GAMMA-RAY COLLIDERS AND MUON COLLIDERS

High-energy physicists have learned much from colliders with beams of protons, antiprotons, electrons and positrons. Now it seems both feasible and useful to build gamma–gamma and muon–muon colliders.

Andrew M. Sessler

The physics of beams is a discipline that has developed over the last 70 years, concerning itself with the manipulation and acceleration of beams of particles and light. Starting with electrostatic accelerators and advancing through cyclotrons and synchrotrons, this science has become ever more sophisticated. Nuclear physics exploits it nowadays in devices such as the continuous-beam superconducting electron accelerator at the Thomas Jefferson National Laboratory in Virginia, the ASTRID cooler ring at Aarhus University in Denmark and the Relativistic Heavy Ion Collider nearing completion at Brookhaven National Laboratory. The modern physics of beams has also made possible the dozens of synchrotron light sources that are currently making significant contributions to physics, chemistry and biology in many countries.

In high-energy physics, almost all of the present accelerators are colliding-beam machines. In recent decades these colliders have produced epochal discoveries: Stanford's SPEAR electron–positron collider unveiled the charmed-quark mesons and the τ lepton in the 1970s. In the realm of high-energy proton–antiproton colliders, the Super Proton Synchrotron at CERN gave us the $W^{\pm}$ and Z^0 vector bosons of electroweak unification in the 1980s, and in the 1990s the Tevatron at Fermilab finally unearthed the top quark, which is almost 200 times heavier than the proton.

Aside from e^+e^- and $p\bar{p}$ colliders, particle physics also has at its disposal the hybrid HERA collider at the DESY laboratory in Hamburg, which brings 800 GeV protons into collision with 30 GeV electrons. What about other particles? Beam physicists are now actively studying schemes for colliding high-energy photons with one another,[1,2] and schemes for colliding a beam of short-lived μ^+ leptons with a beam of their μ^- antiparticles.[3,4] If such schemes can be realized, they will provide extraordinary new opportunities for the investigation of high-energy phenomena. How such new colliders might be realized, and what new physics possibilities their realization would provide, are the subject of this article. (A briefer discussion can be found in the article by Jonathan Wurtele in PHYSICS TODAY, July 1994.)

ANDREW SESSLER, *an accelerator theorist at the Ernest Orlando Lawrence Berkeley National Laboratory, is also the current president of the American Physical Society.*

These exotic collider ideas were first put forward in Russia more than 20 years ago: Muon colliders were proposed by Gersh Budker, Alexander Skrinsky and Vasily Parkhomchuk, and gamma-ray colliders were proposed a few years later by Valery Telnov and Ilya Ginzburg. More recently these ideas have been picked up and significantly advanced—for gamma colliders by Kwang-Je Kim (Lawrence Berkely National Laboratory) and coworkers, and for muon colliders by David Neuffer (Fermilab), Robert Palmer (Brookhaven) and coworkers. Telnov remains the leading advocate for $\gamma\gamma$ colliders. The study of muon colliders is now a very active subject, with Palmer in the vanguard. The $\mu^+\mu^-$ Collider Collaboration brings together more than 100 physicists and engineers at 18 institutions worldwide.

Now, for the first time, particle physicists can effectively explore the origin of mass. How do the quarks and leptons acquire their masses? We have the theoretical understanding to know where to look, and we can build accelerators and detectors powerful enough to reach the requisite energy and sift the rare gold nuggets from the overwhelming dross. The two novel collider concepts described here are admirably suited for this task. One can expect that gamma–gamma and muon colliders would be copious producers of the putative Higgs boson—to which the standard theory attributes the masses of the quarks and leptons. The new colliders should open up many other exciting new channels for exploring particle physics.

Furthermore, the muon collider has the possibility of putting lepton colliders at the high-energy frontier, where nowadays we have only hadron colliders. Whereas the Large Hadron Collider (LHC) now under construction at CERN will have a proton–proton collision energy of about 14 TeV, the e^+e^- colliders foreseen for the next generation have collision energies of only about 1 TeV. Unlike the pointlike leptons, however, the proton has to share its energy out among its component quarks and gluons. Therefore a lepton collider of a given beam energy provides about the same effective collision energy between point particles as a proton collider with six times its beam energy. So the 3 TeV muon collider now under active consideration would join the LHC at the energy frontier.

