IEEE Transactions on Nuclear Science, Vol. NS-28, No. 3, June 1981

THE COMMISSIONING AND PERFORMANCE CHARACTERISTICS OF CESR*

B. D. McDaniel, Cornell University, Newman Laboratory, Ithaca, N.Y. 14853

Introduction

The Cornell Electron Storage Ring, CESR[1,2], came into operation for experimental use in September 1979, just two years after authorization of funds for the conversion from the Cornell Electron Synchrotron Facility. This facility was designed to provide colliding electron and positron beams up to center-of-mass energy of 16 GeV. The storage ring is mounted in the same tunnel with the synchrotron which formerly provided electrons up to an energy of 12 GeV for fixed target experiments. The synchrotron itself is used as the injector for the storage ring. CESR was designed to fill in the luminosity gap which lies between the efficient operating ranges of the e^+-e^- machines at SLAC and DESY. The luminosity of SPEAR and DORIS both fall rapidly before reaching 9.5 GeV in the center-of-mass, while the luminosities of the higher energy machines, PEP and PETRA, are unusably low in the same region. By good fortune the optimum operating energy of the CESR facility falls in the upsilon region, that is at about 10 GeV in center-of-mass, the energy required to create the bound and unbound states of $b\bar{b}$ quarks. The study of these states is one of the exciting activities of experimental high energy physics today. As a consequence, we have found ourselves to be in a unique position to study this important energy region.

Description of the Facility

The Wilson Synchrotron Laboratory which houses the synchrotron and CESR facility is located near the center of the Cornell campus and is shown in the cutaway drawing of Fig. 1. The main laboratory building

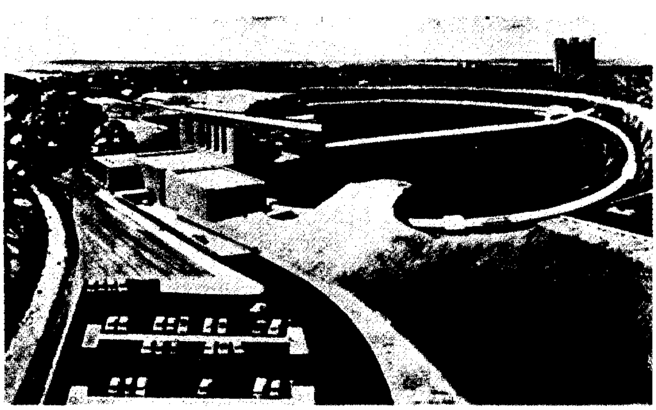

Fig. 1: CESR facility with a cutaway view of the tunnel for the storage ring and synchrotron.

is situated on the side of a deep gorge and the tunnel which houses both the synchrotron and storage ring extends back under an adjacent playing field at a depth of about 50 feet below the surface of the field. The circumference of the tunnel is about a half-mile in length. The storage ring and accelerator pass through a large experimental hall in the main laboratory building called the South Area, and through an underground room called the North Area located diametrically opposite. The intersecting regions of the storage ring are located in these areas. The guide field magnets of the synchrotron and storage ring are mounted side by side in the tunnel as shown in Fig. 2.

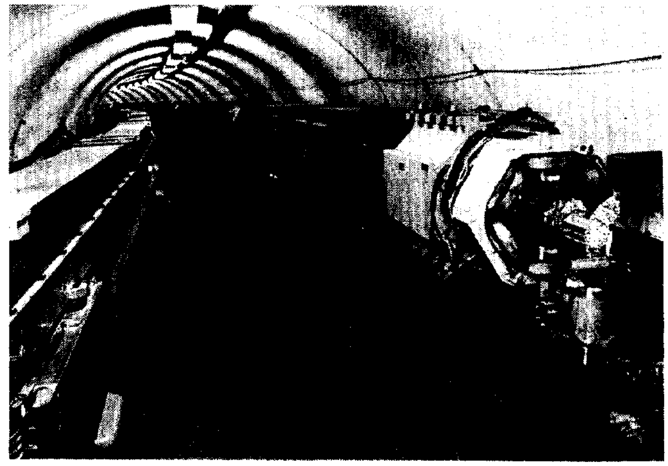

Fig. 2: Interior view of the tunnel with the synchrotron on the left and CESR ring on the right.

