PATH IDENTIFICATION IN STRUCTURAL ACOUSTICS

K. H. Heron

Defence Research Agency
Aerospace Division
RAE Farnborough, England

Abstract

This paper describes an experimental method for the identification of the structure-borne path by which the vibrational energy from a helicopter gearbox is transmitted to the cabin. The method involves taking detailed measurements of the gearbox/cabin interface vibration field as well as of the internal noise field; typically 24 accelerometers and 12 microphones are used.

The method also involves calculating the statistical accuracy of the various predictions and thus being able to display the results as confidence intervals. In many ways this statistical approach is vital since it is necessary to know the quality of any result as well as its numerical value. The heart of the method lies in the way these statistical results are calculated as well as the way statistical confidence tests are used to drive and steer the necessary data fitting processes.

This paper also describes the results from a full scale laboratory validation trial using a grounded Lynx helicopter. A series of known 'flight' conditions were measured and then the method applied. The known results were indeed reproduced; and furthermore the statistics were validated in the sense that a 90% confidence interval for a given result was indeed wrong about 10% of the time!

The method can now be considered to be an engineering tool of proven quality, able to accurately diagnose and quantitatively identify the different paths.

Introduction

Practically all aircraft internal noise fields are made up of contributions from more than one source; for example, engine noise, noise induced by the boundary-layer pressure fluctuations, noise induced by the gearbox on helicopters, etc. These sources fall into two main categories: firstly the broad-band sources such as jet noise, and secondly the narrow-band sources such as propeller noise. Here we are mainly concerned with the latter which are generally the major source of cabin noise for helicopters and for propeller driven aircraft. In these cases we can easily identify the source by the distinctive frequency characteristics of the cabin noise spectrum. However modifying the source to reduce the noise cannot normally be entertained because of overriding considerations associated with aircraft performance. Thus to reduce cabin noise it is not sufficient to identify the source of the noise, we must also identify the mechanisms or paths by which energy from the source is carried to the cabin and converted to noise. This would allow palliative research to become more focused and thus more successful.

Unfortunately there is not just one path to identify. On helicopters there are two main paths: firstly the so called air-borne path associated with acoustic radiation from the gearbox casing, and secondly the structure-borne path where vibrational energy is transmitted via the gearbox feet. Only when we have a good qualitative understanding of these different paths on a given aircraft will we be able to apply and optimise acoustic palliatives scientifically. This problem of path identification has been one of the most important unsolved problems of structural acoustics for many years, and various attempts have been made to solve it.

In this report we describe the laboratory validation of a new method for path identification that involves few restrictions. The method is based on a complete ground simulation of the structure-borne noise field, where instead of attempting a direct ground simulation, a series of ad hoc ground experiments are performed and then the measurements from these are combined to give an accurate representation of the flight situation.

Theory

The principle of the method

Consider the case of a single structure-borne path and a single air-borne path between a given source and a receiver (for example a helicopter gearbox and a cabin microphone). Assume that there exists some interface on the structure through which all the structure-borne energy at a given frequency passes (for example the gearbox feet). Now consider the response at this interface during flight. If we could reproduce this response on the ground, using mechanical exciters for example, the resultant noise levels would be the flight structure-borne noise levels.

That is, if $\underline{a} = \underline{A}\,\underline{\beta}$, (1)

then $\underline{v} = \underline{V}\,\underline{\beta}$, (2)

where $\underline{a}$ is the flight interface response vector, $\underline{A}$ and $\underline{V}$ are the ground interface response matrix and the ground cabin noise matrix respectively, $\underline{\beta}$ is a column vector of regression coefficients, and $\underline{v}$ is the required structure-borne cabin noise vector. All these vectors and matrices are complex to take full account of inter-transducer phase information.

Unfortunately such an approach is impractical, since with too few ground tests $\underline{\beta}$ will in general not exist. It is important to treat the problem as a statistical one, the requirement being to define and find the 'best' $\underline{\beta}$, best in the sense of using a least square fitting procedure on equation (1) with the constraint that we retain statistical significance.

The actual fitting procedure

The procedure can be divided into three parts. Firstly an algorithm is needed to order the ground vectors in terms of their usefulness for inclusion in the fit. Secondly a statistical test is needed to decide how many of these ground vectors to include in the final fit, and finally the expected variance must be calculated and propagated through to the final results. These three stages are briefly described below, and more details can be found in reference 1.

