ometer. Curved mirrors are arranged in such a way that the light "walks" across the mirror surface as it bounces back and forth along the interferometer arm around 50 times. The arms of the Caltech interferometer are gigantic Fabry–Perot cavities—the laser light bounces off the same spot on the mirror on each of its transits. The Fabry–Perot scheme requires lasers with excellent frequency stability but has the advantage of using smaller mirrors. In a 4-km detector the Fabry–Perot interferometer needs 7-inch-diameter mirrors; the delay-line system needs mirrors 20 or 30 inches in diameter. Reduction in mirror size could be important, because the mirrors themselves must be of extremely high quality. They are of a type developed originally for laser gyroscopes and lose only about 1 part in 10^4 of the incident light.

Prototypes of the two detectors are currently operating at MIT and Caltech. Interferometers of both types have been built by groups at the Max Planck Institute, Garching, FRG, and at the University of Glasgow. The Caltech Fabry–Perot interferometer has 40-m arms, making it effectively a $\frac{1}{100}$-scale model of the proposed detector. Light can be stored for as long as milliseconds in the interferometer, making 10 000 trips between the mirrors; the full-scale detector should store light for an astonishing half second. The MIT prototype delay-line interferometer has 1.5-m arms. The mirrors are arranged so that the light traverses the arm 56 times before leaving.

The signal-to-noise ratio of either interferometer increases as a function of the arm length. There are several reasons for this; an important one is simply that whereas the phase delay due to a gravity wave grows with the length of the interferometer arm, seismic noise, which enters only at the mirrors, is independent of arm length. On the other hand, although sophisticated measures such as active damping—which involves canceling sensed vibrations with active controllers—can be taken to isolate the mirrors, seismic vibrations are still an important source of noise at low frequencies.

Another important source of noise is the lasers themselves. Gravity waves will only produce a tiny fraction of a fringe shift in the interferometer. The apparatus is, in a sense, an alternative physical realization of that well-known *gedanken* experiment, the Heisenberg microscope. If the laser light is too faint, the tiny fringe shifts caused by a passing gravity wave will be impossible to resolve because of a type of quantum fluctuation known as shot noise. Increasing the laser power solves this problem but potentially causes another: The hail of photons may jostle the mirrors sufficiently to wipe out the fringe pattern. These two extremes define the quantum limit of detectability. Currently the experimenters are safely below the quantum limit, but they believe a later generation of detectors might be limited by quantum effects.

Practical detectors. The interferometer, regardless of which design is eventually chosen, is by far the least expensive component of the proposed detector. Both interferometers must operate in a vacuum to eliminate noise due to statistical fluctuations in air density. The current design calls for vacuum pipes 48 inches in diameter, which allows for different interferometer designs and choices of laser wavelength and the simultaneous operation of several different interferometers. The motivation for this design is to allow the development of future generations of gravity-wave detectors without interruption of the ongoing search.

The wide diameter of the vacuum tube will also allow the experimenters to play other tricks aimed at increasing the sensitivity of the detectors. One plan is to include a second interferometer half as long as the first in the vacuum tube. The response of this second detector to gravity waves will be exactly half as large as that of the first, but they will respond differently to seismic shocks. This should help the experimenters to discriminate the waves from the noise.

The only sure way to discriminate between the two is to look for coincidences between two or more widely separated detectors. In addition, by measuring the difference between the arrival times of gravity-wave signals at the detectors in California and Maine the experimenters will be able to approximate the general direction of the gravity-wave source. But to really pin a source down, which is necessary if gravity-wave detectors are to become astronomical tools, four widely separated detectors will be needed. With luck, these will be built by the groups now active in Europe.

Second-generation improvements. Although the experimenters believe there is a fair chance that the initial detector system will detect gravity waves, they have thought of improvements that should increase the sensitivity by several orders of magnitude, possibly opening up new, fainter astrophysical sources to their scrutiny.

One such improvement is known as light recycling. The sensitivity of an interferometer depends on the amount of light bouncing between the mirrors: The more photons, the better a fringe can be resolved. One can increase the amount of stored light by using more and more powerful lasers, but the energy costs of running these lasers eventually get prohibitive. Another approach stems from the fact that most of the light that goes into an interferometer is wasted. The light that travels down the arms recombines at a beam splitter, where it exits in two directions. The experiment is adjusted so that the observed pattern is at a dark fringe to avoid overloading the photodiode. This means that most of the light escapes through the other side of the beam splitter. The idea behind light recycling is simply to feed this light back into the interferometer. The difficulty is that the light must be fed back in precisely the right phase. The idea does not give a useful advantage in small systems and awaits testing in the large interferometers.

