Phase TV: Wavefront Analysis of an Optical Resonator

1988 S.U.R.F. Report

by Brian E. Lemoff

Sponsor: Dr. Ronald W.P. Drever

Abstract

In the Caltech gravity wave detector, a laser beam resonates in two high finesse forty meter optical cavities, as well as a smaller, lower finesse mode cleaning cavity. To diagnose errors in allignment of the mirrors, as well as poor mode matching of the incident beam, it is useful to look at the phase front of the back reflected spot. My project was to design and build a "Phase Camera" which would scan the spot, and produce pictures of both phase and intensity on a color video display.

LIGO-T880053-00-R

I. How the Gravity Wave Detector Works

Before discussing the "Phase Camera", we must take a brief look at how the Caltech gravity wave detector works. This is a necessary first step for two reasons. First of all, we must understand the gravity wave detector in order to see why such a phase mapping device is necessary. Secondly, the method by which the "Phase Camera" senses the phase of the light can only be understood once the basic operation of the gravity wave detector has been explained.

In its current configuration, the Caltech gravity wave detector consists of three Fabry-Perot optical cavities. The first cavity, about 1 m in length, is known as the *mode cleaner* has a storage time of about 25 μs, or a bandwidth of 40 kHZ. This cavity serves to stabilize the laser frequency (An argon laser with $\lambda \approx 514$ nm is used) as well as to strip off all but the TEM_{00} mode from the light. The light transmitted through the mode cleaner is inserted into the two larger Fabry-Perot cavities. Each of these cavities is 40 m long and has a storage time of about 1 ms, or a bandwith of 1 kHz.

The two 40 m cavities are oriented at 90 deg to one another, and share the same input light. The phase of the input light is servoed such that one of the arms, call it the *primary* arm is in resonance. The other arm, call it the *secondary* arm, has its length servoed such that it too remains in resonance. The voltage required to keep the secondary arm resonating with the same light that has its phase determined by the first arm is therefore an indication, neglecting DC and low frequency components, of the relative changes in length between the primary and secondary arms.

The three cavities described above all differ as to where the phase correction signal is applied so as to keep them in resonance. For the mode cleaner, the correction is applied to the phase (and hence to the frequency which is the time derivative of phase) of the light emanating from the laser. For the primary arm, the correction is applied to the phase of the light emanating from the mode cleaner. The secondary arm has its phase correction signal applied directly to the length of the cavity. In reality, the primary arm also has a low frequency length correction. This is for stability, and no gravity waves in this frequency domain can be detected.

While the correction signals for the three cavities may be applied in many different ways, they are all derived in exactly the same way. Assume for the moment that the cavity is resonant. If we look at the light coming back from the cavity, it is the superposition of the part of the input beam reflected off the cavity and the part of the resonating light leaking back out of the cavity. The phase of the back reflected light is opposite that of the back leaking light, and thus the two beams interfere destructively. Now, suppose we introduce a 12 MHz modulation into the phase of the incident beam. Because the modulation frequency is orders of magnitude above the bandwidth of the cavity, the back leaking light has no phase modulation (In fact, there is a small modulation, but it is attenuated by a factor which goes like the bandwidth divided by the modulation frequency.). The phase of the back reflected light is of course still modulated. Thus, the interference between the two beams introduces an amplitude modulation of the light intensity.

The intensity of the interfered light depends only on the absolute value of the relative phase between the beams. Thus, if the phase modulation is centered exactly on minus the phase of the resonant light, an amplitude modulation of twice the frequency of the phase modulation, or 24 MHz, is acheived. If the phase modulation is slightly off center, then there will also be a small 12 MHz component in the resulting amplitude modulation. The direction in which the phase modulation is off center determines the sign of the phase of the 12 MHz amplitude modulation with respect to the original phase modulation.

Thus, we have a method for deriving the phase correction signal necessary to keep the cavity in resonance. We apply a 12 MHz phase modulation to the input beam. The superposition of the back reflected and back leaking light is looked at by a photodiode. The voltage coming out of the photodiode has a small 12 MHz component due to the amplitude modulation produced by the interference. The phase of this component relative to the phase modulation signal is determined by mixing the two signals and then averaging. The resulting voltage indicates the overall error in the phase of the incident light (or the error in cavity length, which is equivalent). The appropriate correction signal can then be applied.

The object of any detector is to achieve the maximum possible signal to noise ratio. In the case of the gravity wave detector, one would like the amplitude modulation produced by the interference to be as large as possible. One problem which is often encountered results from errors in mirror allignment. Suppose that one of the mirrors in the cavity has a small tilt from vertical. In this case, the phase front of the back leaking light will also be tilted. The photodiode used in deriving the correction signal spatially averages over the entire cross section of the spot. Thus, if the back leaking spot has a tilted phase front, the top of the spot may contribute a 12 MHz amplitude modulation opposite in phase from the bottom of the spot. As long as there is some net 12 MHz signal, the cavity can be kept resonant, but with different parts of the spot cancelling each other, the signal to noise ratio is reduced.

Tilted mirrors are only one example of ways in which the phase front of the back leaking light can be distorted. Other distortions can come from physical distortions in the mirrors and from errors in mode matching, i.e. if the waist and angle of divergence of the input beam are not properly matched to the mirror curvature and cavity length. All of these distortions reduce the sensitivity of the gravity wave detector. It is therefore desirable to look at the phase variation over the surface of the spot so that one can diagnose and correct the sources of phase distortion.