The vector ordering procedure

Consider a subset $\underline{A}_k$ of $\underline{A}$ made up of just k columns of $\underline{A}$, then the least square fit for this subset is given by

$$\underline{\beta}_k = \underline{C}\ \underline{A}_k^{*}\underline{a}\ , \tag{3}$$

with $$\underline{C} = (\underline{A}_k^{*}\underline{A}_k)^{-1}\ , \tag{4}$$

where $\underline{A}_k^{*}$ denotes the complex conjugate transpose of the matrix $\underline{A}_k$. If σ_k^2 is the fit variance, then $\sigma_k^2\underline{C}$ is known as the covariance matrix, and an unbiased estimate for σ_k^2 is given by

$$\sigma_k^2 = R_k/[\ 2N-2k\ ]\ . \tag{5}$$

where N is the number of interface response transducers, and the fit residue, R_k, is given by

$$R_k = \underline{a}^{*}\underline{a} - \underline{a}^{*}\underline{A}_k\underline{\beta}_k\ . \tag{6}$$

Assuming that we have already ordered and chosen the first k vectors, then by adding to these k vectors every other unused vector in turn and fitting to each of these k+1 vectors we can choose that extra vector that minimises R_{k+1}. This defines the algorithm used for ordering the vectors.

Choosing the fit

We could proceed by using the above formulae for R_k and for σ_k^2 until we found a minimum for σ_k^2. However, when using experimental data, it is unreasonable to use equation (6) for R_k since R_k can then be made as small as we like which is equivalent to assuming perfect data.

We decided to modify equation (6), such that R_k is given by

$$R_k = (1+\varepsilon)\underline{a}^{*}\underline{a} - \underline{a}^{*}\underline{A}_k\underline{\beta}_k\ , \tag{7}$$

where ε can be related to the experimental errors in the measurement of $\underline{a}^{*}\underline{a}$.
The fitting process is continued by adding the next 'best' vector in turn, thus producing N possible fits; assuming enough ground vectors. The fit k is chosen if and only if <u>all</u> subsequent fits are statistically insignificantly better than fit k.

Error estimation

The covariance matrix $\sigma^2\underline{C}$ of the fit, will be rewritten as $\underline{C}_\beta$, since it is an unbiased estimate of the covariance matrix for the regression coefficients $\underline{\beta}$. That is

$$\mathrm{var}(\beta_i) = C_{\beta ii}\ , \tag{8}$$

and $$\mathrm{covar}(\beta_i,\beta_j) = C_{\beta ij}\ . \tag{9}$$

Now the required structure-borne noise vector is given by equation (2), and thus if $\underline{C}_w$ is the covariance matrix for the vector $\underline{w}$, we have

$$\underline{C}_w = \underline{W}\ \underline{C}_\beta\underline{W}^{*}\ , \tag{10}$$

with $$\mathrm{var}(w_i) = C_{wii}\ , \tag{11}$$

and $$\mathrm{covar}(w_i,w_j) = C_{wij}\ . \tag{12}$$

Also if the root mean square of the M structure-borne cabin noise levels is w, given by

$$w^2 = (\underline{w}^{*}\underline{w})/M\ , \tag{13}$$

we have $$w\ \partial w/\partial\underline{w} = \underline{w}^{*}/M\ , \tag{14}$$

and thus

$$\mathrm{var}(w) = [\ \underline{w}^{*}\underline{C}_w\underline{w}\ ]/[\ M\ (\underline{w}^{*}\underline{w})\ ]. \tag{15}$$

Now consider the M air-borne cabin noise levels, $\underline{v}$ say, given by

$$\underline{v} = \underline{m} - \underline{w} , \qquad (16)$$

where $\underline{m}$ is the column vector of the total in-flight cabin noise levels. We must first assign some value to the errors in $\underline{m}$, that is we must set a covariance matrix, $\underline{C}_m$ say, for $\underline{m}$. We assumed that the variance for each microphone is a fixed proportion of its own mean square level, and furthermore we assumed that all of these errors were independent of each other. That is

$$\underline{C}_m = \delta^2 \text{ diag}(mj^* mj), \qquad (17)$$

where δ is an estimate of the measurement error in the flight microphones; so that each microphone is assumed to have a standard error of δ times its root mean square level. Thus

$$\underline{C}_v = \underline{C}_m + \underline{C}_w . \qquad (18)$$

Also if the root mean square of the M air-borne cabin noise levels is v, given by

$$v^2 = (\underline{v}^* \underline{v})/M , \qquad (19)$$

we obtain

$$\text{var}(v) = [\underline{v}^* \underline{C}_v \underline{v}]/[M (\underline{v}^* \underline{v})]. \qquad (20)$$

Finally it should be noted that because the number of degrees of freedom associated with these results is a relatively small number, small sample theory should be used to convert the variances into confidence levels.

The Measurement Programme

A lynx airframe was used. Response blocks were attached to the gearbox feet, close to the gearbox/cabin interface, to allow the mounting of 6 accelerometers at each foot encompassing all the mechanical degrees of freedom. Furthermore, on each gearbox leg and above all the accelerometers, a force block was attached to allow mechanical excitation at 6 positions. 12 microphones were installed at various points within the cabin. These 36 channels of response were amplified and then digitised simultaneously using a 40 channel data acquisition system.

