Established 1845 **SCIENTIFIC AMERICAN** June 1971 Volume 224 Number 6

Fusion by Laser

Experiments indicate that energy-releasing fusion reactions can be initiated and to some extent controlled without a confining magnetic field by focusing a powerful laser pulse on a frozen pellet of fuel

by Moshe J. Lubin and Arthur P. Fraas

The rapid evolution of high-power pulsed lasers over the past decade has made available to workers in numerous fields a new tool whose potential usefulness has just begun to be explored. Prominent among the applications that have been proposed to date are a variety of schemes aimed at exploiting the fact that the focused light from such a laser is capable of heating a small amount of matter to extremely high temperatures: in some cases to more than 50 million degrees Kelvin! To the chemist, for example, such a capability means that chemical reactions can be initiated in less than a nanosecond (a billionth of a second) and can be studied under conditions previously considered unattainable.

To the plasma physicist the unprecedented heating capability of high-power pulsed lasers is particularly exciting. In principle it opens a promising new avenue of attack on the long-standing problem of how to produce the ultrahigh-temperature plasma needed to maintain thermonuclear, or fusion, reactions under controlled conditions. Recent experiments have shown that energy-releasing fusion reactions that may ultimately lead to the production of useful electrical power can be initiated and to some extent controlled within an "inertially confined" plasma obtained by focusing a powerful laser pulse on a dense frozen pellet consisting of a mixture of the heavy hydrogen isotopes deuterium (H^2) and tritium (H^3). In this article we shall review the current status of research on the application of high-power lasers to the generation of fusion power.

The case for fusion power and the progress that has been made over the past 20 years or so by nonlaser technologies toward achieving this goal have been discussed at length in several earlier articles in this magazine, the most recent of which appeared only a few months ago [see "The Prospects of Fusion Power," by William C. Gough and Bernard J. Eastlund; SCIENTIFIC AMERICAN, February]. Suffice it to say here that controlled fusion reactions involving light isotopes such as deuterium and tritium represent a potentially inexhaustible source of inexpensive, efficient, "clean" energy for supplying all mankind's future power requirements.

The fusion-power option, however, has not yet been shown to be technologically feasible. The difficulty is twofold. First, the task of heating the light-element fuel to a sufficiently high temperature at the required density so that it will begin to "burn" slowly has proved to be harder than was anticipated. The temperature required to initiate burning varies between 50 and 100 million degrees K., depending on the fuel. Second, once the fuel is ignited, the hot, dense, gaseous mixture of ions must be held together for a long enough time so that more energy is liberated through burning than is invested in the ignition process. Attempts at confinement by means of intense magnetic fields of various geometries have so far not yielded a practical solution. In spite of the difficulties encountered in heating a gaseous mixture to such temperatures and confining the ionized gas for large fractions of a second, however, the potential return on the investment of time and money in research on this problem is so great that all the major nations of the world have substantial controlled-fusion research programs.

Before describing how high-powered lasers can play a role in the solution of this problem it will be instructive to review briefly (1) the current state of the art in the development of high-power pulsed lasers, (2) the mechanism by which the energy from a laser beam interacts with a dense medium and (3) the requirements demanded of lasers in initiating controlled fusion reactions.

A laser is a device for generating or amplifying a beam of light whose waves are both monochromatic (all the same wavelength) and coherent (all in step). The light beam emitted by a laser can be made almost perfectly parallel, its divergence angle being theoretically limited only by diffraction effects. In principle such a beam can be focused to a spot with a diameter of only a few hundred-millionths of an inch. The laser produces its coherent light beam through the interaction of electromagnetic radiation with the laser medium. By elevating more atoms to an upper energy level than exist at a lower level the absorption of excitation radiation produces an "inverted" atomic population in the laser. The energy stored in the inverted population is then available to amplify a propagating light wave at a particular frequency. This frequency of emission is generally different from the frequency band over which the laser medium has absorbed excitation radiation. For example, a ruby laser absorbs "pumping" light in the green region of the optical spectrum but emits coherent light in the red region.

The small light pulse to be amplified is produced by a laser oscillator. For a laser to work as an oscillator some feedback mechanism must be provided. This is accomplished by the use of mirrors on both ends of the laser medium to provide a means by which waves at the emission frequency may travel back and forth through the laser. In this way oscillations are built up at the emission frequency by feeding on the inverted population until the stored energy in the laser medium is depleted, at which time emission ceases.

The first laser action in a solid material was observed in a ruby rod in 1960 by T. H. Maiman, who was then working at the Hughes Aircraft Company. Numerous other solid-state laser materials were discovered in rapid succession. The most important of these is neodymium-doped barium crown glass, in which laser action was first demonstrated in 1961 by Elias Snitzer of the American Optical Company. Whereas ruby is an expensive crystal to grow in any reasonable size, barium crown glass is comparatively inexpensive and can be manufactured in large sizes and quantities. Size becomes important in dealing with ultrahigh-power laser systems.

The end surface of the laser rod from which the coherent radiation is emitted can withstand only a limited amount of energy per square centimeter. Generally it is felt that impurities in the laser material contribute to significant absorption of the laser light, which ultimately leads to fracture of the laser medium. For example, the energy-handling capability of neodymium glass ranges between 10 and 100 joules per square centimeter, depending on the duration of the laser pulse.

At present solid-state ruby and neodymium-glass laser materials are used to obtain the highest peak-power output in pulsed operation. This is attributable in large part to the higher density of atoms available in solids to participate in the laser action. Liquids that sustain laser action hold promise for the future in high-power pulsed applications for this same reason; in addition they have the advantage of not being vulnerable to permanent structural damage. Unfortunately liquid lasers in the form of dye solutions or liquid glass are still in the initial stages of development.

Solid-state laser systems available at present are capable of delivering energies in excess of 1,000 joules in less than a nanosecond. This coherent energy flux is in the form of a highly collimated

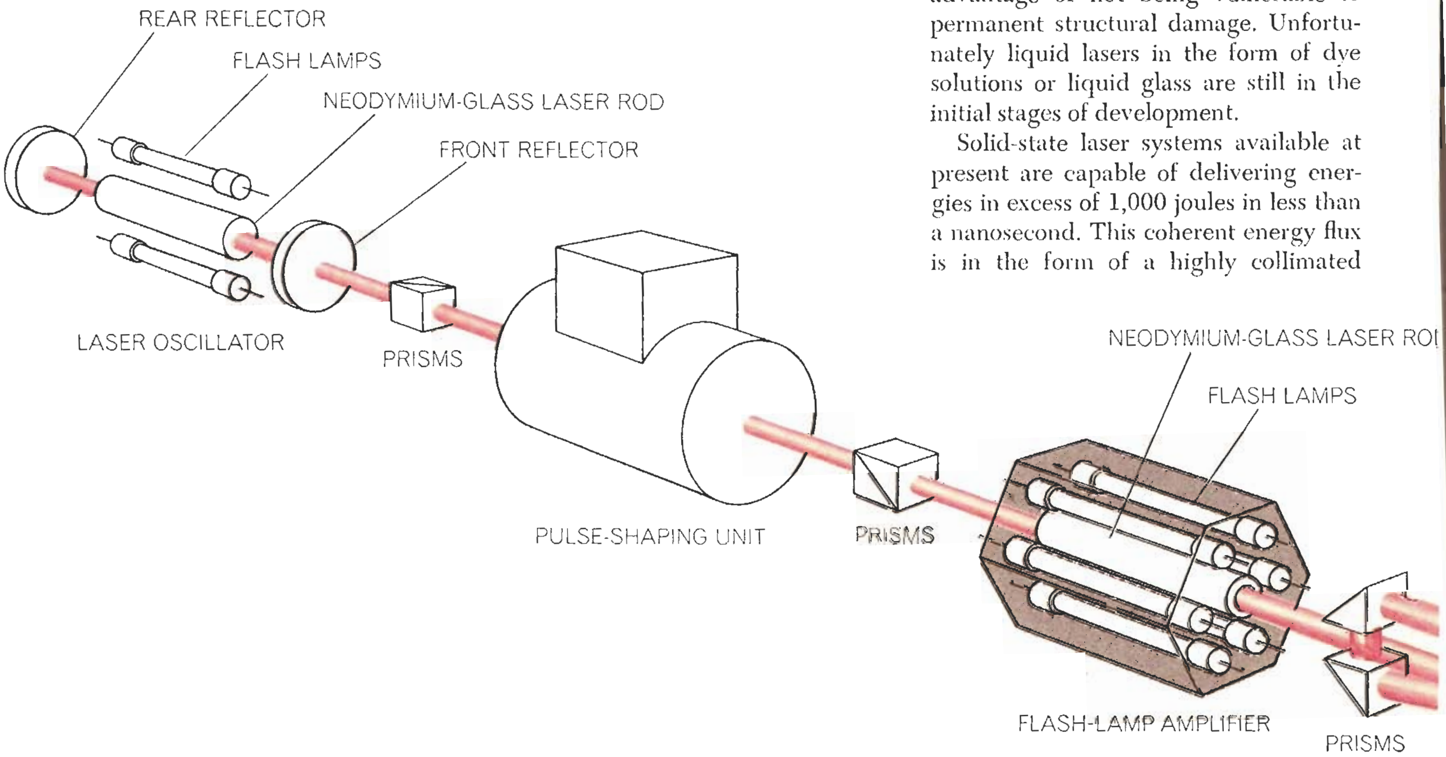

TYPICAL HIGH-POWER LASER SYSTEM shown on these two pages was designed for use in a plasma-heating experiment. The system consists of a string of neodymium-doped barium-crown-glass components, beginning with a laser oscillator (*extreme left*), in which a low-energy laser pulse is formed. The pulse is next "tailored," or shaped both in time and space to suit a particular application, before being amplified by additional laser rods and finally by a large disk amplifier. A suitable combination of prisms sends the pulse through the disk amplifier three times. The resulting beam, some 15 centimeters in diameter, is finally focused on a deuterium droplet inside a vacuum chamber (*extreme right*) to produce a hot plasma, or ionized gas, in which fusion reactions can take place. The fusion reactions are studied by means of the energetic neutrons they emit. The peak power currently available from such a laser system before focusing exceeds 10^{13} watts; when it is focused, it is possible to achieve power densities greater than 10^{17} watts per square centimeter.

