Negative Absolute Temperatures

No substance can be cooled below absolute zero, but some physical systems can have a negative temperature on an absolute scale. Such temperatures are not colder than zero but are hotter than infinity

by Warren G. Proctor

Temperature is a property of matter that seems to have a well-defined range of possible values. If all the heat could be extracted from a body, it would be assigned a temperature of absolute zero, or zero degrees Kelvin, and it could not become any colder. At the other extreme, heat can always be added to a body, at least in principle, and so its temperature can increase without limit. All absolute temperatures therefore seem to fall in the range between zero and positive infinity. (Negative temperatures on the Celsius and Fahrenheit scales arise only because the zero points of those scales do not correspond to absolute zero.)

Temperatures of the kind that are measured with an ordinary thermometer are indeed confined to the range of positive numbers. There are a few physical systems, however, whose temperature when measured on an absolute scale can assume negative values. What may seem even more peculiar, such temperatures are not colder than absolute zero; a system with a negative temperature can give heat to one at an infinite temperature and is therefore hotter. The negative temperatures are not "below zero"; they are "above infinity."

In order to understand how an absolute temperature can be negative, it is first necessary to recognize that a single body can simultaneously include systems of particles that exhibit more than one temperature. In a crystalline solid, for example, the ordinary temperature—the property measured by a mercury thermometer—reflects the average vibratory motion of the atomic nuclei: that temperature cannot be negative. The same nuclei, however, have freedom of orientation as well as freedom of motion, and the set of all the orientations makes up an independent system of particles that can also have a temperature. It is temperatures of this kind that can take on negative values. Because the energy in question is small, a crystal in which the nuclear orientations are at a negative temperature feels neither hot nor cold to the touch. Negative absolute temperatures are nonetheless real; they can readily be measured and can even be manipulated by the methods of calorimetry.

The concept of temperature describes a relation between two quantities, energy and disorder. In fundamental terms what is measured by temperature is the change in the disorder of a system as the energy of the system changes.

In a crystalline solid the major contribution to disorder is the vibratory motion of the atomic nuclei. A crystal of table salt, for example, is often repre-

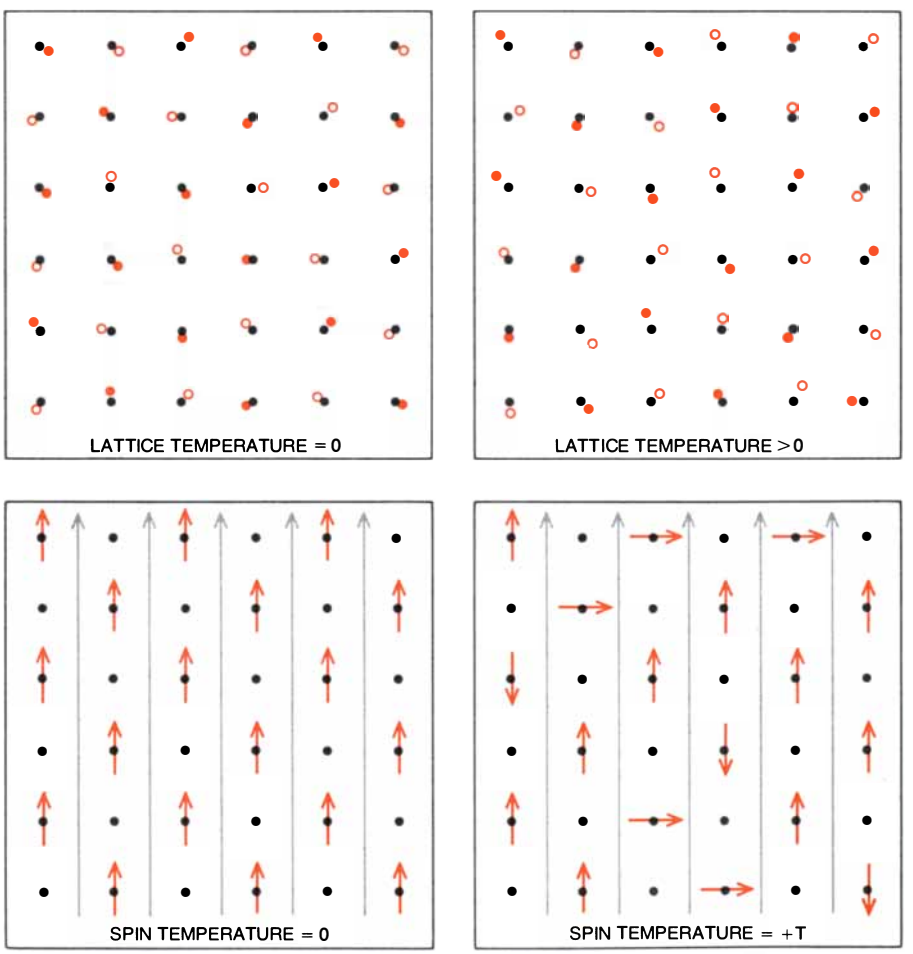

CONCEPT OF TEMPERATURE can be applied to any physical system that has a definite energy and an entropy, or measurable disorder. In a solid the lattice, or ordinary, temperature (*top*) reflects the vibratory motion of the nuclei (*colored dots*) around their average positions (*black dots*). Temperature increases to the right, and so does the average energy of the nuclei. Moreover, every increase in energy is accompanied by an increase in entropy. (The displacements are greatly exaggerated here.) Another system with a defined temperature (*bottom*) is

sented as a perfect three-dimensional lattice in which the nucleus of each sodium and chlorine atom occupies a corner of a cube. Actually the ideal lattice represents only the average or equilibrium positions of the nuclei, around which they vibrate randomly. The amplitude and the frequency of the vibrations vary from moment to moment and from atom to atom; the temperature of the crystal summarizes in a single number the average amplitude and the distribution of frequencies. This number is the ordinary temperature, or what I shall designate the lattice temperature.

As a crystal is cooled the average amplitude and frequency of the nuclear vibrations diminish. Before the development of quantum mechanics it was thought that the vibrations would stop entirely if the crystal were cooled to absolute zero, but in the quantum-mechanical interpretation the vibrations cannot cease: residual vibrations, called zero-point motions, persist even when the temperature reaches absolute zero. Because the zero-point motions cannot be abolished, the residual disorder of the system cannot be eliminated, no more energy can be extracted and the crystal cannot be further cooled, even in principle.

Moving in the other direction on the temperature scale, as heat is added to a solid the average amplitude and frequency of the vibrations increase, but again the motions of the atoms are constrained by the rules of quantum mechanics. Those rules specify that only certain modes of vibration are possible, each with a definite energy, and intermediate energies are forbidden. At absolute zero all the nuclei are in the lowest-energy vibrational state, the state that specifies the zero-point motions. As the crystal absorbs heat some of the atoms jump to higher energy levels, each of which describes a different vibrational mode. As the temperature is raised further, progressively higher levels are populated, but many atoms still remain in the lower levels, so that the population distribution has a pyramidal form. There is no intrinsic limit on the energy that can be acquired by a nucleus, and the number of energy levels available is infinite. It follows that temperature can increase without limit. (Of course any real substance heated indefinitely must melt and then boil, but these complications do not alter the fact that the number of possible energy levels has no upper bound.)