II. The Phase Camera

In order to get some idea of how the phase of the light varied over the surface of the spot, Dr. Ron Drever came up with an idea that he called the "Phase Camera". The basic idea of the "Phase Camera" is quite simple. In the previous section, it was explained that the 12 MHz amplitude modulation of the laser spot is proportional to the relative phase between the input beam and the light resonating in the cavity. The signal output from the servo photodiode is the average of this signal over the spot. This average is of course weighted by the intensity of each part of the spot. Thus, if we could look at the individual elements of area which make up the surface of the spot, we would be able to determine the way in which the convolution of intensity and phase varied over the spot surface.

In the "Phase Camera", the laser spot is scanned across an aperture much smaller than the spot. Behind this aperture is a photodiode, tuned to 12 MHz. As the spot is scanned, e.g. fast horizontally and slow vertically, the amplitude modulation of the area element passing through the aperture is detected. By mixing this signal with the 12 MHz phase modulation signal and then averaging (This procedure is known as *demodulation*) we derive a voltage proportional to the phase and intensity of the area element. This voltage can then be displayed on some type of output device, e.g oscilloscope, chart recorder, video screen. If the horizontal and vertical scanning of the output device is synchronized to the laser scan, then a picture representing the product of phase and intensity versus spatial position is displayed. If we also know the intensity profile of the spot (Ideally, a Gaussian, but the actual intensity profile can be measured using the same scanner and display but looking at the low frequency part of the photodiode output.) we can determine the phase profile of the spot.

Several months ago, Dr. Drever and Dr. Yekta Gursel quickly put together a very crude version of the "Phase Camera". As their scanner, they used a single mirror whose vertical and horizontal orientation were determined by the voltages to two small audio speakers. The demodulated signal was displayed on a storage scope. The horizontal sweep on the scope was synchronized to the horizontal scan of the mirror, and the voltage input to the scope was the sum of the voltage driving the vertical motion of the mirror and the demodulated phase signal. About ten sweeps were accumulated on the scope display, going from the top to the bottom of the screen. The resulting pictures were interesting because they showed that the phase was in fact not constant over the surface of the spot; however, the crudity of the apparatus seriously limited the usefulness of the pictures. First of all, the scanner used was not designed for linear scans, and hence the resulting pictures were spatially quite distorted. Secondly, the fact that only 10 traces were displayed made the resolution quite poor. Finally, the fact that the time taken to accumulate the scans was quite long meant that features in the display could be the result of either spatial variation or time variation.

III. Phase T.V.

My project this summer was to develop a much more sophisticated version of the "Phase Camera" which would be quick enough and accurate enough to be used as a reference while aligning the mirrors in the optical cavities or while adjusting mode matching optics. There are three major components of my system, known collectively as "Phase T.V.". The heart of the system is an Everex System 1800 microcomputer (PC/AT compatible). The computer is responsible for controlling the scanner, accumulating the data, and displaying it on the screen, as well as other data processing features to be described later. The second major component is a NIM standard equipment rack. This rack contains an assortment of high and low power amplifiers as well as a 12 MHz phase shifter and an RF mixer/demodulator. The third major component is the scanning head, which is the only part of the system that actually needs to be in the path of the laser beam.

The scanning head is mounted on a portable aluminum base so as to allow for mobility while ensuring that the allignment of the various components of the scanner are not misalligned. The actual scanning is done by two mirrors, one scanning vertically and one horizontally. Each mirror is mounted on the shaft of a limited rotation galvanometer motor. Each motor is controlled by the computer, via a DAC and a power amplifier. The horizontal scanning mirror oscillates at a frequency of about 70 Hz, while the vertical scanning mirror is driven by a .5 Hz sawtooth. The scanning speed is limited by the A/D conversion rate and not by the mechanical make up of the scanner. In addition to the the two mirror/motor systems, the scanning head also includes an adjustable flat mirror for purposes of beam allignment and a photodiode/amplifier system, responsible for converting light signals to voltage signals. The active area of the photodiode is shuttered by a small pinhole whose diameter can be changed depending on the desired resolution (The smaller the pinhole, the better the resolution and the worse the signal to noise ratio).

Two different signals are output from the photodiode amplifier: a low frequency signal, corresponding to the average intensity of an area element of the spot; and an RF signal, containing the 12 MHz signal resulting from the amplitude modulation of the spot. The low frequency signal goes straight to an amplifier and an A/D converter. The RF signal goes first to a demodulator, to extract the 12 MHz component, and then to an amplifier and an A/D. The computer can read either one or both of these signals, depending on the desired display. The data is stored in a 100 × 100 array, with the indices of the array indicating the vertical and horizontal components of the area element. After each new 100 × 100 array is filled, the data is displayed.

The display uses a VGA card and a NEC MultiSync monitor. During scanning, there are two basic types of display. The first is the two dimensional color plot. In this mode, the voltages of each of the array elements are converted to colors, the highest being violet, and the lowest being red, with the others in order of the rainbow. The data is then plotted on the screen in a 100 pixel × 100 pixel display of 256 possible colors. After each new vertical scan is completed, the display is updated with the new data (unless averaging is enabled). The updates occur about once every five or six seconds. It is possible to have up to three 2-d color plots on the screen at once. The other type of rapid update display is a 14 color, three dimensional contour plot. This display is like a rubber sheet display, except that there is no provision for eliminating hidden lines. The x and y values correspond to the spatial coordinates of the area element, and the z value corresponds to the intensity or phase voltage (depending on which is being displayed). The angle of declension of the xy plane is adjustable, so that the plot can be viewed from any vertical angle. Points plotted above the xy plane are blue, and points below the plane are red. There are 7 intensity levels of red and blue, corresponding to the magnitude of the z coordinate of the point. The colors are redundant on this display, however they make it easier to identify the zero crossing of the phase data. The resolution of these plots is adjustable, and the speed of update varies accordingly. Only one 3-d plot can be on the screen at a time.