The North Area has the dimensions of 9x12 meters and is served by a short tunnel spur which provides space for electronics and by a stairway to the surface where trailers are used for counting rooms and desk space. Fig. 3 shows the layout of the South Area. In

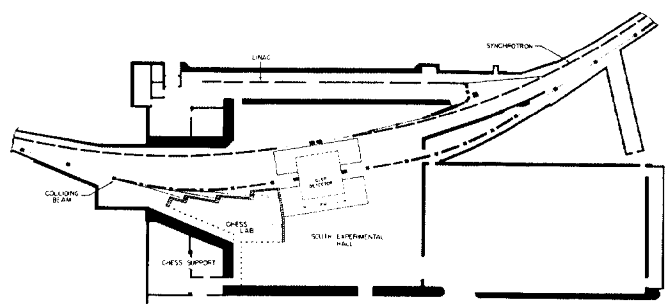

Fig. 3: Plan view of the South Area experimental hall showing the intersection point at the CLEO detector and the synchrotron radiation laboratory, CHESS.

order to provide adequate separation between the synchrotron and the storage ring for the large magnetic detector, CLEO, a bulge is made between the two rings by providing a long drift space on either side of the interaction point. The accelerating cavities of the storage ring are mounted in this drift space. To make up for the lack of bending in this region, high field magnets are required on either side of the interaction area. Immediately adjacent to the interaction, additional very weak field magnets are installed to minimize the effects of synchrotron radiation on the detector. Also shown in the diagram is a synchrotron radiation laboratory called CHESS (Cornell High Energy Synchrotron Source). A total of 3 beam lines, each with several ports, is now set up and operating. This facility is operated by a separate management group of Cornell University.

Table I below gives a few of the design parameters of CESR.

* Supported in part by the National Science Foundation.

0018-9499/81/0600-1984$00.75©1981 IEEE

TABLE I

Maximum design energy, GeV	8
Design luminosity, at 8 GeV, $cm^{-2}sec^{-1}$	$1x10^{32}$
at 5 GeV, $cm^{-2}sec^{-1}$	$4x10^{31}$
Number of interaction points	2
Bunches per beam	1
Particles per bunch at 8 GeV	$1.5x10^{12}$
Current per beam at 8 GeV, mA	100
Bending radius, normal dipoles, m	89
high dipoles, m	32
weak I.R. dipoles, m	180
Magnet power at 8 GeV, MW	1
R.F. frequency, MHz	500
R.F. power, 8 GeV, full current, MW	2
Cost, $ millions	13

Special Features

Positron Vernier Coalescence Scheme

In order to decrease the filling time for positrons, a total of 60 buckets in the synchrotron are populated by the linac and transferred by single turn extraction to the storage ring. A system of vernier phase compression is used to coalesce the multiple bunches of stored positrons into a single bunch. This is accomplished by extracting the bunches from the storage ring one at a time and recirculating them in the synchrotron for a suitable number of turns before returning them to the chosen r.f. bucket in the storage ring. The orbital length of the synchrotron and storage ring differ by one r.f. bucket separation so that in this way, after reinjection into the storage ring, a given bunch may be advanced in azimuth position relative to the other bunches in the storage ring. In this manner all bunches may be transferred, one at a time, into a common bucket. Because of inefficiencies in the transfer, the actual maximum coalescence factor which has been achieved is approximately 20 of a possible 60, that is, an efficiency of 33 percent. After the storage ring is filled with positrons, single electron bunches are accelerated in the synchrotron on successive cycles and deposited in a single bucket of the storage ring until it is filled.

Lattice Restrictions

The guide field lattice of the storage ring is subject to two important restrictions. One is the restriction of the orbital length imposed by the vernier coalescence scheme and the other one is the restriction imposed by the confines of the already existing synchrotron tunnel.