The disorder introduced into a crystal lattice by heating can be described quantitatively in terms of entropy. For the purposes of this article entropy can be defined as a measure of the difficulty of guessing the energy level in which a particular nucleus lies. At absolute zero all the nuclei are in the same energy level (the lowest one possible). Hence there is no uncertainty about their state, and the entropy of the system is at a minimum. As the temperature increases and more levels become available the probability that an atom occupies any given level becomes smaller, and so the entro-

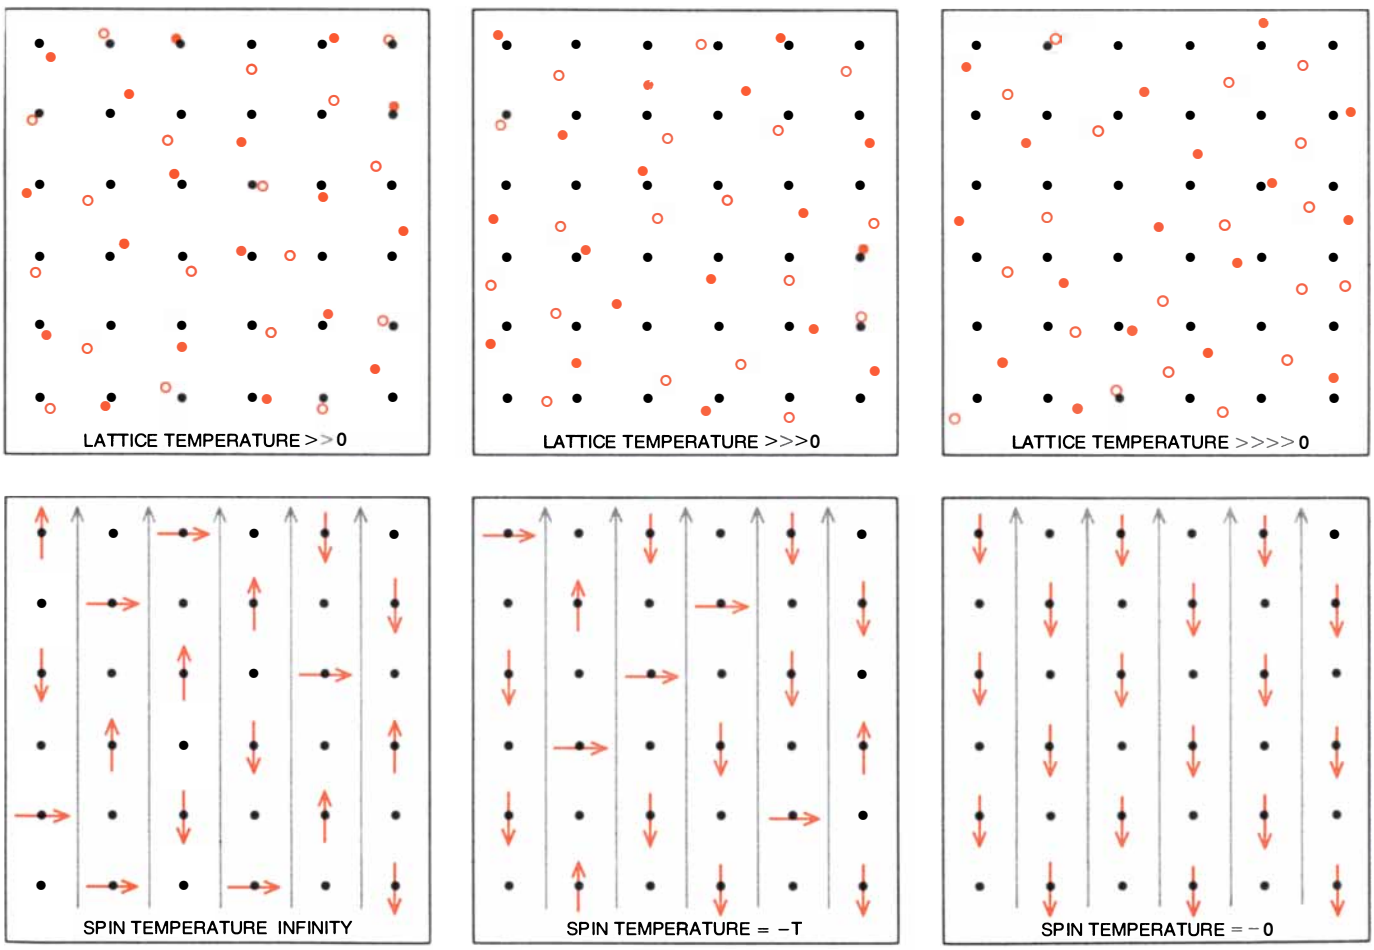

made up of the spins (*colored arrows*) of the nuclei in a crystal. In a magnetic field the spins have only a few possible orientations, in this case three: parallel to the field, antiparallel to it or perpendicular to it. At a spin temperature of absolute zero (*far left*) all the spins line up parallel to the field in a state of minimum energy and entropy. As the spin temperature rises a few spins flip to other configurations, and at infinite temperature (*center*) the spins are equally distributed among the possible orientations; in this state entropy is at a maximum. Adding energy forces more of the spins to assume the antiparallel orientation. As a result entropy is reduced; what is more, the spin temperature becomes negative. When all spins are antiparallel (*far right*), entropy is minimal and spin temperature has reached "minus zero."

py increases. At infinite temperature any atom would have an equal probability of occupying any one of an infinite spectrum of levels, and so the entropy would also be infinite. It is important to note that as energy is added to the system both the entropy and the temperature increase continuously.

Like the vibrations of the nuclei in a crystal, the orientations of the nuclei are governed by the rules of quantum mechanics. In order for an atomic nucleus to have a defined orientation it must first have some detectable asymmetry: it must "point" in some direction. The orientation is defined by the spinning of the nucleus, which can be imagined as being much like the rotation of the earth around its axis. The presence of the axis designates a direction in space along which the nucleus is aligned. Indeed, each spinning nucleus can be associated with an imaginary arrow (the spin vector) that indicates its orientation. The arrow is drawn parallel to the spin axis, and by convention it points in a direction such that an observer looking along the arrow sees the nucleus spinning clockwise.

Not all nuclei spin, and those that do not spin have no definable orientation. Nuclei that do spin, however, are allowed to take on only certain well-defined quantities of angular momentum. When the angular momentum is measured in fundamental units, it can have values of 0, 1/2, 1, 3/2 and so on; intermediate values are forbidden. For the sake of simplicity I shall assume here that all nuclei have the same spin, equal to 1 unit. This assumption is actually false (sodium and chlorine nuclei, for example, each have a spin of 3/2 units), but it makes the geometry of the nuclear spins much easier to visualize and it introduces no serious errors.

The direction of the nuclear spin vector can be detected experimentally only because it corresponds to the direction of the nuclear magnetic field. All nuclei, of course, bear an electric charge (equal to the number of protons in the nucleus); if the nucleus is spinning, the charge is in motion and it generates a magnetic field parallel to the spin vector.

The magnetic field of a nucleus is detected through the interaction of the nuclear field with an artificially imposed external field. Under the rules of quantum mechanics a spinning nucleus can have only a finite number of possible orientations in an external field, the number being proportional to the spin angular momentum of the nucleus. A nucleus with a spin of 1/2 unit has only