Another clever idea is applicable only to the detection of periodic sources. When a gravity wave passes through the interferometer for one half-cycle, one arm is shortened while the other is lengthened. During this time the light in the arms begins to accumulate their phase shift. If, after a half-cycle of the gravity wave, the light in the arms is interchanged, the phase shift will continue to grow. By adjusting the frequency of interchange to that of the gravity wave a significant sensitivity improvement due to resonance should be possible.

The National Science Foundation has provided around $2.5 million this year toward the continuing development of the prototypes at Caltech and MIT. If nature and NSF are kind, sometime in the 1990s astrophysicists and relativists will have a new tool for viewing the universe.

—Bruce Schechter

Superconducting-cavity progress spurs new CEBAF design

What is the best way to provide nuclear physicists with continuous beams of 4-GeV electrons? Three years ago the pulse-stretcher-ring concept embodied in the winning SURA design for the Continuous Electron Beam Accelerator Facility appeared to be the optimal solution (PHYSICS TODAY, July 1983, page 57). But while construction of CEBAF has not yet begun, accelerator technology has not stood still. In the light of recent progress in the development of superconducting rf cavities, a review of CEBAF technology initiated last summer by Hermann Grunder, the

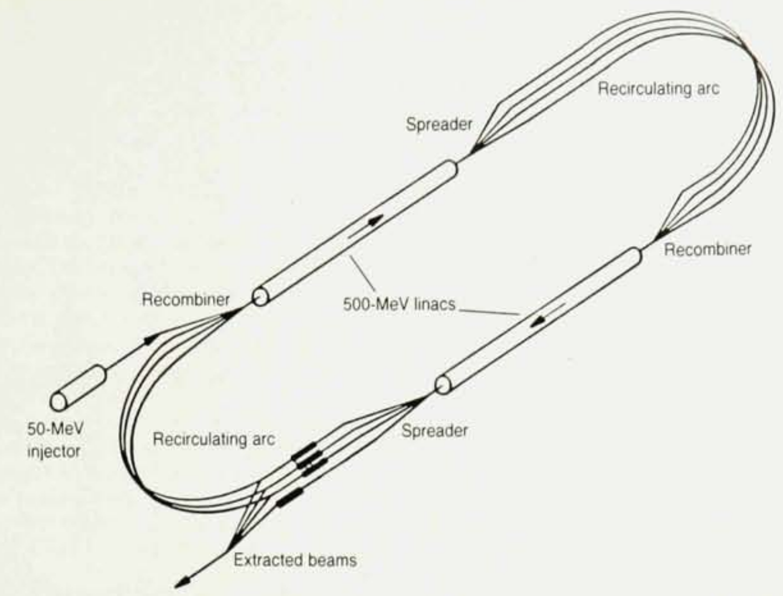

A pair of superconducting linacs are to provide a continuous beam of 4-GeV electrons in this recently revised design for the Continuous Electron Beam Accelerator Facility, planned for Newport News, Virginia. The beam is circulated through the system four times, gaining 500 MeV each time it traverses one of the 235-m-long linacs. Electrons at four different energies emerge continuously from each linac, requring four separate arrays of bending magnets in each recirculating arc. One can also extract continuous beams at these different energies simultaneously for several experiments. Recent progress in the development of superconducting niobium rf cavities makes this an attractive alternative to the original CEBAF design.

project's new director, has concluded that a superconducting-linac design should now replace the original pulse-stretcher-ring concept (see page 51). Looking toward the beginning of construction in fiscal 1987, DOE in December directed CEBAF to prepare a conceptual design for such a machine.

Continuous beams. The basic technological problem is that conventional rf-electron linacs are intrinsically pulsed accelerators. The electric field gradient that accelerates the electrons in such a machine is produced by introducing rf power into a linear array of copper cavities. The electric fields, however, produce enormous resistive heating in the copper walls, so that the linac can be operated only in a pulsed mode. In continuous (as distinguished from pulsed) operation, a conventional rf linac of adequate gradient would generate about 300 kilowatts of heat per meter, and the resulting thermal stresses would exceed the yield strength of copper.

The nuclear experimenters, however, want continuous electron beams, or at least duty cycles close to unity. With the 10^{-3} duty cycles typical of electron linacs, it's very difficult to do the coincidence experiments that are regarded as crucial to the high-energy electron–nucleus scattering program. Coincidence experiments depend on being able to tell whether two particles emerged from the same collision or from different collisions that occurred at about the same time. In pulsed accelerator operation at useful beam intensities the number of spurious coincidences is simply too high when all the collisions occur in very brief bursts. One has to find a way of spreading the beam out in time.