The flight simulation tests

The objective of these tests was to create a series of 'flight' results with known air-borne and structure-borne components. Pure structure-borne noise was generated by a mechanical exciter attached at points high up on the gearbox casing, and pure air-borne noise was generated by an external loudspeaker positioned about 3m away on the port side. Two air-borne noise levels and two structure-borne noise fields were generated. In all cases the same coherent source was used to drive both the loudspeaker and the mechanical exciter. The source was a cyclic computer generated signal comprising 10 frequency components; 280Hz and all its harmonics up to 2800Hz.

The path identification tests

This main phase of testing was performed as if this was a full flight application of the path identification method. A mechanical exciter was attached to each of the 6 allocated positions on each of the 4 gearbox legs in turn; thus a total of 24 'ground' tests were performed.

Results

In all the results that follow the different 'flight cases' are designated by the code FxNx. F1 refers to mechanical excitation at force position 1, and F2 refers to mechanical excitation at force position 2. N1 refers to acoustic excitation at a relatively low level, and N2 refers to acoustic excitation at a higher level.

The prediction process has already been described and we must now choose values for the error factors ε and δ of equations (7) and (17). A value for ε of 20% was chosen based upon the data measurements themselves (see reference 1 for the details).

No similar method for assessing δ directly from the data was developed, perhaps such an approach is impossible. However the main aim of this work was to predict the structure-borne noise levels together with their confidence intervals and none of this involves δ. The prediction of the air-borne noise levels must be considered a bonus, and of course even here the mean predictions are not dependent on δ, it is only when we come to the air-borne predicted confidence intervals that we need a value for δ. We tried various values and eventually chose 2% mainly based on forcing the air-borne confidence levels roughly into line with the known results at least for the lower frequencies.

Figures 1 and 2 show some typical predicted results for the structure-borne and air-borne noise levels plotted against the individual cabin microphone positions for a few frequencies. The predicted levels are depicted as I's and represent 90% confidence intervals. The solid I's are plotted offset to the left and are the structure-borne predictions, and these should be compared with the known results which are plotted as squares. The dotted I's are plotted offset to the right and are the air-borne predictions and these should be compared with the known results which are plotted as crosses. The known total noise level results are plotted as asterisks. All the results are simple magnitude plots in dB, the phase information has been suppressed. The root mean square (rms) results are shown on the right.

The graphs show how we can indeed predict both the structure-borne and air-borne noise levels. Method validation depends on the predictions agreeing with the known results. However, with agreement we mean statistical agreement, for example with our chosen confidence level of 90% we should expect 10% of the predictions to disagree with the known results, but more of this later.

The structure-borne and air-borne noise fields are coherent and interfere with each other, it is quite possible and fairly likely that they will sometimes destructively interfere. Such destructive interference can be seen in figures 1 and 2; a good example is displayed by microphone 7 of figure 1b. Path identification methods that take no account of phase are doomed to fail at such points. Furthermore palliative action against just one of these paths would cause an <u>increase</u> in noise levels when such conditions prevail, it is thus vital that such occurrences are fully identified.

The accuracy of the prediction process and hence the validation of the whole method is displayed in figures 3 and 4. Consider figure 3, on the abscissa is plotted the chosen prediction % confidence level, now varied from 10% to 90%, and on the ordinate is plotted the % of the resultant intervals that encompassed the known result. Each point on the curve is a result of 60 prediction calculations, 5 frequencies for each of the 12 microphones. The results were split into two frequency bands (280Hz-1400Hz and 1680Hz-2800Hz) to bring out any gross frequency trends. The curves exhibit very good agreement with the correct diagonal line, particularly the structure-borne results. At the higher frequencies the air-borne predictions show a consistent bias indicating that δ should perhaps have been considered to be a function of frequency.

The 90% confidence intervals shown on figures 1 and 2 are wider than we would have wished, and although accurate detract from the usefulness of the method. A solution to this problem lies with predicting the averaged or rms structure-borne and air-borne noise levels, averaged over some or all of the cabin microphones. This has a double advantage, firstly the confidence intervals will indeed be reduced (roughly by $\sqrt{M}$), and secondly the final result is an overall cabin noise result which is of more practical use than individual results associated with a particular cabin location.

These rms results are shown in figures 5, and 6. The results are plotted in a similar fashion to those of figures 1 and 2, but here they are plotted against frequency rather than microphone position. Again a 90% prediction confidence value was chosen, and again the structure-borne results are plotted offset to the left with the air-borne results plotted offset to the right. Figure 5 for the cases when the structure-borne path and the air-borne path are similar show excellent agreement for the former but poor agreement for the latter at the higher frequencies. Figure 6 shows excellent results for both paths, but because the air-borne path dominates its prediction close to the total is practically guaranteed regardless of method accuracy.

Conclusions

1) A new method for the identification of a structure-borne noise path in the overall transmission and generation of aircraft internal noise has been fully validated.

2) The method has been revised and improved during this validation process and the predicted statistical confidence intervals for the results have been shown to be accurate.