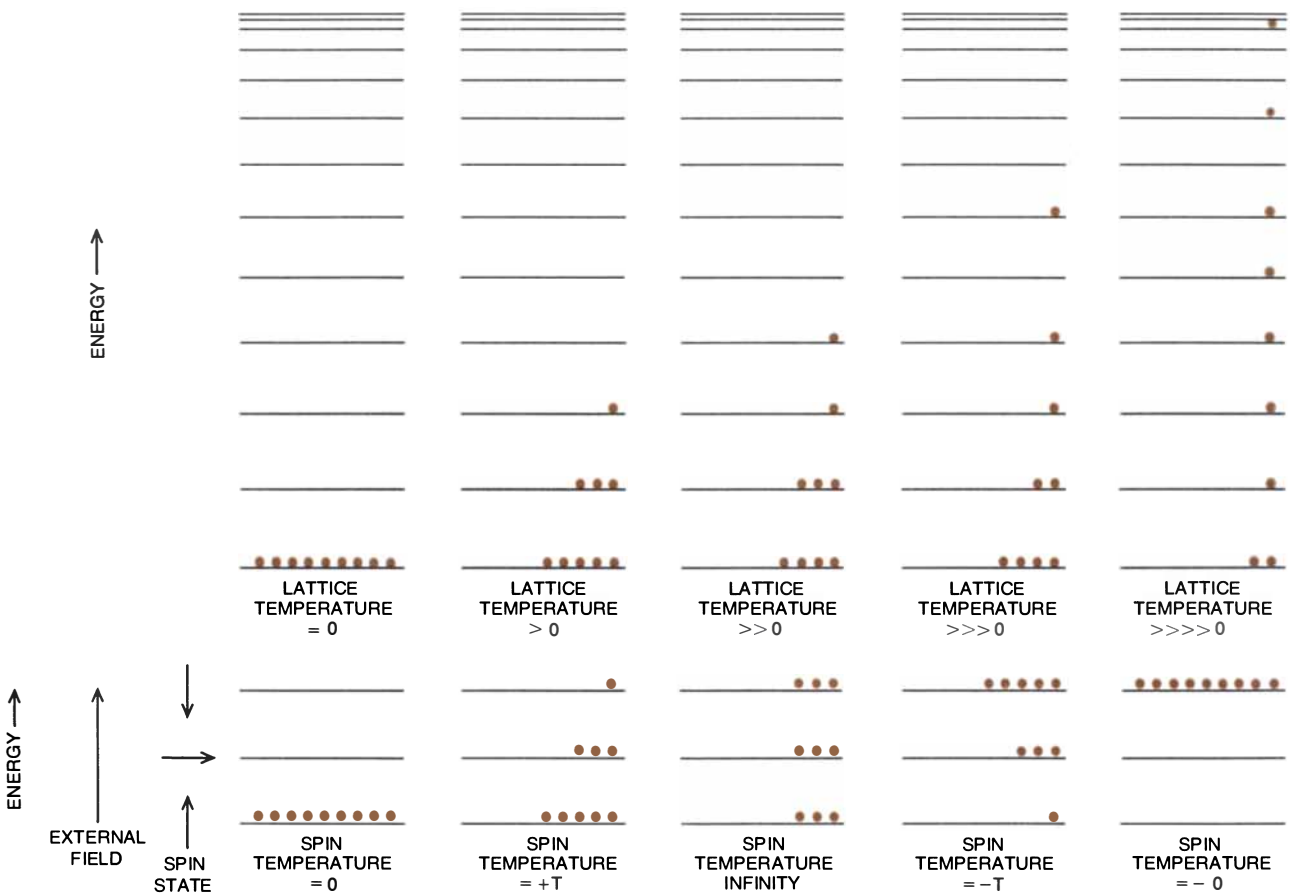

ENERGY LEVELS in a solid correspond to the possible states of vibratory motion (*top*) or to spin orientations (*bottom*). Here the levels are represented as horizontal lines and their energy is proportional to their height above the lowest level; colored dots mark the number of nuclei that occupy each level. For either system absolute zero (*far left*) is a state of minimum energy, since all the nuclei are in the lowest possible energy level; it is also a state of minimum entropy. The entropy can be regarded as a measure of the difficulty of guessing what level a nucleus is in at a given moment; the guessing is obviously easiest when all the nuclei are in the same level. In the vibrational system adding energy forces some nuclei to occupy higher levels, but larger numbers of nuclei remain in the lower levels. As more levels are occupied the entropy increases, since it is more difficult to decide in what level a given nucleus lies. The number of levels available is potentially infinite, and at infinite temperature (*far right*) the nuclei are randomly distributed among them, that is, each level has an equal probability of being occupied. In the spin system there are only as many energy levels as there are possible orientations of the nuclei, in this instance three. As in the vibrational system, adding heat promotes a few nuclei to higher levels, and at infinite temperature (*center*) the nuclei are evenly distributed among the available levels. The entropy is then at its maximum. Because the number of levels in the spin system is finite, however, states of even higher energy are possible; most or all of the nuclei can be driven to the highest, antiparallel energy level. Because the entropy declines as the energy increases in this region (the opposite of the usual relation), the temperature is assigned a negative value. When all the nuclei are in the highest level (*far right*), the entropy is again at a minimum and the temperature is minus zero.

U.S. NAMES RABBIT DIESEL #1 IN ECONOMY.

We always knew we made terrific economy cars.

Now it's official.

Because according to the figures published in the official 1978 EPA Mileage Guide, a Rabbit Diesel gets the highest mileage of any car in America: an incredible 53 MPG on the highway, and 40 MPG in the city.

(Of course, these estimates may vary depending on how and where you drive, optional equipment and your car's condition. That's also official.)

We could go on and on about all the innovations we've built into our newest Rabbit.

Like the fact that it goes on and on (amazingly!) without ever needing a major tune-up. (There are simply no spark plugs, points, condensers, or carburetors to tune.)

Our mechanical wizardry doesn't stop with the engine either.

There's front-wheel drive for better tracking.

There's more room for people than 37 other cars you could buy.

And a Rabbit Diesel goes like a bat out of you-know-where. In fact, it set 31 world records for diesels on a track in Miramas, France.

Come test it for yourself.

For years we've built cars that use very little gas.

Now we've built one that doesn't use any at all.

VOLKSWAGEN DOES IT AGAIN

The cream. The crop.

Olympus introduced the OM-1 and startled the world of photography with the creation of the compact SLR. Today, the OM System is still the cream of the crop.

Because while others have emulated our compact design, OM cameras continue to offer features others can't.

The OM-1 Becomes #1.

Enter the OM-1. Suddenly, the SLR camera is 33% smaller and lighter, yet incredibly rugged to meet the demands of professional wear and tear. Miraculously, the viewfinder is 70% brighter and 30% larger for faster, easier composing and focusing.

And suddenly, the OM-1 became the #1 selling compact SLR. Its metering system is designed to give complete control to professionals and photojournalists. No distractions, blinking lights, or obscured images in the viewfinder.

A Quiet Revolution.

Olympus created a unique shock absorber and air damper system to eliminate noise and vibration, for sharper, unobtrusive photography. Especially vital for long tele shots and macro/micro photos.

The Biggest Smallest System.

More than 280 components, all compact design, include 13 interchangeable screens so you can meet any photographic challenge. Ingeniously designed to change in seconds through the lens mount. And more compact lenses than any other system, each a marvel of optical design and performance.

Olympus "Unlocks" Motor Drive.

OM-1 is still unsurpassed in its continuous-view motor drive capability: 5 frames per second. And a Rapid Winder that fires as fast as 3 shots a second! With no mirror "lock-up," regardless of lens used.

Enter The OM-2. Automatically.

It's the fully automatic OM, with major differences from *all* other automatics! The only SLR with "off-the-film" light measurement for those photographers demanding the ultimate innovation in automatic exposure control. Which means each frame in motor drive or rapid winder sequences is individually exposure-controlled. And it makes possible the unique Olympus 310 Flash whose exposure duration is controlled by the camera's metering system.

And of course, the OM-2 shares every other innovation and system component with the OM-1.

We Wrote The Book On Compact SLR's.

Write for our full color brochure: OLYMPUS, Woodbury, New York 11797. Read it all. Discuss the advantages of an Olympus with your photographer friends.

Visit your camera store. Compare. You'll discover that Olympus is not only the cream of the crop. It's the crème de la crème!

OLYMPUS OM SYSTEM

two possible orientations: the spin vector can lie parallel to the field or it can be antiparallel, opposing the field. No other positions are allowed. The nuclei I shall be considering, with one unit of angular momentum, have three possible orientations: the spin vector can be parallel to the field, antiparallel to it or perpendicular to it.

All the nuclei of the same type in a crystal, or even in a solid made up of many crystals, can be considered as a spin system. Thus a crystal of table salt includes two spin systems, one composed of all the sodium nuclei and the other of all the chlorine nuclei. When the crystal is permeated by an external magnetic field, each of the systems has a measurable energy and a definite entropy. As will be seen, each of the spin systems can also be assigned a temperature, which I shall distinguish from the lattice temperature by calling it the spin temperature.