Because we wished to have a large separation of the synchrotron and storage ring orbits at the intersection point in the South Area experimental hall for the large detector CLEO, it was necessary to make a bulge in the orbit at that point. Since a similar bulge could not be made in the North Area without excessive cost, the ring thus has only one axis of symmetry. It is further made very irregular by the transfer lines and various straight sections distributed throughout the orbit in order to provide for beam handling components such as the fast kicker and orbit bumps. As a result this lattice is perhaps the most irregular of any currently operating machines.

Individual Control of Quadrupole Magnets and Correction Elements

Because of the lattice restrictions referred to above, and because of the desire to maximize the flexibility of the lattice, it was decided to provide independent control of all of the quadrupole elements As a consequence, each quadrupole magnet is driven by an independent chopper-type power supply operating from a common d.c. power bus and controlled by computer. In order to match the high impedance of the chopper supplies, the excitation coils for the magnets were strip wound with water-cooled copper plates providing edge cooling for the coils. The stability specification for each of the supplies was set at one part per 10,000. Similar power supplies are used to drive the dipole, sextupole, and octupole elements, as well as the elements in the various transfer lines.

Performance

On April 1, 1979, just 18 months after initial funding, the first turns of circulating electrons in CESR were obtained. A few days later, after locating and correcting the usual number of misconnected quadrupoles, we obtained the first stored beam with a life time of several seconds. We were pleased to find that the whole system was very manageable, that the independent quadrupole power supplies worked adequately reliably, and that we could in fact store the beam in the very irregular lattice. At that time the positron injection system was not yet complete, and the separator plates had not yet been installed. During the next few months these components were installed and the behavior of the single beams of electrons was studied and improved. By the end of June we had stored positrons, and on August 18, 1979, the first luminosity was detected. By adjusting the chromaticity and focusing, (β^*), in the interaction region, the peak luminosity was quickly raised to more than $10^{30}cm^{-2}sec^{-1}$. At this point the experimental program was initiated. By December 11, 1979, we had observed the ground state and first excited state of the upsilon and verified the existence of the second excited state.[3,4] By mid-January 1980 we had discovered the third excited state of the upsilon[5,6] and showed that its width was greater than that due to the energy width of the beam. This therefore implied we had observed the first unbound state of the upsilon which decays with the emission of B mesons. Our experimental program now has been under way for about 18 months.

In Fig. 4 we see a histogram of the monthly integrated luminosity which has been provided for high energy physics runs. This plot includes periods