3) The method has been extended to include air-borne noise prediction, and the prediction of the spatially averaged noise levels.

4) The method can now be considered to be an engineering tool of proven quality, able to accurately diagnose and quantitatively identify the different paths.

Acknowledgements

This work was sponsored by the CEC under BRITE/EURAM contract number AERO-0028, in particular the work was performed with the collaboration of Agusta, MBB, Westland, Brüel & Kjær, CIRA, and Southampton University.

References

1 K H Heron 'Laboratory Validation of the RAE Path Identification Method', BRITE/EURAM contract AERO-0028, ARW4RA03A, April 1991.

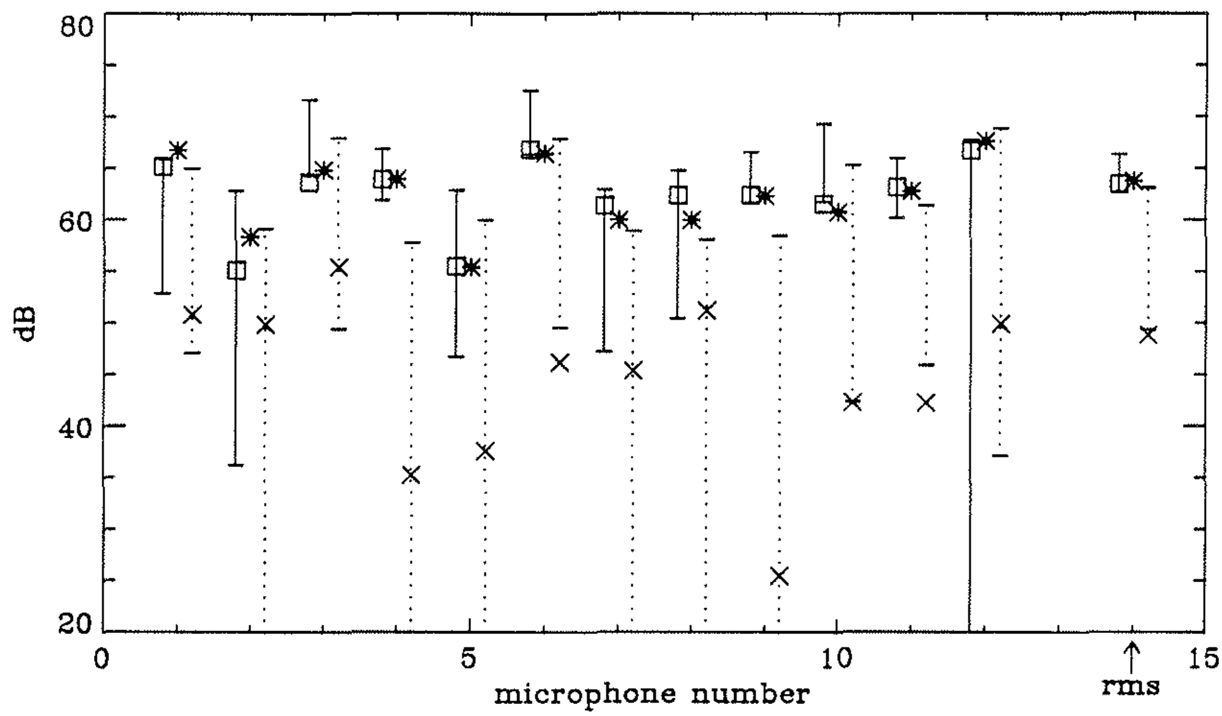

Figure 1a 90% predictions for case F1N1 at 280Hz

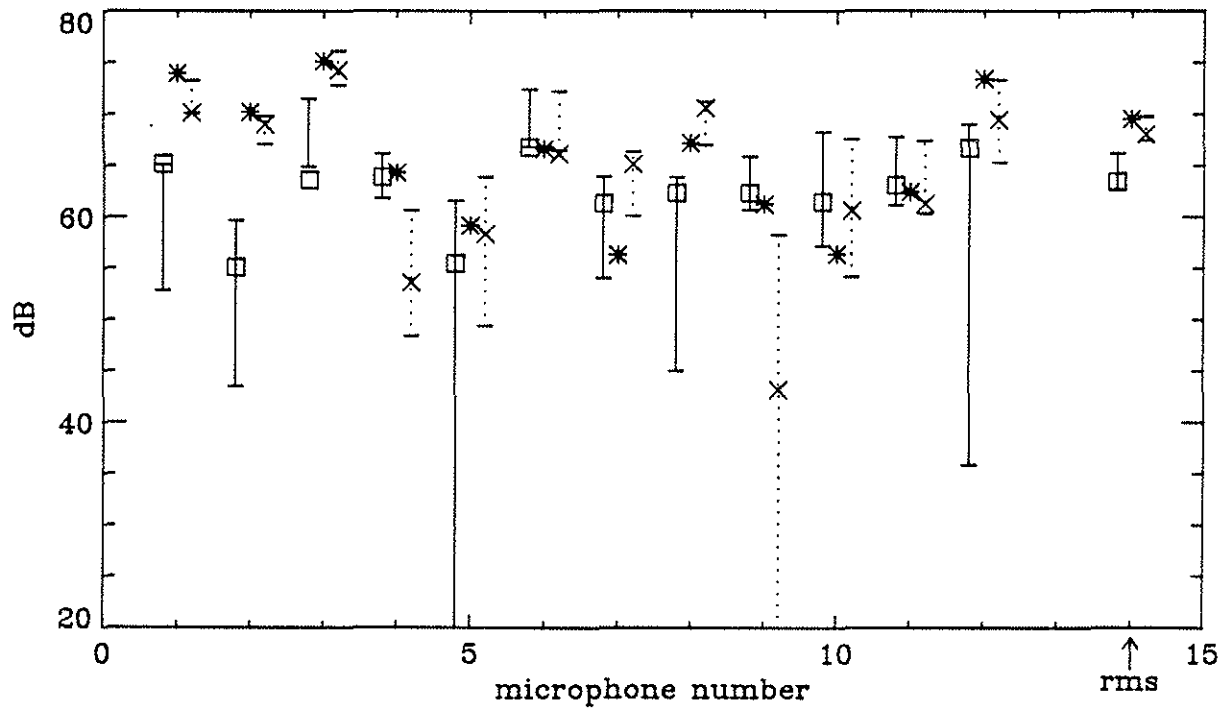

Figure 1b 90% predictions for case F1N2 at 280Hz

