VOLUME 18 NUMBER 4
APRIL 1980

THE PHYSICS TEACHER

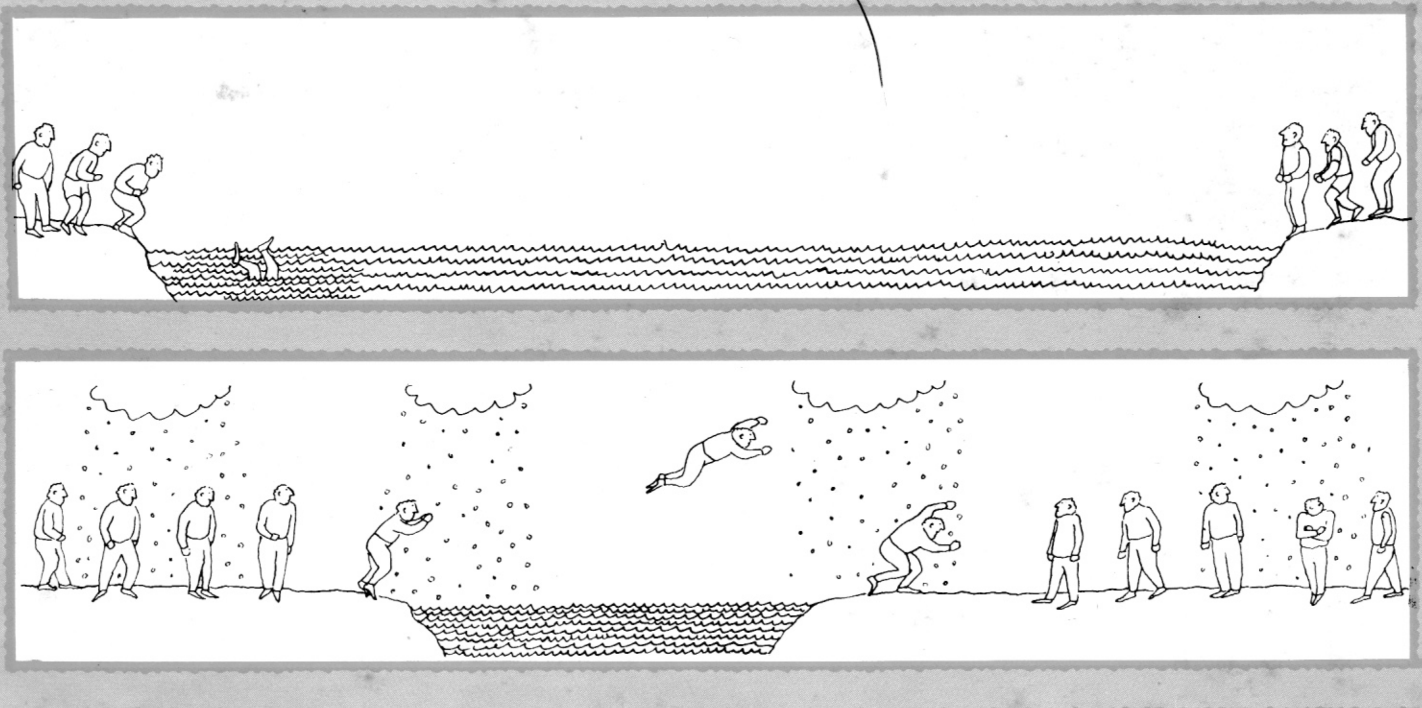

Electric conduction in a granular metal

Size effects in conductivity and superconductivity

(apples, wires, and grains)

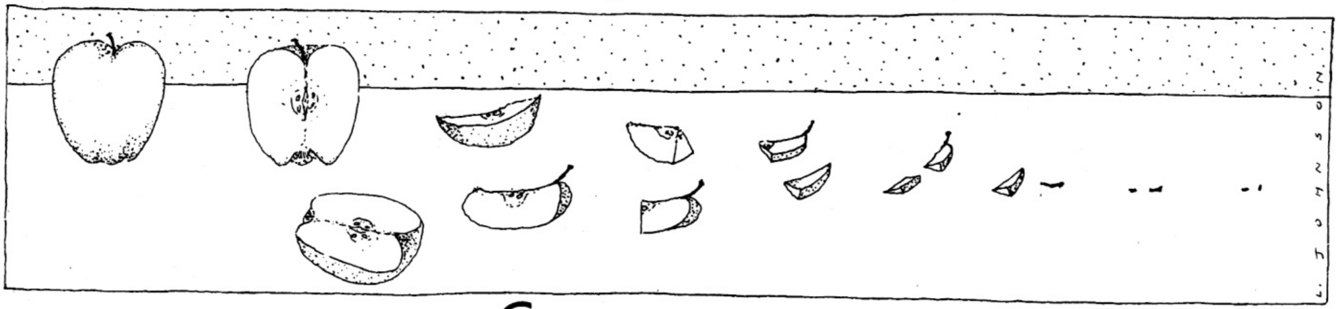

PETER LINDENFELD

The first experiment **Peter Lindenfeld** *remembers was trying to cut a lampcord with a pair of scissors. (Yes, it was plugged in . . . .) He nevertheless went on to study electrical engineering and physics, and for the past 27 years has been teaching at Rutgers University. He has given many courses for teachers, and was responsible for the introduction of the MAT and MST degrees at the University. (Department of Physics, Rutgers University, New Brunswick, New Jersey 08903)*

Cut an apple in half and you have two pieces. But it's still apple. Keep cutting and long before you come to the atoms (the "not-to-be-cut") you can no longer recognize the properties of the material that you started with.

Even the first cut causes changes in some properties: the mass of each piece is different, the volume is different. Both depend upon the particular piece. That's why we like to use their ratio, the density, which characterizes the material regardless of its size or shape.

In the same way we use the resistivity to describe the electrical behavior of a metal, rather than the resistance which depends on the shape of the particular wire which is fashioned from it.

More subtly and interestingly, our first cut has changed the amount of surface. The surface-to-volume ratio has changed, and this is perhaps the most important parameter which characterizes scale changes. It is responsible for such facts as that elephants don't fly and that birds breathe quickly.

Keep cutting the apple. And let's not worry about the peel or the pits. We expect nothing drastic to happen until we come to the basic components, the cells which are the structural units, the molecules which are responsible for taste and smell. If we destroy these then truly we don't have "apple" any more.