Gamma–gamma colliders

The experimental physics. Maxwell's classical electromagnetism is a linear theory. Therefore it does not admit the scattering of light *by light*. But quantum electrody-

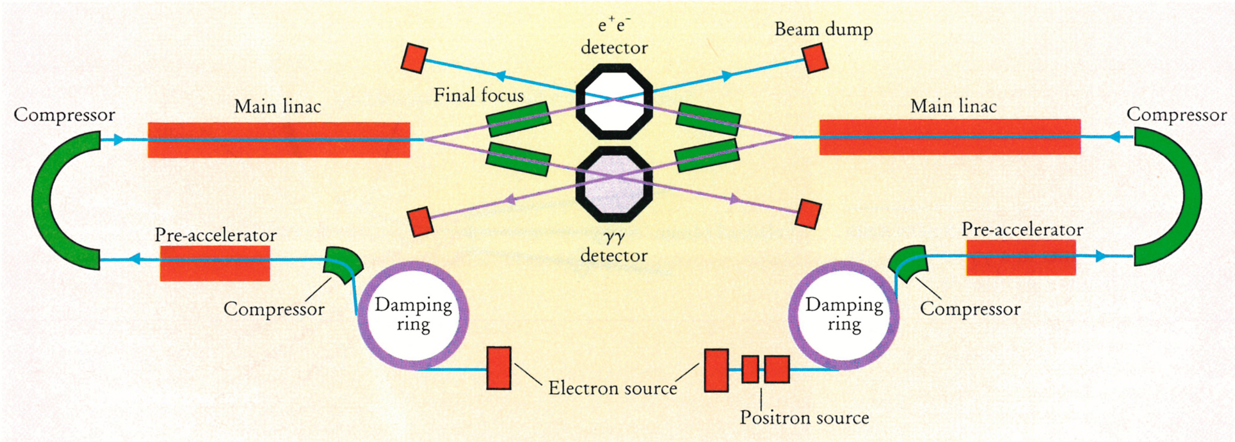

FIGURE 1. TEV ELECTRON-POSITRON LINEAR COLLIDER design with a second interaction region for $\gamma\gamma$ collisions. When running as a $\gamma\gamma$ collider, the machine would accelerate electrons in both of its 10 km linacs. The gamma rays are made by backscattering laser photons off the electron beam near its final focus.

namics (QED) does describe the scattering of photons by photons. (See PHYSICS TODAY, February 1998, page 17.) As the photons become more and more energetic, one eventually gets processes beyond pure QED, such as the photoproduction of hadrons. At that point the theoretical description must incorporate quantum chromodynamics (QCD), the standard theory of the hadronic interactions. And at still higher energies, as yet unexplored, there may well be entirely new phenomena beyond the ken of currently standard theories.

It is in this very high energy range of QCD and beyond that photon–photon colliders would be of considerable interest to particle physicists. Since we are speaking of multi-GeV photons, it is appropriate to call such machines gamma–gamma colliders.

What kind of physics could we explore with a $\gamma\gamma$ collider?[5] The Higgs boson may weigh less than 100 GeV or more than 1 TeV. We've never seen one. The "Higgs mechanism," invoked by the Standard Model of particle physics to break the underlying symmetry between the electromagnetic and weak interactions, may even turn out to be more complicated than just the single neutral boson supposed by the minimal model.

It should be noted that, in a $\gamma\gamma$ collision, the full center-of-mass energy is available for creating a single Higgs boson, whereas, for all practical purposes, an electron–positron collider can produce the Higgs only in pairs. Thus, for a given beam energy, the effective energy reach of a $\gamma\gamma$ collider would significantly exceed that of an e^+e^- collider for doing Higgs physics.

Perhaps the most interesting physics program for a $\gamma\gamma$ collider is the measurement of the decay rate of the Higgs boson into a pair of gammas. This rate is particularly sensitive to new physics beyond the current Standard Model. Its measurement is a sensitive test of various Higgs models that predict the existence of massive new particles too heavy to be produced directly by a first generation $\gamma\gamma$ collider. We would hope to learn about supersymmetric models, which predict a plethora of new heavy particles, the rival Technicolor models and other extensions of the Standard Model.

Another interesting set of reactions is the decay of the Higgs boson to ZZ or pairs of heavy b (for *b*ottom) quarks. A very special opportunity would be the chance to study the *CP* eigenvalue of the Higgs by means of polarized gamma rays. ("*CP*" denotes the combined operations of parity inversion and charge conjugation.) More generally, the ease with which one can polarize gamma beams is a particularly attractive prospect of $\gamma\gamma$ colliders.

