PHYSICS TODAY VOL 22 NO 10

ELECTRONS IN METALS

WALTER A. HARRISON

Easy electrical conduction in metals is usually attributed to the regularity of metallic crystals. Why then does the conductivity remain when a crystal melts? Pseudopotentials provide an answer.

IT IS REMARKABLE how much of the basic description of electrons in metals was made available to us 40 years ago in the doctoral thesis of Felix Bloch. He not only explained how electrons can travel through perfect crystals without colliding with the constituent atoms, but he also discussed the scattering of those electrons when the crystal lattice is vibrating. The origin of this scattering is now called the "electron–phonon interaction," but the calculation came some 20 or 30 years before the word "phonon" was even coined. He also evaluated the electronic contribution to the specific heat.

After looking at this list it is fair to ask, "What on earth have we been doing for the subsequent 40 years?" One answer to that question was provided in the next to last sentence of Bloch's thesis publication. He pointed out that superconductivity was not, at that time, understood. I need not discuss whether it is understood at present; that certainly depends very much on how one translates the German verb "erklären." I will discuss instead the problem that Bloch did address, the motion of electrons in periodic potentials. I will first present Bloch's argument in a novel form and then see how our insight into the problem has evolved through experimental studies of the Fermi surface and through the evolution of the concept of a weak pseudopotential. That concept now appears to be the real physical basis for metallic properties.

It was realized before Bloch's work that, if we *assumed* electrons could run through a crystal lattice without seeing the atoms at all, we could account for a wide range of the properties of metals. On the other hand, it was not at all clear why electrons should not see the atoms; this was the question that was so effectively addressed by Bloch. I will restate his argument as a special case of the conservation of momentum.

Momentum as a symmetry property

First I should talk about conservation of momentum itself. The conservation of momentum is a consequence of the translational invariance of space, which means that in free space all points are equivalent. This invariance implies that a translation of coördinates leaves the form of the Hamiltonian unchanged. The question we would like to ask is: "What can we learn about the states of an electron from the translational invariance alone?" The states of the electron can of course be computed in detail from the time-independent Schrödinger equation

$$H(\mathbf{r})\psi(\mathbf{r}) = E\psi(\mathbf{r})$$

if we know the Hamiltonian H, but now all we want to assume known about the Hamiltonian is its translational invariance.

In order to learn anything, clearly we must translate the equation. This could be a coördinate transformation, but we will think of it as picking up the Hamiltonian and the wave function and moving them a distance $\mathbf{T}$, leaving the coördinate system fixed. Then

$$H(\mathbf{r}-\mathbf{T})\psi(\mathbf{r}-\mathbf{T}) = E\psi(\mathbf{r}-\mathbf{T})$$

Because we move everything together, the equation has to remain true. But the invariance of space tells us that $H(\mathbf{r}-\mathbf{T}) = \mathrm{H}(\mathbf{r})$; so

$$H(\mathbf{r})\psi(\mathbf{r}-\mathbf{T}) = E\psi(\mathbf{r}-\mathbf{T})$$

We have found a new function, the translated wave function $\psi(\mathbf{r}-\mathbf{T})$,

Walter A. Harrison is a professor in the applied physics department at Stanford University. He took his bachelor's degree at Cornell in 1953 and his PhD at the University of Illinois in 1956. For his doctoral thesis he worked with Frederick Seitz on electron–phonon interaction in semiconductors. The nine years prior to his move to Stanford were spent at the General Electric Research and Development Center, where he worked on the theory of metals.

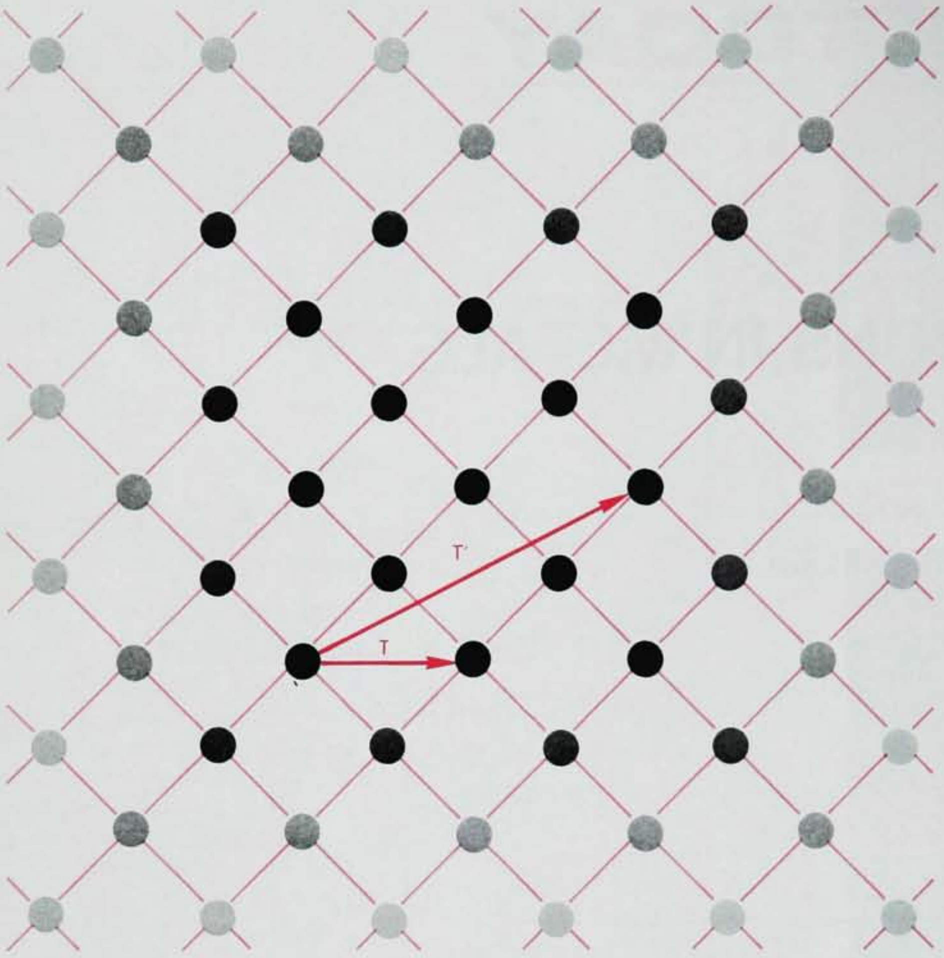

ATOMS IN A METALLIC CRYSTAL are arranged in a regular repeating pattern. Many discrete translations, for example **T** and **T′** shown by the colored arrows, lead to an arrangement in the interior identical to that before translation. —FIG. 1

that is a solution of the Schrödinger equation with precisely the same energy as the wave function we started with. If there is no degeneracy, the new function $\psi(\mathbf{r} - \mathbf{T})$ must be essentially the same as the old function $\psi(\mathbf{r})$. (A little group theory is required to tidy this up, but it is not essential to our discussion.) The two functions need not be identical but can differ at most by a phase factor. We conclude that translating the wave function can at most change it by a constant phase factor, and the phase factor can depend on what translation has been made.