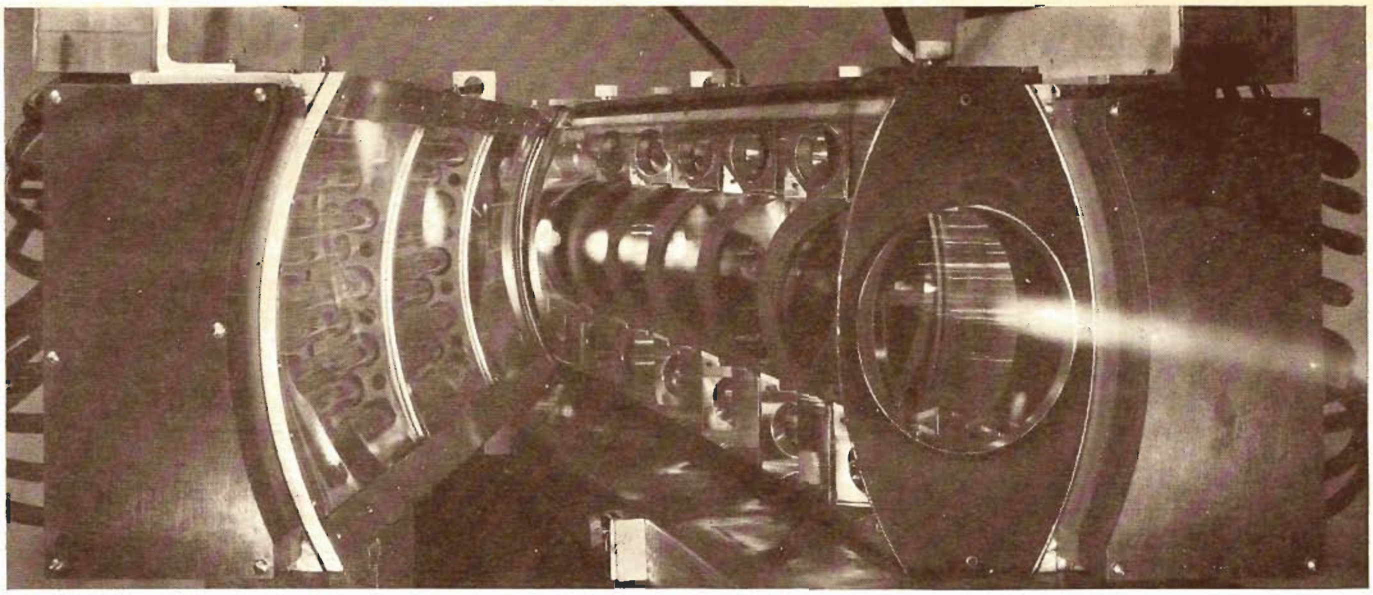

DISK AMPLIFIER is capable of emitting an output pulse with an energy of more than 1,000 joules through an aperture 15 centimeters in diameter. The neodymium-glass laser disks are mounted at an angle to the incoming laser beam in order to minimize reflection at the same time that they are excited efficiently on their faces by the flash lamps located in the reflector tanks on each side of the disks. A hundred such amplifiers arranged in parallel could in principle achieve an energy output greater than 100,000 joules.

beam with a total divergence of less than a hundredth of an angular degree. Such a beam of light starting out six inches in diameter would be less than 18 inches in diameter after traveling a mile.

In the case of a ruby laser the wavelength of the emitted laser radiation is .6943 micron. This wavelength falls in the red region of the spectrum and is therefore visible. The neodymium glass used in high-power lasers emits radiation at 1.06 microns, which is in the infrared portion of the spectrum and is not directly visible to the human eye.

A typical high-power laser system [*see illustration below*] consists of an oscillator in which a laser pulse, appropriately shaped in space and time, is formed at low energy. The pulse is an envelope containing a burst of electromagnetic radiation oscillating at the

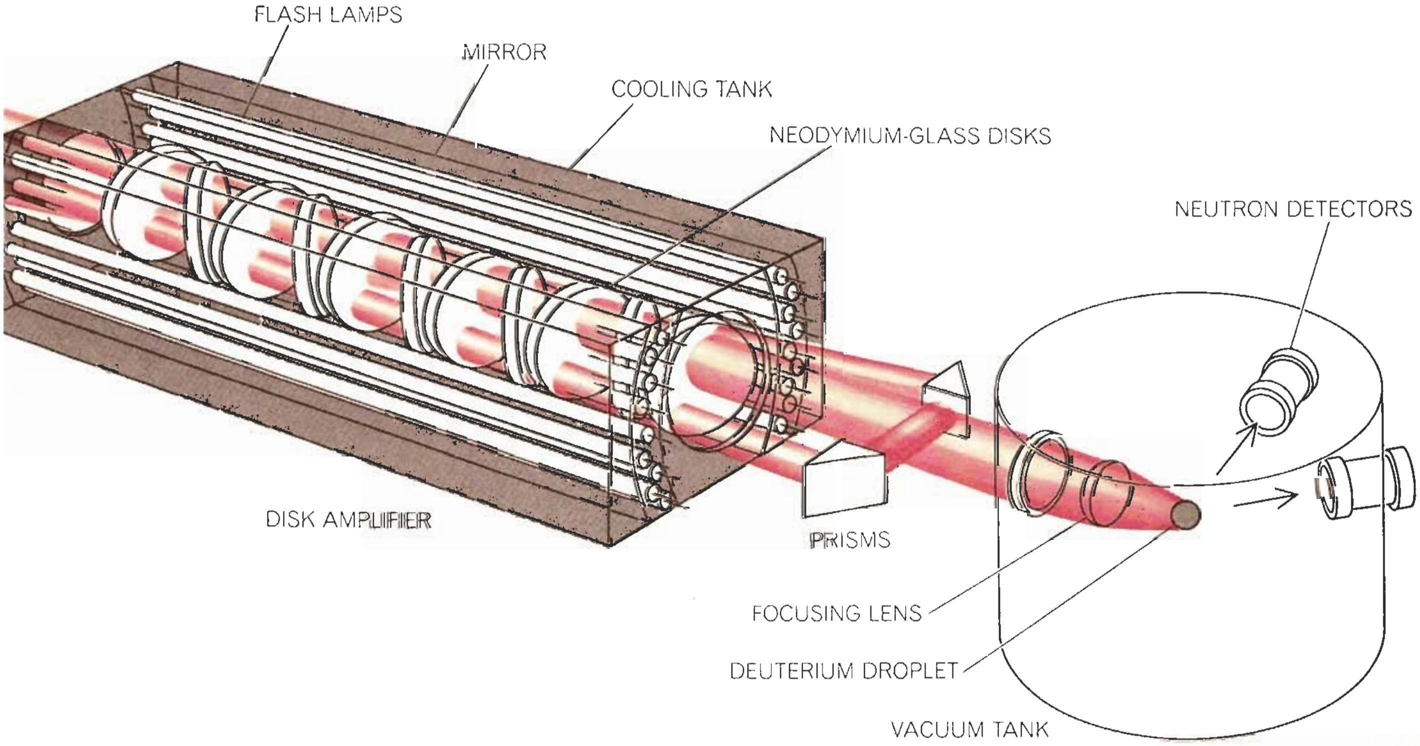

laser frequency. "Tailoring" a pulse to suit a particular application can be accomplished in a number of ways. Pulses lasting on the order of a nanosecond or longer can be reliably produced simply by chopping out a suitable portion of a fatter pulse with a fast electro-optical shutter. The shutter, activated by a high-voltage electrical pulse, is synchronized with the arrival of the longer pulse by monitoring the intensity of the laser radiation as a function of time. Such synchronization requires high-voltage switching on the time scale of a nanosecond.

Producing narrower pulses by this technique is unreliable. Instead a technique known as "mode-locking" is commonly employed [*see illustration below*]. As the energy in the laser cavity is built up and then decays, a string of such mode-locked pulses emerges from the partially transmitting mirror in the front of the cavity. An electro-optical shutter can be used here too, but only to "gate out," or select, a laser pulse with a duration far shorter than the time the shutter is open. One pulse or more can be gated out in this way. Here the duration of the laser pulse is determined solely by the components in the laser cavity. Pulses varying in duration from a billionth to a trillionth of a second have been produced in this fashion.

Once the appropriately shaped laser pulse emerges from the laser, it is directed into a string of laser amplifiers, each amplifier adding energy to the original pulse. The amplifiers usually have successively larger apertures, so that the lasing material can survive the increased energy without damage. The amplifying process itself can be called on to change the shape of the input pulse. If the intensity of the incoming pulse is large enough, most of the amplifier energy is deposited in the pulse's leading edge, thereby narrowing the width of the pulse.

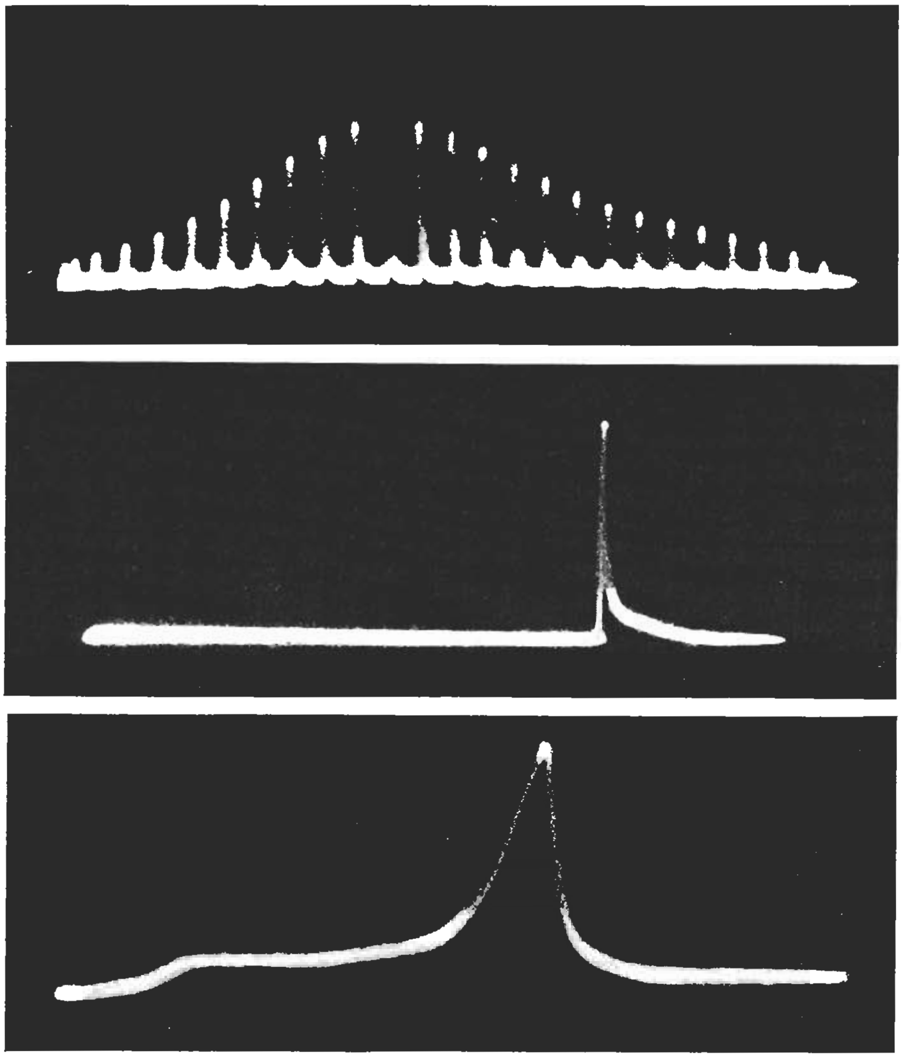

TAILORING OF A LASER PULSE is illustrated by this sequence of oscilloscope traces. A train of very brief "mode-locked" pulses is first produced in the oscillator section of the laser system (*top*). A single pulse is then extracted from the train (it is missing in the top trace) and is amplified (*middle*). The width of the pulse may vary between a billionth (10^{-9}) and a trillionth (10^{-12}) of a second. A low-power leading edge is often added to such a pulse (*bottom*) in order to vaporize the fuel droplet prior to heating by the main pulse.

The atoms in a neodymium-glass laser rod are excited by absorbing light from external flash lamps. Since the light enters the rod from the outside, more atoms near the outer edges of the rod are elevated to an excited energy state than are elevated in the interior. This undesirable situation becomes intolerable for rods more than two inches in diameter. In addition large rods are difficult to cool between pulses. One solution to the problem of exciting large-aperture amplifier stages is the disk amplifier [*see top illustration on preceding page*]. This type of amplifier consists of large neodymium-glass disks that are tilted with respect to the incoming laser beam in order to minimize reflection losses. At the same time the disks have a large projected area of minimal thickness facing the excitation flash lamps, which leads to uniform excitation. One such amplifier with an aperture of six inches is now capable of reaching a power output of 1,000 joules; 100 similar amplifiers could be arranged in parallel to achieve outputs of 100,000 joules.