There is one other method of displaying the data. This is an optional color 3-dimensional hidden line plot. This is not used as a live update display during scanning because it takes about a minute to plot, however it is probably the most informative type of display. This display is a 100×100 perspective rubber sheet plot which clearly shows hills, valleys, ridges, and their relative heights.

The program which runs the scanner and displays the data also has provisions for saving and recalling files. Any 100×100 array can be saved and recalled at a later time. Archived data can then be displayed and compared to current data. For example, suppose that the gravity wave detector is showing lower sensitivity than it had on a previous day. By comparing the phase fronts of the light coming out of each of the three cavities to what they were on a day when the sensitivity was better, it can quickly be determined if the allignment of one of these cavities is at least partially responsible for the decrease in sensitivity.

IV. Conclusion

The "Phase T.V." system has already proven to be quite useful. The first pictures accumulated of the mode cleaner reflected light indicated both allignment and mode matching errors. The allignment is easily adjusted, and it was possible to use the computer display in an interactive loop to correct the allignment errors. Even with the allignment correct, however, the phase display still showed a multiple ring pattern indicative of errors in mode matching. Recently, the mode matching was corrected, and the ring pattern disppeared (Unfortunately, no photographs taken since the mode matching was corrected could be included with this report). One problem with the current system is that there is no good way to calibrate it. The "Phase T.V." is very good at showing where the errors are, but it does not give a very good idea of how large the errors are. Several ideas have been proposed which involve the introduction of deliberate misallignments or phase distortions, but none have yet been made to work.

All of the pictures made to date with the "Phase T.V." system have been of light from the mode cleaning cavity. The two 40 meter cavities of the gravity wave detector have not yet been looked at. This should change in the very near future, and "Phase T.V." pictures of light from the primary and secondary arms should prove quite useful.