$$\psi(\mathbf{r} - \mathbf{T}) = e^{i\varphi(\mathbf{T})}\,\psi(\mathbf{r})$$

We can do the same thing again, making a new translation on the new wave function; of course this translation simply gives an additional phase factor

$$\psi(\mathbf{r} - \mathbf{T} - \mathbf{T}') = e^{i\varphi(\mathbf{T})}e^{i\varphi(\mathbf{T}')}\psi(\mathbf{r})$$

We see that the phases themselves simply add and are therefore linear in the translation. The most general form that is linear in the vector translation **T** is

$$\varphi(\mathbf{T}) = -\mathbf{k}\cdot\mathbf{T}$$

Here **k** is a constant vector that is independent of the translation **T** but may very well depend on the state ψ. Let us then label the state by the **k** that describes its translational properties

$$\psi_k(\mathbf{r} - \mathbf{T}) = e^{-i\mathbf{k}\cdot\mathbf{T}}\psi_k(\mathbf{r}) \qquad (1)$$

This procedure has told us something fairly specific about the transformation properties of the wave function. It still does not quite tell us what the wave function looks like; let us address that question next.

Form of the wave function

We can conveniently start by defining a new function

$$u_k(\mathbf{r}) = e^{-i\mathbf{k}\cdot\mathbf{r}}\psi_k(\mathbf{r}) \qquad (2)$$

In a crystal $u_k(\mathbf{r})$ is of course going to be in the Bloch function, but our discussion here is still directed at electrons in free space. We can turn equation 2 around to write

$$\psi_k(\mathbf{r}) = u_k(\mathbf{r})e^{i\mathbf{k}\cdot\mathbf{r}} \qquad (3)$$

Now the question is: What does this $u_k(\mathbf{r})$ look like? In particular, what are its translational properties? To find out we simply take this form for the wave function and substitute it back into equation 1. We obtain immediately

$$u_k(\mathbf{r} - \mathbf{T}) = u_k(\mathbf{r}) \qquad (4)$$

We have found that translating $u_k(\mathbf{r})$ leaves it unchanged. In free space *any* translation takes the Hamiltonian into itself, and we conclude that any translation will take this function $u_k(\mathbf{r})$ into itself; it must therefore be a constant. We conclude that the electron state $\psi_k(\mathbf{r})$ must be of the form

$$\psi_k(\mathbf{r}) = Ae^{i\mathbf{k}\cdot\mathbf{r}}$$

where A is a constant.

It is remarkable that we have been able to deduce the exact form of the wave function with no knowledge about the Hamiltonian except the complete translational invariance. The resulting form is as well known as the Schrödinger equation, but this derivation may not be.

Electron in a lattice

How does the argument change if the electron is moving in a crystal lattice? A crystal is a regular arrangement of atoms like wallpaper (figure 1). All points in the wallpaper pattern are not the same, but the entire pattern can be displaced so that each flower is replaced by an identical flower, and the resulting pattern is the same as before. Similarly in a crystal we have discrete translations that leave the crystal interior unchanged. In such a translation each plane of atoms is replaced by an identical plane of atoms.

Does this limited translational symmetry tell us anything about the wave functions in a crystal? The argument goes through exactly as before through equations 3 and 4. However, now we cannot say that *any* translation takes u_k into itself, because not all translations take the Hamiltonian into itself. Only the discrete set of translations takes the Bloch function into itself. We conclude as before that the wave function is of the form of equation 3, but now all we can say about u_k is that it has the complete translational periodicity of the lattice. The Bloch function $u_k(\mathbf{r})$ for a given state will be identical in the neighborhood of every atom. It will in fact resemble an atomic wave function near each atom and vary smoothly between them. In the complete wave function

ψ_k, the Bloch function u_k is modulated by a plane wave (figure 2). These are the famous Bloch wave functions that describe the behavior of electrons in metals.

In free space we found that there is a wave number **k** (or a momentum $\hbar\mathbf{k}$) that is a constant of the motion. In the crystal we again find that there is a **k** that is a constant of the motion, and we then call the quantity $\hbar\mathbf{k}$ the *crystal momentum*. The corresponding state carries current, and because it is an eigenstate in the perfect crystal it does not change with time. We have demonstrated that there are current-carrying states in the crystal.

Nature of symmetry arguments

The demonstration is a symmetry argument, and symmetry arguments are peculiar. One of their peculiar features is that they have no real counterpart in everyday life where we make only the most trivial sort of symmetry arguments. The most sophisticated such argument I have been able to find is the search for the men's room in a public building. If I come upon the ladies' room, then I say "Without question the men's room will be in the symmetric position on the opposite side of the building." So I march across the building with almost religious fervor in the knowledge that I am going to find the men's room there.

There are at least two things in common between that everyday example and symmetry arguments in physics. One is the religious fervor; when we conclude something on the basis of symmetry, we feel absolutely certain that it is true. But of course in both cases, in my example and in physics, we may very well be wrong. Frequently the architect will choose to save money on plumbing and put the men's and the women's rooms back-to-back on one side of the building.

There are a number of cases where we have been wrong also in physics. For example the discovery that there was not the expected symmetry of parity came as a thunderous blow. Even within solid-state physics, concerning the existence of the Bloch state, the Great Architect in the Sky saved plumbing when He made antiferromagnets. The electron states, again with a periodic arrangement of atoms, are not the Bloch states that we have deduced here. Also, in the case of the Mott transition, we have a periodic arrangement of the atoms but not Bloch's state for the electrons. In that particular case I am not absolutely sure whether the creator was the Great Architect in the Sky or whether it was Nevill Mott, but recent work seems to favor the former—to the credit of the latter.

There are some other peculiarities about our symmetry argument. One is that, at least in a mathematical sense, if we move one of the atoms out of place, the whole symmetry argument breaks down. Bloch had already dealt with this aspect in 1928, in treating the electron–phonon interaction. He showed that we could describe the states in terms of a perfectly arranged lattice and treat the small displacements as a correction that leads to scattering. There is a further problem: Imagine that we melted the crystal. We then destroy all semblance of translational periodicity and, as it turns out, the resistivity does not change very much. It is fair then to ask whether there is more to the "erklärung" of metallic conductivity than the symmetry argument giving us the Bloch functions. Let us look at that subject next.

After the Bloch theorem

There have been two major areas of activity in the postwar period. One concerns the effects of the Coulomb interactions among the electrons. These are large effects and were not dealt with in a definitive way in any of the earlier work. One of the central results of recent work has been the explanation of why those interactions can be neglected or at least treated in the average way implicit in Bloch's early work. The second area of activity concerned the interaction of the electrons with the ions in the metal. That is the part that I wish to concentrate on here, and, of course, it is specifically the question that Bloch addressed.