Just as a compass needle tends to align itself with the earth's magnetic field and work must be done to turn it from that orientation, nuclear magnets tend to line up parallel to an imposed field and energy must be supplied if they are to assume any other orientation. Hence the three possible orientations of a spin-1 nucleus represent not only different geometric configurations of the nuclear spins but also different energy levels. The difference in energy between the levels is proportional to the strength of the external field. A nucleus with its spin vector parallel to the applied field is in the lowest-energy state; a nucleus oriented perpendicular to the field has a slightly higher energy, and one antiparallel to the field has a still higher energy.

The distribution of the nuclei in a spin system among the three levels is similar to the distribution of nuclei among vibrational levels. At a spin temperature of absolute zero all the nuclei are in the parallel configuration; that is the state of the spin system in which both the energy and the entropy are at a minimum. As the spin temperature increases, some nuclei change their orientation and jump to the higher perpendicular and antiparallel energy levels. As long as the temperature is finite, however, more nuclei remain in the bottom level than in either of the higher levels. At infinite spin temperature all three levels are uniformly populated, a condition in which the system has maximum entropy.

In the presence of an external field any change in the orientation of a nucleus must be accompanied by a change in energy. In many cases these energy transactions are strictly internal to the spin system. For example, when one nucleus "flips" from perpendicular to parallel, giving up a small quantum of energy, a nearby nucleus may make the opposite transition, from parallel to perpendicular, absorbing the same quantum. The energy of the spin system as a whole therefore remains unchanged, since there are always the same number of nuclei in each state. Such exchange interactions take place constantly and very quickly, and they distribute the energy of the spin system uniformly throughout the volume of the crystal.

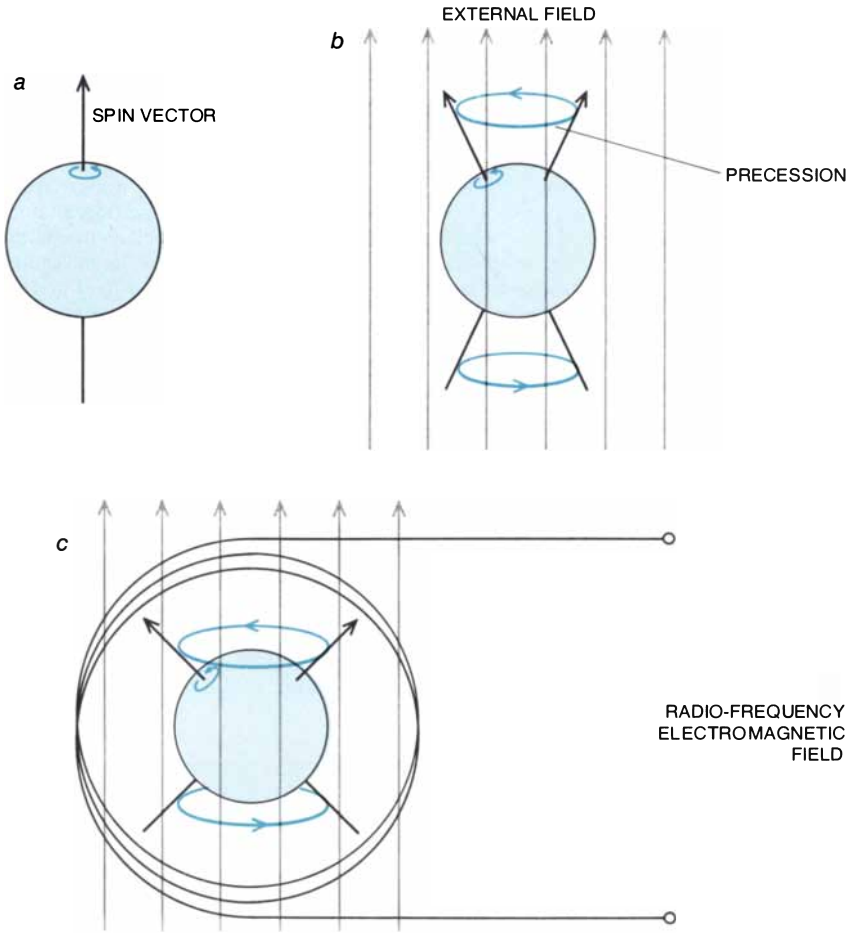

SPIN OF A NUCLEUS defines its orientation. An arrow can be drawn along the axis of spin (*a*); by convention this spin vector points in a direction such that an observer looking along the arrow would see the nucleus spinning clockwise. The spinning of the nucleus generates a magnetic field, which is directed parallel to the spin vector. In an external magnetic field (*b*) the interaction of the two fields causes the nucleus to precess. The orientation of the spin vector can be determined by imposing a high-frequency electromagnetic field transverse to the steady external field. When the frequency of the alternating field matches that of the nuclear precession, the nuclei absorb energy strongly. This technique, called nuclear-magnetic-resonance spectroscopy, can be employed to measure spin temperatures and also to alter the spin directions.

For a nucleus to change its orientation without a compensating transition by a neighboring nucleus, the required quantum of energy must be supplied by or absorbed by some external source. That source is the lattice itself, which can share in the energy available by means of a weak coupling between the energy levels of the spin system and those of the vibrational system. A vibrating nucleus constitutes a moving electric charge, and so it gives rise to a magnetic field that can interact with the spin-generated fields of neighboring nuclei. If the energy of a vibrational mode happens to correspond to the energy absorbed or emitted in a spin transition, then the spin-flip can take place. Such coincidences are rare, but over a period of several minutes they bring the spin system and the lattice vibrations into thermal equilibrium. If the spin temperature is initially the higher one, then most of the spin-flips are to a lower-energy orientation and the spin system gives up energy; it is cooled by contact with the colder lattice. If the spin system is cooler at the outset, then most of the spin transitions represent a gain in energy and the spin system is warmed. In either case the process is called thermal relaxation.

The several minutes required for thermal relaxation in many crystals is a convenience for the experimenter. It is a period brief enough for the spin system to be brought readily into equilibrium with the lattice: the experimenter need only place a crystal between the pole pieces of a magnet and wait for a few minutes; he can then be sure that the spin system

and the lattice are at the same temperature. On the other hand, the relaxation time is long enough for measurements to be made on the spin system as if it were isolated from the lattice.

The state of the spin system is measured by the technique of nuclear-magnetic-resonance spectroscopy. In this technique a crystal is exposed to a strong, steady magnetic field, which tends to align the nuclear magnetic fields and also causes the nuclei to precess, or wobble around their spin axes. The precession has a characteristic frequency, which can be determined by imposing a second, oscillating field perpendicular to the first one. This second field is actually the magnetic component of a radio-frequency electromagnetic wave. When the frequency of the applied field is equal to the precession frequency, a "resonance" is observed: the spin system absorbs energy copiously, and the characteristics of the electromagnetic circuit are thereby altered. By measuring the precession frequency in this way several properties of the spin system can be determined. Among the most important of them is the proportion of the nuclei that are aligned parallel to the external field.

At any temperature below infinity there are always a few more nuclei parallel to an external field than there are antiparallel. This observation in effect states that a crystal in a strong magnetic field becomes slightly magnetized; it is this magnetic polarization that is actually measured by nuclear magnetic resonance. The relation of the magnetic polarization to the external field and to the temperature is given approximately by the principle called Curie's law (after Pierre Curie). The law states that the polarization is directly proportional to the strength of the applied field, customarily designated H, and inversely proportional to the temperature, T, of the spin system. Thus the polarization, M, is proportional to H/T. By simply rearranging the terms of this relation it follows that the spin temperature must be proportional to H/M. Both H and M are quantities that can readily be measured.

The relation expressed by Curie's law suggests a simple experiment. Suppose a crystal of salt is placed in a known magnetic field and the spin system is given time to reach thermal equilibrium with the lattice temperature. The magnetic polarization of the crystal is then measured by nuclear-magnetic-resonance spectroscopy. This stable system can subsequently be disturbed by quickly changing the external field ("quickly" meaning in a short time compared with the thermal-relaxation time). Suppose, for example, the field is reduced to half its original strength. Immediately after the external field is reduced the polarization of the spin system remains essentially unchanged; there has been too little time for the nuclear spins to be reoriented. Since the spin temperature measures the nuclear orientations, intuition suggests that the spin temperature too should be unchanged. Curie's law, however, requires that the temperature be proportional to H/M, and since H, the external field, has been reduced by half, the spin temperature must also be half its former value.