Phase TV: Block Diagram

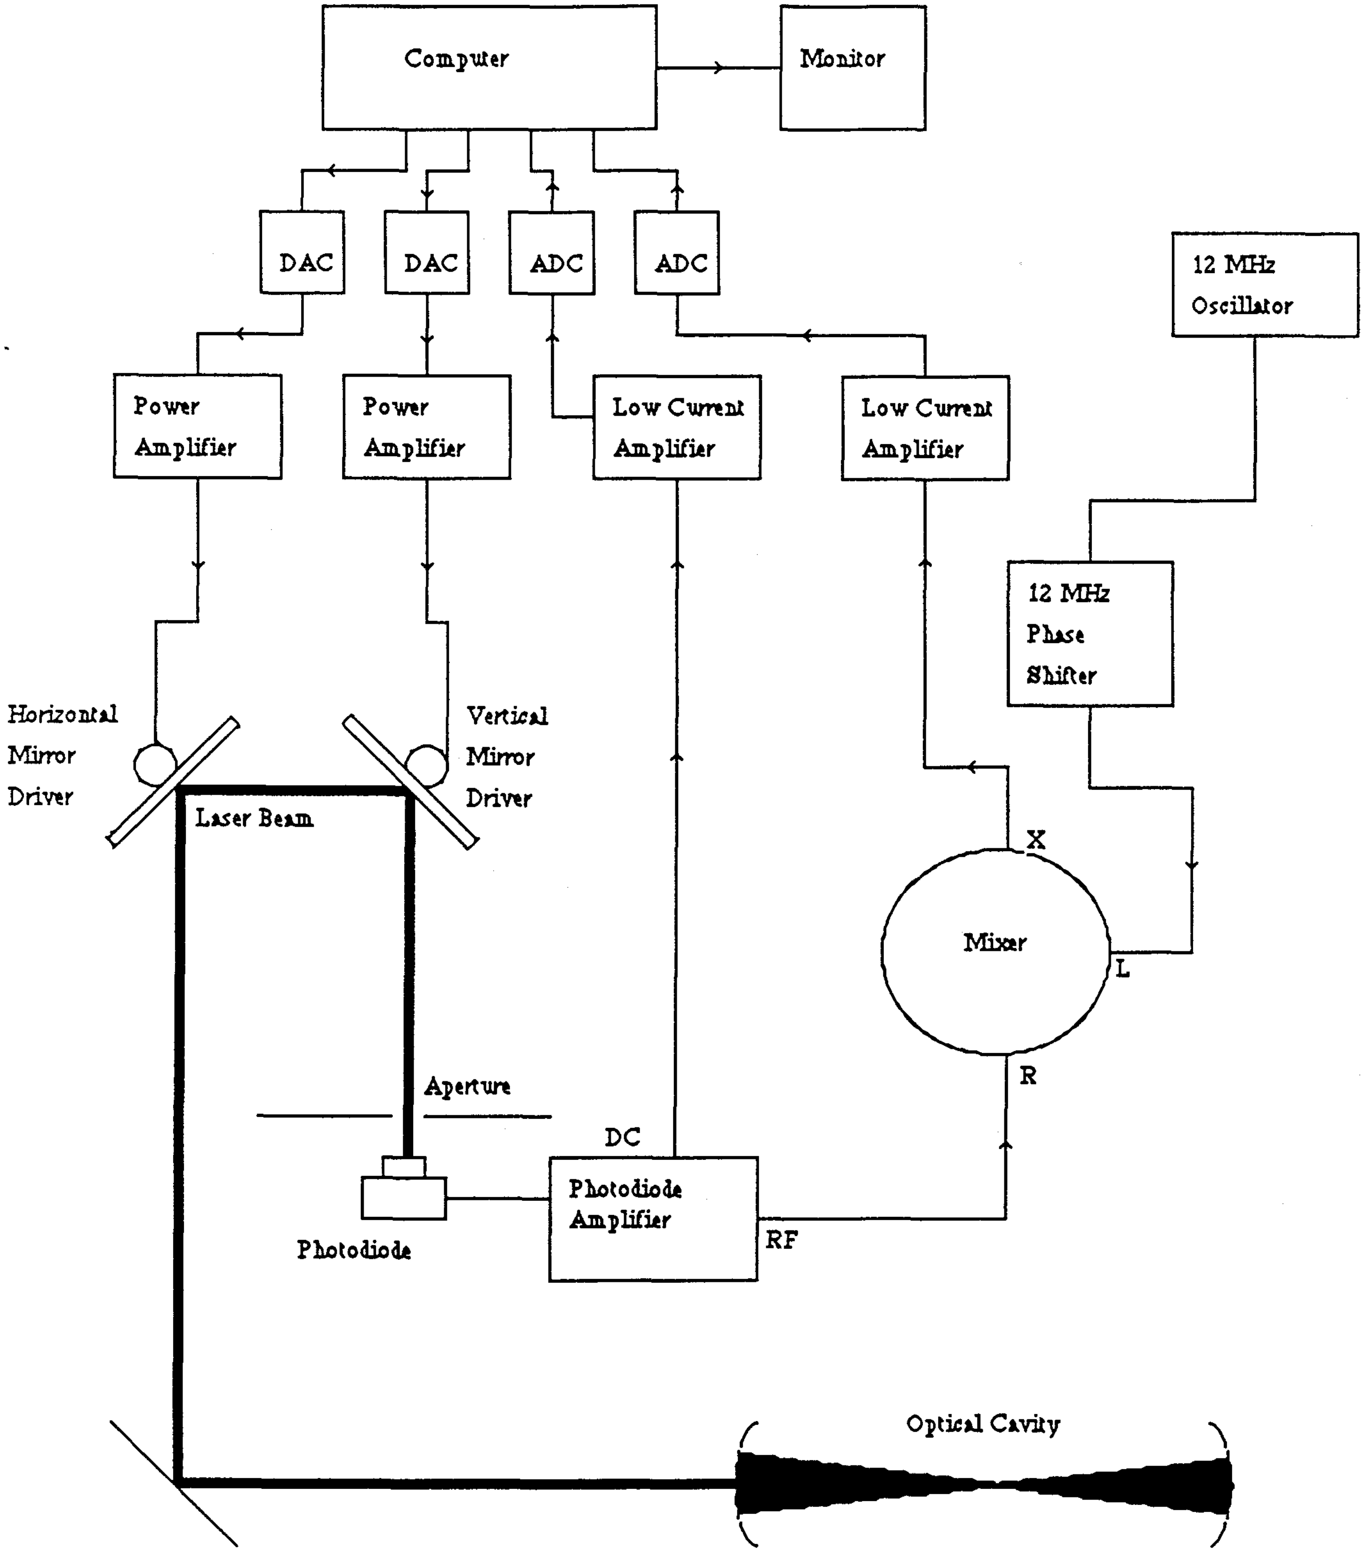

Below: Typical screen displays show adjacent plots of intensity and phase, with a contour plot of intensity superimposed on the phase plot for reference.

