HIGHER ORDER MODE DAMPING IN AN ALS TEST CAVITY

Arne F. Jacob, Glen R. Lambertson, Walter Barry*
Lawrence Berkeley Laboratory, 1 Cyclotron Road, Berkeley, California 94720, USA
*CEBAF, 12000 Jefferson Avenue, Newport News, Virginia 23606, USA

Abstract

The higher order mode (HOM) attenuation scheme proposed for the Advanced Light Source (ALS) accelerating cavities consists of two broadband dampers placed 90° apart on the outer edge. In order to assess the damping efficiency a test assembly was built. The HOM damping was obtained by comparing the peak values of the transmission through the cavity for both the damped and the undamped case. Because of the high number of modes and frequency shifts due to the damping gear, the damping was assessed statistically, by averaging over several modes. In the frequency range from 1.5 to 5.5 GHz, average damping greater than 100 was obtained.

Introduction

The HOMs of the ALS accelerating cavities are a major source of beam impedance and thus a potential cause of coupled bunch instabilities. They must be damped as much as possible using passive damping to reduce the power requirements on an active feedback system, which will still be needed for the full stabilization of the beam. The purpose of this study was to check the working principle of the broadband passive damping gear proposed for the ALS [1], do some development work before commissioning the cavities, and get an approximate quantitative figure of its performance as input data for the feedback system design.

A method that measures relative changes in shunt impedance was used. Following a description of the cavity and the damper configuration(s), this measurement procedure will be discussed in more detail. Finally some results obtained with a test cavity will be presented and discussed.

The Cavity and HOM Damping System

The single-cell ALS cavity is shown in Fig.1. Longitudinally centered ports for a tuner, a monitor, the RF feed, and HOM damping are distributed around the circumference, disturbing its axial rotational symmetry. The attached beam pipe has a bore of 70 mm (cutoff frequency: $f_{cTM} \sim 3.3$ GHz) eventually tapering down to the standard ALS cross section ($f_{cTM} \sim 5$ GHz).

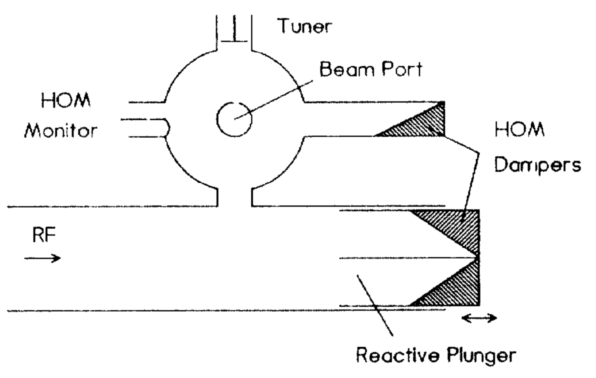

Fig.1: Sketch of cavity with RF feed and HOM damper.

The 500 MHz RF power couples from the fundamental TE_{10} (rectangular) waveguide mode into the TM_{010} cavity mode via a 15 cm diameter port at the bottom of the cavity; this is preferred to the usual loop coupler for power reasons. Its wide band properties are a key feature in the proposed HOM damping scheme. The power coupling is adjusted by means of a reactive plunger in the downstream stub of the waveguide. Two possible arrangements were investigated: the beam axis perpendicular to the waveguide, with the aperture in the narrow wall as sketched, or the beam axis parallel to the waveguide, the porthole being in the broad wall. Because of the available space the first is the arrangement to be used (Fig.1).

The high cutoff frequency of the beam pipe requires the dampers to work over a wide band; the scheme proposed in Ref 1 relies on two devices at 90° to suppress split dipole modes. One damper is connected to the cavity via a port 12 cm in diameter (TE_{11} cutoff ≅ 1.45 GHz) and 10 cm in length. A terminated waveguide of the same diameter is attached to it with the load far enough away to avoid affecting the fundamental mode. The other damper is implemented in the waveguide system which powers the cavity. HOMs generated by the beam couple to the waveguide and excite up- and downstream waves which have to be absorbed without loading the 500 MHz RF signal. The absorbers are backterminated highpass filters which, downstream, are part of the tuning plunger, and, upstream, are located at several waveguide bends. They can be realized with waveguides of reduced cross section, as illustrated in Fig.1. With a metal vane centered along the E-plane of the waveguide, the 500 MHz RF is totally reflected into the waveguide; signals above ~650 MHz are partially admitted and can be absorbed. This scheme has the advantage of being very compact; its drawback, however, is that not all the waveguide modes (e.g. the TE_{10}) are equally well transmitted through the filter. An alternative approach is to taper down the waveguide width until the fundamental mode can no longer propagate at 500 MHz. A smoother coupling can be expected at the expense of greater length.