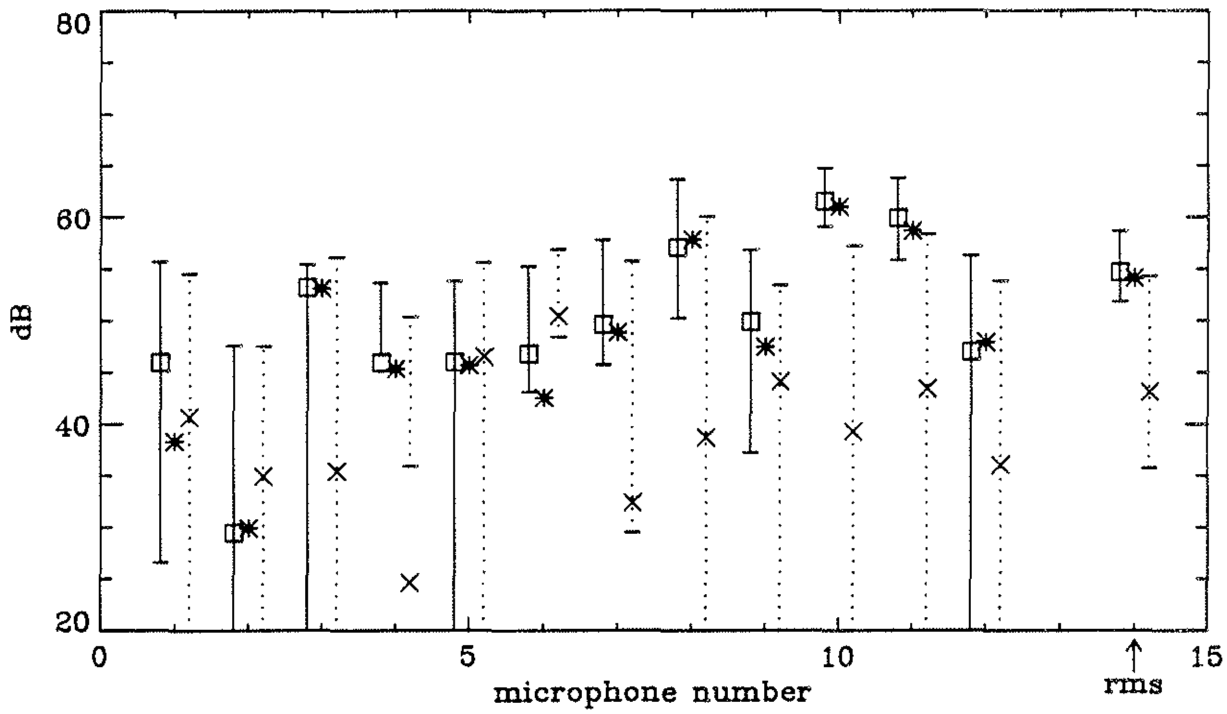

Figure 2a 90% predictions for case F2N1 at 560Hz

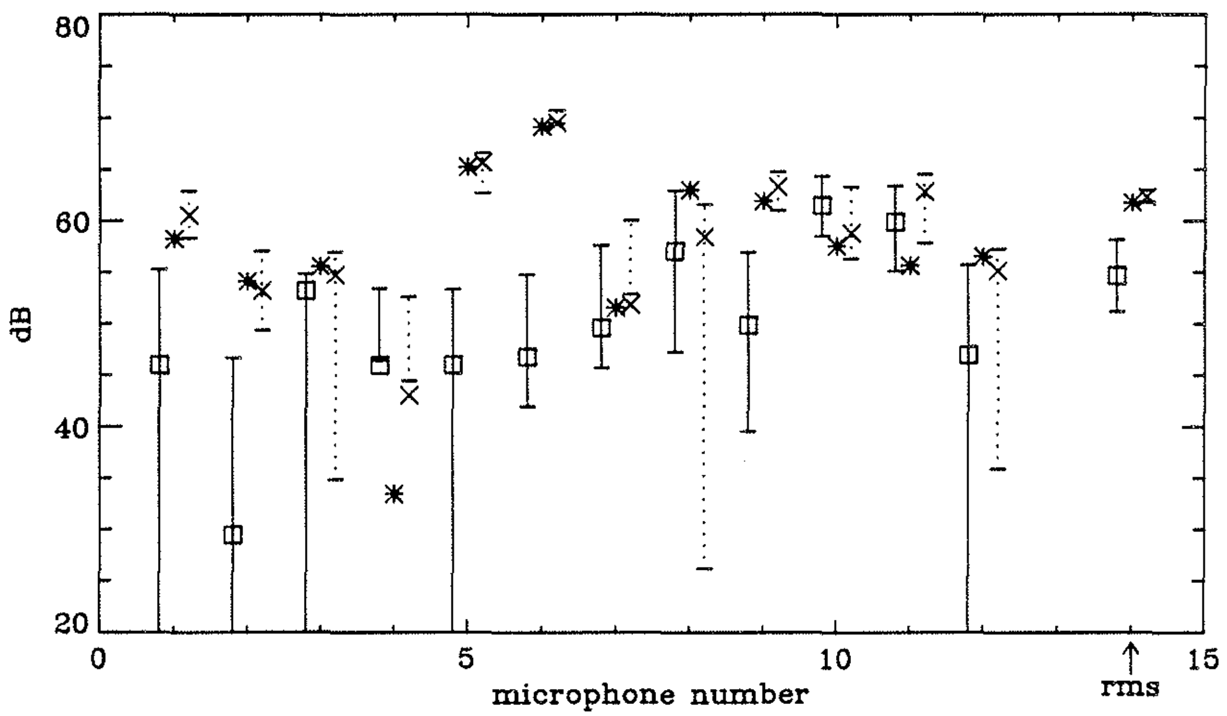

Figure 2b 90% predictions for case F2N2 at 560Hz

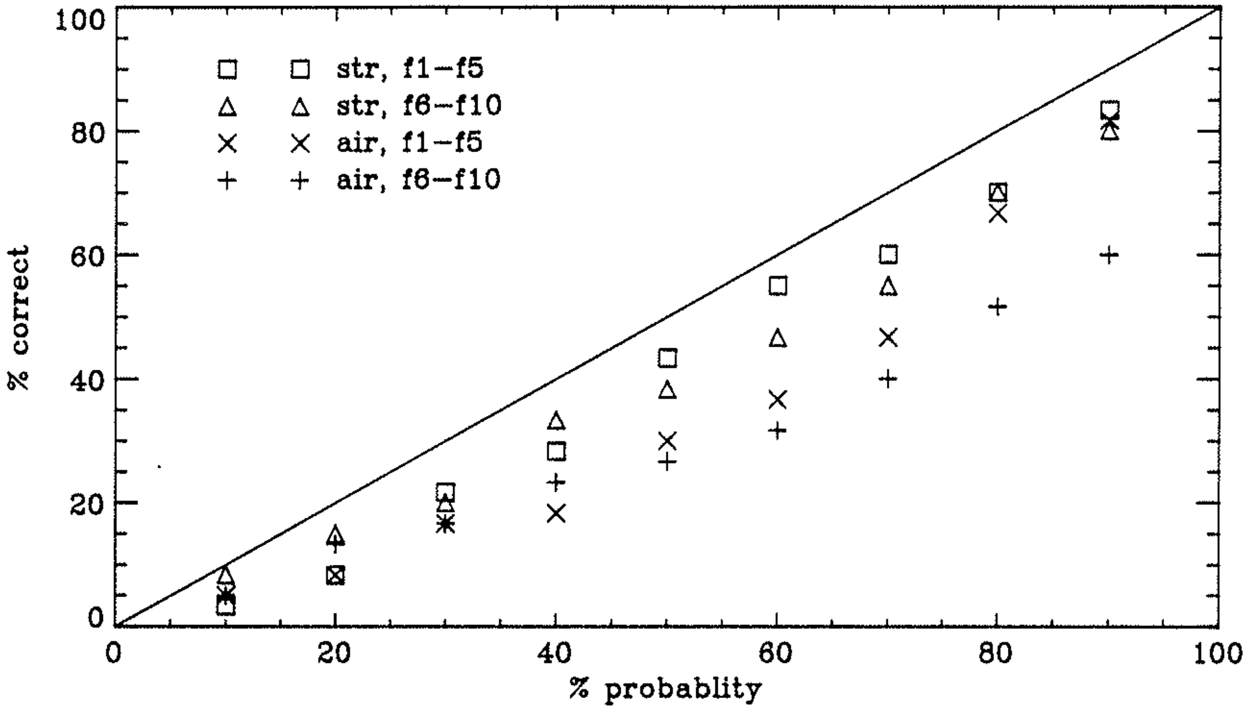

Figure 3a statistical checks for case F1N1