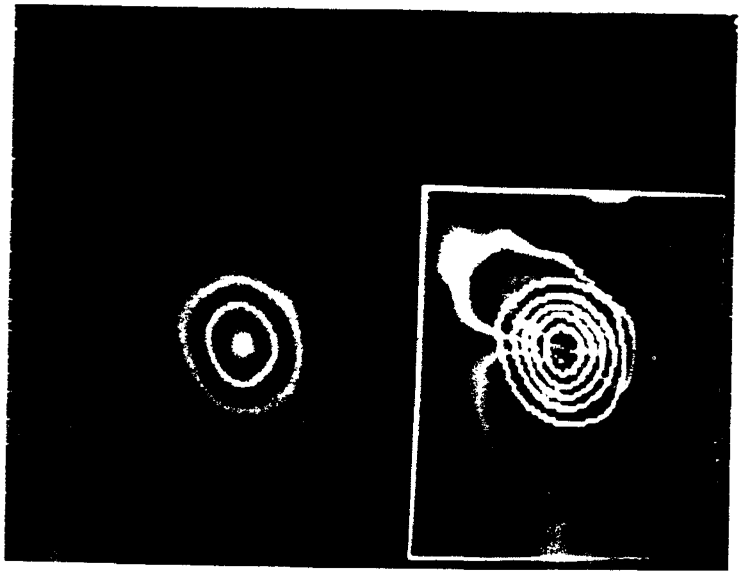

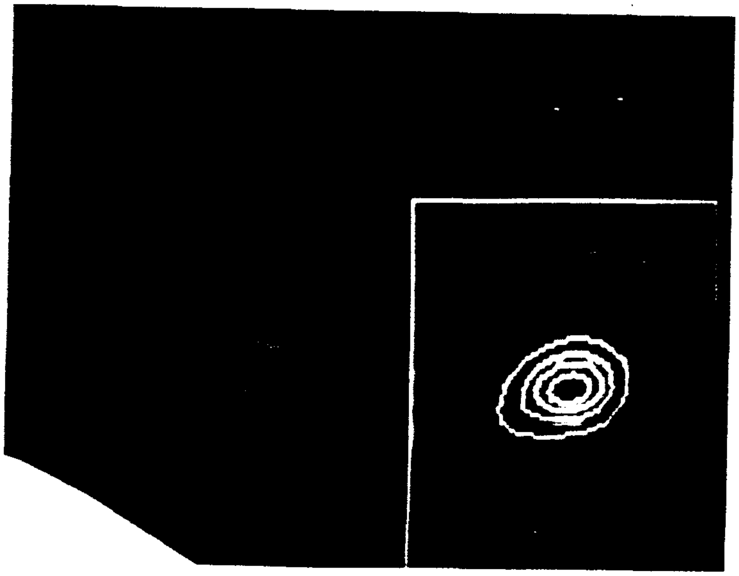

Below: Color and 3-d plots of the same phase data show both a misallignment and a mode matching error. The ring patterns (ridges and moats) indicate poor mode matching. The hill and valley are indicative of a tilt about an axis ~ 135 deg from horizontal. Notice that the overall shape of the distortion is shown more clearly in the color plot, while the relative magnitudes of the distortions are shown more clearly in the 3-d plot.

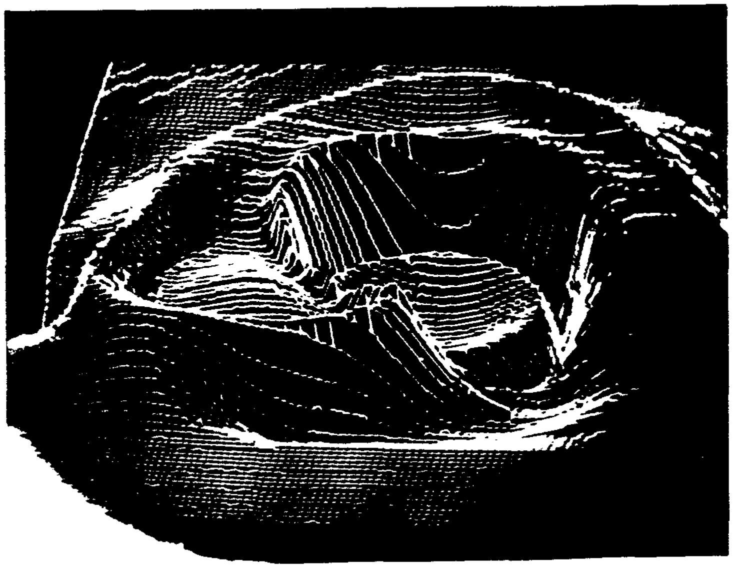

Below: Plots typical of mode cleaner reflected light prior to mode matching corrections. Analogous plots made since the correction are quite circularly symmetric and show no ring patterns.