After the war there were two separate approaches—theoretical and experimental. There are always these two approaches, but seldom are they quite so separate as they were in this case. Bloch had shown us that in a crystal we can associate a wave number with every electron state, and these wave numbers provide a means of indexing the states. Each state has an energy, and that energy is therefore a function of the wave number. This function defines the electronic band structure of the metal, in direct analogy with free atoms. There, by symmetry, we conclude that we can describe the states in terms of angular-momentum quantum numbers, and the electronic structure is the energy of the states as a function of these quantum

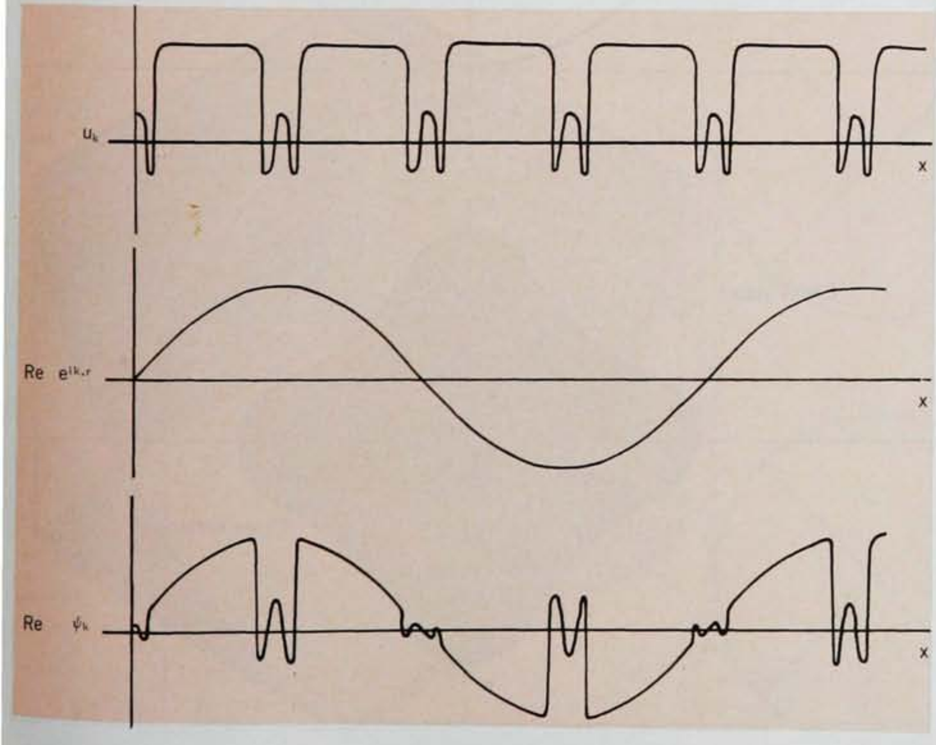

THE BLOCH FUNCTION $u_k(r)$ has the same translational symmetry as the crystal and resembles the atomic wave function near the atoms. The electronic eigenstates take the form $\psi_k = u_k(r)e^{k\cdot r}$, where the states are indexed by the wave number k of the modulating plane wave. —FIG. 2

numbers. In solids the wave numbers take the place of the angular-momentum quantum numbers. The theoretical problem was to make a calculation of the energy as a function of wave number. That turns out to be difficult.

Immediately after the war, it was really out of the question to attempt the calculation of the energy as a function of all wave numbers. What in fact was done was to select a few isolated wave numbers, which a theorist had reason to believe may be important or representative, and compute the energy eigenvalues for those wave numbers. The problem, of course, was that one can not measure such eigenvalues directly, and there was no experimental check on the calculated values. In addition, there was no a priori reason for confidence in the energy-band calculations, largely because of the uncertainty in the treatment of the electron–electron interaction. Theorists were able to develop along their own path and compare their calculations with each other, but there was no contact with experiment.

Experimentalist's difficulties

The situation was not much better with the experimentalists. They set out to make measurements of properties that depend on the electronic structure and, in fact, all properties do. Measurements were made of the resistivity, the Hall effect, and the de Haas–van Alphen effect in order to gain information about the electronic structure. The difficulty was that all of these properties depended on complicated integrals of the energy as a function of wave number, and there was no way to go from the experiments back to a determination of the energy bands. It was necessary to propose a model and adjust parameters. Several people, each studying different properties of a metal such as aluminum, would each propose a different model of the band structure. This diversity clearly led us no closer to a knowledge of the true energy bands.

We can blame the theorists for a further difficulty in the interpretation of experiments. The most complicated form of energy band that a theorist could conveniently use in treating properties corresponded to ellipsoidal constant-energy surfaces, so that virtually all the calculations were done for those surfaces, and the experimentalists naturally used ellipsoidal surfaces to fit their experiments. As it has turned out, these surfaces seldom arise in real metals. The theoretical and experimental approaches proceeding separately were getting nowhere at all.

Fermi surfaces

The big breakthrough came in 1955 when Brian Pippard recognized that the *Fermi surface* is an aspect of the electronic structure that could be fairly directly measured experimentally and that in addition could be calculated by band-calculation techniques. What is the Fermi surface? We have said that the energy of the electronic states is a function of wave number of those states. The wave numbers are three-dimensional vectors lying in a three-dimensional wave-number space. In a metal there are only a certain number of electrons present to occupy these states, two or three electrons from each atom, and in the ground state of the system these electrons will occupy only the lowest lying levels (figure 3). There will be regions of wave-number space where the states are occupied and other regions where they are not. The surface in wave-number space that divides the occupied from the unoccupied states is called the Fermi surface. What Pippard recognized was that the anomalous skin effect depends only on the geometry of the Fermi surface. He then set out to determine the Fermi surface of copper (figure 4).

In fact, it is not possible to go directly from the experiment back to the geometry of the Fermi surface, even by using the anomalous skin effect. It was again necessary to propose a model of the Fermi surface and to adjust it until it fitted. However, the experimental information was more complete and the goal (geometry of the Fermi surface) more limited than with other phenomena. As it turned out Pippard did not get quite the right answer for the Fermi surface of copper. Bulges have been found that he did not discover. Interestingly enough that error is really of very little importance.