Numerous available sources are capable of delivering such bursts of energy in a short time. None, however, equals the laser in the ease with which the output can be concentrated in a small volume. It is this ability to focus high-power radiation using optical elements such as lenses and mirrors that makes the laser uniquely attractive for the purpose of initiating fusion. The peak power in such a beam before focusing is more than 10^{13} watts, and when the beam is focused, it is possible to achieve power densities greater than 10^{17} watts per square centimeter. The focused energy results in a local electric field in the immediate vicinity of the focus on the order of 10^{10} volts per centimeter! It is difficult to grasp the magnitude of such enormous electric-field intensities; an example of their strength is that a free electron accelerated in a field of 10^{10} volts would reach a velocity comparable to the velocity of light in a fraction of an optical wavelength.

It is well known that nuclear-fusion fuels such as deuterium and tritium liberate more energy per nucleon (neutron or proton) than nuclear-fission fuels. Unfortunately these hydrogen isotopes only burn efficiently at high temperature. Below 10 million degrees K. the burning is too slow to be useful in an energy-producing process here on the earth. The sun can burn hydrogen efficiently at lower temperatures because of its large volume. In the limited size of a laboratory plasma one must consider burning isotopes at temperatures greater than 50 million degrees K.

The traditional approach to controlled fusion has been to begin by producing a plasma that is dilute and rather cool. The plasma is next trapped in a confining magnetic field. One then proceeds to heat the plasma to a higher temperature while holding it in place with a magnetic field that is made proportionately stronger. Our present experience and understanding only allow us to heat the plasma rather gently so as not to upset the delicate balance between the hot ionized gas and the magnetic field. Although our knowledge of the forces and conditions controlling this balance has increased greatly over the past 10 years, we are not yet sure of the proper magnetic-field configuration required to contain the hot plasma long enough for significant burning to occur.

When high-power pulsed lasers began to emerge as working systems, one of their first applications in the controlled-fusion program was the production of extremely clean, well-behaved plasmas to fill existing magnetic-confinement devices [*see top illustration at right*]. The question that was asked at the time was: "Can a plasma produced by the laser heating of a small droplet of thermonuclear fuel in a magnetic field serve as the plasma required for nearly steady-state, continuous fusion reactions?"

More recently, as the energy-producing capability of high-power pulsed laser systems has increased, a second question has arisen: "Is it possible to utilize the intense burst of focused coherent radiation from the laser to raise the temperature of a suitable deuterium-tritium target high enough to produce significant burning *without* the necessity of a confining magnetic field?"

To answer these questions one must first examine how high-frequency electromagnetic radiation interacts with a dense plasma made up of electrons and positive ions. The oscillating electric field in a light beam may be absorbed, reflected, scattered and/or refracted by a material surface. Our prime interest is in the absorption of this radiation and the conversion of the electromagnetic energy of the laser beam to the thermal energy of the fusion fuel. Once a few unbound electrons are produced on the surface of the solid hydrogen isotope, these free electrons rapidly pick up energy from the incident oscillating electric field. Their energy is then trans-

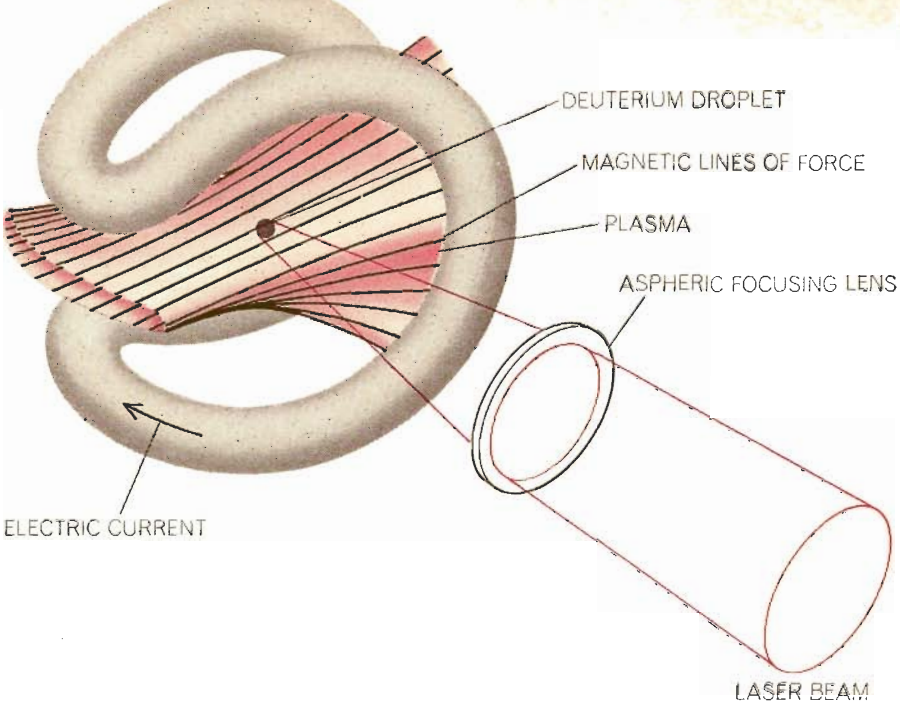

FIRST USE OF HIGH-POWER LASERS in the controlled-fusion program was to produce a plasma to fill existing magnetic-confinement devices such as the one shown in this schematic drawing. An electric current flowing in the "baseball seam" winding induces a magnetic field whose energy density is lowest at the center and increases outward in all directions, thus creating a magnetic "well" in which to confine the plasma, which is generated by laser-heating a fuel droplet at the center. In an experiment carried out at the University of Rochester in 1968 a device of this type was used to produce a confined plasma with a temperature of 10 million degrees Kelvin and a density of 10^{13} ions per cubic centimeter.

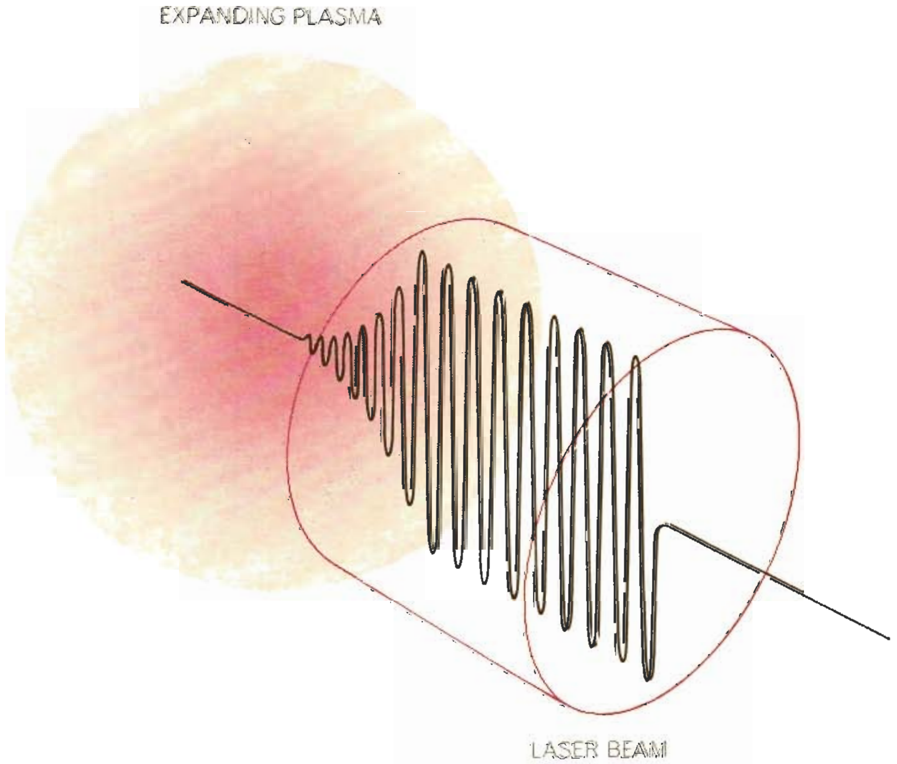

LASER PULSE IS ABSORBED very efficiently by a dense, expanding sphere of plasma, with most of the absorption taking place in a thin surface layer of plasma only about a hundredth of an inch thick. As the energy of the laser beam is absorbed by the plasma, the amplitude of the oscillating electric field (*black curve*) decreases rapidly within this layer.

ferred to the ions in the plasma medium by means of long-range attractive interactions. Thus the heavier ions also “feel” the effect of the laser beam through the mediation of the oscillating electrons. The transfer of this energy from the electrons to the ions can be much slower than the heating of the electrons by the laser beam. In fact, it is doubtful that in a free deuterium-tritium plasma effective transfer can occur on time scales of less than 10^{-11} second. This absorption process is called collisional absorption.

If the incident radiation is to be effectively absorbed, the plasma should be opaque to it. On the other hand, it is desirable that the laser radiation not be wasted by reflection. In more quantitative terms, it is desirable that the incident light wave lose all its energy to the electrons in their first few collisions with neighboring ions. This model enables us to estimate how thick a dense plasma must be to absorb all the incident laser light. When a plasma is surrounded by a vacuum, the density of electrons and ions increases rapidly from zero at the outer surface to a maximum in the bulk of the plasma. Most of the laser radiation is absorbed in this layer of increasing density [*see bottom illustration on preceding page*]. For complete absorption in a dense plasma at thermonuclear temperatures this layer need be only .005 inch thick.

The answer to the first question has been provided by the theoretical and experimental work of laboratories throughout the world. In the U.S. a group at the United Aircraft Research Laboratories under Alan F. Haught is