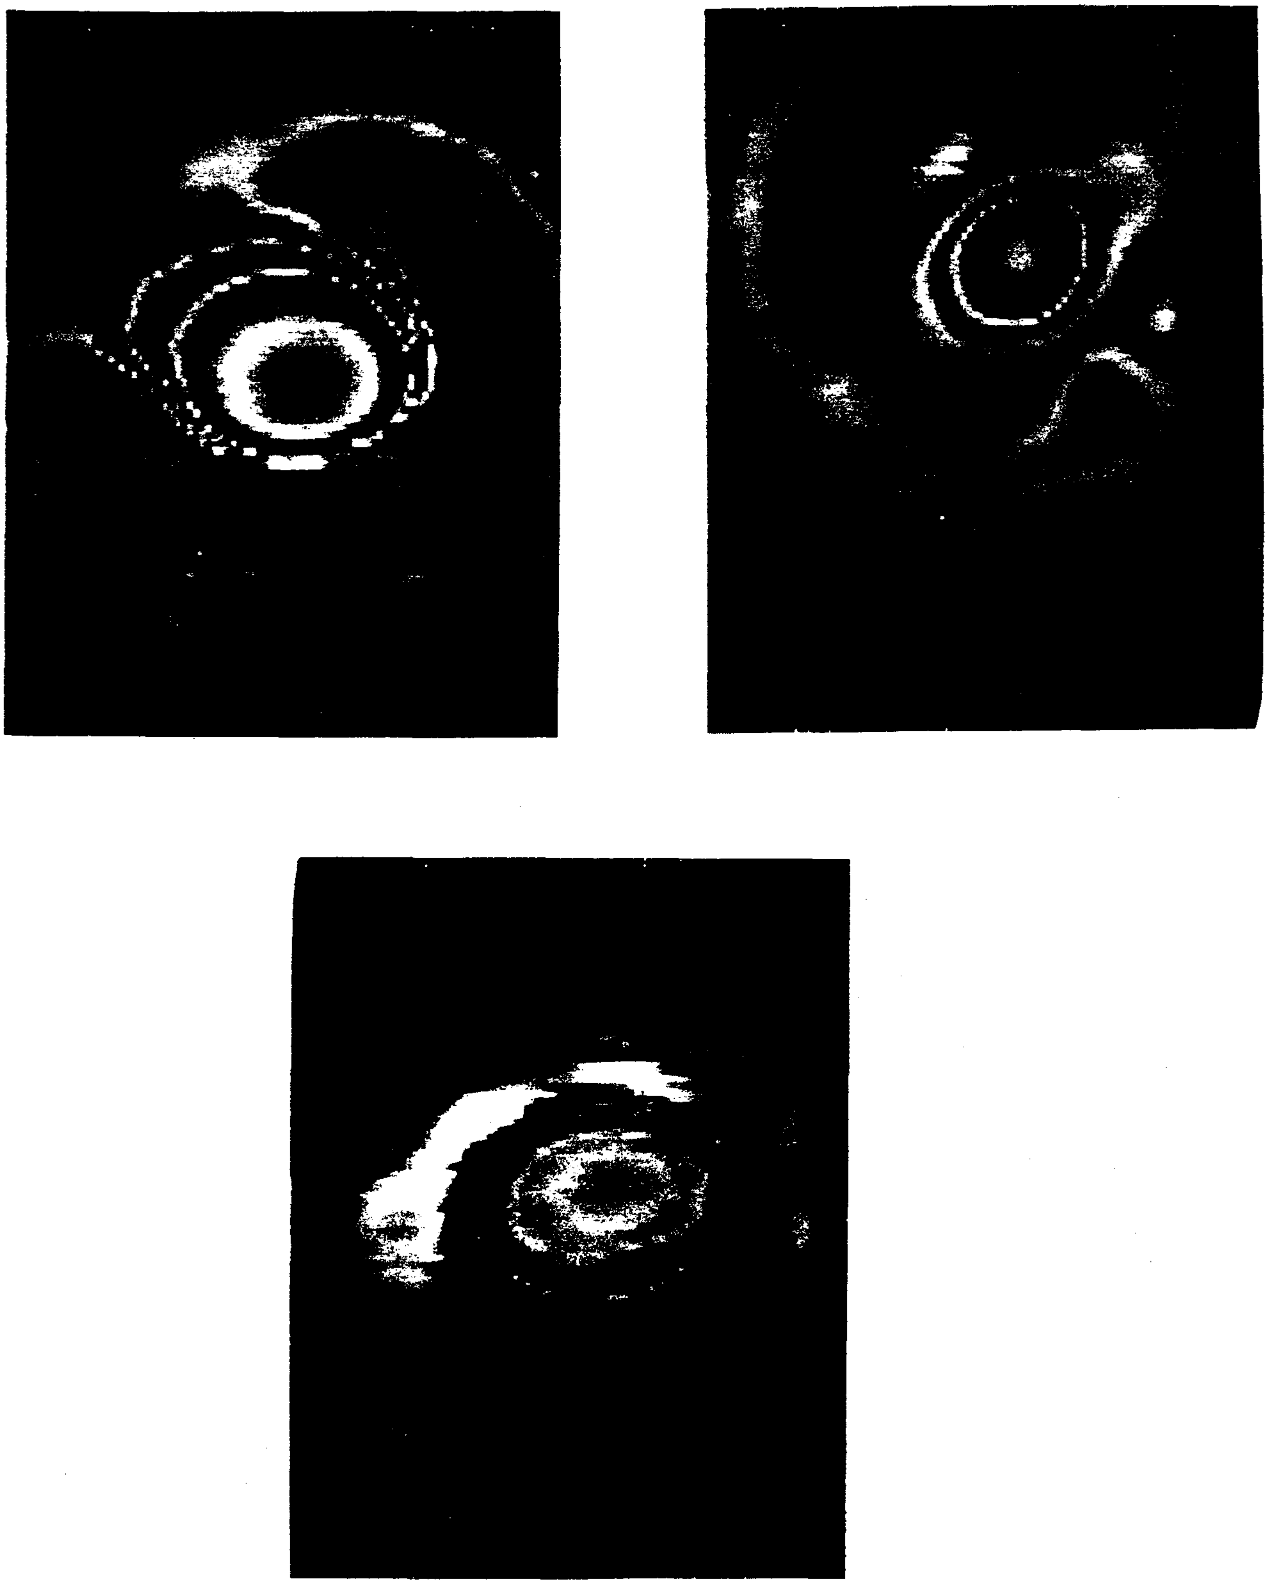

Below: Phase profiles of light reflected from the mode cleaner while it is locked in modes other than TEM_{00}. In the TEM_{00} mode, only one spot would be visible.