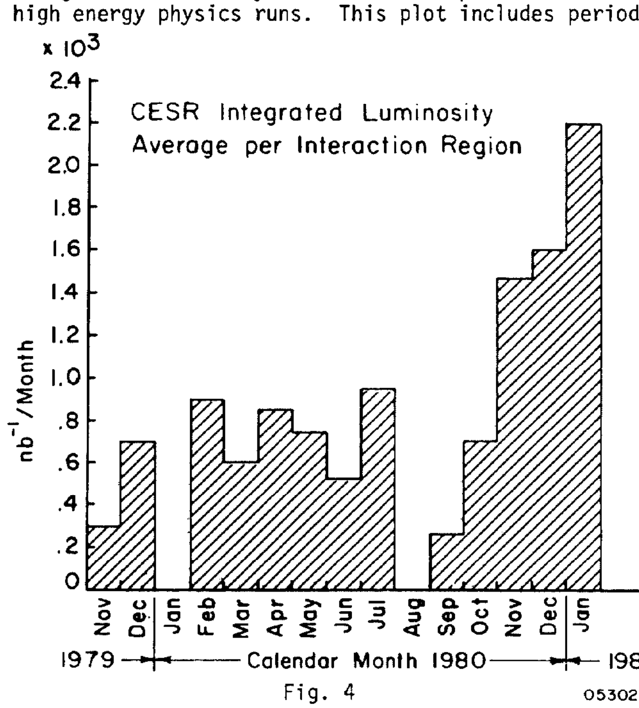

Fig. 4

0530281

of major shutdown for maintenance and installation of experimental equipment, as well as the 25 percent of time devoted to machine studies. As you can see from the graph, the integrated monthly luminosity has increased from a few hundred to 2200 nb^{-1}. This latter rate corresponds to a daily average of about 100 nb^{-1} per day during periods of high energy physics running or an average luminosity of about $1.2x10^{30}cm^{-2}sec^{-1}$ averaging over fill times and interruptions of the operation for maintenance of the accelerator or the experimental equipment. The maximum initial luminosity which has been observed is $3x10^{30}$ $cm^{-2}sec^{-1}$. The luminosity lifetime is typically 4 hours and the filling time between running cycles varies from 45 to 75 minutes. Because of the great interest in the upsilon physics the range of operating energies during the past year has been limited to the interval between 4 and 6 GeV per beam. All of this operation has been accomplished using only one klystron driving a single R.F. cavity. Only during the last month has a second cavity been installed in the facility. We have not yet attempted to operate at any energy in excess of 6 GeV per beam.

The whole synchrotron and storage ring facility is controlled by a net of three PDP 11-34 units backed up by access to a PDP-10. All controls are available at a single console and the standard mode of operation requires only one operator of the senior technician level.

Operational Sequence

We have been operating the storage ring using the following sequence while filling. Prior to each fill, the magnets are run through a hysteresis cycle to restore them to a standard field condition corresponding to an energy near to the desired operating energy. Positrons are then injected into the storage ring according to the vernier coalescence scheme described earlier. When the coalescence system and synchrotron are well tuned up, we can fill at the rate of 2 mA per minute and complete the fill using 6 coalescence cycles. The separator plates are then turned on, and electrons are injected. During the positron and electron filling period, the lattice is adjusted to have a relatively large beta function at the interaction region in order to make it easier to fill and store. After filling with both positrons and electrons, the lattice tune is maintained while the energy of the stored beam is adjusted to the desired operating energy. To accomplish this, the computer-controlled power supplies gently adjust all of the magnet currents in unison to the desired values. After arrival at the operating energy, the proper low beta lattice is then called. At this point the separator plates are turned off and luminosity is obtained. During the entire process great care is taken to monitor and maintain the betatron and synchrotron oscillation frequencies at the proper value to avoid crossing through a resonance. One of the greatest difficulties which we have, when trying to store large beams, comes from the sudden and sometimes inexplicable loss of either the positron or electron beam, or both, at some point in the filling process. We have taken considerable time to study these effects and to map out the disastrous canyons in the tune plane. As we have become more experienced and better understand the peculiarities of the tune plot, our efficiency has has greatly improved.

Results of Accelerator Studies

Single Beam Studies

1. Maximum current. We have stored electron beams up to 36 mA and positrons up to 27 mA at energies of 5.5 GeV. There seems to be no real limitation to the current that can be stored except for the increase in the pressure in the cavity which trips the r.f. power. It is believed that this effect is due to heating of the cavity walls due to higher mode losses. Certain portions of the wall cannot be outgassed by the normal processing mode because such processing only produces the normal mode fields. The present limit is about 60 percent of the design current at this energy, but is well beyond the maximum beam current which can currently be stored with two colliding beams; hence, it is not presently a limitation to the luminosity. At these high currents, the average vacuum deteriorates by about a factor of 10 from about 10^{-8} to 10^{-7} Torr. The single beam lifetime correspondingly is reduced from 5 to 2 hours. We have not felt it necessary thus far to activate our provisions for baking out the chambers.

2. Betatron frequency as a function of stored current. We find that there is a significant variation of the betatron frequency with current and that the dependence is different for electrons and positrons. The frequency difference persists even at zero beam current limit. The variation with beam current for vertical oscillations is $-4x10^{-4}$ units of Q per mA. The effect is much weaker for horizontal oscillations. The e^+-e^- difference at zero current is about 10^{-3} units of Q. It is suspected that the reason for the frequency difference at zero beam current may be due to using only one cavity in an asymmetric position and that this effect will disappear when two balanced cavities are used.

3. Aperture measurements. Measurements of the working aperture of the machine have been made using both a beam pinger and a resonant shaker. These measurements indicate that the aperture in both horizontal and vertical are approximately 80 percent of the design values.

4. Verification of lattice parameters. The independent quadrupole controls make it very easy to measure the beta function of a given lattice. By varying individual quadrupole strengths and observing the corresponding tune shift which is observed, it is a matter of just a few minutes to measure the beta function around the whole machine, including the measurements of β^* at the interaction region. After the absolute strengths of the quadrupole system were well calibrated we found that a given lattice configuration could be accurately predicted.