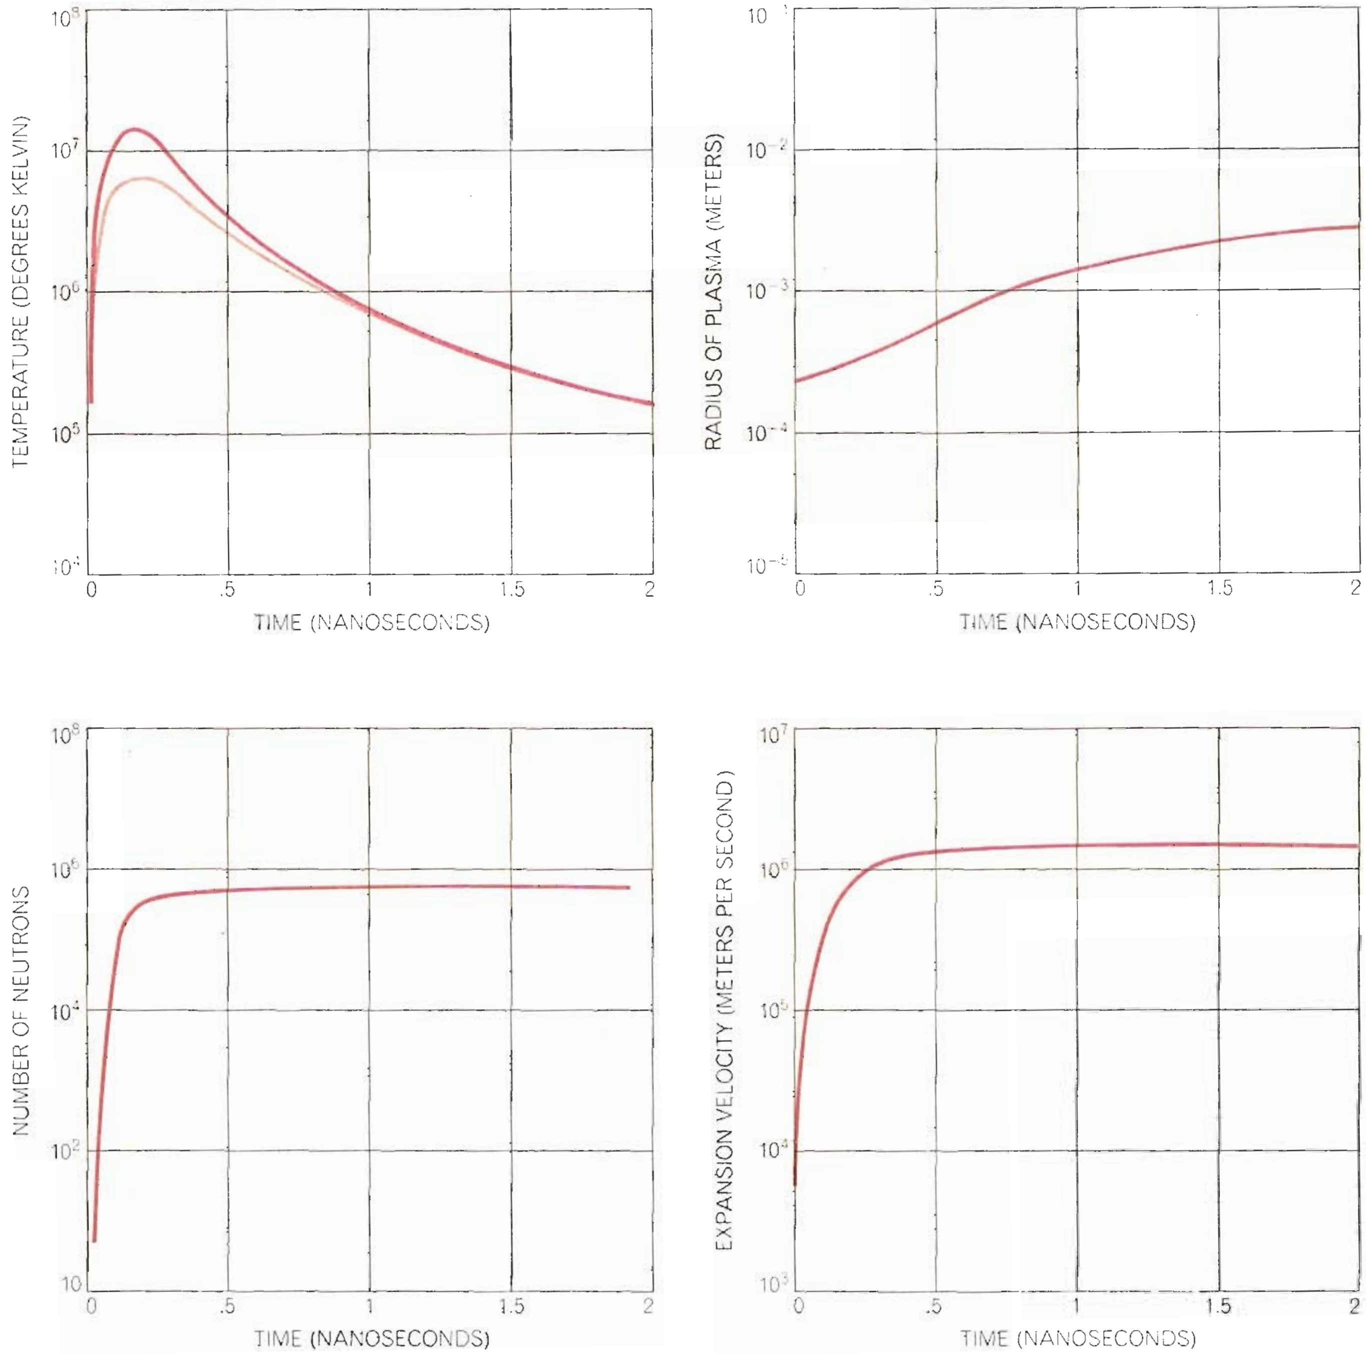

TYPICAL ABSORPTION HISTORY for a laser-heated deuterium fuel pellet is represented by these four graphs. The pellet, originally a tenth of a millimeter in diameter, was heated by a tailored laser pulse that lasted 10^{-10} second and had a peak amplitude of 10^{13} watts. Darker curve in first graph is electron temperature; lighter curve, ion temperature. The incoming pulse was absorbed solely by nonrelativistic processes and hence resulting neutron yield, a measure of success of burning, is a conservative estimate.

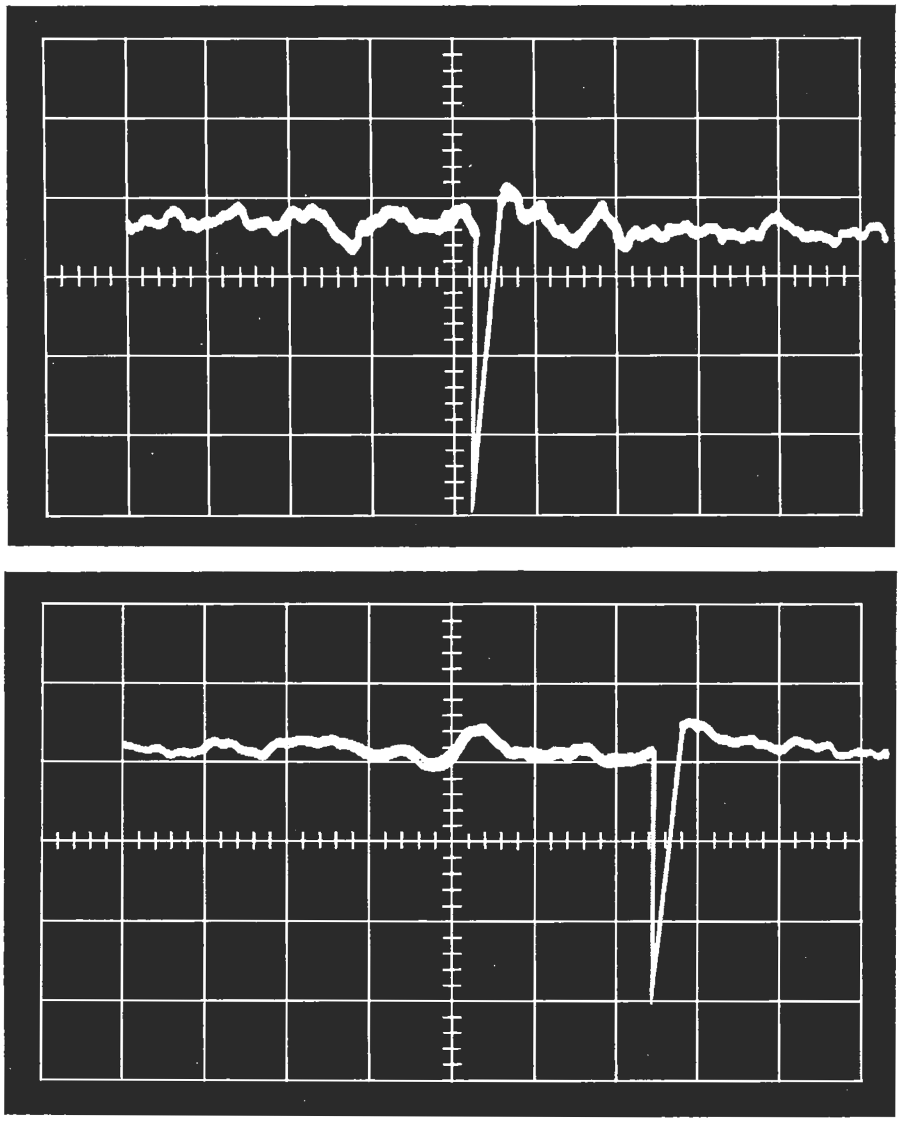

BURST OF NEUTRONS arriving successively at two detectors spaced 225 centimeters apart is signaled by the sharp deflections in these two oscilloscope traces. The difference in the arrival times is a measure of the energy of the neutrons and indicates that these particular neutrons were produced by deuterium-deuterium reactions in the experimental plasma. A similar experiment carried out with a pure hydrogen plasma yielded no neutron output.

demonstrating that plasmas produced by the laser heating of small pellets can be confined in existing magnetic-field geometries over a wide range of temperatures. The resulting quasi-steady plasma, although not hot enough for significant controlled fusion reactions, can then be augmented and heated to higher temperatures by other methods. Thus clean, laser-produced plasmas can serve as the first step in the buildup to controlled fusion conditions. Since the laser plasma is only the catalyst that serves to turn on the steady-state fusion process, the low efficiency of the laser system is of no particular significance.

The second question—regarding the feasibility of laser-heating a small dense plasma to thermonuclear conditions without the necessity of a confining magnetic field—is receiving increased attention. The conditions required to achieve this result can best be understood by taking a step-by-step look at the heating process in a laser-produced plasma. In the typical experimental arrangement pulsed laser radiation is focused on a small fuel pellet of solid deuterium-tritium. The solid fuel pellet is converted to a dense plasma by the leading edge of the incident laser pulse, which is of sufficient duration to vaporize (but not significantly heat) the droplet. Once the droplet is in gaseous form the main heating pulse is applied.

The plasma does not remain stationary as it is heated. Detailed consideration of the absorption process shows that the dynamics of the expanding plasma heated by the incident radiation are vitally important to the time-dependence of the absorption. A typical absorption-time history for an initially solid droplet [*see illustration on opposite page*] shows that there is a time beyond which the plasma has expanded to a point of transparency to the incident laser light. The duration of usable laser energy for heating dense, freely expanding, small plasmas is bounded by two characteristic times. The first is the time it takes for the plasma to expand to the point where it is too dilute to continue to absorb a significant portion of the incident laser energy (a few nanoseconds). The second is the time associated with the electron-ion energy transfer. As we have mentioned, this time scale is on the order of 10^{-11} second. Hence a laser pulse with a duration of 10^{-10} second, bracketed by these two characteristic times, is well suited to this purpose. The laser energy coupled efficiently into the plasma appears initially as thermal energy, and this energy, coupled with the energy from fusion, is ultimately divided up among the fusion reaction products, radiation losses and expansion energy.

One could in principle reduce the amount of energy that goes into expansion by using a strong containing magnetic field. To do so, however, would call for magnetic fields capable of exerting pressures of tens of millions of pounds per square inch. Such fields are not available at present.

Once the density and temperature of the ions are known as a function of time, the thermonuclear yield is easily determined. In particular, rigorous calculations on such typical configurations indicate that one can achieve a ratio of liberated energy to input energy that is greater than 1 at input energies of around 10^5 to 10^6 joules. The assumptions that form the basis of these calculations are currently being studied experimentally at a number of laboratories in the U.S. and abroad.