For mechanical convenience the test system for the measurements was simpler. First the cavity had a pillbox instead of a re-entrant shape, and its frequency was scaled up by a factor of 1.2. These changes will of course affect individual modes but not significantly alter the average behavior, except for a global frequency shift. Second the HOM coupling port at 90° was the same diameter as the other one, which is not an exact scaling. Finally the upstream waveguide dampers were replaced by a plunger-type filter, which should exhibit the same behavior.

Measurement Procedure

A thorough way to characterize the efficiency of the damping system would be to compare one by one the measured shunt impedance of each mode in the isolated cavity and in the damped cavity. In principle this can be done by pulling a bead through the cavity. In the frequency range up to cutoff there are hundreds of modes and this method becomes impractical; also in the case of our test cavity mode identification may not be needed. We therefore have made an alternative measurement which will be described now.

For the longitudinal measurement electric probes were mounted in the end walls of the cavity and aligned on the beam axis. In the transverse case magnetic loops were used instead. They will, however, also excite a family of TE modes, thus limiting the validity of the results for transverse damping. The alternative use of two electric probes symmetrically placed about the beam axis and driven in opposite phase was not possible because the required 180° hybrids were not available in the entire band. At resonance the transmission coefficient S_{21} between the antennas is a measure of the shunt impedance. Thus by comparing the transmission data for the undamped and the damped cases the reduction in shunt impedance can be assessed. Under the assumption that R/Q is not significantly altered this also measures the decrease in Q, and, for unchanged transit time factor, also the achieved reduction in beam impedance. At resonance the damping factor is given by

$$\frac{R_u}{R_d} = \frac{Q_u}{Q_d} = \frac{S_{21u}\,[S_{21d}\,(n^2 + m^2) - 2nm]}{S_{21d}\,[S_{21u}\,(n^2 + m^2) - 2nm]} \qquad (1)$$

The subscripts d and u refer to the damped and undamped case respectively. The input and output coupling coefficients (n, m) are different for each mode, and unknown. However, for like antennas, i.e. n=m, eq.(1) simplifies to

$$\frac{R_u}{R_d} = \frac{Q_u}{Q_d} = \frac{S_{21u}\,(S_{21d} - 1)}{S_{21d}\,(S_{21u} - 1)}, \qquad (2)$$

which can readily be evaluated if the antennas couple strongly enough to detect the peaks with good accuracy in the damped case.

In order to find the peaks, the frequency response of the cavity must be carefully measured. To keep the error on S_{21} below 10%, the data points should not be spaced by more than f_r / (2 Q_L), where f_r is the resonant frequency and Q_L the loaded Q. The average error on S_{21} is then less than 4%. Because of Q's on the order of 10^4 a large number of data points is required to measure S_{21} accurately for each mode. To cover the frequency range from 0.5 GHz to 1.0 GHz, for instance, approximately 200 intervals with 401 points each (a standard option on the HP8510B Network Analyzer [NWA]) are needed. This number can be reduced as frequency increases because, on the average, Q grows slower than f_r. Nevertheless, for the total frequency range we have considered here (0.5 - to 5.5 GHz), this represents a huge amount of data. The time needed for the data acquisition is on the order of an hour, depending on the number of averages during the measurement. The data acquisition and analysis, i.e. the search for resonant peaks, is handled by a computer program. A simple through calibration accounts for the losses in the cables.