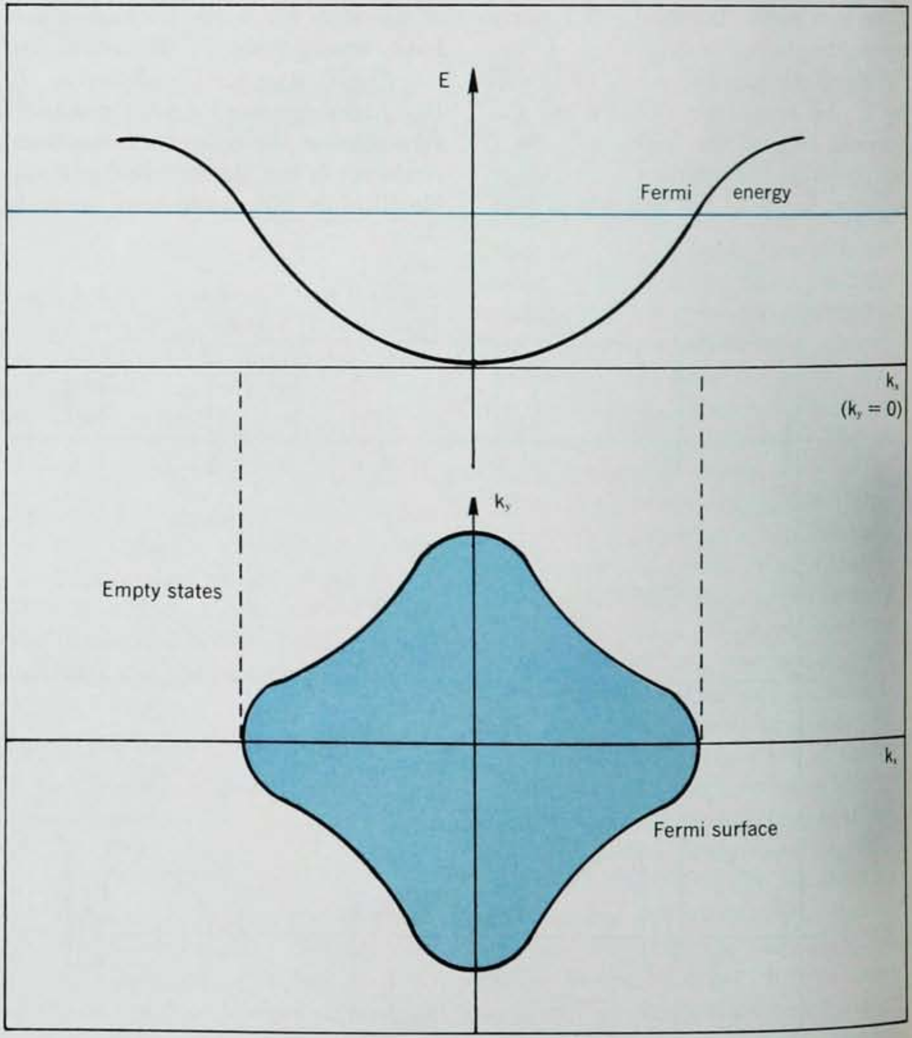

THE ELECTRONIC STRUCTURE is the energy of the one-electron states as a function of wave number k. In the ground state of the metal, states in some regions of wave-number space are occupied by electrons; states in other regions are not. The Fermi surface divides these regions in wave-number space. —FIG. 3

The real point was that Pippard had focused attention on one aspect of the electronic structure, which could be proven right or be proven wrong. If he had the right solution, all other experiments should be consistent with exactly the same Fermi surface. Band calculations could be done to see if they were also yielding the correct Fermi surface.

Pippard's puzzle

It is interesting in passing to recall how Pippard arrived at the Fermi surface he proposed, even though it turned out to be in error. Once he had collected the data and began puzzling over the shape of surface that would fit his data, he actually found the solution in his pocket.

He was at the time amused by a novel brain teaser. The puzzle was: What object has three identical perpendicular projections, as in a mechanical drawing, each of which is a circle with two crossed diagonal diameters? Pippard carried a model of this object in his pocket. A little thought will lead to the solution: the intersection of three perpendicular cylinders (figure 5). That shape had the essential features that were needed to explain his data. He began with such a shape, rounded it off a little, bulged it here and there, and finally fitted the data.

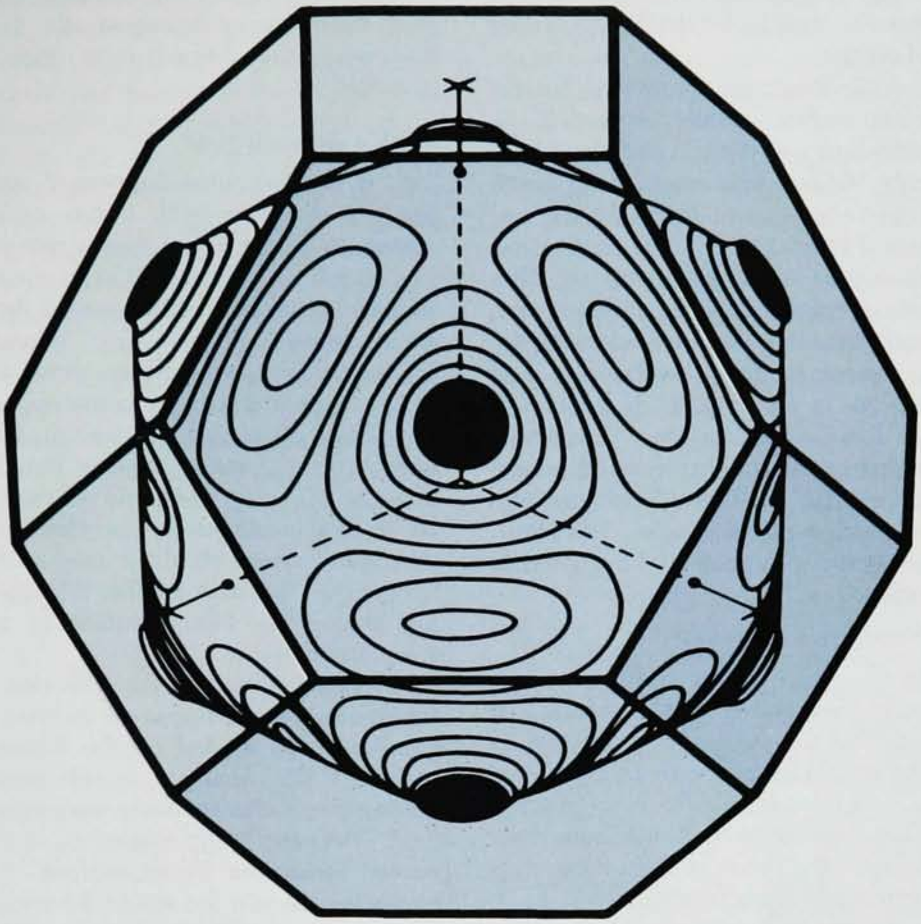

PIPPARD'S FERMI SURFACE FOR COPPER. Pippard noted that the anomalous skin effect depended only on the geometry of the Fermi surface; on the basis of his experiments he proposed this shape as the Fermi surface of copper. The polygon is made up of Bragg reflection planes. An electron accelerated into one face is diffracted to the opposite face. —FIG. 4