5. Chromaticity. Using a two-family sextupole configuration, we have studied the head tail instability for the haromonics 1, 2 and 3, and have measured the damping time as a function of current. If the chromaticity is too great, the beam is destroyed. Though we have the flexibility to set up a more complex set of sextupole corrections, these have not been studied.

6. Dispersion. Since we wish to control η_v^* in the interaction regions, we have studied the behavior of η_v as a function of a number of variables by looking at the orbit distortions produced by a variation in the cavity frequency. We found a significant dependance on beam position in the accelerating cavity, presumably due to higher mode excitation. Skew quadrupoles also have a strong effect on the value of η_v and provide some measure of control.

7. Higher mode losses. We have studied higher mode excitation losses by measuring the power input to the cavity as a function of beam current at 4.6 and 5.5 GeV beam energies. From this data we can obtain a figure for the total higher order mode loss. We have calculated independently the higher order mode loss in the R.F. cavity and have deduced the losses in the separator plates and vacuum chamber walls from wire

impulse measurements made before assembly. We have found fortuitously good agreement between these calculations and the measurements using the stored beam. Approximately two thirds of the losses are in the cavity and the remainder is in the vacuum system and separater plates. The total obtained from the beam measurement gives the value 6.2 V/pc for the loss parameter.

8. Synchro-betatron resonances. We have studied the various synchro-betatron resonances as a function of beam current. We find a broadening of these resonances with an increase in beam current. We suspect that the space for operating in our lattice is somewhat restricted by the presence of a number of lines and stop bands which are less apparent in machines which have a higher degree of symmetry.

9. Octupole suppression of bunch-to-bunch coupling. While filling with positrons, we store the beam in 60 uniformly spaced buckets before coalescing. In this mode we find that as the charge is built up there is a threshold for a "frying" sort of oscillation that builds up due to an interaction between the filled buckets. By applying a weak octupole correction we are able to raise the threshold for such frying by a large amount, i.e. to a level such that the limitation to the storage of positrons at this part of the filling cycle is no longer limited by this effect but is limited by the ability to recirculate the large bunches through the synchrotron without loss during the coalescence process.

10. Tune stability. Because each of the quadrupoles is controlled individually, one may be concerned that the aggregate effects of their instabilities would cause a serious wandering of the betatron tune. We have observed a wandering in the operating point amounting to an rms variation in Q of about 3×10^{-4}. This is somewhat greater than would be expected due to the normal rms variation of the individual supplies and is probably due to nonstatistical fluctuations in a few individual supplies. We have found variations of this magnitude to be acceptable.

Beam-Beam Studies

1. Measurement of luminosity as a function of beam current. In Fig. 5 we see the general behavior of the luminosity as a function of beam current. Though the exact operating point will affect the luminosity curve, varying the specific luminosity and other parameters, the general characteristics are as shown. For small currents, the luminosity increases linearly as the square of the beam current, for equal beams of positrons and electrons. As the current increases, the luminosity curve breaks away from the square law dependence and begins to increase at a more linear rate. The calculated tune shift as a function of beam current shows a corresponding saturation type behavior at a value of about 0.025. We have made an attempt to parameterize the dependence and find some agreement with the following model.

At low currents, the vertical beam size at the crossing point, σ_v^* increases. If we write $(\sigma_v^*)^2 = (\sigma_{vo}^*)^2 + (gI)^2$ and use this in the standard tune shift formulas, we can obtain a satisfactory fit of the data for the luminosity versus current curve. σ_{vo}^* is the beam height at zero current. The constant g can be empirically obtained by fitting the curves.

We find that at a certain beam current, the beams become unstable and are lost. We have found empirically that the peak luminosity obtainable varies approximately as E^4 and the maximum attainable calculated tune shift and peak beam current increase with energy, while the vertical beam size at maximum luminosity decreases with energy. Attempts to increase the vertical beam size by the use of skew quadrupoles do give some increase in attainable luminosity, but excessive amounts of correction result in loss of the beams.