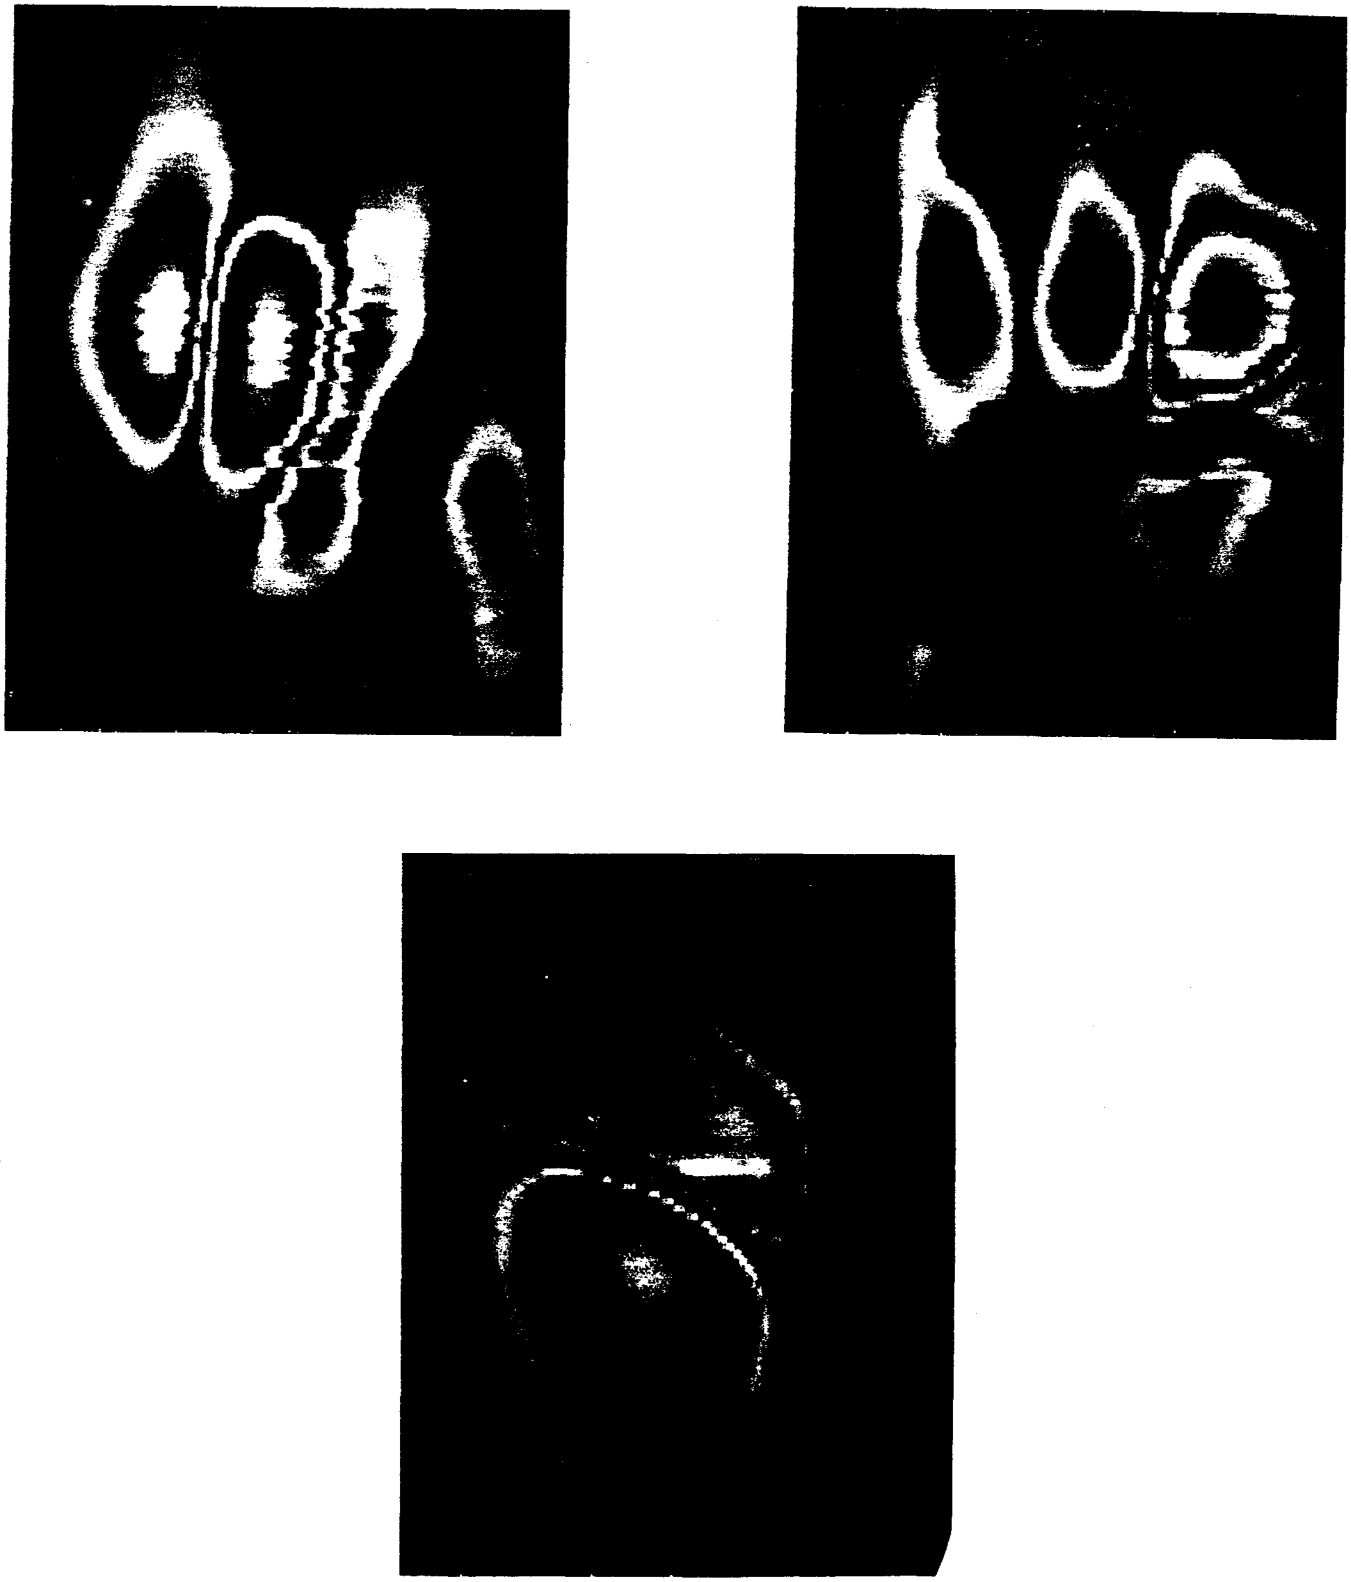

Below: Intensity profile of light transmitted through the mode cleaner while it is resonating in the wrong mode. In the TEM_{00} mode, only one Gaussian spot is transmitted; here, three spots are transmitted.