Hold on. Haven't we forgotten something? Are there not changes between those which characterize the first cut and those which finally destroy the substance of our apple? How about the interaction between cells, the transport of fluid from cell to cell, dynamic "happenings" which are altered before the structure itself is affected?

Resistivity and electronic mean free path

Let's leave our apple and take an example of a simpler system. The electrical resistivity is supposed to be independent of the size of a wire. And it is, but only while the wire is sufficiently large. When it gets below a certain diameter, which may be hundreds or thousands of times as large as the distance between the atoms of the wire, the resistivity increases.

We have chosen a parameter (the resistivity) which is independent of large-scale shape changes, and we are nowhere near the scale of the structural units, the atoms of the metal. Yet we observe a size effect.

The explanation is of great physical and historical interest. It has to do with the microscopic variable which the resistivity tells us about. That is the electronic mean free path, the average distance between collisions as the "free" electrons move through a metal.[1]

The product of the resistivity and the electronic mean free path is a constant for a given metal, independent of temperature and of small concentrations of impurities or imperfections. Normally the mean free path depends almost entirely on the collisions which the electrons make in the bulk of the

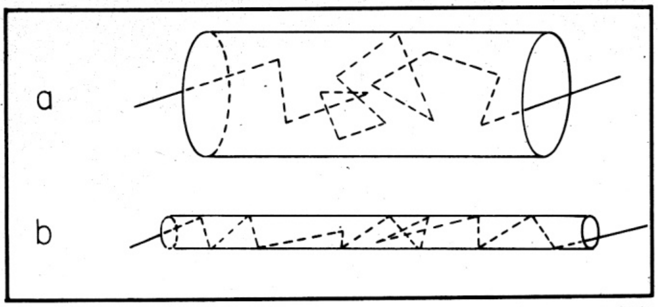

Fig. 1. Part a illustrates the path of an electron in a bulk material. Most of the collisions take place in the interior of the specimen. The average distance between collisions (the "bulk" mean free path, ℓ_b) depends on impurities, inhomogeneities, and other departures from the regularity of the crystal lattice.

In part b the wire diameter is so small that a major fraction of the collisions take place at the specimen surface. The wire diameter is smaller than the bulk mean free path ℓ_b, so that the actual mean free path, ℓ, is smaller than ℓ_b.

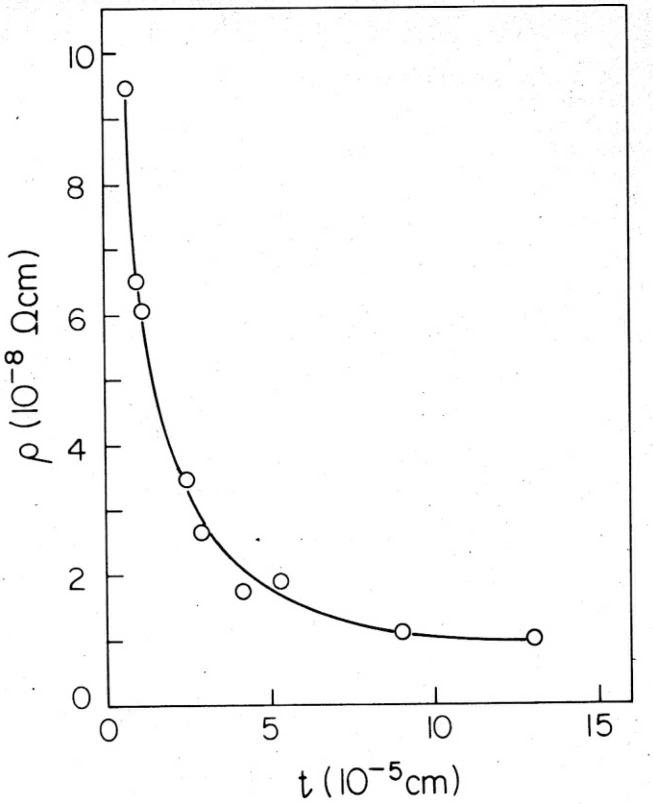

Fig. 3. In thin films the resistivity rises when the thickness becomes small compared to the bulk mean free path. This graph shows the results of measurements by Larson and Boiko on silver films. [D. C. Larson and B. T. Boiko, Appl. Phys. Lett. 5, 155 (1964)]. A comprehensive review of the field may be found in "Size-dependent electrical conduction in thin metal films and wires" by D. C. Larson, in *Physics of Thin Films*, Vol. 6 (Academic Press, New York, 1971) p. 81.

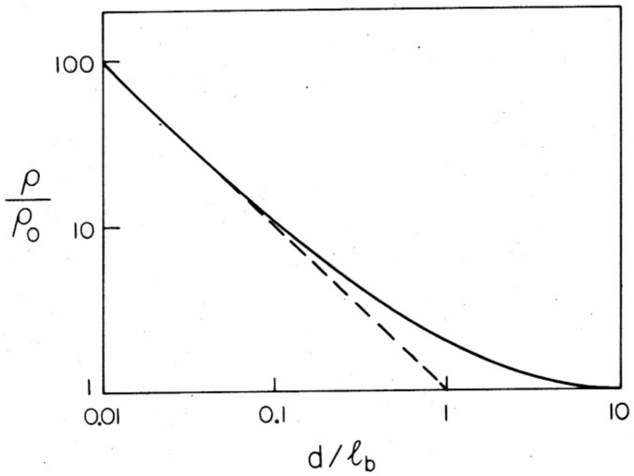

Fig. 2. Theoretical relationship between ρ/ρ_0 and d/ℓ_b for small wires. Here ρ_0 is the resistivity of the bulk material, corresponding to the mean free path ℓ_b. As the wire diameter, d, gets smaller the actual mean free path, ℓ, becomes smaller than ℓ_b, and the resistivity, ρ, becomes larger. A measurement of ρ for a given d then allows ℓ_b to be determined from this graph.

wire (Fig. 1a). But if the wire diameter is so small that the scattering at the surface is significant, the mean free path will be smaller and the resistivity larger than in a bigger wire (Fig. 1b).