If nature is supersymmetric, a $\gamma\gamma$ collider should be able to produce charged superparticles at useful rates. The main source of background would be W^+W^- pairs. If some of the particles in the Standard Model's inventory are, in fact, composites made of more fundamental objects, the composites should have excited states that decay to the ground state by emitting gammas or Z bosons. Or their finite spatial extensions might manifest themselves as anomalous interactions. If, for example, the W boson is composite, it may have an anomalous magnetic moment (or electric quadrupole moment) that could be measured at a $\gamma\gamma$ collider, with its expected copious W production.

In short, then, there are plenty of good particle physics reasons for considering the construction of a $\gamma\gamma$ collider.

The machine. Given all this motivation for constructing a $\gamma\gamma$ collider, we must consider how to make sufficiently energetic and intense gamma beams. The method of choice is Compton backscattering of laser photons from energetic electrons. The first $\gamma\gamma$ collider would most likely be a hybrid facility: a TeV e^+e^- collider with an extra pair of final beam lines for backscattering directed into laser beams near the focus. (See figure 1.)

The electron–photon backscattering cross section is close to the Thompson-scattering cross section, which is $(8/3)\,\pi r_0{}^2$, where $r_0 = 2.8 \times 10^{-13}$ cm is the classical radius of the electron. The momentum of the (effectively very massive) electron is largely transferred to the optical photon, turning it into a very energetic gamma ray. In fact, the top energy of the resulting broad gamma spectrum extends up to about 80% of the incident electron energy.

The intensity of the backscattered gamma beam can be quite high if the incoming laser beam is sufficiently intense. Suppose, not unreasonably, that we could produce one gamma ray for every incident electron. That would require about 10^{19} laser photons for a pulse of 10^{10} electrons. Then, if the gammas are as tightly focused as the electrons, the luminosity of the $\gamma\gamma$ collider would be the same as that of the electron–positron collider.

The incident laser photons must not be too energetic, lest they interact inelastically with the backscattered gammas to produce electron pairs. The criterion for avoiding this undesirable pair production is that incident laser

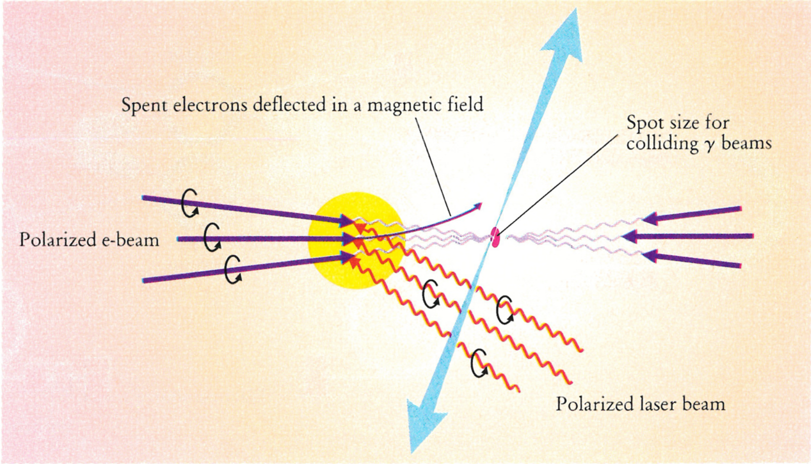

FIGURE 2. CONVERSION AND COLLISION points in a $\gamma\gamma$ collider. At each conversion point (one is shown in yellow) a converging high-energy electron beam runs into a laser beam, backscattering its photons to form a focused gamma-ray beam. Then the two γ beams, from the two electron linacs, meet at the collision point.

photon energy must be less than about $4 \times 10^{-12} E$, where E is the electron beam energy. For a $\gamma\gamma$ collider that starts with a pair of 250 GeV electron beams, this means that the laser photon energy must not exceed 1 eV. A pulse of 10^{19} such photons would have a total energy on the order of a joule. That's quite managable for a pulsed infrared laser.

The high luminosity foreseen for a TeV e^+e^- collider is attributed to the very small spot size at the collision point—typically tens of nanometers. If the conversion point at which the Compton backscattering takes place is too far from the final $\gamma\gamma$ collision point, natural spreading of the gamma beam will render the final collision point too diffuse. On the other hand, the greater the distance between the conversion point and the final $\gamma\gamma$ collision point, the more desirably monochromatic will be the gamma-ray spectrum. The trade-off yields collider designs in which the gammas typically travel less than a centimeter from creation to collision.