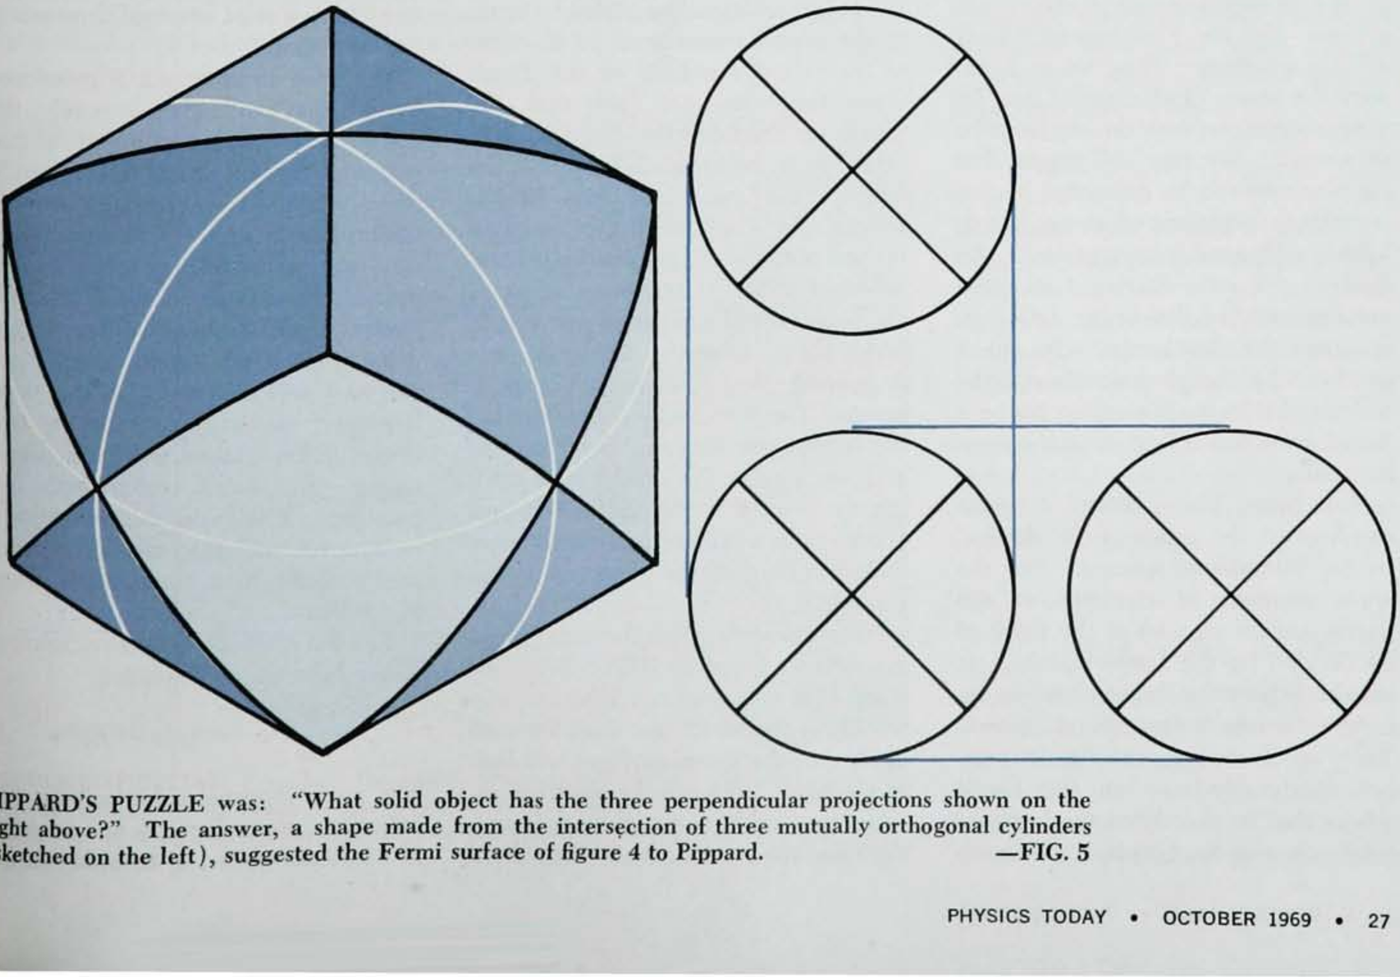

PIPPARD'S PUZZLE was: "What solid object has the three perpendicular projections shown on the right above?" The answer, a shape made from the intersection of three mutually orthogonal cylinders (sketched on the left), suggested the Fermi surface of figure 4 to Pippard. —FIG. 5

It gave a specific model, which he thought then to be the Fermi surface of copper.

After this time, the activity in the Fermi-surface studies expanded rapidly and so did the calculation of energy bands. The energy-band calculators were in fact obtaining the correct Fermi surfaces; the approximations they had used to treat the electron–electron interactions were sound, and therefore the methods and approximations they used were good enough to give a good description of the electronic structure. A mass of experimental activity went into documenting in detail the Fermi surfaces of a wide range of metals. This effort, now called "Fermiology," continues to expand.

Nearly free electrons

In the course of this study an exceedingly remarkable fact emerged. It was that the energy-band structure in the simple metals, such as aluminum, is very close to that for free electrons. That conclusion was not immediate because the Fermi surfaces themselves were exceedingly complicated.

To see that such complicated surfaces are consistent with nearly-free-electron behavior we imagine first a gas of electrons with no perturbing potentials from the atoms. The energy in a free-electron gas increases monotonically with the wave number; so the constant-energy surfaces are spheres, and the Fermi surface itself will be a sphere. Now let us introduce the atoms of the crystal, but let their interactions with the electrons be very weak. We may still expect that the electrons can be diffracted by this periodic arrangement of atoms, just as light is diffracted from a grating. An electron will only diffract if its wave number satisfies the Bragg reflection condition for this lattice. To put it another way, the electron can only be diffracted if its wave number lies on a Bragg reflection plane in wave-number space.

The Bragg planes occur at wave numbers of the order of 2π divided by the interatomic spacing. But the wave numbers of electrons at the Fermi surface are also of the order of 2π divided by the lattice spacing; so we can expect the Bragg planes to intersect the free-electron Fermi sphere. There is, of course, a negligible number of the electrons on the Fermi sphere that lie exactly on the intersection with the Bragg plane. Thus almost all of the electrons will proceed as if there was no lattice at all. In that sense the diffraction is unimportant.

With a magnetic field

Let us imagine what happens if we apply a magnetic field. Then each electron with the Fermi energy moves in a circular or helical orbit as its wave number moves in a circle around the Fermi sphere (figure 6). When, however, the wave number meets a Bragg plane it will jump to the opposite Bragg plane and continue on its way. When it meets another Bragg plane, it will jump again and continue on its way. Although the electron propagates almost all of the time as if it were free, the wave number is jumping around the Fermi surface in a complicated way.

It is interesting to ask how this electron is moving in *real* space as its wave number jumps around on the Fermi surface. We imagine an electron moving perpendicular to the magnetic field. Between Bragg reflections, as it moves across the Fermi surface, it moves in a circular arc within the crystal. When it diffracts, it abruptly changes direction and proceeds along another circular arc. Thus the orbit in real space is made up of a collection of circular arcs and may be quite complicated.

As it turns out, when we represent these free-electron-like states in terms of the crystal momentum, or in terms of the wave numbers of the Bloch wave functions, we find that the Fermi surface takes the shape of these orbits in a magnetic field. It is no longer the original free-electron sphere, but a sphere that is chopped up and reassembled into surfaces composed of spherical segments. Part of the business of Fermiology was assembling these spherical segments and comparing the results with experiments. The remarkable thing was that the interaction between the electrons and the ions in the crystal *was* sufficiently weak that this diffraction approach gave a rather good quantitative understanding of the observed Fermi surfaces.