Two slightly different methods of data analysis are employed. In the first, the noise floor is suppressed, and peaks are identified by comparing each data point with its two direct neighbors. In order to reduce residual errors due to noise, an alternative method has also been used: here more than three (typically seven) points have to match the pattern of a resonant curve (Fig.2). For instance, the value P_4 is a resonant peak if the increments $D_i = P_{i+1} - P_i$ (i = 1,2,...,6) satisfy the following criteria: D_1, D_2, $-D_5$, $-D_6 > T$, $D_3 D_4 < 0$, $D_3 > 0$, where the threshold T is determined empirically to reduce errors. The choice of the frequency spacing depends on Q. In the damped cavity both high- and low-Q resonances exist, so that the analysis has to be repeated with varying intervals.

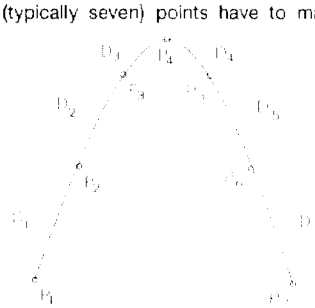

Fig.2: Data points and intervals on a resonant curve.

Results

Qualitative Results

Using the first peak search method various configurations were investigated for both longitudinal and transverse coupling in order to define which structure to study in more detail. However the damping was assessed only qualitatively by visually comparing the peak spectra. Then only obvious changes could be noticed. This saved some time at the expense of accuracy.

Two different absorber materials, NZ-51 ferrite and carbon impregnated cardboard, were evaluated. The ferrite was placed on the walls of the waveguide and auxiliary port whereas the resistive material was positioned in high E-field regions. Both materials provided the same damping.

In comparing the coupling through the narrow wall and the broad wall of the waveguide no difference was detected.

The partitioned tuning plunger(s) causes reflections even at higher frequencies. In an attempt to improve the match a tapered plunger was tried. A gradual 11° transition from the full width down to half its value provided a smooth transmission to an attached load at higher frequencies while retaining the ability to tune the fundamental frequency signal. However, no significant difference in damping could be observed.

Although these results are not precise they suggest that the achievable amount of damping is limited by the coupling through the apertures of the cavity, rather than by the quality of the loads, which with some development could certainly be further optimized. For coupling apertures of a given size, one could improve the average damping efficiency by displacing the holes off center, so as to couple to more modes.

Quantitative Results

In order to obtain a more quantitative measure of the damping efficiency, the peak data must be analyzed with more precision. At first the alternative peak search method described above was used to characterize the spectrum of the cavity more accurately. Transverse coupling was not investigated in detail, because of the inaccuracy inherent to the magnetic loops. A very large number (855) of modes was found in the undamped test cavity. This may be compared with a calculated 52 longitudinal modes in an ideally symmetric cavity. Owing to the perturbation of the symmetry by the various apertures, many modes with nonzero azimuthal dependence have longitudinal electric field components at the cavity centerline. Some modes were identified as of basically TE-type. They can be excited by imperfect antennas - a parasitic effect - or because of longitudinal electric field components due to distortions in the apertures - a situation the beam would face as well. In the damped cavity still 267 modes have been found above the threshold. The spectral distribution is shown in Fig.3. The ordinate is the normalized shunt impedance in dB, so that for a given mode the difference between the spectra gives the damping efficiency in dB, i.e. 20 times its logarithmic value. A visual inspection of the two results gives some information on the change in the envelope. However the interpretation of the reduced mode density is not apparent. Therefore more comprehensive ways to analyze the data were employed. Because of the high number of modes and the frequency shifts occurring with damping an exact mode by mode comparison is not possible, except at the lowest frequencies.