Nevertheless, at the backscattering conversion point, the electron beam is about 100 times wider than it would be at the e^+e^- crossing point. Thus one might expect that the gamma rays, at their collision point would have the same large radial spread of the electron beam at the conversion point. If that were so, the luminosity of the $\gamma\gamma$ collider would be very much less than that of the corresponding e^+e^- collider.

But that isn't what happens! The gamma rays are, in fact, focused to the spot size the electron beams would have had at the collision point. This happens automatically, because the electrons have much more momentum than the laser photons. Therefore the gamma produced in a backscattering collision proceeds along the direction of the incident electron. Thus, to make a $\gamma\gamma$ collider, one has only to focus the electron beams, which is what one does anyway in an e^+e^- collider.

The lasers. The incident joule laser pulse would be a few picoseconds long, corresponding to a peak power of a few terrawatts. With a repetition rate of 160 Hz, the average laser would be about 16 kW. To explore the particle physics adequately would require a variety of laser polarization states. Solid-state lasers already yield peak powers that satisfy most of these requirements by means of chirped pulse amplification. Only the average power falls short; it is currently not much more than a few watts.

The two crucial technologies for constructing the requisite high-average-power lasers are high-power diode lasers for pumping and lasing materials that can handle high thermal loading. There are major military and civilian efforts seeking to make more powerful diode lasers. There are also active efforts to develop advanced materials such as athermal glass hosts or new crystals specifically engineered for diode pumping. Therefore it seems quite possible to achieve the requisite average power levels with perhaps a dozen lasers working together. If, on the other hand, the laser pulses could be reused—for example, by storing them in an optical cavity—the average-power requirement could be significantly relaxed.

Another possibility is a free-electron laser. So far, FELs have generated neither the peak power nor the average power one would need for a $\gamma\gamma$ collider. On the other hand, FELs have considerable promise: Chirped pulse amplification could be applied in an induction-linac free electron laser to produce adequate optical pulses.

In devising suitable optics for the intense laser beam one must consider several key factors: Transmissive optics are, for the most part, not feasible; they would absorb too much energy from an intense beam. The overlap between the electron and photon beams must be good; this dictates that the two axes must be closely aligned and that the depth of field be adequately long. Also, because of the high peak and average power, the spent laser beam must be transported to an external dump. Furthermore, one cannot avoid having the used laser beam from one side intersect the optics of the other side. Because two laser pulses will therefore fall on the same mirror, it is necessary to locate the mirrors so that the two pulses will not hit them simultaneously. All of this must be accomplished in the crowded space of masks, tracking chamber and the quadrupole focusing magnets surrounding the collision point. (see figure 2.)

The necessary R&D to make real a $\gamma\gamma$ collider will consist primarily of work on laser development and optical elements. Much of this will be "table top" work. That is to say, it does not require high-energy beams. But actual experience with a $\gamma\gamma$ collider will eventually be necessary for investigation of backgrounds, detector issues and the lifetimes of optical elements near the collision point.

Muon colliders

As difficult as it may be to make intense enough light for a $\gamma\gamma$ collider, the idea of a $\mu^+\mu^-$ collider is even more exotic. How would you make sufficiently intense beams of muons? Muons are short-lived charged leptons. A charged π meson almost always decays into a muon and a neutrino. But even an intense pion beam produces a rather meager and diffuse decay-muon beam. Furthermore, a muon at rest decays in 2 microseconds. Even a TeV muon lives