This was very much the same problem as was faced in 1928. If we assume that the electrons interact only weakly with the lattice, then we conclude that the Fermi surfaces will look as they do. How can the electrons respond so weakly to the potentials arising from the metallic ions when those potentials are known to be very strong? Even with quantum theory and the periodic lattice assumed by Bloch, we would expect a much stronger interaction than that observed.

Pseudopotential theory

The simplest way to resolve this difficulty is in terms of pseudopotentials. The concept of a pseudopotential is quite old and originated with Fermi, who was interested in the scattering of strongly interacting particles. Consider, for example, the scattering of a neutron by a proton. These particles interact sufficiently strongly that there exists a bound state of the two, the deuteron. In terms of the scattering theory, that means that the phase shifts are greater than π, which in turn means that we can not use perturbation theory in treating the interaction.

Fermi suggested that we construct an artificial potential, a pseudopotential, which gives exactly the same phase shifts as the real potential but with any integral multiples of π dropped. Such a pseudopotential will yield precisely the correct scattering because adding multiples of π to the phase shift does nothing to the scattering cross section. Yet it is a potential for which the phase shifts are smaller than π, and so perturbation theory can give meaningful results. Pseudopotentials in a solid accomplish precisely the same goal.

We wish to construct a pseudopotential that will give precisely the same energy-band structure as the true potential but will reflect the inherently weak observed interaction between electrons and ions. The most direct way to describe the corresponding pseudopotentials is in the framework of the calculation of the energy eigenvalues themselves, a band calculation, and so I will describe it in that way. However, we will see that the real virtue of the method will be in eliminating the band calculations altogether. I will do this elimination in a novel way that can be generalized directly to a much wider range of problems.

For the moment we seek a solution of the Schrödinger equation,

$$-\frac{\hbar^2}{2m}\nabla^2\psi + V(\mathbf{r})\psi = E\psi \qquad (5)$$

The usual approach for solving this equation, particularly if one is going to use a computer, is to make an expansion of the wave function in a complete set, such as plane waves

$$\psi = \Sigma_k a_k |\mathbf{k}\rangle$$

where the a_k are coefficients and the $|\mathbf{k}\rangle$ are plane waves. We can then substitute this form for the wave function in the Schrödinger equation, multiply on the left by a plane wave and integrate. This procedure converts the partial differential equation into a set of simultaneous linear algebraic equations in the coefficients a_k with as many unknowns as there are terms in our expansion. It turns out that plane waves are a poor set for expansion of states in the solid because of the need to duplicate the sharp oscillations of the atomic wave functions that persist in the Bloch wave function in the neighborhood of the ions. We can not obtain a reasonable description of the state without using many hundred plane waves, and we are left with the problem of solving many hundred simultaneous equations in many hundred unknowns.

We seek an improvement by supplementing our expansion with terms that may duplicate these atomic-like oscillations and hoping to get by with many fewer plane waves.

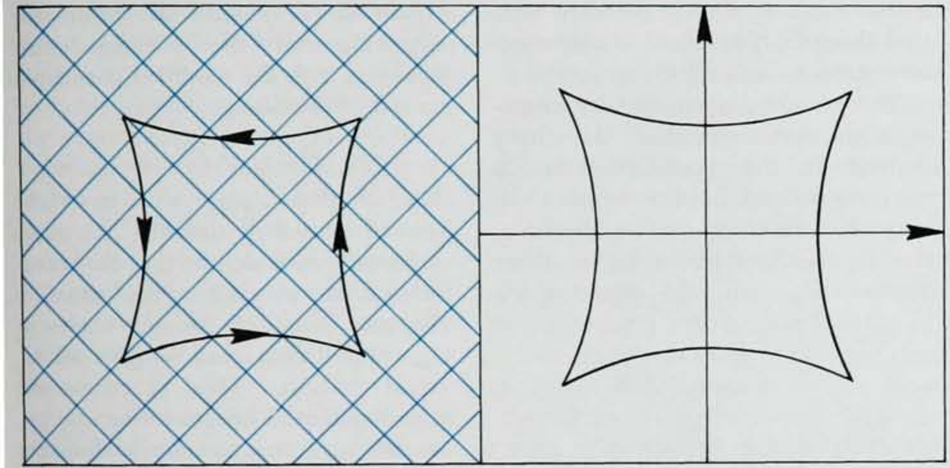

WEAK INTERACTION. **If electrons interact sufficiently weakly with the metallic atoms, the interaction can be treated as if it gives rise to simple diffraction. Then an electron in the crystal, moving in a magnetic field, will move in a circular arc between diffraction events. In such a case the Fermi surface consists of a sphere chopped up and reassembled; it has cross sections of the same shape as the electron orbits. —FIG. 6**

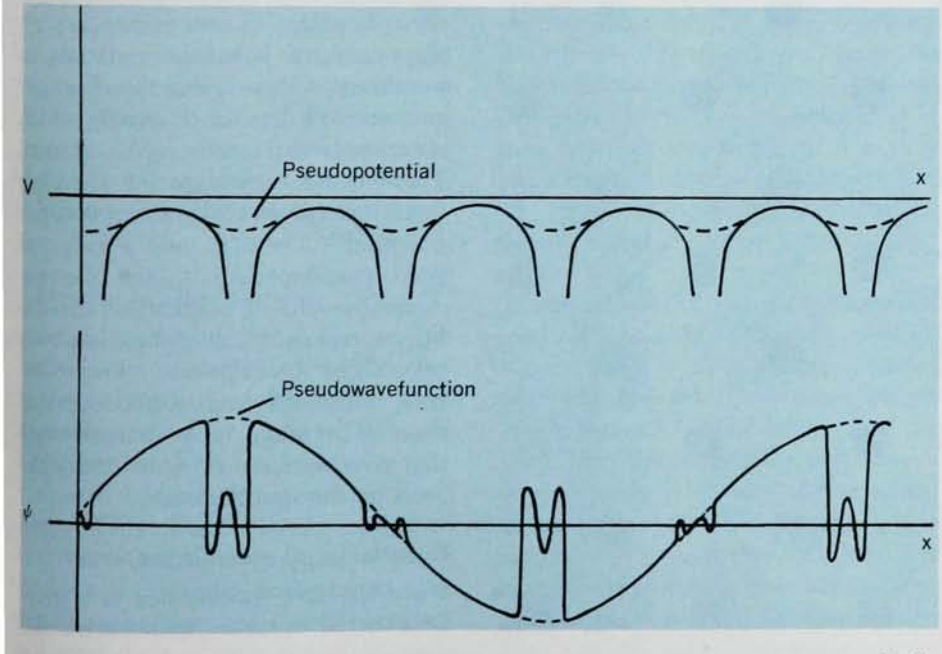

PSEUDOPOTENTIAL **describes the weak interaction between electrons and the metallic lattice. It is much weaker than the true potential. However, it yields eigenstates of the same energy and pseudowavefunctions identical to the true wavefunction between metallic atoms. —FIG. 7**

In supplementing our expansion, we are using an over-complete set because plane waves alone are complete. However, the expansion remains legitimate.