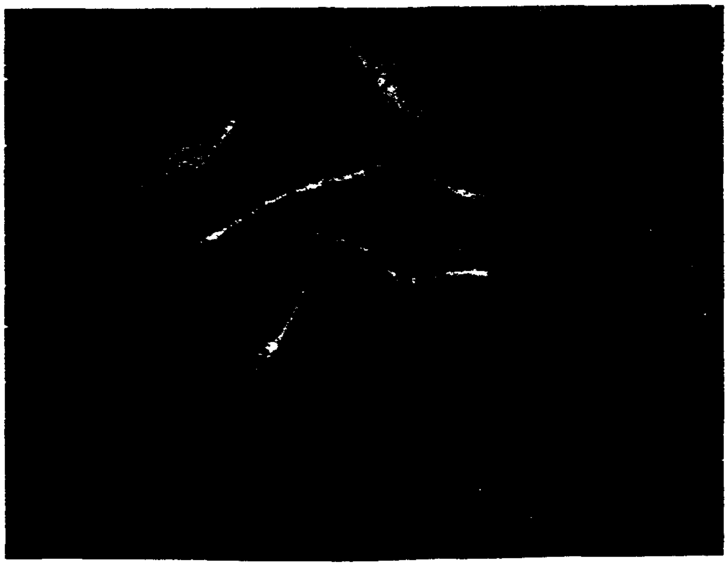

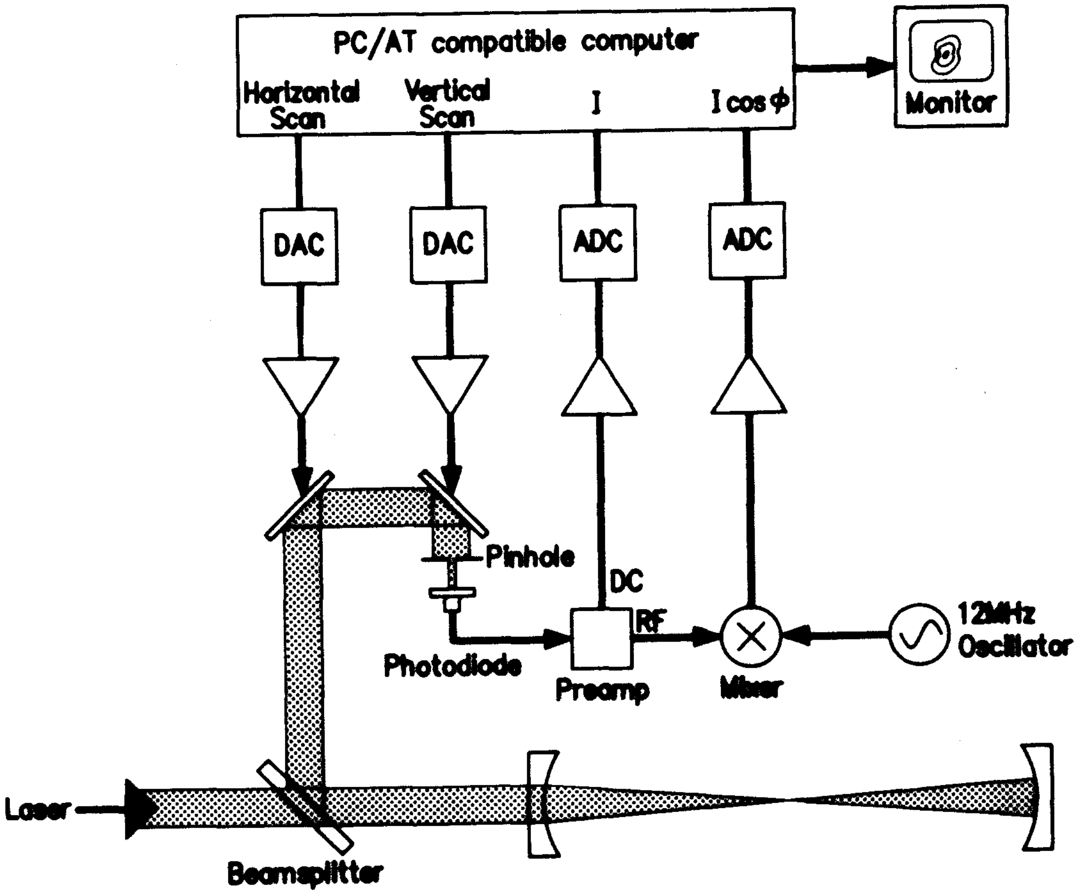
PC/AT compatible computer
Horizontal Scan
Vertical Scan
I
I cos φ
Monitor
DAC
DAC
ADC
ADC
Pinhole
DC
RF
Photodiode
Preamp
Mixer
12MHz Oscillator
Laser
Beamsplitter