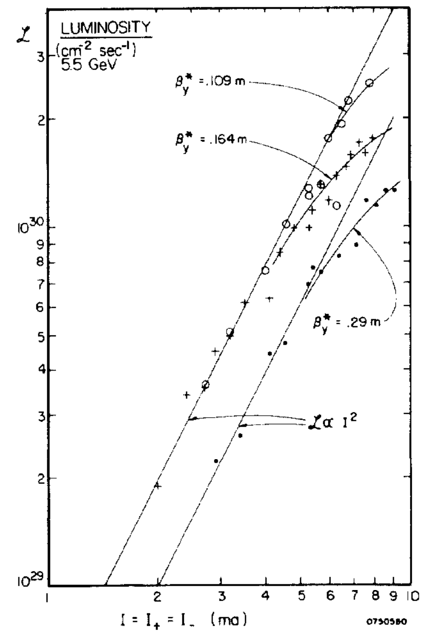

Fig. 5: Log-log plot of CESR luminosity as a function of beam current for some typical cases.

2. Luminosity limits. Our luminosity is limited by our maximum attainable tune shift, 0.025. This is similar to the values that have been obtained for maximum tune shift both at PETRA and PEP. The three machines (PEP, PETRA and CESR) were all designed on the assumption that a tune shift of 0.05 could be achieved. This is twice the value that has actually been obtained. In our machine this has caused a reduction in luminosity by about a factor of 2. Instability of the beam also limits the maximum beam current to about 10 mA. This is about a factor of 6 lower than the design current for the machine at our operating energy. This limitation in current is not well understood; however, recent work by Talman and Peggs may be shedding some light on this question. This will be discussed below. The combined effect of these two factors is a reduction of luminosity by a factor of 12 below the design luminosity.

In order to increase the luminosity without increasing the tune shift, we have attempted to increase the emittance of the beam and add more current. The attempts to improve the luminosity in this way have yielded no significant improvement.

As a result of these difficulties, all of the machines that were to have produced luminosities of about 10^{32} $cm^{-2}sec^{-1}$ are now operating at luminosities lower by a factor of 10 to 20. To compensate for this loss in luminosity PETRA has recently installed a "low beta" intersection region which lowers the value of β^* at the intersection region. This has accomplished the desired result of increasing the luminosity at the same beam current and leads to the expectation of an overall improvement in the luminosity of more than a factor of two. PEP is now installing similar insertion units for the same purpose. We, too, are planning the installation of low beta sections.

3. Talman-Peggs "Swamp" Diagrams. In order to better understand the beam-beam interaction problems R. Talman and S. Peggs of our group, undertook a theoretical analysis of this problem. They have generated

a self-consistent beam-beam model, and have made extensive calculations to find the regions of the tune plane which would produce the maximum luminosity under stable operating conditions.[7] This model assumes that as long as the horizontal width of the beam at the interaction point is large compared to the vertical size of the beam, the luminosity is determined almost entirely by the parametric amplification of vertical oscillations by the horizontal betatron oscillations. Their analysis makes use of a beam tracking program running through several thousand turns of the machine.

The initial calculations by Talman and Peggs were made with η^*, the dispersion at the intersection region, set equal to zero and they obtained a contour map of achievable luminosity for many points in the tune plane. This so-called "swamp diagram" is shown in Fig. 6.

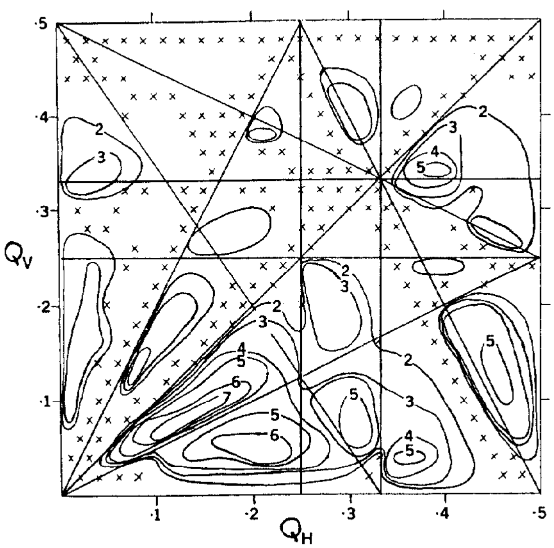

Fig. 6: Talman-Peggs "Swamp" diagram. A plot of relative luminosity contours in the tune plane plotted against fractional tunes of Q_V and Q_H for dispersion, η^*, equal to zero at the interaction point. The crosses indicate bad lifetime regions.