If an experiment is made to measure the resistivity as a function of wire size it is possible to find how the resistivity goes up, and to determine the electronic mean free path. In this way this elusive quantity is measured directly, as if with calipers (Figs. 2 and 3).

Measurements of the size effect of the electrical resistivity show that the mean free path can be, in sufficiently pure metals, as large as several millimeters. Classical physics is quite unable to account for the fact that electrons can go that far in a solid without any collisions. It is only the wave nature of electrons which allows us to understand this phenomenon.

It is sometimes thought that quantum mechanics affects only the description of systems of atomic size. Here we have a macroscopic manifestation of the wave and quantum mechanical behavior of electrons.

Thermodynamic fluctuations

Another size effect, also "dynamical," yet of quite different origin, provides a limitation to an entire science.

The laws of thermodynamics are thought to be absolute, in the sense that they are entirely reliable, without known exceptions. And indeed they are, as is emphasized by the long history of the search for perpetual-motion machines.

Yet the science applies only to systems with a sufficiently large number of components. The condition of thermodynamic equilibrium which is fundamental to a thermodynamic description is not one of static stability. It is subject to departures from the equilibrium state, i.e., to fluctuations, which become more important as the number of participating particles decreases.

Systems of a hundred or even a thousand particles,

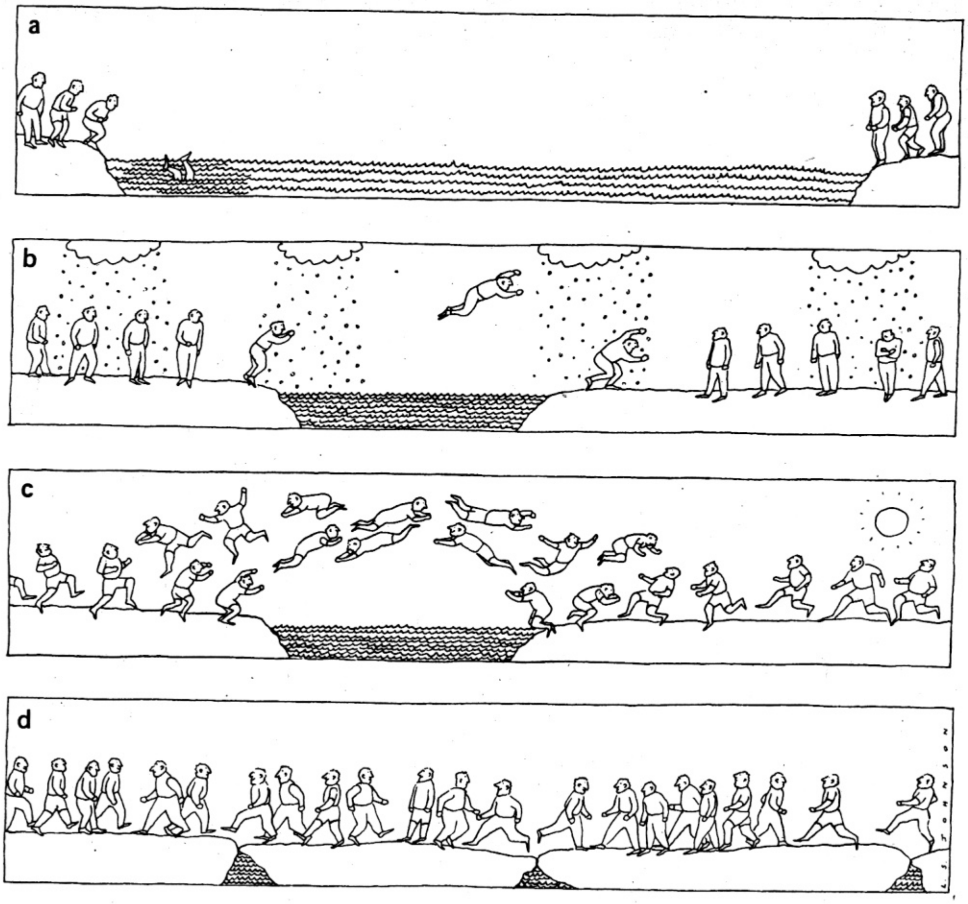

Fig. 4. A free interpretation of a granular metal. In part a the grains are so far apart that no electrons can cross from one to the other. Electric conduction is not possible and the material is an insulator.

In parts b and c some electron transfer is possible, but only for those electrons which have sufficient energy. In part b the temperature is so low that only a few electrons take part in the conduction process. In part c the temperature is higher, more electrons hop from grain to grain, and the conductivity is larger. The increase in conductivity which accompanies the larger number of charge carriers is characteristic of semiconductors.

In part d the grains touch, or are so close together that electron transfer is possible without any "activation energy." Electron motion is not as unimpeded as in a bulk metal, so that the resistivity is higher, but, as in a metal, the electrons are not localized. In quantum-mechanical language, the electron wave functions extend throughout the material, leading to all the properties associated with metallic behavior.

which might show the effects of thermodynamic fluctuations, are not common. They are, however, interesting precisely because the study of their fluctuations can shed light on fundamental thermodynamic behavior.

Granular aluminum

When aluminum is heated under vacuum it melts and then evaporates, only to deposit on a cold surface as a shiny layer of pure metal. If, instead of a vacuum, there is a small amount of oxygen, the aluminum deposits in the form of small metallic grains, surrounded by aluminum oxide. It is possible in this way to make layers of rather uniform grain size, about 30 Å in diameter.

The amount of oxide, and therefore the separation between grains, can be varied by varying the oxygen pressure during the evaporation. With small amounts of oxide the grains touch and the whole specimen behaves like a piece of metal (Fig. 4d). Because of its granular constitution the resistivity can, however, be much larger (ten or a hundred times larger) than that of pure aluminum.

With larger amounts of oxide the grains are entirely separate, electrons cannot travel throughout the specimen, and it is then an electrical insulator (Fig. 4a).

There is also an intermediate regime, where the grains are close, and where electrons of sufficient energy can hop from grain to grain. Here the resistivity is a thousand or more times as great as in the pure metal. The higher the temperature, the more likely it is that an electron will have enough energy to go from one grain to another. The electrical resistivity will get lower as the temperature rises, and in that respect the specimen behaves like a semiconductor (Fig. 4b, c).