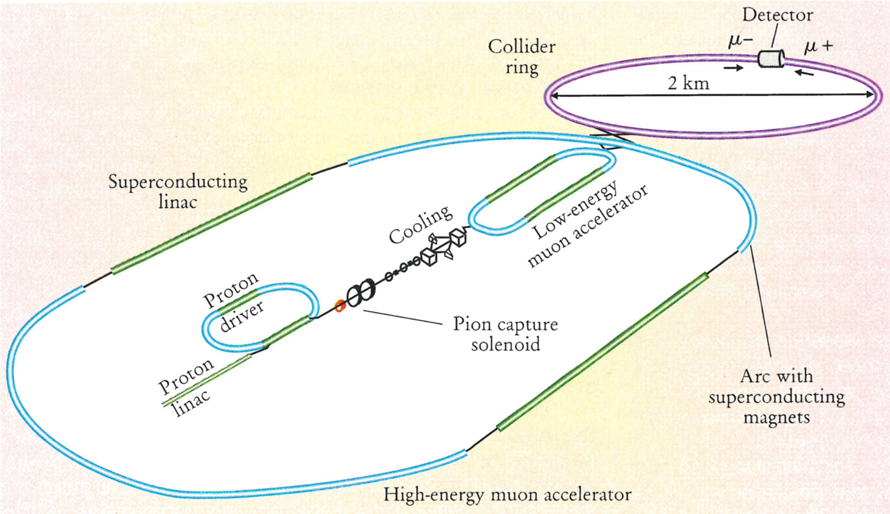

FIGURE 3. A 3 TEV MUON COLLIDER COMPLEX. After the proton driver come sections for making and capturing pions and cooling their decay muons, followed by low- and high-energy muon accelerator rings with linacs and return arcs. Finally, the μ^+ and μ^- bunches countercirculate in the collider ring, meeting again and again inside the detector.

for only a few milliseconds. So everything—capture, cooling, acceleration and collision—must be done in very short order. Clearly, making a $\mu^+\mu^-$ collider will not be easy. So first we must address the question of why one would want to undertake so daunting a task.

First of all, the muon is essentially a very heavy electron. Its larger mass (more than 200 m_e) does away with the synchrotron radiation problem that makes it impractical to build circular electron machines at high energy. A next-generation TeV e^+e^- collider has to be a pair of 10-kilometer linacs. But, for a given radius of curvature, synchrotron radiation loss falls as the inverse fourth power of the mass of the circulating beam particle. That's why the Fermilab Tevatron proton–antiproton collider ring is only two kilometers across, and it's why one can consider building a circular TeV $\mu^+\mu^-$ collider of comparably modest dimensions.

Of course, building a muon collider will require successful solutions of many difficult problems. But the result might well be a machine that is less expensive than an e^+e^- linear collider with the same final energy, though a TeV muon collider would still be a billion-dollar undertaking. Eliminating synchrotron radiation has another important consequence. It makes for a very well defined initial state of the colliding particles, which opens up interesting opportunities for the experimenter.[6,7]

The experimental physics. Perhaps the most important attraction is the possibility of constructing a very-high-energy lepton collider—say 1.5 TeV on 1.5 TeV. The current plans for e^+e^- linear colliders are largely limited to 0.5 TeV beams by available accelerating gradients. A circular machine, like a sling, can of course make do with more modest accelerating gradients, because the beam circulates repeatedly through the structure.

Lepton colliders suffer much less than proton machines from obscuring hadronic debris. Furthermore, a lepton collider of a given beam energy provides the same point-particle collision energy as a proton collider with six times its beam energy. That's important for Higgs physics and other issues at the high-energy frontier.

Like the proposed $\gamma\gamma$ collider, a $\mu^+\mu^-$ collider could form Higgs bosons one at a time. In a TeV e^+e^- collider, by contrast, this process would be negligibly rare, essentially because the coupling of the Higgs to any (pointlike) particle–antiparticle pair is simply proportional to the square of the particle's mass.

The experimenters want to operate a collider with the smallest possible energy spread in the beam. That increases the sought-after signal relative to background, and it provides a well defined initial-state energy. In a muon collider, with its negligible "beamstrahlung" (synchrotron-like radiation where the crossing beams perturb each other electromagnetically) one could hope for $\Delta E/E \approx 10^{-5}$.

A colliding $\mu^+\mu^-$ pair can form a Higgs boson as a resonance in the "s channel," that is to say,

$$\mu^+ + \mu^- \to \mathrm{h} \to \mathrm{X}\overline{\mathrm{X}} \ ,$$

where X can be any lepton or quark that's light enough. (A quark would manifest itself as a collimated "jet" of final-state hadrons.) A first muon collider would, therefore, spend much of its time sitting right on the Higgs resonance. That is to say, its beam energies would be tuned to precisely half the Higgs mass. If m_h turns out to be as little as 100 GeV, then a rather low-energy machine (say 50 GeV on 50 GeV), might be the first muon collider one seeks to build.

A second goal might be to operate the collider near the energy threshold for making a Higgs together with a Z^0. That would allow the determination of many Higgs boson properties, as well as a careful measurements of the top-quark mass. If the Higgs mechanism turns out to involve more than just a single minimal-standard-model Higgs particle, it will be important to measure the quantum numbers, masses and widths of any newly discovered particles. The muon collider would be excellent for studying leptoquarks and supersymmetric particles, if they exist.