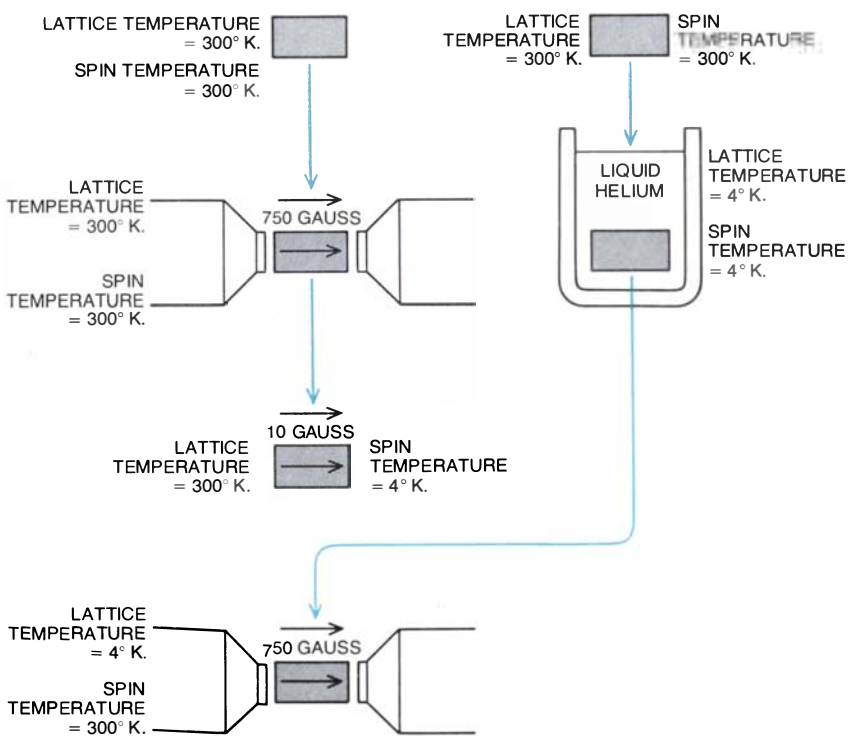

EQUIVALENCE OF SPIN TEMPERATURE and lattice temperature was demonstrated in an experiment carried out by Anatole Abragam and the author. The experiment begins with two identical crystals at room temperature, say 300 degrees Kelvin. One crystal is placed in a magnetic field of 750 gauss, which tends to align the nuclear spins with the external fields. After several minutes the spin temperature reaches equilibrium with the lattice temperature; the magnetic polarization, or the proportion of the spins parallel to the field, is then measured by nuclear magnetic resonance. The second crystal is immersed in liquid helium until the spin and lattice temperatures reach equilibrium; the temperature of equilibrium must be the same as that of the helium, about four degrees K. Removing the first crystal from the magnet reduces the field by a factor of about 75 and reduces the temperature by the same amount, to four degrees K. If either crystal is then placed in the field, the magnetic polarization of the two crystals is found to be identical. Thus there is no difference between the spin temperature of a crystal cooled by removing it from a magnetic field and that of a crystal cooled by contact with a cold lattice.

This behavior is so startlingly unlike the behavior of more familiar thermal systems that it calls into question the very notion of assigning a temperature to a spin system. It could be argued, for example, that the concept of temperature should be employed only where knowledge of a system (such as the distribution of vibrational states in a crystal) is incomplete and cannot be specified any more precisely. In the nuclear-magnetic-resonance experiment, however, the temperature is in fact calculated from a more detailed knowledge of the underlying state of the system. Perhaps it would be better in this case to ignore Curie's law and say that the temperature is undefined.

These issues can be further explored in another experiment. Again a crystal is placed in a magnetic field and the polarization is measured after it has reached its equilibrium value. This time the external field is briefly reduced all the way to zero, and then it is restored to its original strength and the polarization is measured again. A convenient way to produce these abrupt changes in the external field is simply to remove the crystal from between the pole pieces of the magnet, hold it a few seconds and replace it. When such an experiment is undertaken, it turns out that the magnetic polarization is unchanged.

When the external field surrounding a nucleus is abolished, the spin energy levels immediately collapse. There is no longer any energy difference between the three possible orientations; indeed, there is no longer a field to define a set of allowed orientations. The persistence of the magnetic polarization after an interval outside the imposed field might therefore be interpreted as evidence for the validity of the spin temperature. In this argument it is the spin temperature

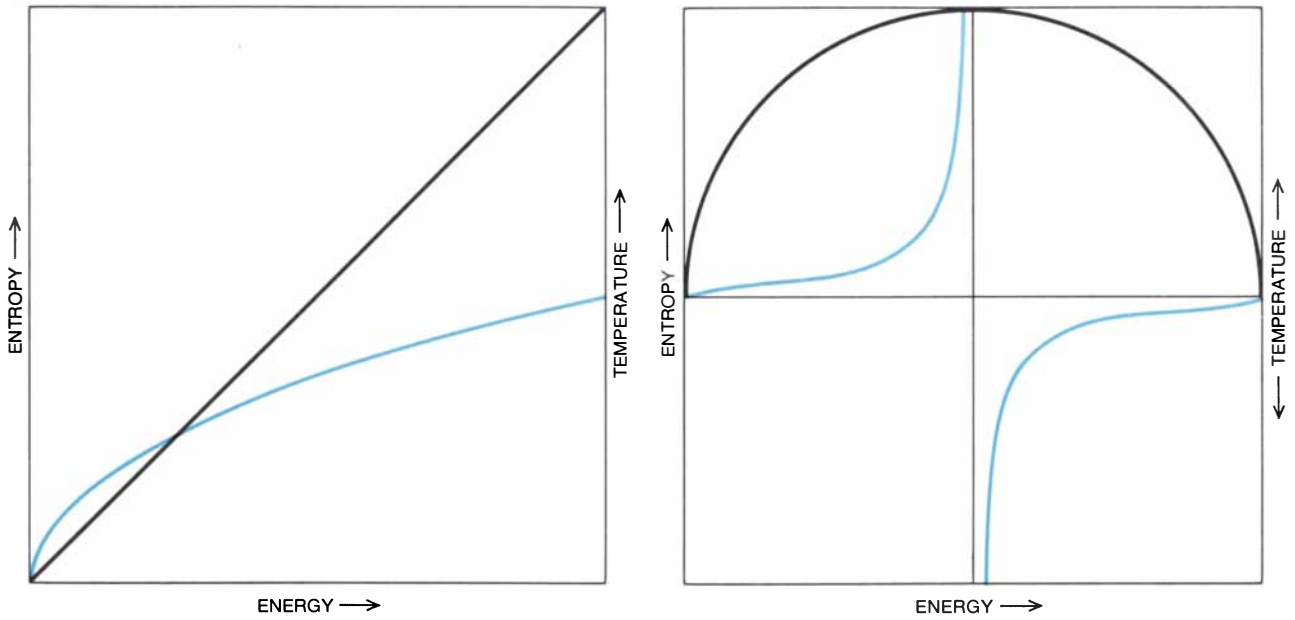

ENERGY AND ENTROPY are the basic elements in a definition of temperature in which negative values arise naturally: Temperature measures the amount of energy that must be added to a system to yield a given change in entropy. Here entropy (*black curves*) and temperature (*colored curves*) are both graphed as functions of energy. In a vibrational system (*left*) an increase in energy invariably brings an increase in entropy, and so the temperature is always positive. In the spin system, however (*right*), the entropy has a maximum possible value; at that point the change in the entropy is zero, and so the temperature is infinite. With each further increase in energy the entropy is reduced, and so the sign of the relation changes: the temperature becomes negative and at the maximum energy reaches minus zero.

that is "remembered" by the system and that specifies the correct polarization when the field is restored.