The trick envisioned in the original SURA design is to inject the brief output pulse of a conventional linac into a racetrack-shaped "pulse-stretcher ring" 380 meters in circumference—just about long enough to be filled by the 1.2-microsec beam pulse of relativistic electrons. During the following millisecond of recovery time before the next linac pulse, the beam circulating in the ring would be slowly and continously extracted to provide an almost continuous electron beam for experimenters; the pulse-stretcher ring would raise the effective duty factor of the linac from 10^{-3} to about 80%.

Superconducting linacs. Although it is generally agreed that the pulse-stretcher-ring scheme would work adequately, there now seems to be a better way. If the rf cavities in the linac were superconducting rather than normal metal, the resistive heating would be all but eliminated and one could run the linac itself in a cw (continuous wave) mode. Superconducting accelerator cavities have in fact been under development for more than 20 years, but only recently have they been able to demonstrate the minimum performance levels necessary for an accelerator like CEBAF—an accelerating gradient of 5 MeV/m and a Q of 3×10^9. Last year, in fact, three major high-energy-physics laboratories decided to incorporate some superconducting rf cavities in accelerators now under construction (LEP at CERN, HERA at DESY in Hamburg, and Tristan at KEK in Japan).

The "pre-conceptual" design report issued by CEBAF in December points out a number of advantages a cw superconducting linac would have over the original pulse-stretcher-ring scheme: conceptual simplicity, inherently continuous beam with 100% duty factor, better beam quality, considerable power savings during operation, potential for significant energy upgrading, and the ability to deliver simultaneous continuous beams at three different energies.

Experimenters would have simultaneous access to three (or even four) different beam energies because the cw superconducting design, as outlined in the pre-conceptual design report, involves four circuits of the electron beam through a pair of 0.5-GeV linacs on its way to the 4-GeV maximum energy. The two superconducting linacs, each 235 meters long, would lie side by side, with their accelerating gradients pointed in opposite directions. The beam would circulate from one linac to the other through the bending fields of recirculator arcs at both ends. The two linacs are kept 100 m apart for fear that the tighter bending radius required by closer proximity might excite destabilizing quantum excitations in the beam. (This becomes an even more serious consideration when one anticipates upgrading the system to 16 GeV at some future date.) Starting with an injection energy of 50 MeV, the electrons would gain 1 GeV for each complete circuit through both linacs, reaching the maximum energy after four round trips.

Thus, in continuous operation, electrons with four different, discrete energies are continuously emerging from the downstream end of each linac, depending on how many passes an individual electron has completed. Because the electrons are already fully relativistic at the initial injection energy, they traverse the straight linac sections in lockstep, irrespective of energy. But in the curved recirculating arcs between the linacs, the different momenta require different magnetic field strengths to bend them around. Thus the beam emerging from each linac is spread out into four discrete beams of different energy to follow four separate magnet arrays. At this point one can also extract beams at the various emerging energies for different experiments.

Ednor Rowe (University of Wisconsin) reminds us that this new CEBAF design is a recirculating system not unlike the Argonne three-sided microtron proposal that emerged as the chief rival of the original CEBAF design. Rowe was a member of the NSAC panel that—by a split decision—rejected the Argonne proposal in favor of the old CEBAF.

Niobium rf cavities. The acceleration in each of the 235-m-long linacs would be accomplished by a string of 200 superconducting niobium cavities, each about half a meter long. These rf cavities, designed to operate at 1.5 GHz (20 cm wavelength), have been developed by Ronald Sundelin and his colleagues at Cornell, where four prototype cavities have been built, each exceeding the gradient and Q performance criteria that have been specified for a superconducting CEBAF.

The Cornell group has been in the superconducting-cavity business since the late 1960s, when plans were afoot to build a 35-GeV electron synchrotron at Cornell. That project was eventually dropped in favor of the 8-GeV CESR e^+e^- storage ring, but superconducting-cavity development continued in anticipation of a 100-GeV e^+e^- storage ring Cornell hoped to build (PHYSICS TODAY, August 1981, page 20). With the shelving of that project in 1983, the Cornell group continued its work on superconducting rf cavities with an eye toward the eventual construction of a TeV e^+e^- linear collider.