It was expected that this diagram would give the general characteristics of regions of luminosity even with η^* not equal to zero. Certainly the general features, including stop bands and resonances, are correct. We compared attainable luminosities in CESR for the various operating points and found that the results were in approximate agreement with the calculations. On the basis of a prediction of the calculation, we then moved our standard operating point to a new position in the tune plane and found an increase in the maximum luminosity by nearly a factor of two, as well as a very significant improvement in stability.

The normal integral tune values used in CESR in the past have been $Q_H = Q_V = 9$. Because of the geometry of CESR, the strong dipoles near the interaction region make it difficult to achieve low values of η^*. By increasing the horizontal integer tune to $Q_H = 12$, the emittance can be lowered and this permits η^* to be lowered more nearly to zero. We have found that we can operate in this mode and are planning to study this region of the plane. Peggs has recently discovered that in some regions, the stability changes rapidly with small changes in η^* near zero values of the dispersion. He has now calculated new diagrams for higher values of η^* which we will use in our future attempts to improve the luminosity.

Future Improvements

1. Low Beta Insertion for the Interaction Region

We are currently studying the design of low beta sections for each of our interaction areas. In the North Area, where the Columbia-Stony Brook detector is located, it appears rather easy to provide the insertion, since their apparatus does not require a long free space between interaction quadrupoles. It appears that the free length can be reduced to less than 4 meters, and that a decrease in β^* of at least a factor of 3 can be obtained. Thus we can hope for an increase in luminosity at that interaction area amounting to a factor of 3.

In the South Area where the CLEO detector is located, the design of a satisfactory low beta section is much more difficult. This is because the CLEO detector is very large and contains a solenoid of 2-meter length which requires additional space for compensating coils. In order to make the low beta section, a quadrupole must be moved close into the interaction area. It appears difficult to do this without removing the compensating solenoids or moving them to a point outside of the nearest quadrupoles. We are investigating the possibility of eliminating the compensating solenoids and using skew quadrupoles to correct for the cross coupling induced by the main solenoid of the detector. A number of additional constraints are imposed on the system by the separator plates, luminosity monitor and vacuum pumps. We do not yet have an acceptable solution to this problem, but believe that one can be found.

2. Increase in Energy

We have thus far operated with one cavity to about 6 GeV. We have recently installed a second cavity so that with each cavity operating with one tube we should easily reach 7 GeV. Arrangements already exist for the addition of two more power tubes. With this addition we will be able to operate at 8 GeV. Because of the great interest in operating in the upsilon region we have not pressed this program as fast as we otherwise would have done.

3. Positron Accumulator

In order to improve the filling rate for positrons, we are currently working on a plan to use a small intermediate storage ring which will allow us to dispense with the vernier coalescence scheme and permit rapid injection of single bunches of positrons into the storage ring. If this can be done, we can save much time in the injection process, and improve reliability. We should also be able to use the "topping off" mode of injection which will further improve the average luminosity.

REFERENCES

1. M. Tigner, IEEE Trans. on Nuclear Science NS-24, 1849 (1977).
2. Design Report, Cornell Electron Storage Ring, CLNS-360, April 1977, Floyd R. Newman Laboratory of Nuclear Studies, Cornell Univ., Ithaca, NY 14853.
3. D. Andrews et al., Phys. Rev. Lett. 44, 1108 (1980).
4. T. Bohringer et al., Phys. Rev. Lett. 44, 1111 (1980)
5. D. Andrews et al., Phys. Rev. Lett. 45, 219 (1980).
6. G. Finocchiaro, Phys. Rev. Lett. 45, 222 (1980).
7. S. Peggs and R. Talman, CLNS-80/463, July 1980, Floyd R. Newman Laboratory of Nuclear Studies, Ithaca, NY 14853.