Superconductivity

Below 1.2 K aluminum is a superconductor. That it is so only at this low temperature is the consequence of a battle between almost equal forces. By themselves the electrons repel each other. In the presence of the positive ions in a metal they can, through a subtle dynamical interaction, also attract. When the attraction wins we have a superconductor, but this special attraction has to overcome not only the electrostatic repulsion, but also the disruptive effect of the ceaseless thermal motion.

The attractive forces need all the help they can get, and only when the thermal motion is quieted down at low temperatures do they sometimes win. It is because of the precarious balance between opposing forces that superconductivity is particularly sensitive to thermodynamic fluctuations.

Figure 5 illustrates the effect of fluctuations on the way in which the resistance goes to zero when the material becomes superconducting. In a large piece of metallic material all of it is either in the normal state or in the superconducting state and the resistance goes abruptly to zero at the transition temperature (Fig. 5a). In a single small grain, on the other hand, the thermodynamic fluctuations can have a significant effect. Even at temperatures where the bulk material would be in the normal state, the shifting, fluctuating electron configuration can produce

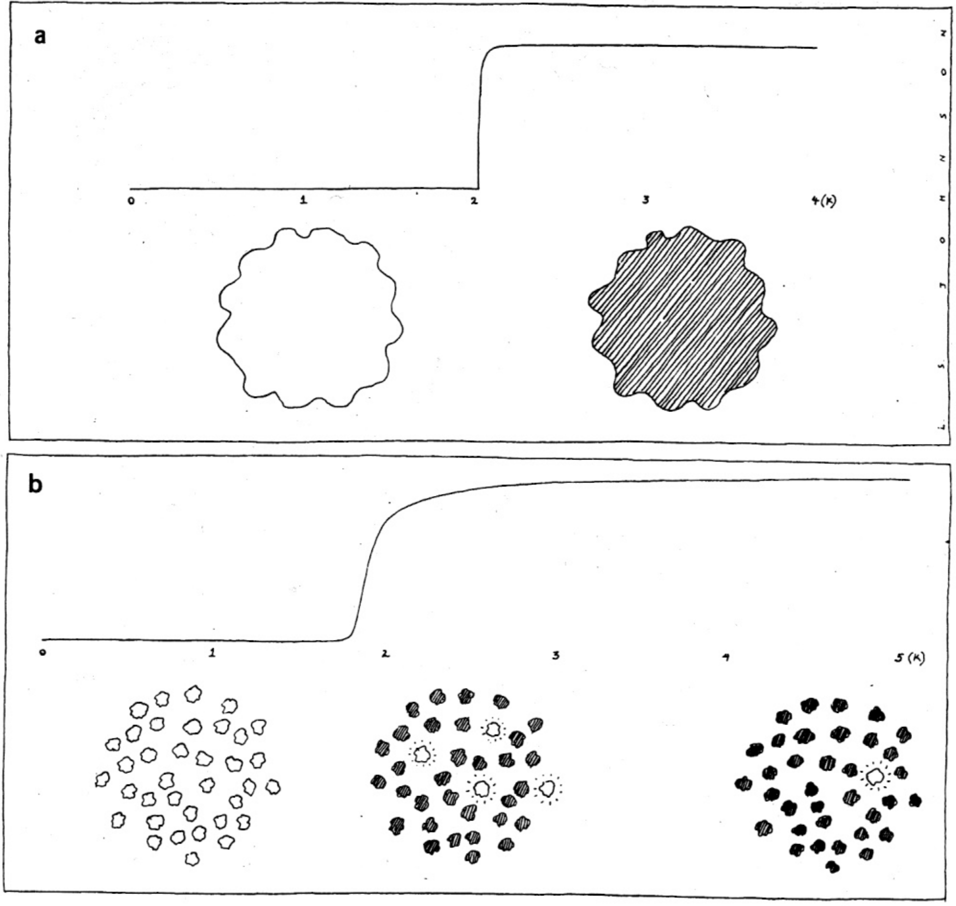

Fig. 5. Part a shows a curve of resistance against temperature for a bulk superconductor. The resistance drops abruptly to zero at the transition temperature which is characteristic for the material. At low temperatures the whole piece is superconducting; above the transition temperature all of it is in the normal state.

Part b shows a curve for a granular material, with a transition which is broadened by fluctuations. Because of the small number of electrons on each grain the thermodynamic fluctuations have a much greater effect than in the bulk material. Even at temperatures so high that the equilibrium state is normal, the fluctuations cause the grains to be superconducting some fraction of the time.

brief moments during which the grain is in the superconducting state. The experimental result is that the transition is no longer abrupt, but rather occurs gradually over an extended temperature range (Fig. 5b).

Threads and bulk

In its granular form aluminum can also be superconducting, but only under certain conditions. To investigate these conditions we measured the heat capacity of a series of specimens.

Why didn't we just measure the property which gave superconductivity its name, the resistivity, and its reciprocal, the conductivity? Actually we did that too, because they are so much easier to measure. But you learn much less that way. You can see at what temperature the resistance goes to zero, but even that information is often in doubt. The reason is that if even a single thread, from one end of the specimen to the other, goes superconducting, it will short out the rest, and the measured resistance will be zero. There is no way of knowing whether the whole specimen is superconducting, or whether, because of inhomogeneities, it is only some small fraction of it.

The heat capacity, by contrast, is a "bulk" property, i.e., it is proportional to the mass of material which is involved.

Heat capacity

Why is the heat capacity (or the heat capacity per unit mass, i.e., the specific heat capacity or specific heat) one of the physicist's most treasured and fundamental quantities? You let some heat flow to a chunk of matter and its temperature rises. The ratio of the added heat energy to the temperature rise is the heat capacity. What's so interesting about that, unless you're making a frying pan?

The reason becomes clearer if we look in more detail at what we are doing when we heat something. We are giving it some energy. The temperature rise depends on what happens to that energy. From it we learn something about the way in which the added energy is distributed over the possible kinds of internal motion.

There is hardly anything more fundamental about a material than its energy-level structure, and that's just what the specific heat tells us about.

The energy-level structure of a superconductor is different from that of a normal metal. At the temperature at which a metal makes the phase transition to become a superconductor the heat capacity changes discontinuously. The size of the discontinuity and the way in which the heat capacity decreases at lower temperatures are precisely described by calculations which follow from the theory of superconductivity. Any disagreement with the theoretical variation indicates that the specimen is not behaving like a "proper" superconductor.