Neutrons resulting from fusion reactions are ideal "thermometers" to use in evaluating the success of a particular high-power heating experiment. For example, for every deuterium-tritium reaction a neutron with an energy of 14.1 million electron volts (MeV) is produced. The number of such reactions is strongly temperature-dependent. Neutrons produced by the interaction of laser light and a plasma were first reported by N. G. Basov at the Lebedev Institute in Moscow in 1963. The Russian work was followed by high-power laser-plasma neutron experiments at the Limeil Laboratory in France and at several places in the U.S., specifically the Livermore Laboratory of the University of California, the Los Alamos National Labora-

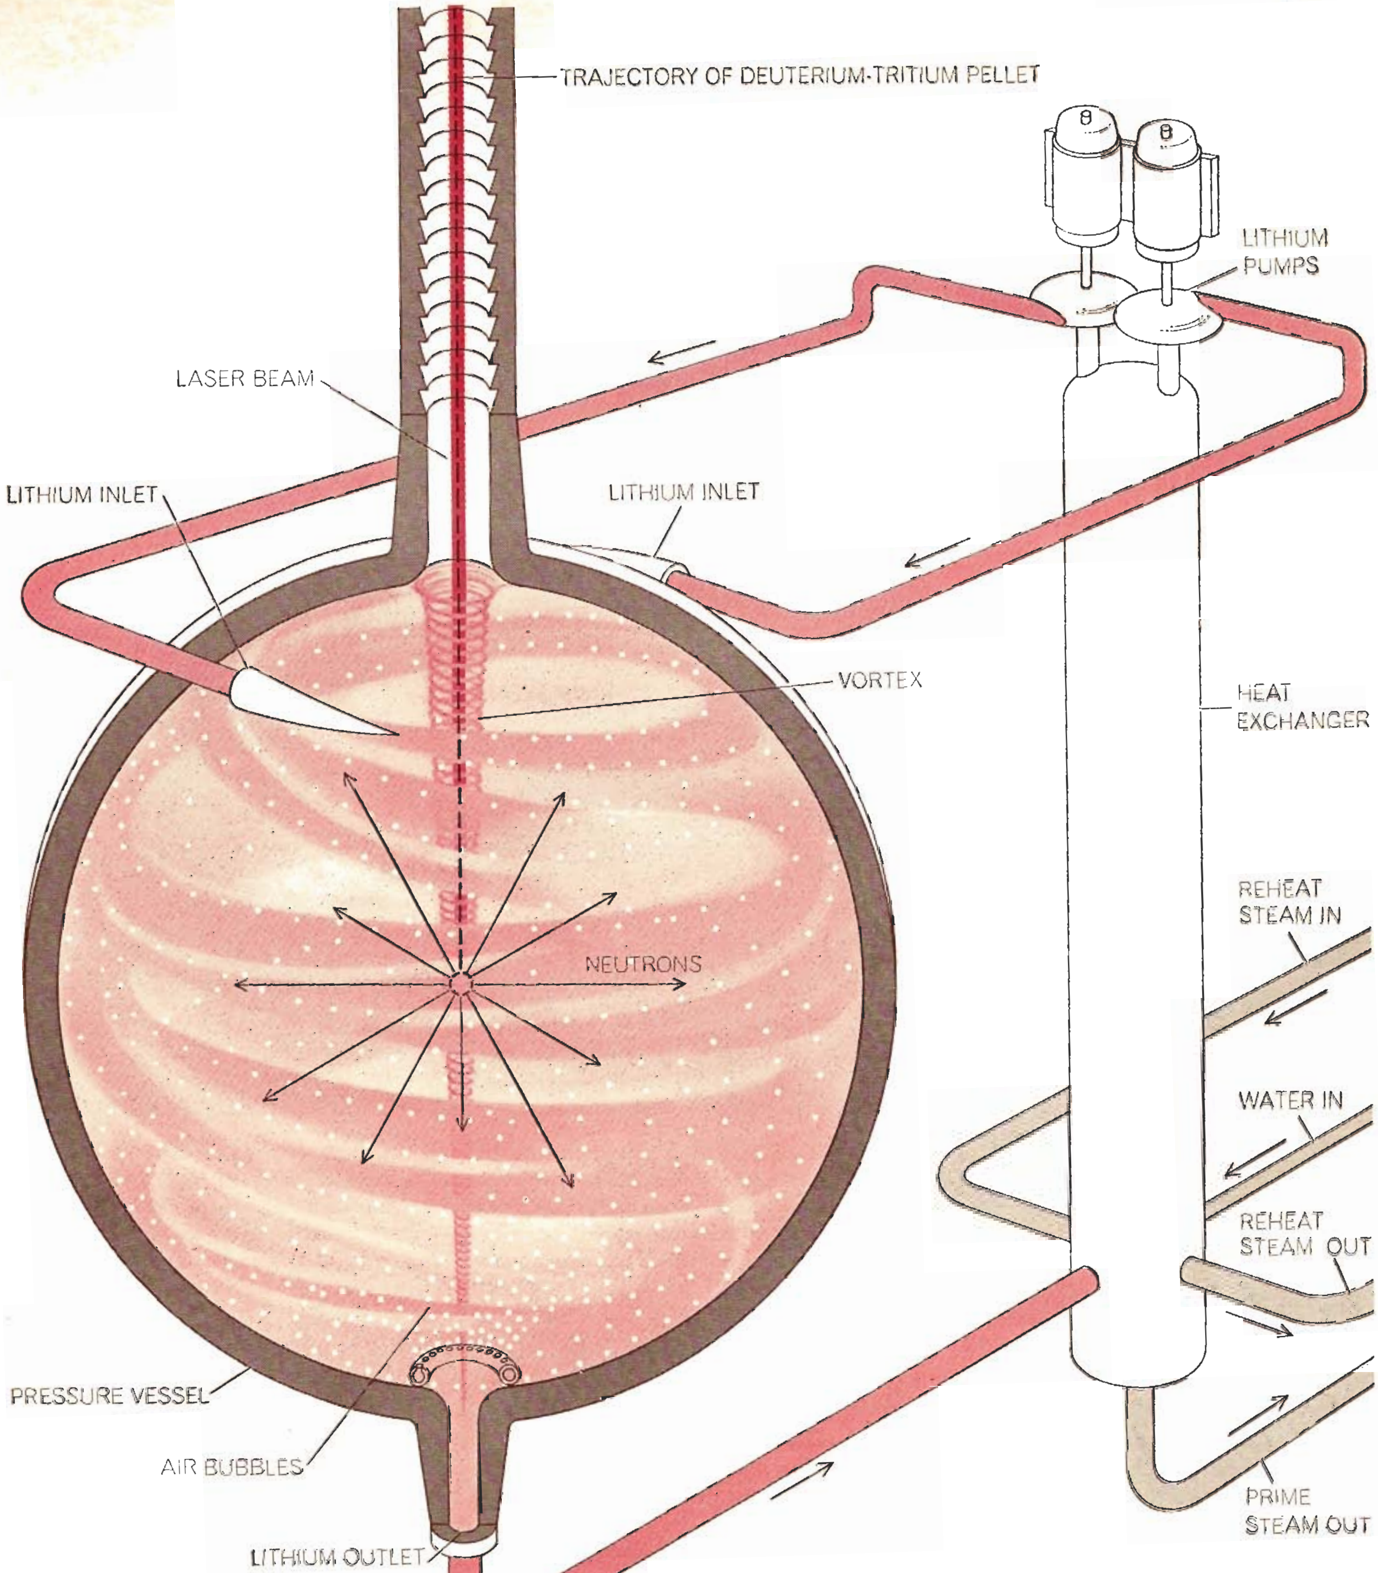

PROPOSED FUSION REACTOR was designed by workers at the Oak Ridge National Laboratory as a method for converting the fusion energy from laser-ignited deuterium-tritium pellets into useful electric power. The fusion energy would be absorbed in a pool of lithium, which would be contained in a cylindrical or spherical pressure vessel some 10 to 15 feet in diameter. A free-standing vortex would be maintained around a vertical axis by swirling the lithium at a sufficiently high velocity. A frozen pellet of deuterium-tritium would be injected into the center of the vortex cavity and would be ignited with a laser pulse when it reached the midplane. The energy deposited in the lithium as heat in the form of energetic neutrons (*black arrows*) would be removed by drawing off lithium from the bottom of the vessel, circulating it through heat exchangers and returning it through centrifugal pumps to tangential nozzles in the perimeter of the vessel. In this particular design the blast waves created by the explosion of the pellet would be attenuated to protect the pressure vessel by introducing shock-absorbing gas bubbles into the lithium pool through a perforated ring in the bottom of the vessel. The pellet-injection and laser systems would be protected by making the port for pellet injection quite long with a wall profile that would break the normal shock wave into many oblique shock waves, thus attenuating the primary blast wave.

tory, the Sandia Corporation and the University of Rochester. Laboratories around the world have been improving on those initial results. Total neutron yields in excess of 10^5 from the deuterium-deuterium reaction have been measured using targets of deuterium with incident laser-beam energies of between 50 and 250 joules. The interpretation of these results indicates a slightly better absorption than simple collisional theory predicts.

Of special interest is the detailed study of the absorption of the incident radiation. There are a number of important contributions to the absorption that may be more significant than the simple collisional process described above and that may hence affect the energy release. Nature, however, seems to favor a freely expanding dense plasma, since the majority of these special absorption processes increase the effectiveness of the incident laser beam.

At present lasers with a properly shaped pulse can deliver 10^3 joules, whereas at least 10^5 joules is needed for significant amounts of laser-produced fusion. This laser energy can be made available using present technology, but as yet no such working system exists. Although the experimental results are few, they appear to justify optimism.

Meaningful use can be made of the energy released from a dense centimeter-size fuel pellet heated by focused laser radiation in a number of different energy-conversion applications. For example, a central station generating electric power by means of controlled thermonuclear reactions from a periodic vaporizing of deuterium-tritium pellets would look quite different from a configuration designed for use in, say, space propulsion. As an illustration let us consider the feasibility of this general scheme for a controlled thermonuclear power plant.

A method for converting the fusion energy from laser-ignited deuterium-tritium pellets into electrical power was evolved at the Oak Ridge National Laboratory early in 1969 in conjunction with fusion-power feasibility studies that have been under way there since 1967. The Oak Ridge approach to fusion entails absorption of the energy from fusion in a pool of lithium, which in turn delivers the energy as heat to a thermodynamic cycle.

The lithium pool would be contained in a cylindrical or spherical pressure vessel 10 to 15 feet in diameter and would be swirled at a sufficiently high velocity to form a free vortex around its vertical axis [*see illustration on opposite page*]. By adjusting the swirl velocity distribution properly it should be possible to obtain a central cavity with a fairly uniform diameter of perhaps five centimeters through the region from the top of the vessel to well below the midplane. A frozen pellet of deuterium-tritium would be injected into the center of the vortex cavity and would be ignited with a laser pulse when it reached the midplane. The energy deposited in the lithium as heat would be removed by drawing off lithium from the bottom of the pressure vessel, circulating it through heat exchangers and returning it to the pressure vessel. The process would be repeated perhaps every 10 seconds. The large thermal inertia in the lithium circuit would act to maintain an essentially constant flow of heat to the thermodynamic cycle.

Drawing off the lithium from the bottom of the pressure vessel would help to stabilize the vortex. After circulating through the heat exchangers the lithium would be returned through pumps to tangential nozzles in the perimeter of the pressure vessel, thereby maintaining the desired vortex.

An obvious objection to this concept is that it would not serve to yield useful amounts of power because the amount of energy one might release in each explosion would have to be kept quite small or the pressure vessel would burst. A relatively simple analysis discloses that this is indeed the case unless some measures are taken to attenuate the blast wave before it strikes the wall of the pressure vessel. In fact, several steps can be taken to increase the rate at which the energy in the shock wave is degraded. One step would be to introduce a large number of gas bubbles in order to reduce the mean density of the liquid by about 5 percent. As the shock wave traversed a gas bubble some of its energy would go into spalling liquid from the inner surface of the bubble and projecting it through the void to strike the bubble wall at the other side, where much of the kinetic energy of the projected material would be converted to thermal energy. The extent to which this would attenuate the main shock wave is difficult to estimate, but simple tests indicate that a sufficiently high bubble density could be obtained to allow an increase in the energy released per explosion by as much as a factor of 10 over the corresponding value for a pool of bubble-free lithium.

The steps outlined here are designed to reduce the intensity of the blast wave before it impinges on the pressure vessel, but they would have little or no effect on the blast wave progressing upward through the cavity in the liquid lithium toward the laser and the pellet-injection system. Fortunately the low density of the lithium vapor will diminish the intensity of the blast wave. In addition the port for pellet injection and the laser beam can be made quite long, with a wall profile that would break the normal shock wave up into a host of oblique shock waves and thus dissipate the bulk of its energy.