In this experiment too, however, there is a counterargument based on what seems to be a more fundamental description of the spin system. The net magnetic field acting on any nucleus in the interior of a crystal is made up not only of the externally imposed field but also of the fields generated by neighboring nuclei. Although the external field is removed, the quantum states of the nuclei may be preserved by interactions with these internal fields. The situation is complex, but it appears that the experimental result could be explained without reference to spin temperature.

The question of whether or not spin systems can have a temperature was resolved through an experiment suggested by Anatole Abragam of the Collège de France and carried out by Abragam and me at the Saclay Nuclear Research Center near Paris. The experiment tested the equivalence of spin temperatures created by magnetic cooling and those achieved by more conventional methods of refrigeration.

A salt crystal at room temperature (roughly 300 degrees Kelvin, or degrees Celsius above absolute zero) was placed in a magnetic field of 750 gauss, and the spin temperature was allowed to reach equilibrium with the lattice temperature. The crystal was then removed from the magnet, so that the field acting on the nuclei was reduced from 750 gauss to whatever residual field was provided by the local, internal fields; these fields are estimated to be about 10 gauss. The intensity of the field was therefore reduced by a factor of 75, and if Curie's law makes correct predictions under such circumstances, the spin temperature should have been reduced by the same factor, from 300 degrees K. to four degrees K.

It had already been shown that if the same crystal were replaced in the field, the magnetic polarization would be found to be unchanged. Instead of following this procedure another crystal was substituted for the first one; the second crystal was identical with the first in all respects except that it had been kept for some time in a bath of liquid helium at four degrees K. and in the absence of a magnetic field. It could therefore be assumed that the lattice temperature and the spin temperature of the second crystal were in equilibrium at four degrees K. When the second crystal was introduced into the 750-gauss field, the measured polarization was the same as that in the first crystal.

The interpretation of the experiment is straightforward: there is no measurable difference between a spin system cooled to four degrees K. by removing it from a magnetic field and an identical system cooled by thermal contact with a lattice at the same temperature. Although the behavior of the spin temperature in a rapidly changing magnetic field seems bizarre, it is indistinguishable from the behavior of a spin temperature set by conventional means.

Once the validity of spin temperatures has been established, it is a short step to demonstrating the existence of negative spin temperatures. It was pointed out above that as temperature increases, progressively higher energy levels are occupied. At absolute zero all the nuclei are in the lowest level; at temperatures between zero and infinity there is a pyramidal distribution, with more nuclei in the lower levels than in the higher ones; at infinite temperature all energy levels have an equal probability of being occupied. Up to this point the relation between temperature and the distribution of occupied states holds as well for the spin temperature as it does for the lattice temperature. A difference emerges, however, if an attempt is made to add more energy to a system that is already at infinite temperature.

For a lattice system to reach an infinite temperature would require an infinite source of energy, and there is no conceivable state of the system with higher energy. A spin system at infinite temperature, on the other hand, clearly has a finite energy: it is simply the sum of the energies needed to keep a third of the nuclei in each of the three possible spin states. Moreover, it is easy to imagine configurations of the spin system that would have higher energy. All that is necessary is to invert the population

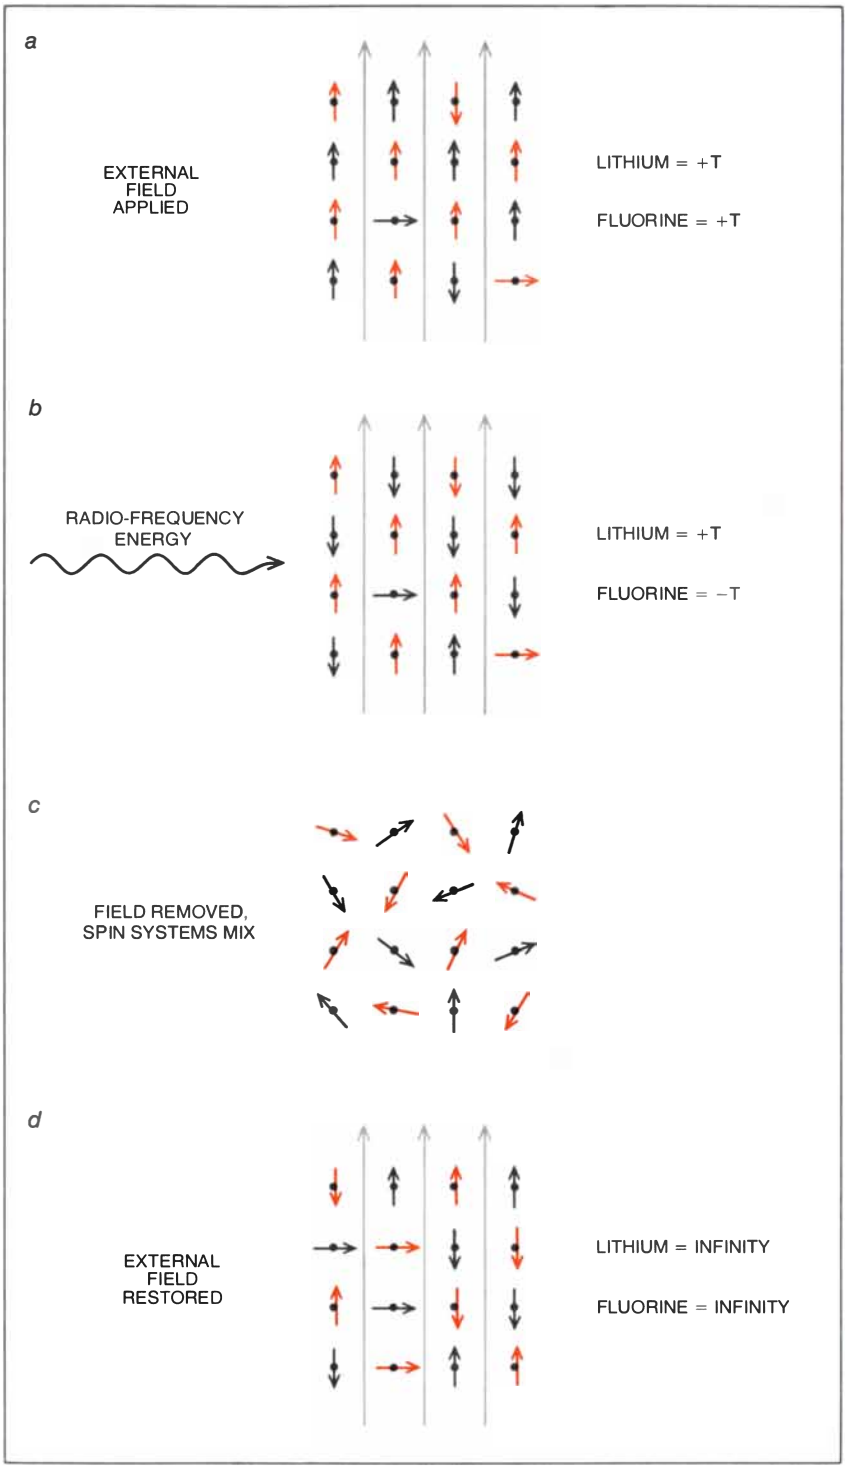

SPIN CALORIMETRY measures the equilibrium temperature reached when two spin systems are "mixed." Each spin system is made up of many identical nuclei; in a crystal of lithium fluoride, for example, there are two spin systems, one consisting of all the lithium nuclei (*color*) and the other of all the fluorine nuclei (*black*). Because the energy levels in these nuclei have different spacings the lithium and fluorine spins do not interact as long as a magnetic field is maintained. When the field is first imposed (*a*), both systems reach thermal equilibrium with the lattice at some positive temperature $+T$, meaning that there is an excess of spins parallel to the field. A burst of radio-frequency energy of the appropriate frequency causes all the fluorine nuclei to reverse their orientation (*b*), but it has no effect on the lithium nuclei. Hence the lithium system retains the temperature $+T$, whereas the temperature of the fluorine system is converted to $-T$. When the magnetic field is removed, all the energy levels collapse (*c*) and the two systems can interact freely. If the field is restored (*d*), each system is found to have a polarization of zero and an infinite spin temperature, the value midway between $+T$ and $-T$.

distribution observed at lower temperatures, so that there are more nuclei in the higher-energy levels than in the lower-energy ones. The state of highest possible energy would have all the nuclei in the antiparallel orientation.