A muon collider could also be used as a fixed-target machine, as distinguished from a collider. One would bombard fixed targets not only with the machine's powerful muon beams, but also with the neutrino beams resulting from their decay. Such opportunities far exceed anything we have now, or anything else that's planned in this arena.

The machine. To realize a $\mu^+\mu^-$ collider, we would have to produce muons, cool them, accelerate them and bring them into collision. A possible accelerator complex designed to accomplish all of this is shown in figure 3. Most of the difficulties are consequences of the muon's 2.2-μs lifetime.

A muon beam, created by the decays in a pion beam of the same charge, starts out with a very large phase-space dispersion, or "emittance." So, first off, the muon beam has to be "cooled" to much lower emittance. To get useful $\mu^+\mu^-$ collision rates, we will require very intense intersecting beams. Unfortunately, none of the traditional cooling methods—stochastic, radiation, laser or electron cooling—are fast enough for such ephemeral beam particles. But there is a new method, called ionization cooling, that would seem to meet our needs.

Accelerating the cooled muon beams to the final energy must also be done very quickly. A conventional synchrotron would be too slow, and a full-energy linac would be too expensive. One needs something like a recirculating linac or a very rapidly cycling synchrotron. Similarly, the traditional radio-frequency techniques for injecting and extracting beams are too slow for the short-lived muons.

The muon lifetime limits the number of times the beam can circumnavigate the collider ring to something like 1000. Clearly, one wants the smallest possible collider-ring circumference. That calls for very strong, presumably superconducting, bending magnets.

A muon decays into an electron (or positron) and two neutrinos. As the decay electrons and positrons spiral in toward the collider ring's inner wall, they bombard the outer wall with x rays. The ring's superconducting magnets would have to be shielded against this unwelcome radiation.

The particle detectors would also have to function in this harsh muon-decay environment. It is proposed, therefore, to surround each inner detector with a tungsten cone pointing towards the vertex. That does shield the detector, but a significant flux of background track remains and the cone impinges on the detector's viewing solid angle.

The muon-decay neutrinos can be a health hazard! It is not the neutrinos that hurt; but when they go through matter they produce hadrons that can hurt. At these neutrino energies, the resulting hadron flux reaches equilibrium after a few tens of meters of material. With increasing muon energy, the decay neutrino beam becomes narrower and the neutrino scattering cross section increases. A low-energy "Higgs factory" would pose no hazard. But a 1.5 TeV × 1.5 TeV muon collider would have to be buried under about 250 meters of earth. The 27-km-circumference LHC, by comparison, is more than 300 meters under ground.

To be interesting to experimenters, a muon collider would have to attain not only some specified high energy, but also some specified "luminosity." The luminosity of a collider is defined as its event rate for any process per unit reaction cross section for that process. A 50 GeV × 50 GeV muon collider would have to have a luminosity of 10^{31} events per second per square centimeter of reaction cross section. Because interesting point-particle cross sections fall rapidly with increasing energy, a 1.5 TeV × 1.5 TeV muon collider would need a luminosity of $5 \times 10^{34}/(\mathrm{s\ cm}^2)$. (See the table of machine parameters[8] below).

Working backward through the design complex from the required luminosity at the collision point, one ends up finally with the required machine pulse rate and proton bombardment intensity at the target where the pions are generated. One must consider the polarization of the muons; they are polarized in production. Therefore the final beams can be polarized to some degree by restricting the muon capture solid angle. But that exacts a cost in luminosity.

Ionization cooling, the only thing we know of that's fast enough for muon beams, reduces the emittance by making the beam traverse and ionize some suitable material and then replenishing the lost energy in an RF cavity. One can't do anything like the conventional radiation cooling of electron beams in damping rings, because the circulating muons emit so little synchrotron radiation.

In going through material, however, the muon beam does experience some "heating" as a result of multiple scattering. For hadrons, this multiple scattering would be so severe as to preclude the use of ionization cooling. The muon beam's emittance is balanced between reduction by ionization cooling and increase by multiple scattering. For the most suitable materials, one finds an equilibrium between these two opposing effects when the normalized Lorentz invariant emittance

$$\varepsilon \equiv \beta\gamma(\Delta\theta)(\Delta x) \approx 85\ \pi\ \text{mm-mrad},$$

where the two Δs are, respectively, the rms beam spreads in angle and transverse dimension. That's not as good as one can do with electron beams in damping rings, but it's good enough for a muon collider. Of course, if the equilibrium emittance were lower, we would need fewer muons for the requisite luminosity, and the collider would be much simpler to construct and operate.