At this point we distinguish two categories of states in the metal. First are the core states, the lowest-lying states in the free atom. (In aluminum they are the 1s, 2s and 2p states.) Those states remain essentially the same in the solid as in the free atom, and are therefore known at the outset. The conduction-band states, which derive from the free-atom valence states (3s and 3p in aluminum) differ appreciably from the free-atom states and must be described in our expansion. However, they will be orthogonal to the core states, and it turns out that an expansion in plane waves *and* core wave functions requires only a small number of plane waves for description of the conduction-band states. The addition of core wave functions orthogonalizes the smooth plane-wave terms to the core states and duplicates the sharp oscillations near the metallic ions. This procedure is also the essence of the orthogonalized plane-wave method for band calculation developed by Conyers Herring in the 1930's. He orthogonalized each plane wave to the core states and then expanded the conduction-band states in the resulting orthogonalized plane waves.

In our case we include core wave functions, which we write $|\alpha\rangle$, as part of the overcomplete set and expand the wave function

$$\psi = \Sigma_k a_k |\mathbf{k}\rangle + \Sigma_\alpha a_\alpha |\alpha\rangle \qquad (6)$$

This form is substituted in the Schrödinger equation, equation 5, as before. But from there on the tack is entirely different. Instead of obtaining algebraic equations to be solved by machine, we seek an equation that can be treated in perturbation theory. We can do it by calling the sum over plane waves a "pseudowavefunction," $\varphi = \Sigma_k a_k |\mathbf{k}\rangle$. Then the Schrödinger equation becomes

$$-\frac{\hbar^2}{2m}\nabla^2\varphi + V\varphi + \Sigma_\alpha (E_\alpha - E) a_\alpha |\alpha\rangle = E\varphi \qquad (7)$$

We have noted that the Hamiltonian operating on a core state $|\alpha\rangle$ gives simply the core-state energy E_α times $|\alpha\rangle$ and have taken the term in $E|\alpha\rangle$ to the left. The coefficients a_α are obtained by multiplying on the left by a particular $\langle\alpha|$ and by letting the first two terms operate to the left. We obtain $a_\alpha = -\langle\alpha|\varphi\rangle$. Substituting this form in equation 7 we obtain the *pseudopotential equation*

$$-\frac{\hbar^2\nabla^2}{2m}\varphi + W\varphi = E\varphi \qquad (8)$$

where the pseudopotential operator W is given by

$$W = V + \Sigma_\alpha(E - E_\alpha)|\alpha\rangle\langle\alpha| \qquad (9)$$

The pseudopotential equation becomes the basis for the calculation and is remarkable in a number of ways.

First the pseudopotential is small and can reasonably be treated by perturbation theory. The reason is that the pseudowavefunction is smooth, consisting of a small number of plane waves; it is obtainable with only small corrections to the plane-wave solutions obtained from equation 8 if W is dropped altogether. At the same time the pseudowavefunction must equal the true wave function between atoms where the core wave functions vanish as may be seen from equation 6. Thus the pseudowavefunction is simply the true wave function with the atomic-like oscillations removed (figure 7). We may also see more directly from equation 9 that the pseudopotential will be small. The energies E of interest are much higher than the core energy; so the second term in equation 9 is positive and will tend to cancel the negative first term in the same equation, which describes the attraction of the electrons to the ions.

An exact equation

The second important aspect of this equation is that it is *exact*. We have not made any approximations on the initial Schrödinger equation, except for the assumption that the core states remain unchanged. (If we were worried about that, we could of course use core states calculated in the crystal.)

There is also something very mysterious about this equation. As I have derived it the pseudopotential is uniquely defined, but it turns out to be somewhat arbitrary. For example, $E - E_\alpha$ can be replaced by any other constant C_α, and the equation remains exact. That result is most easily seen by imagining we have solved this equation with the new pseudopotential

$$-\frac{\hbar^2}{2m}\nabla^2|\varphi\rangle + V|\varphi\rangle + \Sigma_\alpha C_\alpha|\alpha\rangle\langle\alpha|\varphi\rangle = E'|\varphi\rangle$$

We multiply on the left by the corresponding true wave function (that is, the same wave number) and integrate. In the first two terms we let the Hamiltonian operate to the left to obtain $E\langle\psi|\varphi\rangle$, where E is the true energy of the state. In the third term, we note that the true wave function is orthogonal to all core states; so these terms go out whatever the value of the C_α. On the right we obtain $E'\langle\psi|\varphi\rangle$ with E' again the eigenvalue of the pseudopotential equation with the modified pseudopotential. We obtain

$$E\langle\psi|\varphi\rangle = E'\langle\psi|\varphi\rangle$$

The overlap $\langle\psi|\varphi\rangle$ will never be zero if ψ and φ describe the same state, and we conclude that the eigenvalue E' we obtain here is identical to the true energy E for any choice of C_α. One trivial case is with all C_α equal to zero. Then of course the pseudopotential becomes the true potential, and the pseudowavefunction becomes the true wave function. The pseudopotential equation obviously leads to the correct eigenvalues. This choice is consistent, but it does not yield a useful pseudopotential. It does not matter which pseudopotential we use as long as we solve the equations exactly. However, we are always going to solve the equations in perturbation theory, and therefore our answers will depend on exactly which pseudopotential we have chosen. There is no *correct* pseudopotential; there is not even really a best pseudopotential. There is only a range of *good* pseudopotentials, and the numerical results of calculations inevitably depend on which of these has been selected by the calculator. Fortunately it turns out from experience that there is a range of pseudopotentials that give meaningful results in calculations on the simple metals.