Let us try to visualize in detail the sequence of events in the course of one explosion cycle, beginning with the injection of the deuterium-tritium pellet into the vortex. Perhaps 10 percent of the pellet would be vaporized as it moved along its trajectory as a consequence of the absorption of heat radiated to it by the hot lithium. The remainder of the pellet would be vaporized and ionized when it was struck by the laser beam. Some of the deuterium and tritium ions would fuse, and the fusion energy released would increase the temperature of the rest of the pellet material so that about 5 percent of the ions would fuse before plasma expansion terminated the reaction. About 75 percent of the energy from the fusion reaction would go into energetic neutrons, and this energy would be absorbed as those neutrons were slowed down in the inner region of the lithium. The absorbed energy would raise the temperature of the central lithium region a few hundred degrees in about a millionth of a second and would cause the lithium to expand; this expansion, however, should be largely absorbed by contraction of the bubbles so that no severe blast wave would be generated. About 25 percent of the energy released by the fusion reaction would appear as alpha particles; these would produce ion heating and, subsequently, X rays. The energy in the X rays would be absorbed in the first millimeter or two of the lithium layer, causing the lithium in that thin layer to vaporize. The vaporized lithium together with the expanding thermonuclear plasma would induce a weak blast wave in the lithium.

The velocity of sound in hot lithium is about 15,000 feet per second. The liquid would move outward from the explosion center at a lower velocity to form a spherical cavity. The displaced liquid would move into the upper portion of the vortex cavity. This would entail a good deal of sloshing, particularly in the region of the ports for the injection of the fuel pellet and the laser beam. The geometry of the injection passage de-

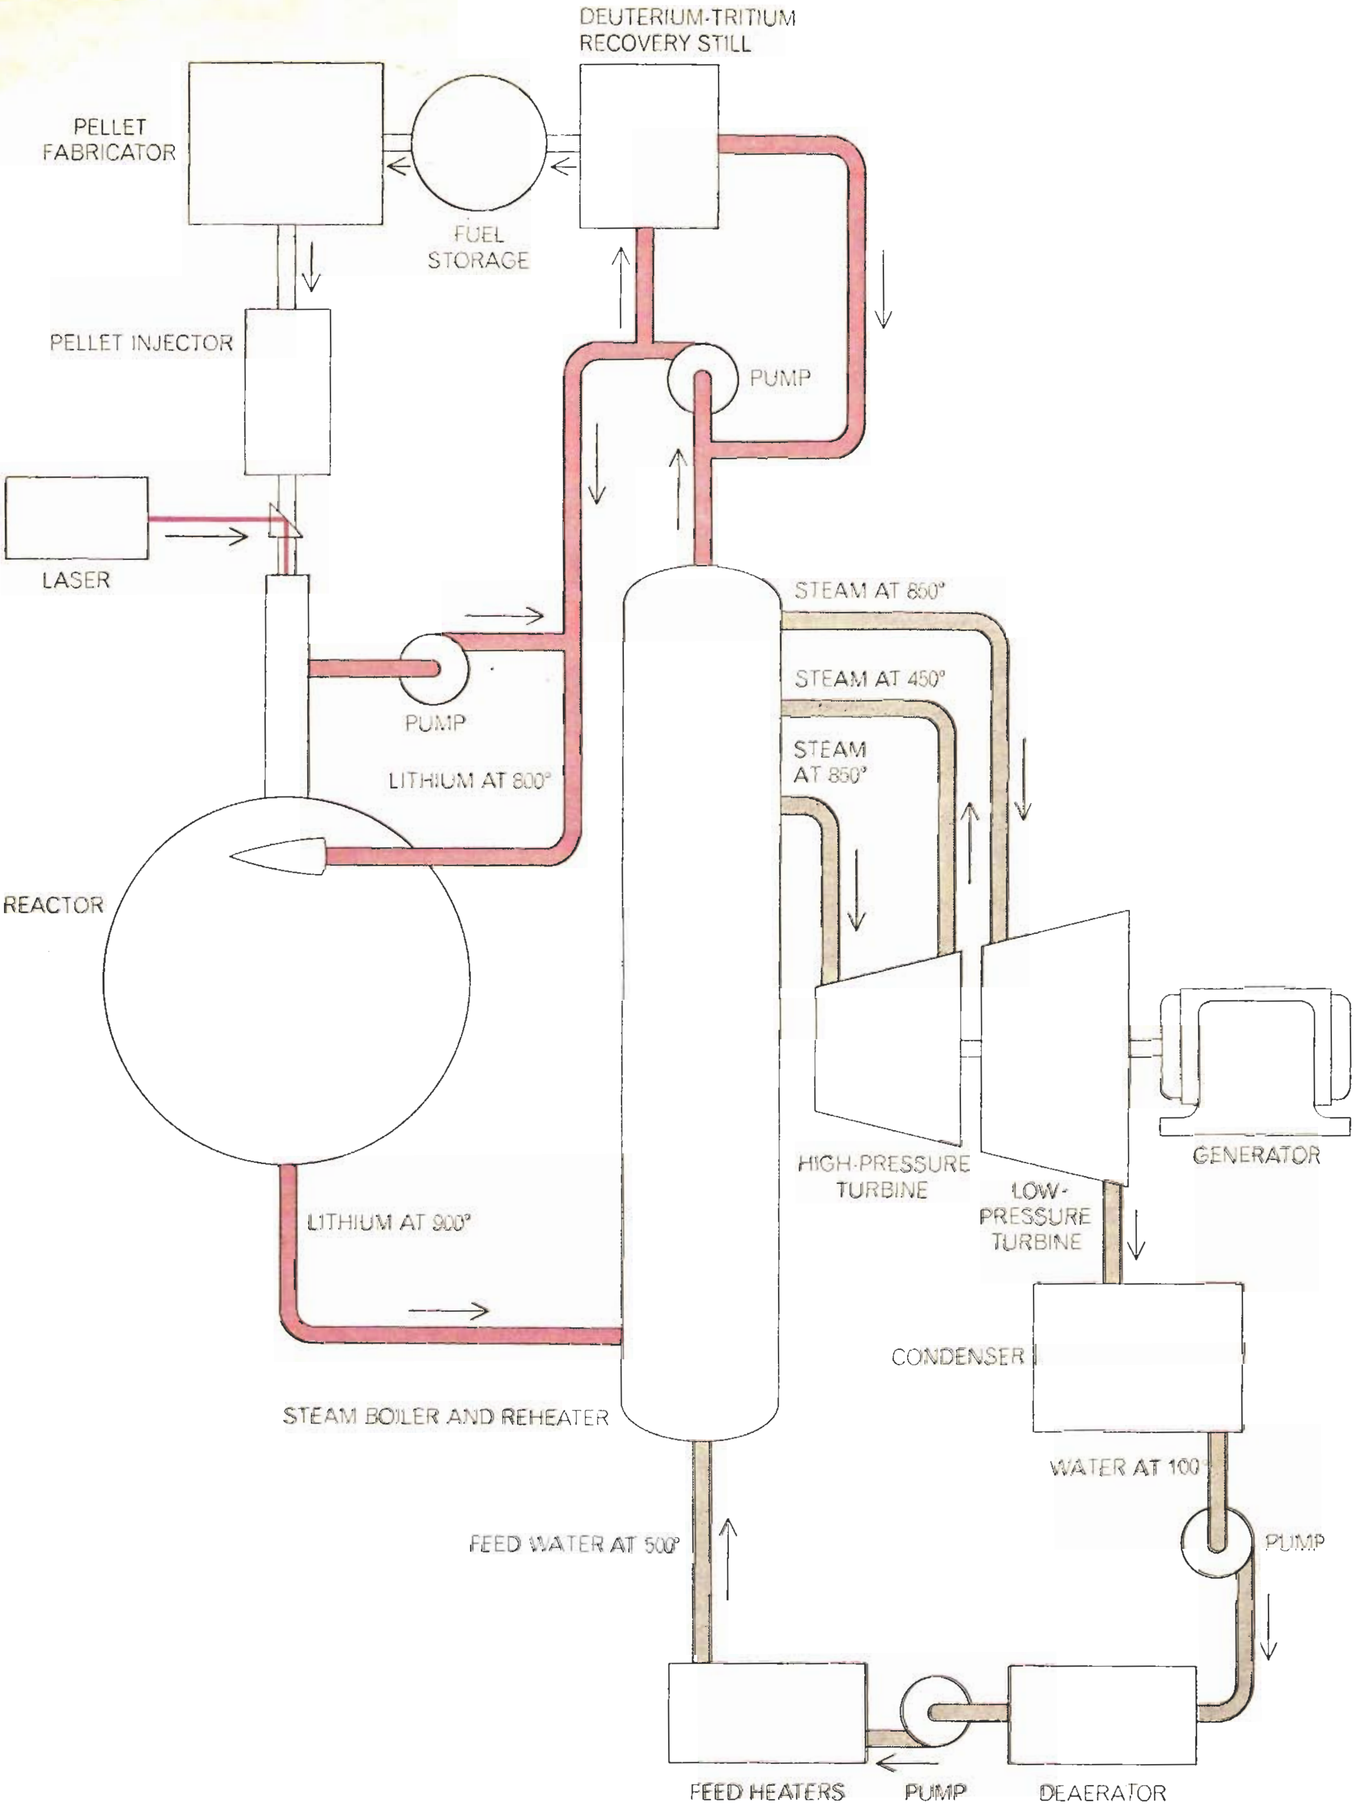

FULL-SCALE POWER PLANT incorporating many of the concepts discussed in this article might be based on a single fusion reactor in which deuterium-tritium pellets would be ignited by high-power laser pulses. The energy of the explosions would be absorbed in a tritium-breeding lithium blanket and used to generate 150 megawatts of electric power by means of a conventional steam cycle.

tory, the Sandia Corporation and the University of Rochester. Laboratories around the world have been improving on those initial results. Total neutron yields in excess of 10^5 from the deuterium-deuterium reaction have been measured using targets of deuterium with incident laser-beam energies of between 50 and 250 joules. The interpretation of these results indicates a slightly better absorption than simple collisional theory predicts.

Of special interest is the detailed study of the absorption of the incident radiation. There are a number of important contributions to the absorption that may be more significant than the simple collisional process described above and that may hence affect the energy release. Nature, however, seems to favor a freely expanding dense plasma, since the majority of these special absorption processes increase the effectiveness of the incident laser beam.

At present lasers with a properly shaped pulse can deliver 10^3 joules, whereas at least 10^5 joules is needed for significant amounts of laser-produced fusion. This laser energy can be made available using present technology, but as yet no such working system exists. Although the experimental results are few, they appear to justify optimism.

Meaningful use can be made of the energy released from a dense centimeter-size fuel pellet heated by focused laser radiation in a number of different energy-conversion applications. For example, a central station generating electric power by means of controlled thermonuclear reactions from a periodic vaporizing of deuterium-tritium pellets would look quite different from a configuration designed for use in, say, space propulsion. As an illustration let us consider the feasibility of this general scheme for a controlled thermonuclear power plant.

A method for converting the fusion energy from laser-ignited deuterium-tritium pellets into electrical power was evolved at the Oak Ridge National Laboratory early in 1969 in conjunction with fusion-power feasibility studies that have been under way there since 1967. The Oak Ridge approach to fusion entails absorption of the energy from fusion in a pool of lithium, which in turn delivers the energy as heat to a thermodynamic cycle.

The lithium pool would be contained in a cylindrical or spherical pressure vessel 10 to 15 feet in diameter and would be swirled at a sufficiently high velocity to form a free vortex around its vertical axis [*see illustration on opposite page*]. By adjusting the swirl velocity distribution properly it should be possible to obtain a central cavity with a fairly uniform diameter of perhaps five centimeters through the region from the top of the vessel to well below the midplane. A frozen pellet of deuterium-tritium would be injected into the center of the vortex cavity and would be ignited with a laser pulse when it reached the midplane. The energy deposited in the lithium as heat would be removed by drawing off lithium from the bottom of the pressure vessel, circulating it through heat exchangers and returning it to the pressure vessel. The process would be repeated perhaps every 10 seconds. The large thermal inertia in the lithium circuit would act to maintain an essentially constant flow of heat to the thermodynamic cycle.