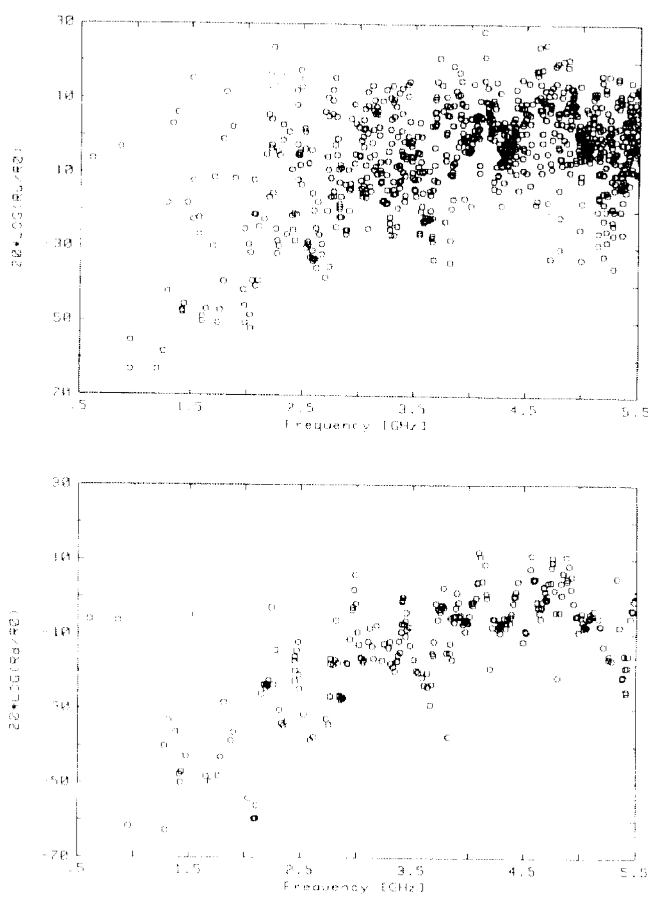

Fig.3: Spectrum of (a) the undamped and (b) the damped test cavity.

One method is to compare each mode of the damped cavity and relate it to the one with the closest resonance frequency in the undamped cavity. For all the attenuated modes which are not detected, a minimum peak value corresponding to the noise floor is assumed, a conservative approach, especially for the weak modes. We worked with S_{21min} ~ - 60 dB. Fig.4 shows the result of such an analysis. It already reveals much more information than the visual inspection of Fig.3. Significant damping starts at about 1.25 GHz. Many modes, which presumably are strongly excited in the undamped cavity, exhibit an attenuation between 40 and 80 dB. Some modes, however, appear to be amplified, a result of the assumed correspondence between modes.

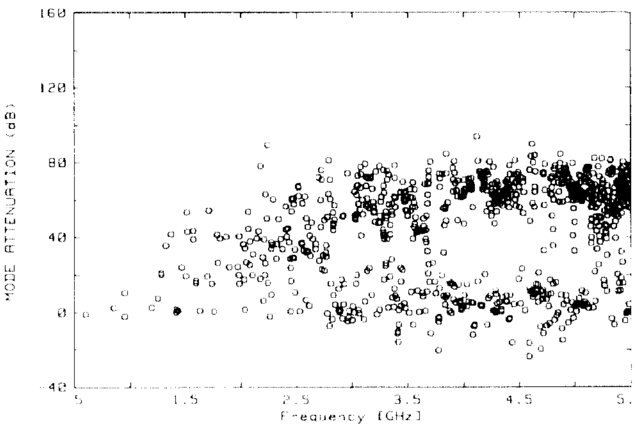

Fig 4: Comparing peaks with closest resonant frequencies.

A means to smooth the errors out is to take averages over parts of the frequency interval. The uncertainty does not so much lie in the mode-to-mode variation of the attenuation as in the erratic mating of the modes. An arithmetic average of the individual damping values would emphasize the larger numbers. If their reciprocal values are averaged one introduces the opposite bias. Therefore the shunt impedances were averaged over a given frequency range before computing the mode attenuation. Two different averages were used, arithmetic and geometric. Figs.5 and 6 depict the results obtained in 100 MHz intervals. While both results agree quite well up to 2.8 GHz some divergence appears towards higher frequencies, where the arithmetic averaging yields values in the 20-to-30 dB range, and the geometric one gives numbers in excess of 40 dB.