Measuring the heat capacity

The measurement of heat capacities is notoriously full of pitfalls. There are two main kinds of difficulties. One is that it is necessary to measure small temperature differences and that's a lot harder than to measure small voltages. The other has to do with the heat energy. It is easy enough to measure heat input electrically, but the only part that we are interested in is that which goes to the specimen, not that which goes to the wires, the support, the heater, the thermometer, or to any other part of the surroundings.

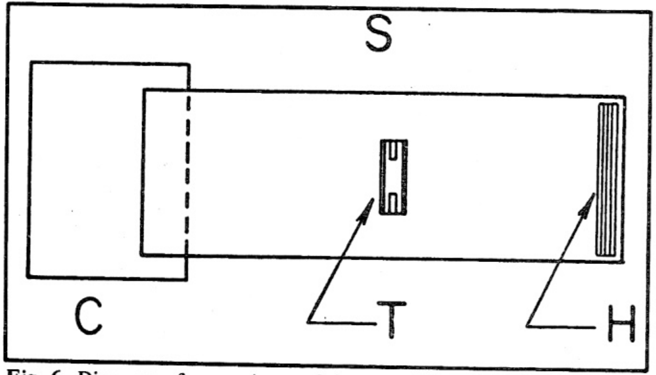

Fig. 6. Diagram of a specimen. C is a copper tab which is held in a clamp attached to a block held at a constant temperature. S is a glass substrate made of microscope cover glass. The granular aluminum is evaporated onto the substrate and is a few microns thick. H is the heater made of constantan. T is the thermometer. It consists of a 1/10-W Allen-Bradley carbon composition resistor which is sanded down until it is thinner than the substrate. At liquid helium temperatures such a resistor has a resistance which is from one to three orders of magnitude higher than at room temperature, and increases rapidly as the temperature decreases.

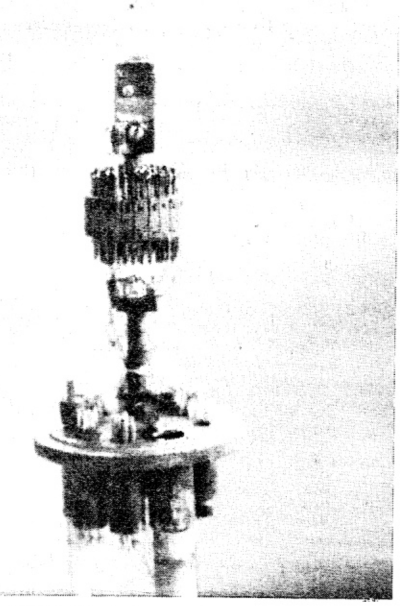

Fig. 7. A specimen in the apparatus.

In our case there is another problem. The mass of the evaporated granular aluminum films is typically only a few milligrams. The usual method of detecting the temperature change which results from a certain amount of heating is quite hopeless.

The first task, therefore, was to find a suitable method. We adapted a procedure which had been described, but never seriously used. It involves the production of a short pulse of heat at one end of a specimen. A temperature pulse travels along the specimen toward the other end which is held at a constant temperature. On its way it passes a thermometer whose temperature is recorded as a function of time. From the temperature variation the heat capacity of the specimen can be calculated.

The result is, of course, not only the heat capacity of the specimen, but includes that of the heater, thermometer, the glass substrate on which the film is deposited, and some unknown fraction of the wires to the heater and thermometer. We therefore make one measurement with the specimen on the substrate, one using the substrate without the specimen and we subtract the two.

Figure 6 shows a diagram of the specimen. Figure 7 is a photograph of a specimen in the apparatus. The wires to the heater and thermometer are very thin so as to minimize the heat flow along them. The whole specimen assembly is in an evacuated enclosure (which was removed for the photograph) so that the heat generated in the heater

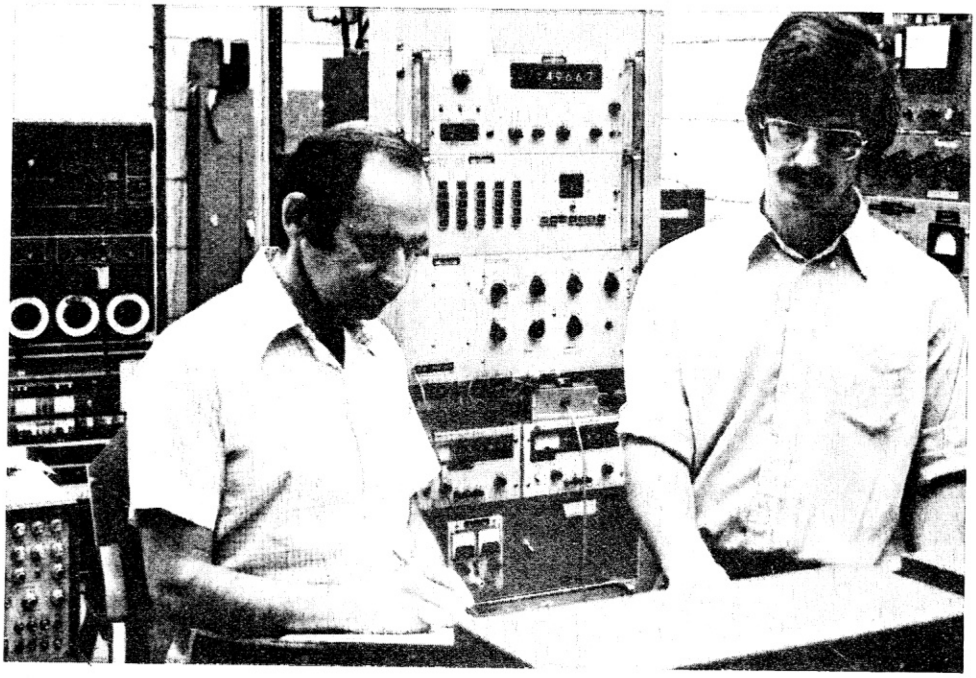

Fig. 8. Ray Filler and Peter Lindenfeld trying to look like serious physicists.