Now for some comments on design issues. Starting at the beginning, the proton driver should supply something like 5×10^{13} protons per machine pulse to the pion production target. (See the table at left.) That has already been achieved at the Brookhaven Alternating Gradient Synchrotron, but a muon collider would require shorter bunches and a higher repetition rate (5 Hz) than the AGS provides. The resulting megawatt power on target is comparable to what is being considered for spallation neutron sources.

The capture of the resulting charged pions would require a 20-tesla solenoidal magnetic field. Magnets that strong have been built at the University of Florida's National Magnet Laboratory. Then there has to be a phase-rotation section for reducing the longitudinal spreading of the muons due to their large range of longitudinal velocities. These upstream components of the muon collider complex are shown schematically in figure 4.

The cooling system requires both longitudinal and transverse cooling—by about a factor of 100 in each of

Parameters for three muon colliders. The collision energy is twice the muon energy. The muons are decay products of pions made by protons hitting a target. For a 100 GeV collider, are considered two different beam momentum spreads dp/p. Most of the parameters quoted in the text are those in the second column, with the larger momentum spread.

Collision energy (GeV)	3000	100	100
Proton energy (GeV)	16	16	16
Protons per bunch (10^{13})	2.5	5	5
Number of proton bunches × rep rate (Hz)	4 × 15	2 × 15	2 × 15
Power on proton target (MW)	4	4	4
Muons per bunch (10^{12})	2	4	4
Collider circumference (m)	6000	300	300
dp/p (%)	0.16	0.12	0.01
Emittance (π mm-mrad)	50	85	195
Muon rms bunch length (cm)	0.3	4	9
Muon rms bunch width (μm)	3.2	82	187
Luminosity ($\mathrm{cm}^{-2}\mathrm{sec}^{-1}$)	5×10^{34}	10^{32}	2×10^{31}

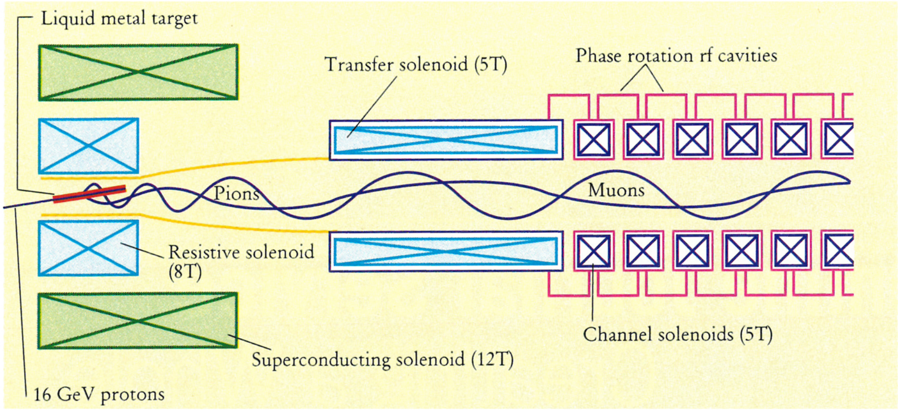

FIGURE 4. MAKING AND CAPTURING PIONS for a muon collider begins, in this schematic, with a 16 Gev proton beam hitting a liquid metal target. A complex of coaxial solenoids (their coils indicated in cross section) captures and directs the charged pions to the decay channel, after which the RF phase-rotation section longitudinally bunches the decay muons.

three directions; that means an overall factor of 10^6, which is quite large. But remember that the Tevatron's antiproton beam is cooled by a factor of 10^9 (albeit by stochastic means that are too slow for our purposes). Proper design of a muon cooling channel has thus far only been partially addressed. One will need lithium lenses at the end of the channel; they can supply both very strong focusing and ionization cooling. There will have to be an experimental facility for testing various ionization-cooling strategies. Such a facility has already been designed.

The accelerating system, on which most of the money would probably be spent, is relatively straightforward. Pulsed magnets in the arc regions of the recirculator would reduce the cost and complexity of multichannel recirculator arcs, but it is not clear whether suitable pulsed magnets can be made.

