is very small when the positive and negative carriers cancel. Thus some people think that this, in conjunction with the lattice vibrations, may have something to do with superconductivity. Of course, that makes the problem more complicated, because it would mean that if Frohlich and Bardeen could solve their model exactly, they still would not find superconductivity, since it would still involve only negative carriers.

As long as one makes a list of such things, one can always add a few things which might have something to do with it. In case they do, at the end then people can say, "Oh, he thought of it," but I don't believe the one I will mention has anything to do with it. The closed shells have been treated adiabatically and forgotten about and they might have something to do with superconductivity. I just do not think that the closed shells have much to do with it.

I also tried the same thing as Frohlich and Bardeen, which is to solve the Frohlich and Bardeen model with such precision that the estimated errors in the calculation are less than the difference in energy one is looking for. I made diagrams and loops. When one is working with partition functions, it turns out there are all kinds of wonderful tricks that one can do with the diagrams. I calculated the specific heat to enormous precision, and the specific heat was still proportional to T . However, a student of mine, Michael Cohen,

pointed out that the series expansion which I was carrying out may not lead to the lowest state at all. Suppose, for example, that one has a harmonic oscillator and considers the effect of a small perturbation, ϵx^3 , on the ground-state energy. Regardless of how small ϵ is, one can never derive the answer from perturbation theory even though the series appears to be nicely convergent. The reason, of course, is that if one goes far enough from the origin the potential becomes arbitrarily large and negative, and hence no lowest state exists. In other words, what happens is that one starts with a function which is qualitatively different from the correct function and the right answer lies in a qualitatively different place. Thus when one starts with a simple kind of wave function to make a succession of approximations, one cannot produce the qualitative features of the isolated lowest state of the real system with the original wave functions. Nor can one be sure that the qualitative difference does not exist just because one computes so accurately. Thus, even if one treats this model very carefully and accurately, as long as it is done by a series method one can never be certain that the model does not lead to superconductivity. That is, we may simply be calculating the normal state specific heat very carefully. It is going to be very hard

to show that the phenomenon does not exist in this model. It is also going to be hard to show that it *does* exist.

I am now brought to the same position as Casimir, who first told me about this problem. He said, "There is only one way to go about working this out. It is simply to *guess* the quality of the answer." I think I have come unfortunately to the same position, what we must do is not compute anything, but simply *guess* what makes the ground state isolated at a lower energy. That is, guess what kind of correlation exists at long distances. Why haven't we theoretical physicists solved this problem yet? We have no excuse that there are not enough experiments, it has nothing to do with experiments. Our situation is unlike the field, say, of mesons, where we say, perhaps there aren't yet enough clues for even a human mind to figure out what is the pattern. We should not even have to look at the experiments. Every time we look at another experiment, we make the problem easier. It is like looking in the back of the book for the answer, which is slowly being unveiled by the details of the various experiments. There is no reason to require the experiments. The only reason that we cannot do this problem of superconductivity is that we haven't got enough imagination.