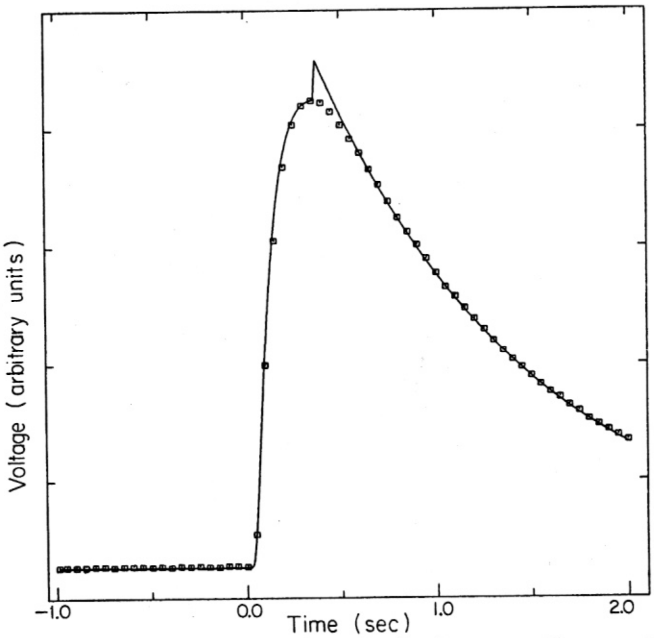

Fig. 9. Typical output pulse from the signal averager. The squares are the averaged experimental points. The full lines are computer fits. To the right of the maximum the fit is exponential, to the left it follows the form expected for an infinitely long specimen. The vertical height at the maximum corresponds to a temperature rise at the thermometer of about 15 millidegrees.

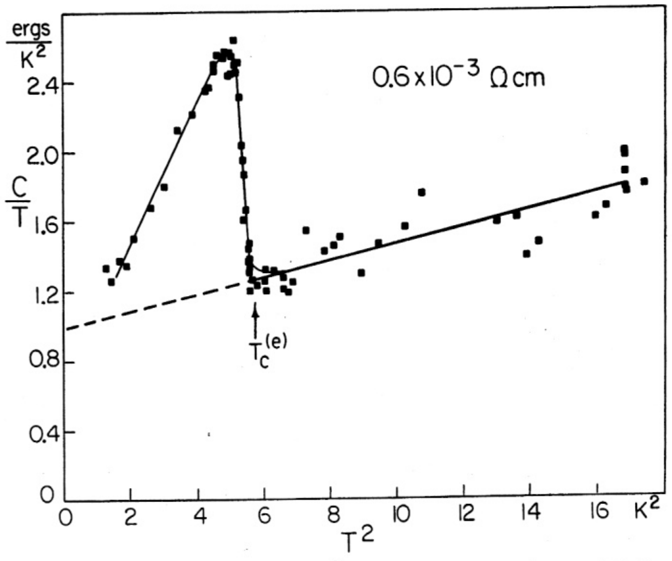

Fig. 10. Graph of C/T against T^2 for the specimen whose resistivity is 0.6 x 10^{-3} Ωcm. Although this specimen has the properties of a metal, the granular structure causes its resistivity to be several hundred times that of pure aluminum. Here C is the heat capacity and T the temperature in degrees Kelvin. $T_c^{(e)}$ marks the temperature at which the electrical resistance of the specimen goes to zero. The squares are the experimental points. The full line shows the variation expected for a bulk superconductor with the same transition temperature.

actually travels along the specimen and not through the air (or other gas) surrounding it. Outside the enclosure is a bath of liquid helium which provides an environment at the temperature at which the measurements are to be made, namely between about 1.2 and 4.2 degrees Kelvin.

The wires lead to the measuring equipment which is shown (with its operators) on Fig. 8. The photograph is staged (someone wanted a picture of "physicists at work") but the scene is otherwise genuine. The rack at right holds a vacuum-gauge control and a temperature regulator. The one in the middle has power supplies, switching boxes, and a digital voltmeter which is used for measuring the energy of a heat pulse. At the left is the heart of the equipment, a signal averager, which stores the response of the specimen thermometer as a function of time in 400 digital channels. We usually repeated the measurement about 40 times, after which the averaged pulse is fed to the minicomputer at the back for the numerical analysis.

Figure 9 shows an example of the output of the signal averager. The heat pulse occurs at the time marked "0.0" and lasts a tenth of a millisecond. The squares are the averaged experimental points (not all are shown). The continuous line to the right of the maximum is an exponential fitted by the computer. The line to the left of the maximum is also a computer fit, using the theoretically expected formula for the shape of the beginning part of the pulse. The computer produces these fits within seconds of the measurement, so that problems and difficulties become apparent right away.

Lattice and electronic heat capacity

The only part of the heat capacity that we are interested in is that of the conduction electrons. The measured amount includes the effect of the vibrational energy of the ion lattice. Fortunately the different temperature dependence of the two components makes it easy to separate them.

At the low temperatures of our experiment the lattice heat capacity varies as T^3 where T is the absolute temperature. In the normal state the electronic part varies as T, so that the total can be written as $C_n = \gamma T + AT^3$, where γ and A are appropriate coefficients. We can also write

$$\frac{C_n}{T} = \gamma + AT^2$$

which suggests a plot of C_n/T against T^2. We expect a straight line, with an intercept on the vertical axis equal to γ, and a slope equal to A.

From such a plot we can find the electronic heat capacity $C_{en} = \gamma T$ and the lattice heat capacity AT^3 Below the transition temperature the lattice part is given by the same relation, and we can subtract it from the measured heat capacity to find the electronic heat capacity in the superconducting state, C_{es}.

Experimental results

Finally we are ready to describe the results. We started with a specimen in the metallic region. Figure 10 shows C/T against T^2 to illustrate the separation of the two components of the heat capacity. Figure 11 shows C_{es}/C_{en} against T. On both figures the curve expected for a bulk superconductor is shown by the full line, and the experimental values are seen to follow the theoretical relation very closely. The agreement encouraged us to think that we were measuring correctly, and that any departures from this behavior in other specimens should be taken seriously.