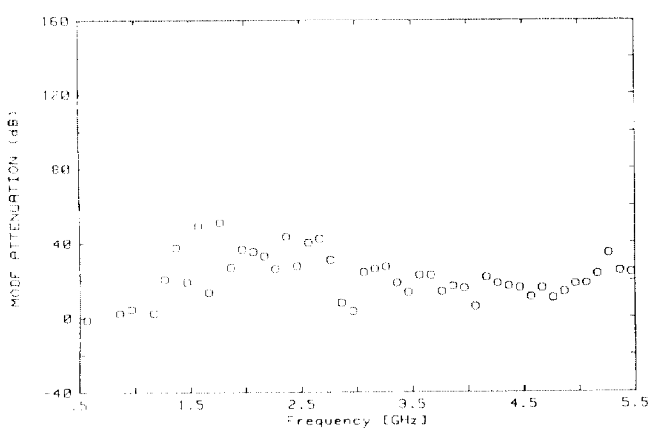

Fig.5: Arithmetic averages of peaks in 100 MHz bands.

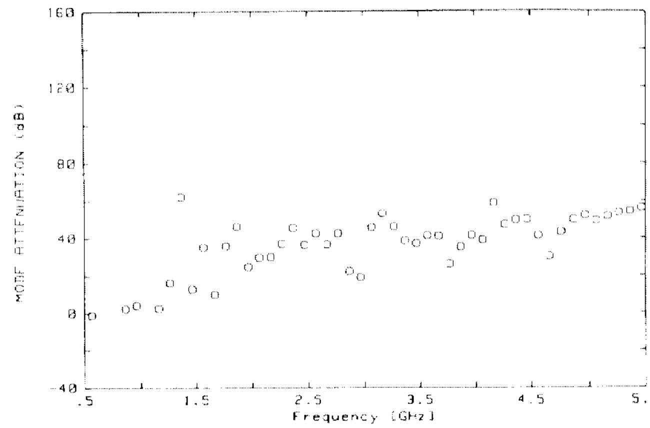

Fig.6: Geometric averages of peaks in 100 MHz bands.

The arithmetic averaging emphasizes the larger peak values in both the undamped and the damped case. Thus one large number in the damped cavity may already spoil the average. The reduced attenuation at higher frequencies also does not fit the intuitive understanding of the dampers: their efficiency should improve at higher frequencies, as the coupling to the loads gets stronger. However, in both analyses errors occur when modes overlap, which is more likely to occur in the damped cavity, and at higher frequencies. Two circumstances are imaginable: (a) The peaks are still distinct; then the detected resistance is higher than the actual one and both methods yield too-low values for the mode attenuation. (b) The peaks cannot be distinguished because they are too close together; the detected resistance then approximately equals the sum of the true values, and the arithmetic average is not affected. However the damping value resulting from the geometric average is too high.

We conclude that arithmetic averaging gives a lower limit for the damping efficiency, but because of the emphasis on the peak resistances of the strong modes, the results seem to be too conservative here. In the case of geometric averaging, mode overlapping at higher frequencies may result in an overestimation of the damping figure. However, a correction due to the limited sensitivity of the measurement would boost the damping effectiveness in both cases, so that 40 dB (~ 100) probably gives the right order of magnitude. It should be pointed out that in the ALS the beam pipe attached to the cavity will be tapered making it likely that the HOMs are much less a problem above 3.3 GHz in the longitudinal case. The transverse modes will even couple to the TE mode of the beam pipe, which has a cutoff frequency of ~ 2.5 GHz.

Conclusions

The performance of the HOM attenuation scheme proposed for the ALS has been investigated on a test cavity. Because of their large number, the modes could not be characterized individually. Thus a global approach has been used, in which the resonant peaks were detected in both the isolated and the damped cavity and their average values compared within frequency intervals. A geometric average yields a damping effectiveness in excess of 100 in the frequency range from 1.5 to 5.5 GHz. Some modes escape attenuation. The damping is limited by the coupling factor of the apertures rather than the quality of the loads.

References

[1] B. Taylor, K. Baptiste, H. Lancaster, C.C. Lo: Advanced Light Source Storage Ring RF System, Proc. IEEE PAC, Chicago, IL, 1989, pp.124-125.

This work was supported by the Director, Office of Energy Research, Office of Basic Energy Sciences, Materials Sciences Division, of the U.S. Department of Energy, under Contract No. DE-AC03-76SF00098.