The collider ring must be made almost isochronous; that is to say, particles of slightly different energy must have almost the same circulation frequency. In that case the muon bunches could be kept short—about 4 cm. There would be severe space-charge problems in the ring, but simulation studies indicate that they could be controlled in an almost-isochronous ring. The ring's high-field superconducting bending magnets, shielded against muon decay products, would be novel; but presumably they can be built. Because the muons would make about 1000 trips around the ring before they decay, the collider's luminosity would be 1000 times what you'd get with the same beam in a single-pass collider. Finally, the detector that records the interesting physics coming from the collisions must be able to operate in the very severe background created by the muon decays. But we expect this background to be less troublesome than what experimenters at the LHC will have to deal with.

In sum, then, a muon collider appears to be feasible, but a great deal of R&D will be needed to determine whether it really is possible. A goodly number of physicists are eager to tackle the subject.

Reprise

A $\gamma\gamma$ collider would seem to be a very natural addition to a TeV e^+e^- linear collider. The additional cost and complexity would be relatively small. A second interaction region, with the capability of doing $e\gamma$ as well as $\gamma\gamma$ experiments would seem to be a small and appropriate investment on top of the billions required for the basic e^+e^- machine.[9]

A $\mu^+\mu^-$ collider is a new concept, with more uncertainty that a $\gamma\gamma$ machine. But it offers the possibility of putting lepton colliders back on the energy frontier. That's essentially because its accelerator and collider will be rings of modest dimension, and therefore presumably less expensive than a comparable linear collider. The greatest uncertainties attach to the production, capture and cooling of the muons, and to the operation of detectors in the severe backgrounds caused by muon decay.

At present, Europe has the lead in electron colliders (LEP), hadron colliders (LHC) and hadron–electron colliders (HERA). Stanford and Japan's High Energy Research Organization (KEK) are jointly working on a TeV e^+e^- collider design, as is DESY. (See PHYSICS TODAY, November 1997, page 21.) Japan and/or Germany seem to be the most likely locations for the next-generation e^+e^- machine. Looking broadly, and also contemplating what the US will do in high-energy physics, one may imagine a $\mu^+\mu^-$ collider in the US, early in the next century.

Acknowledgments

I am indebted to many colleagues who helped in preparing this article. In particular I have drawn heavily upon the work of Robert Palmer and Kwang-Je Kim. My work was supported by the US Department of Energy under contract no. DE-AC03-SF-00098 at the Lawrence Berkeley National Laboratory.

References

1. I. Ginzburg *et al.*, JETP Lett. **34**, 491 (1982). I. Ginzburg *et al.*, Nucl. Instr. and Meth. **205**, 47 (1983).
2. V. Telnov, Nucl. Instr. and Meth. A **294**, 72 (1990).
3. A. N. Skrinsky, V. V. Parkhomchuk, Sov. J. Nucl. Phys. **12**, 3 (1981). V. V. Parkhomchuk, A. N. Skrinsky, in *Proc. International Conf. on High Energy Accelerators*, F. Close, R. Donaldson, eds., Fermilab, Batavia, Ill. (1983), p. 485. D. Neuffer, Particle Accelerators **14**, 75 (1983); in *Proc. International Conf. on High Energy Accelerators*, F. Close, R. Donaldson, eds., Fermilab, Batavia, Ill. (1983), p. 481; in *Advanced Accelerator Concepts*, AIP Conf. Proc. **156**, AIP, New York (1987).
4. D. Neuffer, R. Palmer, in *Proc. European Particle Accelerator Conf.*, V. Suller, C. Petit-Jean-Genaz, eds., World Scientific, Singapore (1994), p. 52.
5. *Proc. Workshop on Gamma–Gamma Colliders*, Nucl. Instr. and Meth. A **356** (1995).
6. *Physics Potential and Development of Muon Colliders, AIP Conf. Proc.* **352**, D. Cline, AIP, New York (1996). *Beam Dynamics and Technological Issues for Muon Colliders,* J. Gallardo, ed., AIP Conf. Proc. **372**, AIP, New York (1996).
7. *Muon Colliders: Proc. Symp. Physics Potential and Development of Muon Colliders*, Nucl. Phys. B (proc. suppl.) **51A** (1996).
8. Muon Collider Collaboration, *Muon Collider Feasibility Study*, BNL-52503, Brookhaven National Laboratory, Upton, New York (1996).
9. NLC Design Group, *Zeroth-order Design Report for the Next Linear Collider*, SLAC Report 474, Stanford Linear Accelerator Center, Menlo Park, California (1996). ■