The feasibility of a superconducting linac with a recirculating electron beam was demonstrated in the 1970s by Alan Schwettman and his colleagues at the Stanford High Energy Physics Laboratory. Three passes through the small HEPL linac produced a 150-MeV electron beam with a high duty factor and low phase-space spread—two essential requirements for CEBAF. But the accelerating gradients of the HEPL cavities were limited to about 2 MeV/m—less than half of what one needs for CEBAF. Furthermore, the external couplers that remove unwanted higher-harmonic power from the HEPL cavities were a hundred times too weak for the high-current operation of CEBAF.

It was eventually understood that the gradient limitation of the HEPL cavities was principally due to "multipacting": Stray electrons generated and accelerated in the cavity strike the walls and knock out secondary electrons, which can do the same after they, in turn, are accelerated. If, on average, each wall collision produces more than one secondary electron, this undesirable process grows exponentially. Multipacting can also occur in normal-metal cavities, but it presents a much more serious problem in a superconductor, where any deposition of energy threatens to defeat the cryogenic system that keeps the metal below its critical temperature.

Multipacting, Sundelin told us, was the most severe limiting factor militating against the improvement of superconducting-cavity performance in the 1970s. In 1979 a group at the University of Genoa discovered, quite by accident, that one could eliminate multipacting by employing a spherically shaped cavity. This geometrical effect was soon explained by Udo Klein and Dieter Proch at the University of Wuppertal. The present elliptical geometry of the Cornell cavities, developed by Peter Kneisel, preserves this elimination of multipacting while permitting a reduction of peak electric fields and improving mechanical rigidity.

Surface defects on the niobium cavity walls have also limited the performance capabilities of the superconducting cavities. These defects become normal conducting spots in an otherwise superconducting structure. The problem is to keep this normal conductivity local, to prevent the quenching of the entire cavity. One therefore wants the thermal conductivity of the superconducting walls to be as large as possible, to carry away the ohmic heat generated at the normally conducting surface imperfections lest it drive the entire structure normal by raising it above 9 K, the critical temperature of niobium.

The last two years have seen significant improvement in the thermal conductivity of the niobium. Beginning in 1983, a number of manufacturers have been able to achieve as much as an order of magnitude improvement in the thermal conductivity of the niobium they supply by employing a method of slow and repeated melting that drives out the interstitial gas that impedes heat transport. Hasan Padamsee of Cornell has also discovered that yttrium sublimated on the surface of the niobium improves its thermal conductivity by an additional factor of 3—by capturing oxygen from the bulk niobium.

Heat transport is also impeded by vacuum voids inadvertently produced in the niobium walls by standard electron-beam welding techniques. These vacuum voids appear to be responsible for breakdowns at rather modest accelerating gradients, and they also cause undesirable surface roughness. In 1983 Joseph Kirchgessner and Sundelin developed a remedy for this serious limitation on field intensity—replacing the sharply focused electron-beam weld with a raster-scan welding technique that moves the welding beam around so fast that not enough heat is generated in any one place to produce vacuum voids.

Another geometric issue that has in the past limited the performance of the superconducting accelerator cavities was the placement of the couplers that introduce the rf power into the cavity at its fundamental frequency and damp out undesirable higher modes. In the early designs, the couplers penetrated directly into the cavity through holes in its walls. Surface currents had to go around these holes, increasing the local magnetic field and hence the danger that local surface defects would quench the entire cavity. The coupler holes also complicated the field geometry in such a way as to exacerbate the problem of multipacting.

Beginning in 1975, Kirchgessner and Maury Tigner at Cornell developed a new coupling system for introducing rf power at the fundamental frequency—moving the coupling holes from the cavity wall to the beam pipes at the ends of the cavity. But the accelerating gradient was still limited to 3.5 MeV/m by the higher-mode couplers, which still penetrated the cavity walls. When Sundelin and Joseph Amato moved the higher-mode couplers to the external beam pipes in 1983, they raised this field-gradient limit while obtaining a thousandfold improvement over the earlier HEPL design in the damping of the higher modes. These higher modes, if they are inadequately damped, would cause serious instabilities at the electron-current intensities of interest for CEBAF.

"Superconducting rf technology," the pre-conceptual design report tells us, "is now sufficiently mature that a high-efficiency cw superconducting linac is practical." This contrasts sharply with the prevailing attitude in the early 1980s that spawned the proliferation of room-temperature cw accelerator designs, Rowe points out. "In those days people thought that a superconducting linac would be lots of trouble—the worst thing you could build."

—BERTRAM SCHWARZSCHILD □