Actually the significance of this first measurement goes even further. The fact that the sharp discontinuity at the transition temperature is well reproduced shows that all portions of the specimen make their transition together, so that the specimen is homogeneous. This was not the case with earlier results in our and in other laboratories. In fact

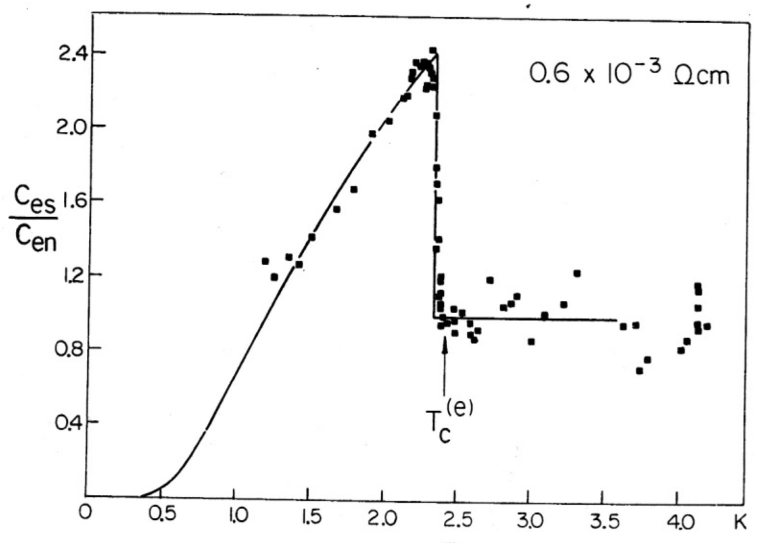

Fig. 11. C_{es}/C_{en} as a function of the temperature for the same specimen as for Fig. 10. C_{es} is the electronic heat capacity in the superconducting state, C_{en} is the electronic heat capacity in the normal state. Below the transition temperature C_{en} is assumed to continue to be proportional to T, just as it is above. Again the squares are the experimental points and the full line is the theoretical curve.

it was common to scoff at results on evaporated films because it was not thought that such specimens could exhibit the behavior characteristic of large, "well-behaved" and homogeneous samples. We are delighted to have been able to calm some of these doubts.

The next specimen had a higher resistivity and showed no evidence of a discontinuity at all. We could tell from the normal-state specific heat that almost half of it was metallic aluminum. Why was there no evidence of a transition?

There we have our main result. We concluded that the grains were now quite far from each other, and were so small that thermodynamic fluctuations would have a dominating effect. The specimen fluctuates between being superconducting and being normal. For a specimen of this size the time it spends being a superconductor is too small to show up in the measurement.

This result had been discussed theoretically, and detailed calculations of the expected variation of the specific heat with specimen size had been made. Our experimental result was in accord with the calculations.

In one sense we had now finished. In another we had a lot of work left. To see no transition at all may have been the result we had hoped for, but it was also somewhat uncomfortable. The measurements are sufficiently hard so that all kinds of things can go wrong. Like any experiment this one had to demonstrate its reliability by its internal consistency.

It was clear what we had to do. We had to make specimens intermediate between those which I have described. Unfortunately we are here in a region where a minute change in the oxygen pressure during the evaporation of the aluminum causes a large change in the resistivity. Just a little more oxide and the electrons have a much tougher time getting from one grain to the next.

It took time, hard work, and some luck, but finally we had enough intermediate specimens to have confidence in our conclusion. Figure 12 shows the gradual decrease in the discontinuity at the transition.

Is that the end of the story? Not quite. In fact, the demonstration of the effect of fluctuations was not primarily what we were after. It is true that we succeeded to do this more clearly than others who had tried in different ways before us. But, as an unkind reviewer of our paper pointed out, the result was entirely expected, and therefore less than exciting.

Physics describes the real world, and no amount of calculation takes the place of an experiment. Still, the conclusion we have described was hardly enough to justify the time and effort we had expended.

What we really tried to do requires us to turn the argument around. We wanted to investigate the structure of the material that we were dealing with, and the interactions responsible for its behavior.

We were interested in two aspects: one concerns the grains themselves, the other the coupling between the grains which allows the whole specimen to conduct electricity and to become superconducting.

First of all, our results show that granular aluminum really consists of grains. That, in other words, the metallic part is in small pieces, so small that when they are sufficiently separated they do not undergo a "bulk" transition to the superconducting state, because the fluctuations prevent them from doing so.

There had been a controversy on this point. There seemed to be a simple reason why a granular system should not be able to become a superconductor at all when the grains were at least partly isolated. The reason is that the transfer of an electron from one neutral grain to another causes both to become electrically charged, and therefore represents an increase in electrostatic potential energy. A calculation seemed to indicate that the size of this charging energy was such as to prevent superconductivity. The only way out of the dilemma seemed to be to say that the specimens did not really consist of separate metallic grains.

Our experiment shows that they do consist of grains after all, and that superconductivity exists where it had been thought to be impossible.

That's interesting, but it leaves a problem: why isn't the electrostatic energy there and doing its thing?

Gradually, under the impetus of the experimental results, the dilemma was resolved. The charging energy is there, but its destructive effect is not as great as had been thought. The first calculation had just used too simple a model.

It is necessary to take into account that there is a whole array of grains, and not just two at a time. Furthermore the actual specimen does not consist of perfectly regular grains, all at the same distance from each other. With these refinements the agreement between theory and experiment becomes excellent.

The calculations take into account the competition between the disruptive effect of the electrostatic energy and the attractive effect of the superconducting coupling energy (the "Josephson energy"). Because of the earlier discrepancy there had also been doubt and uncertainty about the coupling between the grains which causes the whole specimen to become superconducting and to act like a single homogeneous piece of material.