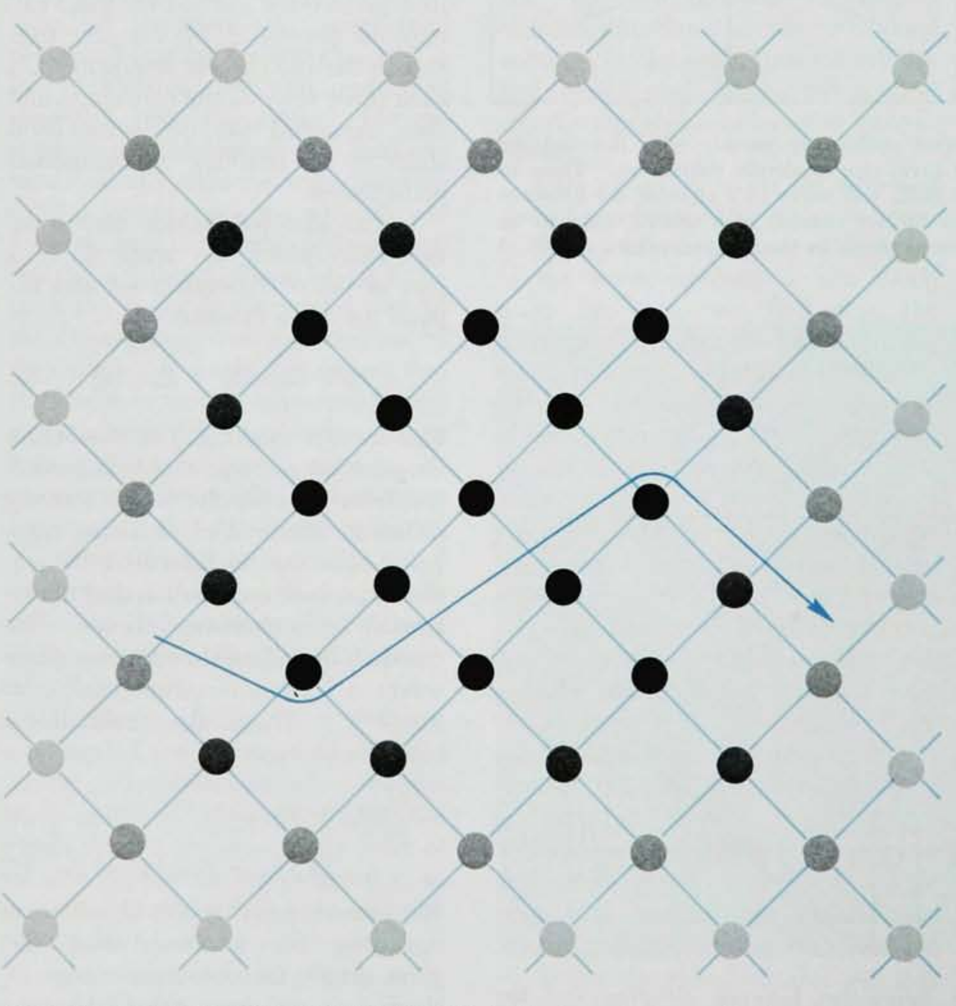

SUCCESSIVE ELECTRON SCATTERING is the important contribution to the interaction between atoms in a metal, after direct Coulomb interaction. It arises by scattering, through a pseudopotential interaction, from pairs of atoms. —FIG. 8

Calculation of metallic properties

Once we have transformed to a pseudopotential equation, with a term that we can treat as a perturbation, we are in a position to forget about the energy-band calculation itself. We calculate directly any property that may be of interest. For example, the resistivity of the liquid metal that we mentioned earlier becomes a very simple calculation. We have, in zero order, plane-wave pseudowavefunctions. The perturbation depends upon the structure of the liquid, which we can obtain from experiment, and upon calculable atomic pseudopotentials. Then the calculation of the resistivity of the liquid metal to lowest order becomes trivial.

We find that the results are reasonably in accord with the experiment, but this case turns out to be one where the result is rather sensitive to the choice of pseudopotential. By thinking through the argument for arbitrariness of the pseudopotential we see that the sensitivity means precisely that this is a case where lowest order perturbation theory is not very good. A particular pseudopotential that gives the experimental answer in a lowest-order calculation does not yield small higher-order terms; rather it yields higher-order terms that cancel in the calculation of this particular property. In many properties the calculations are more reliable, and in all cases the answers appear to be within a reasonable range.

We can even do calculations that were hard to imagine in terms of energy-band calculations themselves. We can sum the total energy of the occupied states to second order in the pseudopotential and thus compute the total energy of the crystal. We then obtain directly the energy changes when the atoms are rearranged. In this way we can calculate, for example, the vibrational spectrum. In doing the pseudopotential calculation of the energy to second order we are including interactions between atoms that arise from successive electron scattering from the two atoms (figure 8).

Extension to other systems

Perturbation theory has made possible a much more complete understanding of simple metals than of any other state of condensed matter. It is unfortunate that this possibility is limited to the simple metals, which are free-electron-like. We can not directly extend it even to the noble and transition metals because of the states in the metal arising from the atomic d states. We can not treat these as core states, because they change as we go from the free atom to the solid. At the same time it is not useful to treat them in terms of a pseudowavefunction; they are sufficiently well localized that we need a very large number of plane waves in their expansion.

However, in terms of the way we have developed the pseudopotential method here it may be obvious how we should proceed. We can again supplement our plane-wave expansion with core wave functions, but now also with atomic d states. They are not eigenstates of the Hamiltonian in the metal, but they are not very far from being so. We can keep track of the corrections, and these extra terms will be small. They, as well as the pseudopotentials, are treated as perturbations. Although it is more complicated, the formulation yields again a fairly simple pseudopotential, which we can use in the calculation of electronic properties. It is even possible again to sum the total energy of the crystal, at least in the noble metals, to obtain the total energy as a function of the configuration of atoms.

The important point about this generalization is not so much that we have been able to treat the transition metals, but that it makes clear that the essence of the method is not free-electron-like behavior; we do not have that in the transition metals. The essence is something different that can be described as a general mathematical method. Again, in terms of an eigenvalue problem, we obtain a Hamiltonian matrix by expansion in some complete set of states. The eigenvalues are obtained by diagonalizing the matrix. If the off-diagonal matrix elements are initially very small we can use perturbation theory; these are the problems with which we have traditionally been very successful. If they are not very small, or if a very large number of them are appreciable, we cannot directly use perturbation theory, and the solution becomes difficult.

The essence of the pseudopotential method is to make a transformation on this matrix to a basis set that includes not only the initial complete set of states but also other terms that we think for physical reasons will allow a more concise description (fewer terms) of the true states. This ordinarily is a nonunitary transformation of the Hamiltonian matrix; it is also a nonsquare transformation because we expand to a more numerous basis. However, if our physical intuition is good and we have included in our expansion terms that make a good approximation to the eigenstates, the off-diagonal matrix elements will be reduced and concentrated near the diagonal. Perturbation theory then becomes applicable. The trouble in perturbation theory is ordinarily not so much the size of the matrix elements but their failure to become sufficiently small far from the diagonal. Good intuition enables us to reduce these distant matrix elements to proceed with perturbation theory.

Insulators and molecules

In this context the method for generalization to insulators and to molecules is clear. We again use an over-complete set. In the case of an insulator the plane-wave expansion is supplemented by a Bloch sum of atomic orbitals. In the case of molecules, the plane-wave expansion is supplemented by linear combination of atomic orbitals. In this method the zeroth-order approximation for molecules becomes the LCAO (linear combination of atomic orbitals) method, which has been so successful in theoretical chemistry. The approach is then to use perturbation theory to improve upon this zeroth-order approximation. My hope is not really for more accurate calculations but that the use of perturbation theory will allow a simpler and more complete description of complicated molecules. Perhaps it will provide a systematic treatment of bonds and, in higher order, the interaction between bonds.

Obviously in describing this current effort I am going well beyond the discussion of electrons in metals. In so doing I am expressing the hope that the very circuitous path that began with Bloch's thesis in 1928 and led through the intricacies of Fermi surfaces and through the mysterious pseudopotential, may take us in a meaningful way into molecular physics, and possibly some day even to organic systems.

* * *

This article is based on a talk given at the 25,26 Oct. 1969 symposium at Stanford University, held in honor of the 40th anniversary of Bloch's thesis publication.

Bibliography

F. Bloch, Z. Physik **52,** 555 (1928).

The Fermi Surface (W. A. Harrison, M. B. Webb, eds.) Wiley, New York (1960).

W. A. Harrison, *Pseudopotentials in the Theory of Metals,* Benjamin, New York (1966).

W. A. Harrison, Phys. Rev. **181,** 1036 (1969). ☐