Drawing off the lithium from the bottom of the pressure vessel would help to stabilize the vortex. After circulating through the heat exchangers the lithium would be returned through pumps to tangential nozzles in the perimeter of the pressure vessel, thereby maintaining the desired vortex.

An obvious objection to this concept is that it would not serve to yield useful amounts of power because the amount of energy one might release in each explosion would have to be kept quite small or the pressure vessel would burst. A relatively simple analysis discloses that this is indeed the case unless some measures are taken to attenuate the blast wave before it strikes the wall of the pressure vessel. In fact, several steps can be taken to increase the rate at which the energy in the shock wave is degraded. One step would be to introduce a large number of gas bubbles in order to reduce the mean density of the liquid by about 5 percent. As the shock wave traversed a gas bubble some of its energy would go into spalling liquid from the inner surface of the bubble and projecting it through the void to strike the bubble wall at the other side, where much of the kinetic energy of the projected material would be converted to thermal energy. The extent to which this would attenuate the main shock wave is difficult to estimate, but simple tests indicate that a sufficiently high bubble density could be obtained to allow an increase in the energy released per explosion by as much as a factor of 10 over the corresponding value for a pool of bubble-free lithium.

The steps outlined here are designed to reduce the intensity of the blast wave before it impinges on the pressure vessel, but they would have little or no effect on the blast wave progressing upward through the cavity in the liquid lithium toward the laser and the pellet-injection system. Fortunately the low density of the lithium vapor will diminish the intensity of the blast wave. In addition the port for pellet injection and the laser beam can be made quite long, with a wall profile that would break the normal shock wave up into a host of oblique shock waves and thus dissipate the bulk of its energy.

Let us try to visualize in detail the sequence of events in the course of one explosion cycle, beginning with the injection of the deuterium-tritium pellet into the vortex. Perhaps 10 percent of the pellet would be vaporized as it moved along its trajectory as a consequence of the absorption of heat radiated to it by the hot lithium. The remainder of the pellet would be vaporized and ionized when it was struck by the laser beam. Some of the deuterium and tritium ions would fuse, and the fusion energy released would increase the temperature of the rest of the pellet material so that about 5 percent of the ions would fuse before plasma expansion terminated the reaction. About 75 percent of the energy from the fusion reaction would go into energetic neutrons, and this energy would be absorbed as those neutrons were slowed down in the inner region of the lithium. The absorbed energy would raise the temperature of the central lithium region a few hundred degrees in about a millionth of a second and would cause the lithium to expand; this expansion, however, should be largely absorbed by contraction of the bubbles so that no severe blast wave would be generated. About 25 percent of the energy released by the fusion reaction would appear as alpha particles; these would produce ion heating and, subsequently, X rays. The energy in the X rays would be absorbed in the first millimeter or two of the lithium layer, causing the lithium in that thin layer to vaporize. The vaporized lithium together with the expanding thermonuclear plasma would induce a weak blast wave in the lithium.

The velocity of sound in hot lithium is about 15,000 feet per second. The liquid would move outward from the explosion center at a lower velocity to form a spherical cavity. The displaced liquid would move into the upper portion of the vortex cavity. This would entail a good deal of sloshing, particularly in the region of the ports for the injection of the fuel pellet and the laser beam. The geometry of the injection passage designed to keep the blast wave from moving up through the ports, however, would also be effective in inhibiting the liquid from sloshing up the port.

The rotational momentum of the lithium in the pressure vessel should help to reestablish the vortex. In addition the lithium through-flow in a typical case would be so large that about 50 percent of the lithium in the vessel would be replaced each cycle.

The pressure-vessel stress problems were examined by C. V. Chester and Lawrence Dresner at Oak Ridge. They quickly reached the conclusion that the proposed system presents so many analytical complexities, particularly if the effects of gas bubbles in the liquid are included, that it is doubtful an analytical model can be developed on purely theoretical grounds. Nonetheless, it does appear that all the phenomena involved are such that the scaling law commonly used for analysis of blast effects would apply. Hence test work has proceeded with small models.

Chester and Dresner designed a series of model tests using small steel vessels and ordinary explosive charges. The first tests were run with the model pressure vessel filled with bubble-free water. The amount of explosive was increased in small increments from about 1.5 grams to about 10 grams. The dilation of the vessel at the midplane was measured with a micrometer following each test, and the results were plotted as a function of the weight of charge detonated. A line drawn through the scatter band of points was passed through the zero dilation axis to define the maximum charge that could be detonated without stressing the vessel beyond the elastic limit. The tests indicated that this value was consistent with the value determined analytically, and that it corresponded to about two grams of explosive. Thus the simplified analytical model was deemed sound and suitable for scaling to full-size vessels.

A second series of tests was run with a thin layer of sponge rubber lining the inner wall of the pressure vessel. This cushioned the blast wave enough to increase to more than four grams the amount of explosive that could be used without stressing the vessel beyond the elastic limit.

A third series of tests was run with a perforated ring placed in the bottom of the model pressure vessel. Air was admitted to the ring to provide an annular curtain of bubbles in the region near the wall of the pressure vessel. A strain gage mounted on the outside of the model in the same horizontal plane as the explosive provided a good measure of the stress in the vessel initially with no air bubbles in the vessel and later with air bubbles filling about 5 percent of the vessel.

In extrapolating these results to a full-scale system—a pressure vessel 25 times larger, with a diameter of 12.5 feet and a wall 10 inches thick—the energy released from the explosion would be increased by a factor of 25^3, or 15,625. Allowing a safety factor of two, one could employ 15 kilograms of explosive in a full-scale system using bubble-free water at room temperature. This corresponds to an energy release of 15×10^5 calories per explosion. If the explosion were repeated at a frequency of one every 10 seconds, the power output would be 6.3 megawatts.

This value of 6.3 megawatts would be the energy in the blast wave for a system in which TNT would be exploded in water with no vortex or bubbles to cushion the explosion. With TNT all the energy would go into the blast wave, but in a fusion reactor about 80 percent of the energy would appear in the form of heat in the liquid as a consequence of the slowing down and absorption of fast neutrons. Most of this energy would simply go into heating the liquid and compressing the bubbles rather than a blast wave; hence the energy output of the reactor could be approximately five times the energy producing the blast wave, that is, about 30 megawatts. The model tests indicate that the cushioning effect of the bubbles ought to reduce the intensity of the blast wave by at least another factor of 2.5, and that the use of a spherical rather than a cylindrical vessel would give another factor of two, thus allowing an increase in the rated power output to perhaps 150 megawatts. Some further increase in power may be possible as a consequence of differences in compressibility between water and high-temperature lithium, but data are not in hand to estimate the extent of this effect.

The design of a full-scale power plant was derived from the concept outlined here within the limitations imposed by basic design considerations. For example, although deuterium is cheap and readily available, tritium is not. As a result, for a deuterium-tritium fusion reactor to be economically attractive it must breed tritium. This can be done because the energetic neutrons have a substantial probability of colliding with other neutrons in the course of slowing down in a natural lithium blanket. When a neutron is absorbed in the lithium, it yields tritium plus an alpha particle. Thus a breeding ratio of perhaps 1.3 would be obtained. If the lithium pool had a radius of at least a meter, there would be virtually no neutron losses from escape or absorption in structural material, and there would be no neutron damage to the pressure vessel. Hence lithium appears to be a particularly attractive choice for the fluid to absorb the energy of the explosion.

In the choice of material for the pressure vessel, extensive compatibility tests indicate that the two most promising candidates for use with lithium are Croloy (a chrome-molybdenum steel) and a niobium alloy containing 1 percent zirconium. The operating temperature with Croloy is limited by corrosion considerations to about 900 degrees F., whereas the niobium alloy could be employed with lithium up to at least 2,000 degrees F. Niobium is very expensive, however, and requires protection from oxidation; Croloy is preferable, at least for initial studies and experiments.

In choosing a thermodynamic cycle the first logical choice is a steam Rankine cycle. Croloy is widely used in existing fossil-fueled steam plants for steam temperatures up to 1,050 degrees F. Thus if an upper temperature limit of 900 degrees is adopted for the lithium system and a 50-degree temperature drop between the lithium and the steam circuit is accepted, the reactor could be coupled to a quite conventional steam system [*see illustration on opposite page*].

It might be argued that an intermediate fluid such as an inert salt should be used between the lithium and the steam circuit, because of the large energy release associated with the lithium-water reaction that would occur if there were a leak from one system into the other. The consequences of such a leak, however, appear no more serious than those associated with the heat exchangers between the sodium-potassium and the steam system employed in fast sodium-cooled fission reactor plants. The amount of radioactivity in the lithium would be small and the loss of tritium in such an accident would not be serious.

In order to keep the thermal stresses to a modest level in the vessels and piping of the lithium system it seems desirable to keep the average temperature rise in the lithium to about 100 degrees F. This could be accomplished by using a flow rate of 2,300 gallons per minute with two pumps operating in parallel.