The discrepancy has now been removed, and the doubt resolved. We now understand the behavior of our material much better, both when its grains are isolated, and

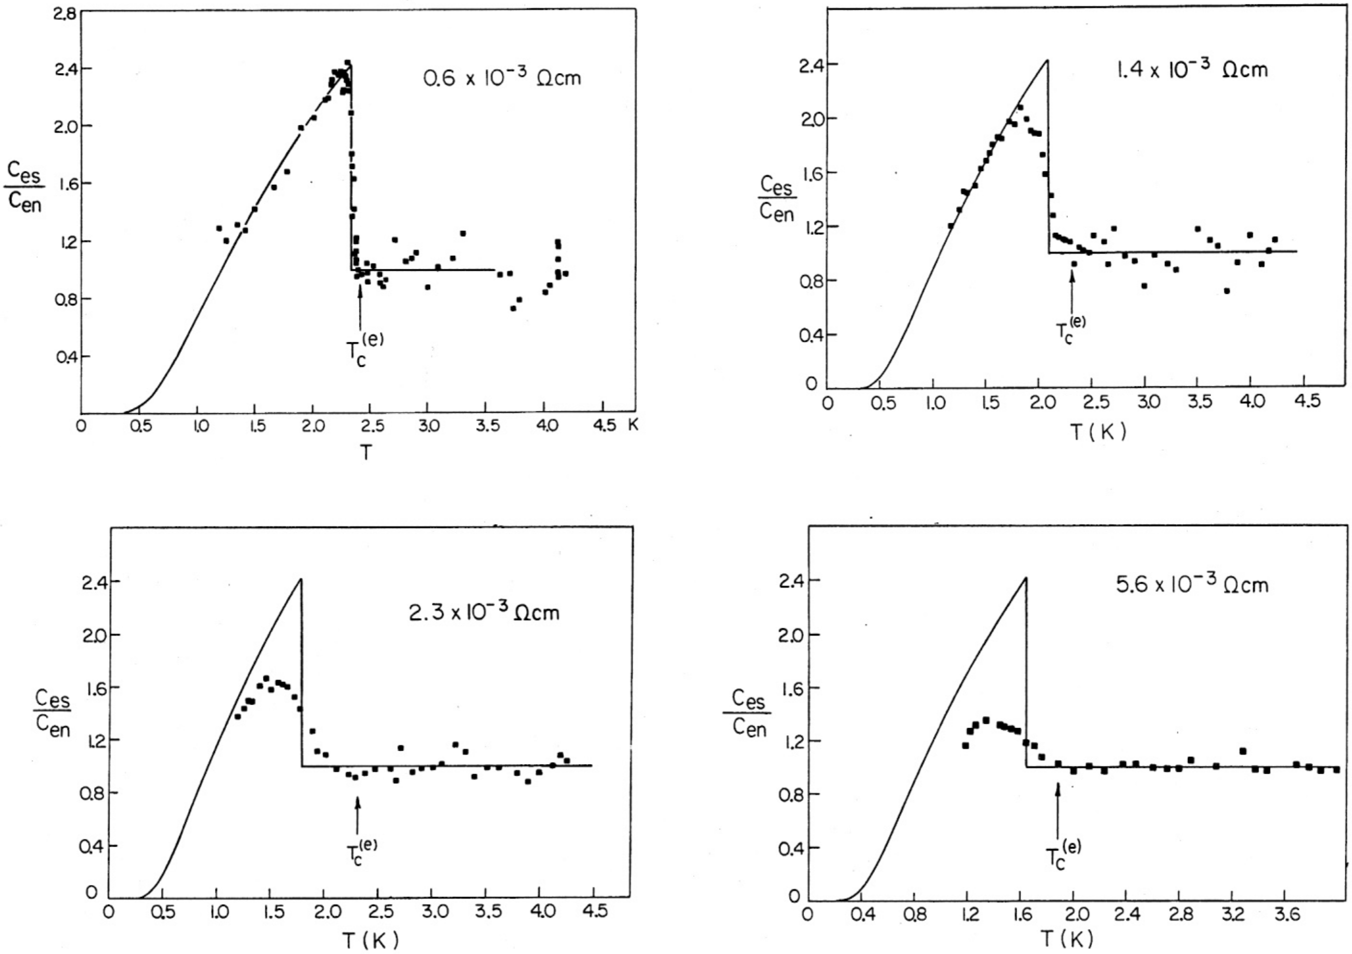

Fig. 12. C_{es}/C_{en} as a function of T for four specimens with increasing resistivity. The discontinuity at the transition temperature becomes smaller as the amount of oxide between the grains is made larger and the grains become more and more isolated from one another. In a completely isolated grain of this size the effect of the thermodynamic fluctuations becomes so large that the discontinuity in the heat capacity is completely suppressed.

when the coupling between the grains gives rise to superconductivity.

I had started the experiment during a sabbatical year as the guest of Guy Deutscher, who was then in the solid-state physics laboratory of the University of Paris in Orsay. Back home at Rutgers the work was continued by Ray Filler who wrote his Ph.D. thesis about the development of the method and some preliminary results. The measurements were finally done by Tom Worthington who was then a post-doctoral fellow. Guy Deutscher, who had in the meantime moved to Tel-Aviv University, came for several summers with help and encouragement.

It's been a long time since the beginning of our effort, and progress was sometimes painfully slow. But at the same time we got to know quite a bit about granular aluminum and found it to be an unusually interesting material. Its grains allow us to study size effects, as well as the coupling and the transport processes between them. We are now investigating the conduction processes of granular aluminum in the normal state, and continuing our study of the limits on the occurrence of superconductivity. The richness and variety of the phenomena exhibited by this material are so great that we will probably keep working with it for quite a while.

Acknowledgments:

My debt to my collaborators is evident. Many others have helped with illuminating conversations and discussions. I would like particularly to thank my colleague W. L. McLean, with whom I have discussed all aspects of this work.

The details of the work described in this article may be found in a letter by T. Worthington, P. Lindenfeld and G. Deutscher [*Physical Review Letters* **41**, 316 (1978)] and an article by R. L. Filler, P. Lindenfeld, T. Worthington, and G. Deutscher (to appear in the *Physical Review*). These papers cite references on related measurements on this and other granular systems. Ray Filler is now at the U.S. Army Electronics, Technology and Devices Laboratory, Fort Monmouth, N.J., and Tom Worthington is at the IBM Thomas J. Watson Research Laboratory in Yorktown Heights, N.Y.

The research was supported by the National Science Foundation under grant No. DMR 76-22962, by the Rutgers University Research Council, and by the U.S.-Israel Bi-national Science Foundation. Helium gas was supplied under a grant by the Office of Naval Research.

The drawings for Figs. 4 and 5, as well as that of the apples were made by Lonni Sue Johnson.

Reference

1. A survey of the theory of electrical conduction can be found in the September 1979 issue of *The Physics Teacher*: Philip B. Allen, "Electrical conductivity," Phys. Teach. **17**, 362 (1979).