The crucial difference between the lattice temperature and the spin temperature is that the number of energy levels for the vibratory motions of atoms is infinite, but the number of orientational levels is always finite. In the lattice system there is no intrinsic limit to the energy of an individual nucleus, and adding energy merely promotes some nuclei to ever higher levels. In the spin system no nucleus can have an energy higher than that of the antiparallel orientation, and adding energy can only drive more and more of the nuclei into that configuration.

What is the temperature of a spin system in which the majority of the nuclei are in the highest energy level? The measurement can again be derived from Curie's law. At temperatures below infinity the preponderance of nuclei in the lowest-energy, parallel orientation gave rise to a magnetic polarization, M, parallel to the external field, H. The temperature was then found to be proportional to H/M. At temperatures higher than infinity the majority of the nuclei would be lined up antiparallel to the field. It follows that the magnetic polarization would also be opposed to the external field and algebraically would have the value $-M$. Applying Curie's law, the temperature is proportional to $H/-M$, and it must have a negative value.

A better understanding of why negative numbers must be included in the absolute temperature scale can be gained by returning to the fundamental definition of temperature. For a lattice system any increase in energy is accompanied by an increase in entropy. If entropy is related to the difficulty of deciding what state a nucleus is in, then this "monotonic" increase is readily accounted for: with every increase in energy there are more energy levels among which to choose. The temperature can be related to the amount of energy that must be added to the system to yield a given increase in entropy; since both quantities must always have the same sign, the temperature must always be positive.

Up to the point of infinite temperature the relation of energy to entropy is similar in the spin system. Each increase in energy brings a corresponding increase in entropy. At infinite temperature, however, the entropy reaches a maximum: it cannot be more difficult to predict the orientation of a nucleus than it is when there is an equal likelihood of its being in any one of the possible orientations. Further increases in energy now cause a decrease in entropy as the system becomes more highly ordered.

When all the nuclei are opposed to the external field, the system has reached a temperature of "minus zero" and the entropy is at a minimum; every nucleus is certain to be in the highest energy state.

The temperature of the spin system becomes negative as soon as the entropy begins to decrease with increasing energy. Mathematically the temperature curve is said to approach an asymptote at the point of maximum entropy. If the distribution of spin states is nearly uniform but slightly biased toward the parallel configuration, then the temperature is numerically large and positive; with a slight bias in the other direction the temperature abruptly assumes a large negative value. Toward either extreme the temperature approaches zero, but at maximum energy the approach is from the negative side. A more formal analysis of these relations has been carried out by Norman F. Ramsey of Harvard University and others.

Although negative absolute temperatures may be intellectually troubling, they are quite easy to create in the laboratory. In an early demonstration Edward M. Purcell and Robert V. Pound of Harvard converted a positive spin temperature to the numerically equal negative one by suddenly reversing the direction of the applied magnetic field. All the nuclei that were originally parallel to the field then became antiparallel, and vice versa. The energy of the spin system increased (since more of the nuclei were in the high-energy antiparallel state) but the entropy of the system remained constant.

Today the most convenient method for producing negative absolute temperatures makes use of the nuclear-magnetic-resonance apparatus that is also employed to measure the temperatures. The specimen is initially polarized by a magnetic field; then it is irradiated with a short but intense burst of radio-frequency energy. If the appropriate frequency is chosen, the population of occupied states can be inverted almost instantaneously. All the nuclei originally parallel to the field reverse their direction to become antiparallel and the initially antiparallel ones become parallel; the number of nuclei with a perpendicular orientation is not changed.

Through nuclear magnetic resonance an infinite spin temperature can be generated with equal ease. The radio-frequency field is merely left on for a few seconds, so that each nucleus flips many times. The result is that the nuclei distribute themselves randomly, and hence for a large number of nuclei uniformly, in all the possible spin states.

If spin temperatures extend from zero to positive infinity and then from negative infinity back to zero, the temperatures should obey a peculiar law of addition. In particular, two temperatures of equal magnitude but opposite sign should add to yield an infinite temperature. That prediction can be tested by an experiment in "spin calorimetry," analogous to the ordinary calorimetry experiments in which two liquids at different temperatures are mixed and the equilibrium temperature of the product is measured.

Abragam and I conducted such experiments in which the two "substances" to be mixed were the separate spin systems of lithium and fluorine in a crystal of lithium fluoride. With nuclear-magnetic-resonance apparatus each of the spin systems could be separately prepared at a positive temperature, $+T$, a negative temperature $-T$ or an infinite temperature. Of course, the two kinds of nuclei are already intimately mixed in the crystal, but as long as the specimen remains in a magnetic field the two spin systems do not interfere with each other. They are effectively isolated because at any given external field the spacing between the energy levels is different in lithium and fluorine. When the crystal is taken out of the field, however, the energy levels of both nuclear species collapse; the two systems can then freely communicate, and they quickly reach a common equilibrium temperature. The crystal can then be replaced in the magnetic field and the polarization of each system can be measured.

A series of such calorimetry measurements showed close agreement with theory. When the lithium system at a temperature of $+T$ was mixed with a fluorine system at $-T$, both systems reached a common temperature near infinity. When one system was at a large negative or positive temperature and the other was at infinite temperature, both reached a common intermediate temperature. Discrepancies from the ideal results could be accounted for by differences in the specific heat of the two spin systems and by relaxation of the spin systems toward thermal equilibrium during the experiment.

The introduction of negative absolute temperatures may seem to accomplish little more than to make obscure and confusing one of the properties of matter that had seemed most simple and most familiar. The experimental results run counter to a strong sentiment that temperature, like volume, is something intrinsically positive. It might even be supposed that the temperature scale was determined too early in the history of science and that some revision of it might encompass all temperatures within a single range of positive numbers. That is not possible; two infinite ranges of temperature are needed. The scale cannot be simplified because the negative temperatures it includes represent real states of any system with a finite number of energy levels.

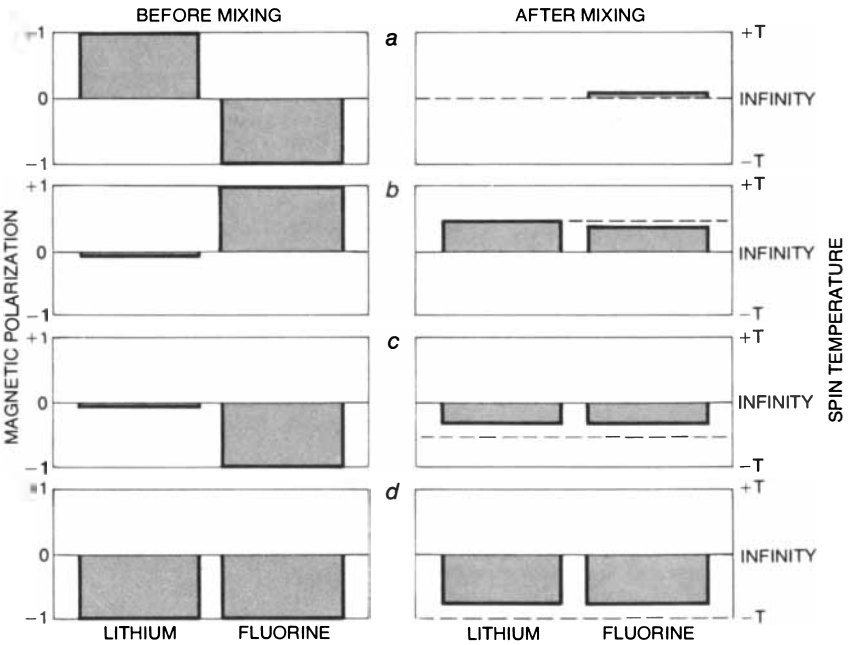

RESULTS OF SPIN-CALORIMETRY EXPERIMENTS show that spin temperatures behave much like ordinary temperatures when two systems are mixed. If the lithium and fluorine systems are oppositely polarized before mixing (*a*), as in the example on the opposite page, the final polarization is near zero and the final temperature is near infinity. Mixing one strongly polarized system with one at an infinite temperature (*b, c*) yields an intermediate temperature for both systems. If the lithium and fluorine nuclei have the same polarization (*d*), no change is expected as a result of mixing. Discrepancies between theoretical expectations (*broken lines*) and measured values (*gray bars*) can be attributed to differences in the heat capacity of the two spin systems and to the thermal relaxation of both systems toward the lattice temperature.