A physically small but economically

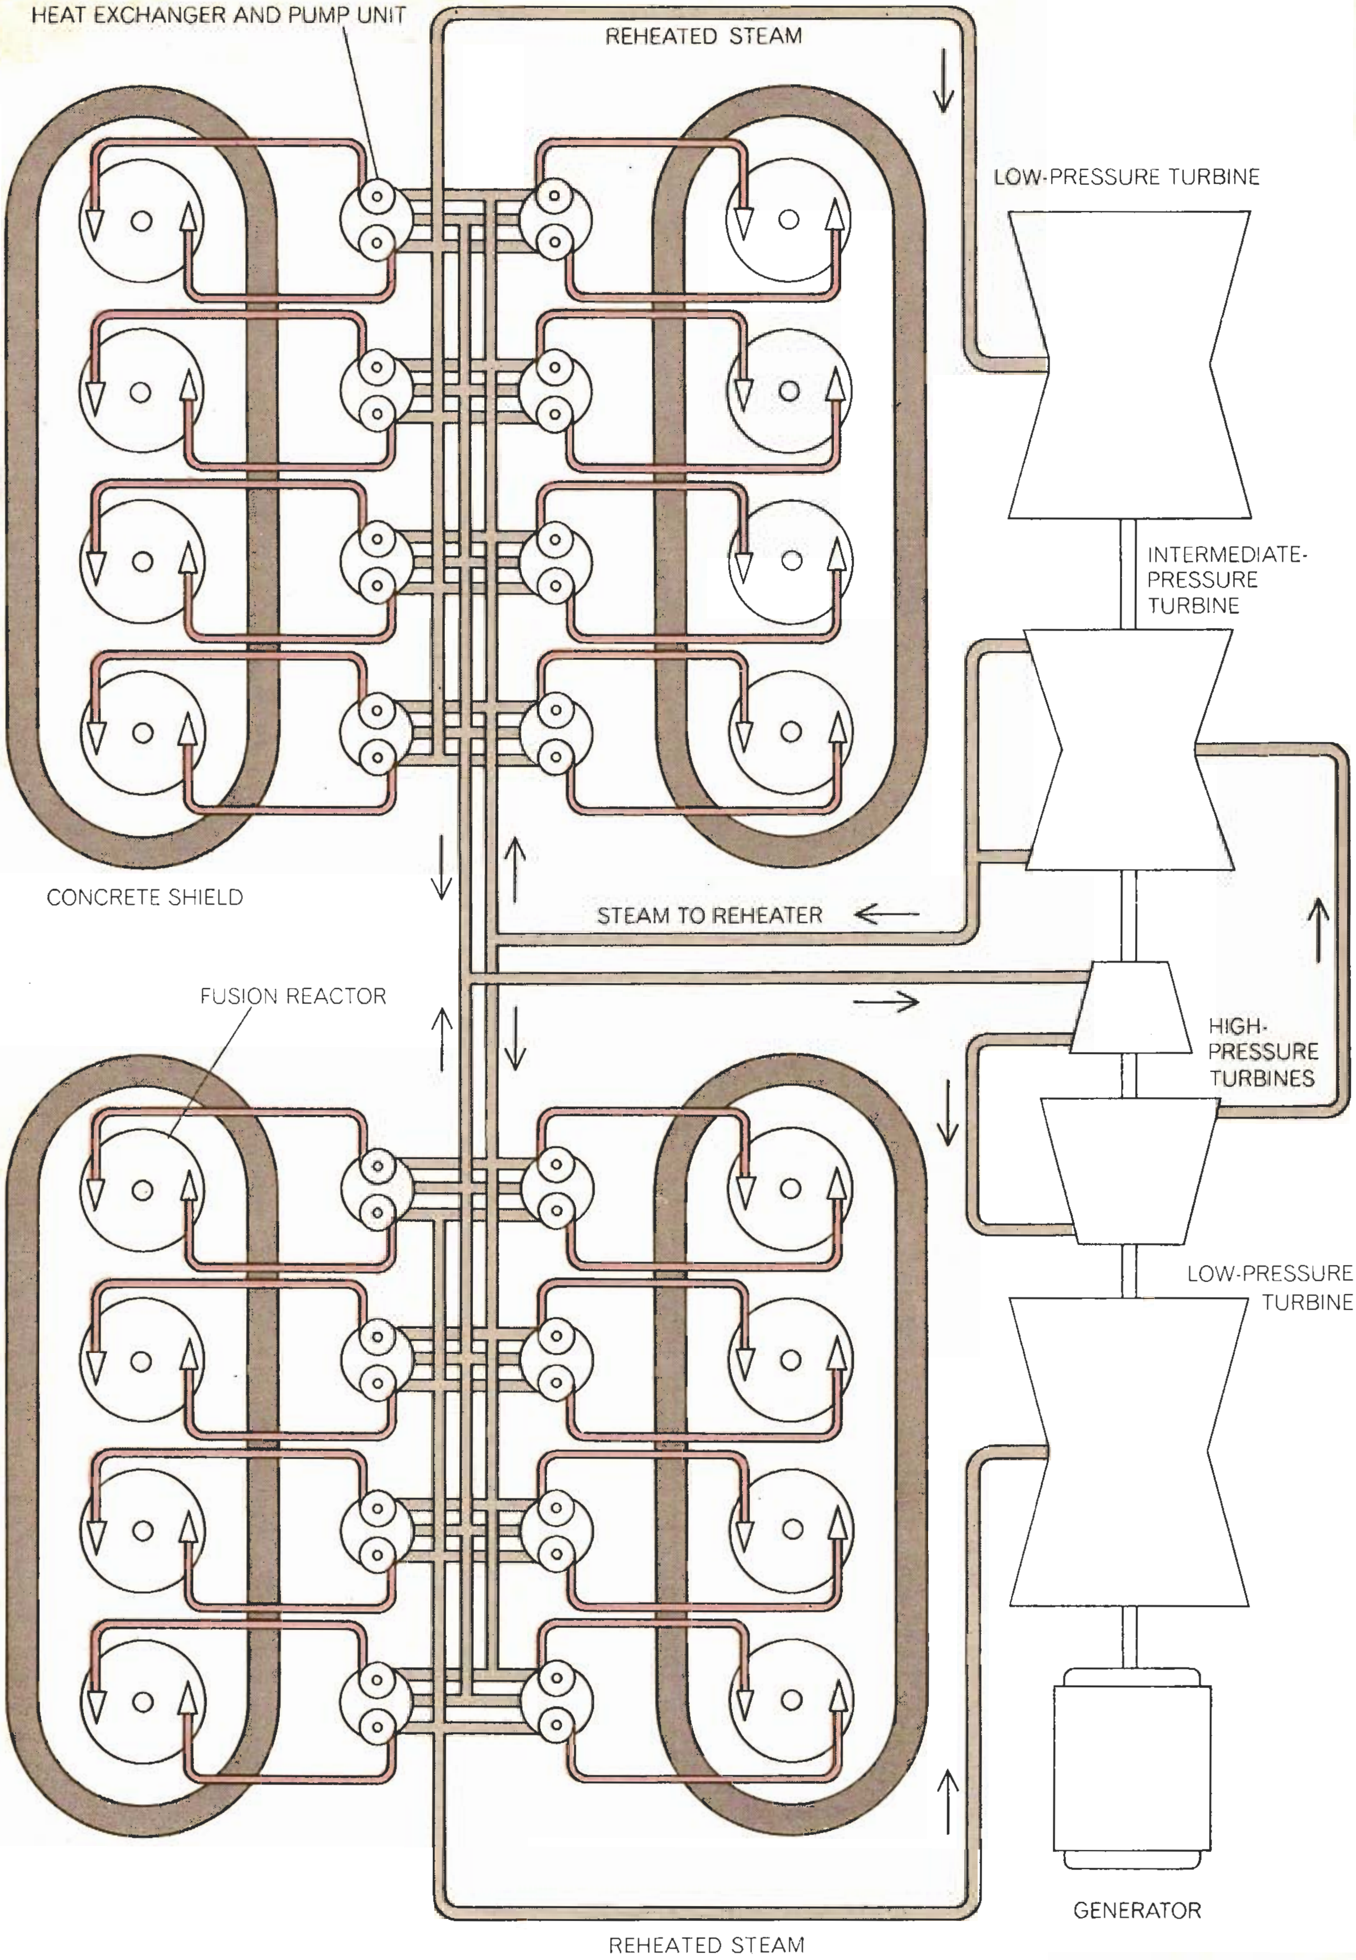
HEAT EXCHANGER AND PUMP UNIT
REHEATED STEAM
LOW-PRESSURE TURBINE
INTERMEDIATE-PRESSURE TURBINE
CONCRETE SHIELD
STEAM TO REHEATER
FUSION REACTOR
HIGH-PRESSURE TURBINES
LOW-PRESSURE TURBINE
GENERATOR
REHEATED STEAM

vital portion of the plant is the fuel-recovery and fuel-reprocessing system. It has been shown that both tritium and deuterium can be removed from lithium at low cost with a small, essentially conventional distillation system. After being compressed, the deuterium-tritium mixture can be liquefied with a cryogenic system. Fuel pellets can then be manufactured by allowing liquid droplets to fall through a vacuum chamber filled with cold helium gas, which would chill the droplet and freeze it. The electric power required for the cryogenic system appears to be less than a watt per kilowatt of thermal output from the reactor; hence it should not detract appreciably from the overall thermal efficiency of the power plant. It would, of course, be necessary to add deuterium and take off a deuterium-tritium mixture to maintain the proper composition of the fuel pellets.

For plants with an output greater than 150 megawatts it might be possible to make the cushion of bubbles more effective and thus increase the power output from a vessel of a given size by a factor of two or more. Power outputs of more than about 500 megawatts per vessel will require that the size and/or thickness of the vessel be increased. This will increase the difficulties of fabrication and consequently the unit cost in dollars per pound and dollars per kilowatt of output. The alternative is to use a number of relatively small vessels operated in parallel [*see illustration on opposite page*]. The latter approach is attractive because it would make possible modular construction with one standard size or more of shop-fabricated reactor vessels coupled with appropriately sized pumps, heat exchangers and connecting pipes. This approach is particularly attractive because large steam turbines are commonly supplied with steam through many steam pipes, since provisions for thermal expansion make it advantageous to keep the diameter—and hence the capacity—of individual steam pipes to a modest level.

It should be mentioned that the type of fusion reactor proposed here differs from both fission reactors and other types of fusion reactors in that there appears to be no theoretical advantage to the use of very large units. This constitutes an important advantage for the concept.

A laser-initiated fusion-power plant built of Croloy and operated to produce steam at 850 degrees F. would have an overall thermal efficiency of about 40 percent, which is approximately that of the better fossil-fueled plants in current operation. The efficiency could be increased dramatically if the system were built of the niobium alloy and operated at a temperature of about 1,800 degrees F. That temperature would be too high for use in a steam cycle because there would be enough dissociation of the hydrogen and oxygen in the steam to cause serious attack of any structural metal that would be economically attractive. It should be possible, however, to employ a binary-vapor cycle with a potassium-vapor Rankine cycle taking heat from the lithium at perhaps 1,800 degrees F. Such a system would have an overall thermal efficiency for the cycle of about 58 percent. This would cut the waste heat rejected per electrical kilowatt to about half that for conventional steam plants, and it opens the possibility of even higher efficiencies of heat utilization by integration with industrial and urban heating systems.

There is little point in developing the proposed power plant unless it looks attractive economically. If a suitable laser system can be built at a reasonable cost—a major question that cannot be resolved at this stage—it appears that the capital cost of the rest of the system should be no higher than the capital cost of more conventional plants. Certainly the cost of the steam system should be the same as that of similar steam systems in current use. The lithium circuit and steam boiler should cost about the same as the corresponding equipment projected for the circuit of the sodium-cooled fast-breeder fission reactor. Moreover, the pressure vessel would be relatively simple compared with those currently employed in water-cooled fission reactors. The weight of a spherical pressure vessel with an inside diameter of 12.5 feet and a wall thickness of 10 inches would be about 200,000 pounds. Assuming a 40 percent thermal efficiency, such a plant would yield an electrical output of about 60,000 kilowatts. Assuming a unit cost of $3 per pound, the vessel would cost $600,000, or about $10 per kilowatt. The cost of other items, such as the cryogenic systems for fuel recovery and pellet fabrication, appears to be a small fraction of the overall cost of the plant.

The operating cost for fuel should be low indeed. The cost of the deuterium and lithium would be about $200 per pound and $15 per pound respectively. This would yield a cost of only about three cents per million BTU (British thermal units) compared with current prices of about 40 cents per million BTU for fossil fuels. Inasmuch as capital charges are roughly the same as fuel costs in conventional fossil-fuel plants, this in turn indicates that the proposed laser-initiated fusion-power plant would be economically attractive even if the capital cost were twice as high.

One of the major problems associated with fission reactors stems from the large inventory of radioactive material inherently present and the potential hazard to the public that it represents. Because of this factor as much as 30 percent of the capital cost of a fission reactor plant may stem from elaborate provisions to prevent or contain any conceivable accident that might release radioactive material to the environment. The only radioactive material of consequence in the proposed fusion-power plant would be tritium, and estimates indicate that it would represent a total hazard potential that would be lower than that of a comparable fission reactor by a factor of about a million. This should reduce costs and greatly ease siting problems.

The design studies of full-scale laser-initiated fusion-power plants make the concept look attractive and therefore raise questions with respect to the development problems. Clearly the most vital are concerned with the laser and the pellet-ignition process. Can a sufficiently powerful laser be built and pulsed to ignite the pellet? If so, will the yield of fusion energy be many times more than the energy input to the laser? How effectively can the blast wave be attenuated by entraining a substantial fraction of gas bubbles in a swirling pool of lithium? These are the principal questions that must be answered experimentally before one can say whether or not the proposed concept is really feasible. There are, of course, many other difficult development problems, but none appears so difficult as to raise doubts about the feasibility of the concept.

THOUSAND-MEGAWATT POWER PLANT might be composed of 16 fusion-reactor modules, each with its associated heat exchanger and pump unit. The system would produce steam at a temperature of 850 degrees Fahrenheit and a pressure of about 1,200 pounds per square inch to drive a combination of high-pressure and intermediate-pressure turbines with one reheat to 850 degrees for a pair of low-pressure turbines. With the addition of a potassium-vapor "topping" cycle operating at a turbine-inlet temperature of 1,800 degrees the overall thermal efficiency of such a system could be brought as high as 58 percent.