HP measurement and computer advances

HP-33E Scientific
49 lines of program memory
8 conditionals
3 levels of subroutines
8 addressable storage registers
Go To, Single Step, and Back Step commands for easy editing

HP-38E Advanced Financial
Powerful cash flow analysis
Easy time and money calculations
Up to 99 lines of program memory
2000-year calendar
Advanced statistics capabilities

HP's two new Series E programmables are without equal and within your reach.

If you are an aspiring professional, it makes sense to step up to HP programmability. On the basis of power, capability, and ease of use, the new HP-33E scientific programmable at $100* and the new HP-38E advanced financial programmable at $120* are surprisingly affordable.

For student and educator alike, an HP Series E programmable is the logical choice. Consider: Each has a specialized set of advanced, preprogrammed functions, plus the convenience of programming.

Consider further: Series E calculators make excellence available at a surprisingly low price. Besides the two programmables, Series E includes three preprogrammed models: the HP-31E and HP-32E scientific and advanced scientific calculators, at $60* and $80* respectively, and the HP-37E business management calculator at $75*.

Like their predecessors, they have no "equal." The key to HP's simplified computer logic is ENTER↑, making possible fast, efficient problem solving with no =,

literally and figuratively. And, like their predecessors, Series E calculators have the feel and reliability born of quality design and construction.

Programming the HP-33E and 37E is easy. Just switch to the Program mode and key in the series of operations you'd normally use to solve a problem. Rather than executing the keystrokes, the calculator stores them in its memory. Switch to the Run mode, key in your data, and press the run/stop key. You can repeat the program with different data as many times as you like. There's no complicated programming language to learn, and no elaborate start-up procedures to memorize.

Series E offers some other helpful advances: Display messages tell you when you've made a procedural error, and what kind. And if you suspect the calculator is at fault, it will perform a diagnostic self-check at the touch of a key. A larger, brighter LED display automatically separates thousands with commas for easy reading. To help you get started, we've prepared modular handbooks for owners, free with the calculator, that let you skip what you already know, and give you what information you *do* need, clearly and simply.

HP has a 16 mm film, *Calculators, Classrooms... Careers*, available on a free-loan basis to teachers. The 14-minute film examines the development of advanced personal calculators and how they help people in a variety of educational and work environments. Narrator is Wally Schirra, astronaut and business executive. To obtain the film for classroom showing, write to Rick Baker, H.P., 1000 N.E. Circle Boulevard, Corvallis, OR 97330.

extend your possibilities.

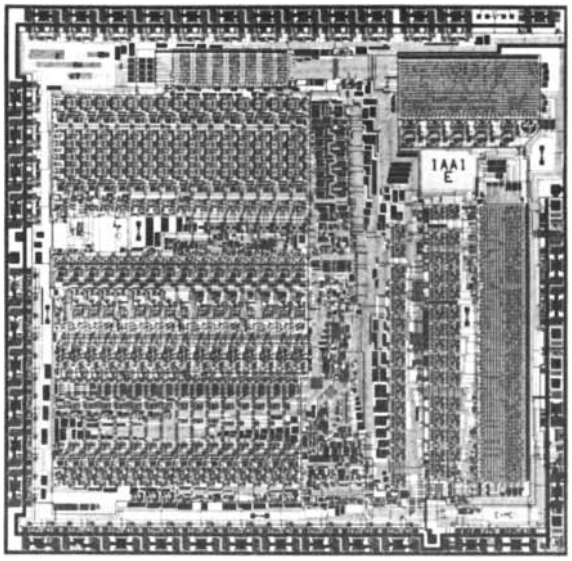

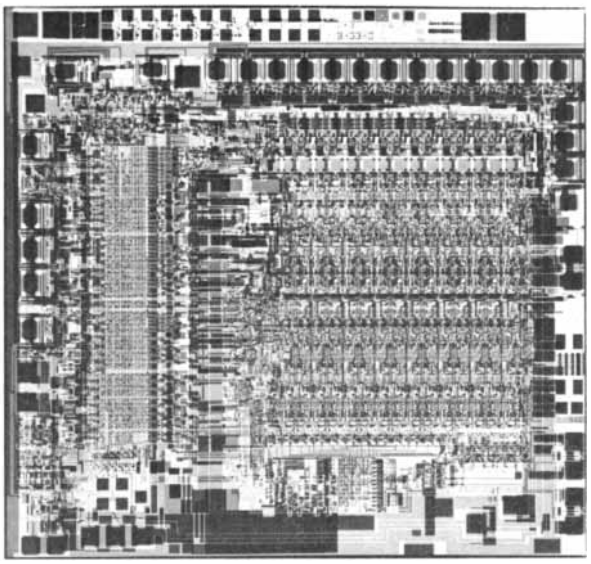

These photographic enlargements of two HP microprocessor chips—one with NMOS II large-scale integrated circuits at right and one with CMOS/SOS at left—illustrate a major reason for the rapid growth of digital circuits: increasing functional density at a decreasing cost. These chips contain as many as 10,000 transistors and are about this big [■].

How to operate more confidently in an increasingly digital world.

While analog circuits generally convey data by a continuum of voltages, the digital world of microprocessors and LSI circuits relies on two voltage levels, conveying information by switching between them. Because of this fundamental difference, digital circuits pose a special set of problems for both designers and troubleshooters. And HP has some excellent solutions.

As the solid state revolution increases their functional density at a decreasing cost, digital circuits are infiltrating areas that previously relied on analog technology, and are cropping up in an endless procession of products that include microwave ovens, automobiles, home computers, process control (analog from its inception), telecommunications, and a range of industrial controls.

In short, the technical, industrial, and consumer worlds are turning digital. While this greatly expands possibilities, it also introduces complications—most notably in the design, debugging, software development, and diagnosis of these functional leviathans in minnow-sized packages. While it may not be immediately obvious, this digital data domain requires special measurement tools. And here Hewlett-Packard is making significant strides, continually introducing new measurement approaches and instruments designed to make the digital revolution manageable, and thus enable technical people to operate more confidently.

If you are wrestling with any aspect of digital circuitry, you will be interested in reading a recent issue of the *Hewlett-Packard Journal*, which explores HP's growing family of digital logic analysis test instruments at some length, and offers a few viewpoints of possible interest. Just mail the coupon, and we will be pleased to send you a copy.

HEWLETT hp PACKARD

1503 Page Mill Road, Palo Alto, California 94304

For assistance call: Washington (301) 948-6370, Chicago (312) 255-9800, Atlanta (404) 955-1500, Los Angeles (213) 970-7500

Mail to: Hewlett-Packard, 1503 Page Mill Road, Palo Alto, CA 94304.
Please send me further information on

( ) HP-33E scientific calculator
( ) HP-38E advanced financial calculator
( ) HP digital logic analysis

Name _____
Company _____
Address _____
City _____ State _____ Zip _____

*Domestic U.S. prices only.

00848