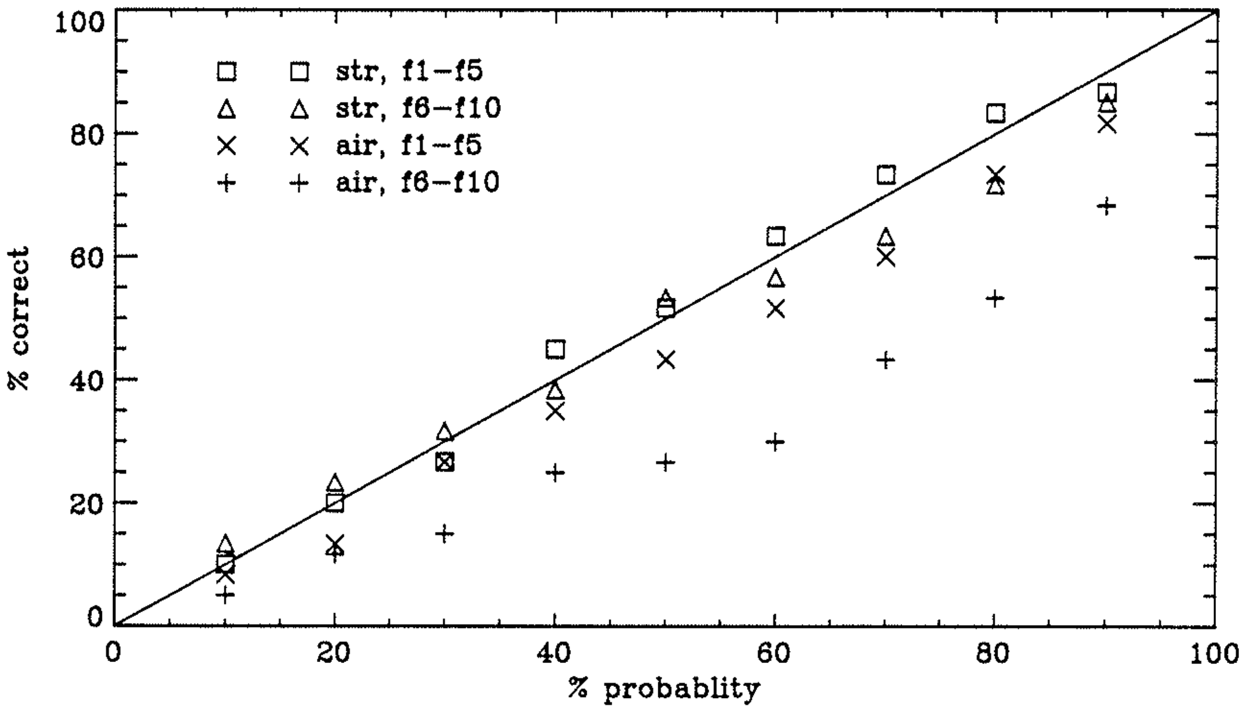

Figure 3b statistical checks for case F2N1

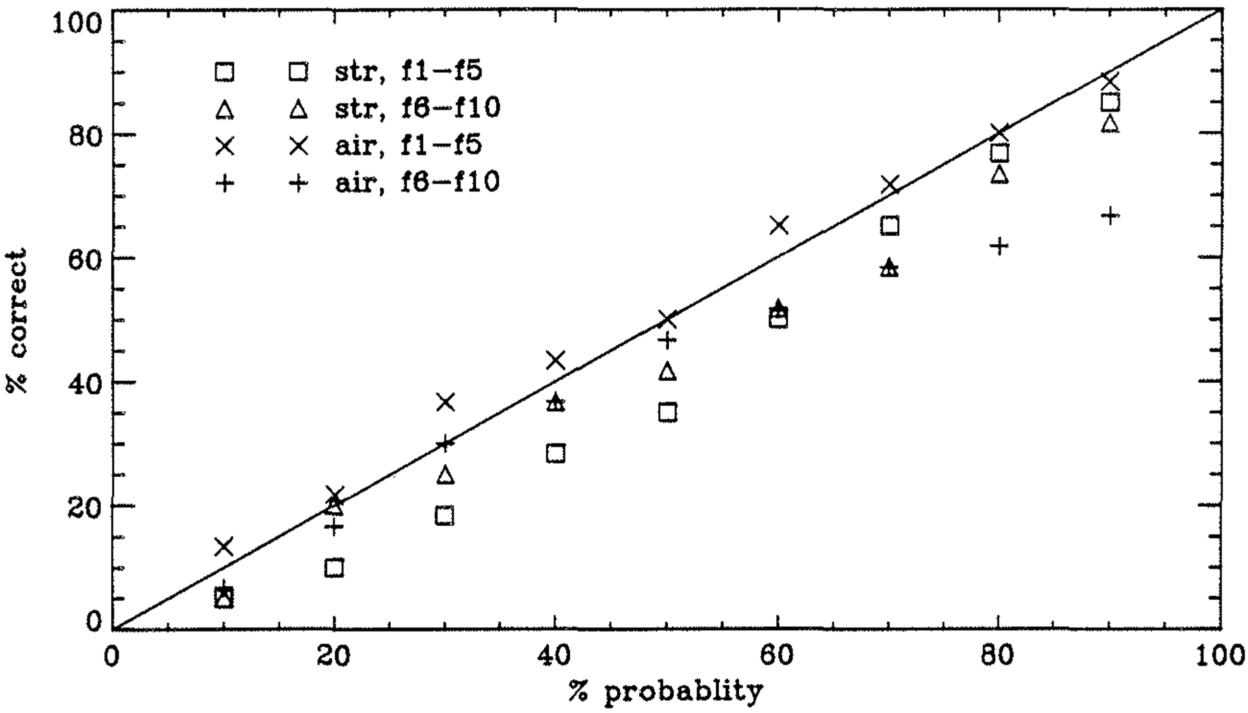

Figure 4a statistical checks for case F1N2

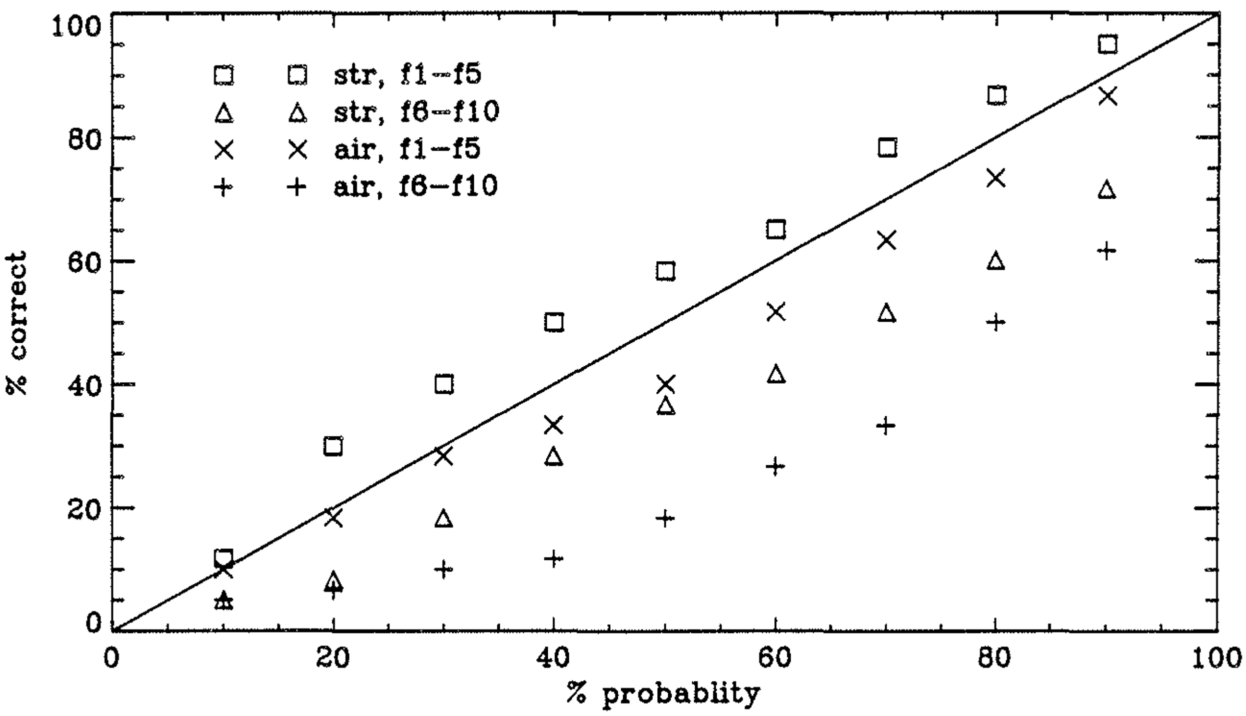

Figure 4b statistical checks for case F2N2

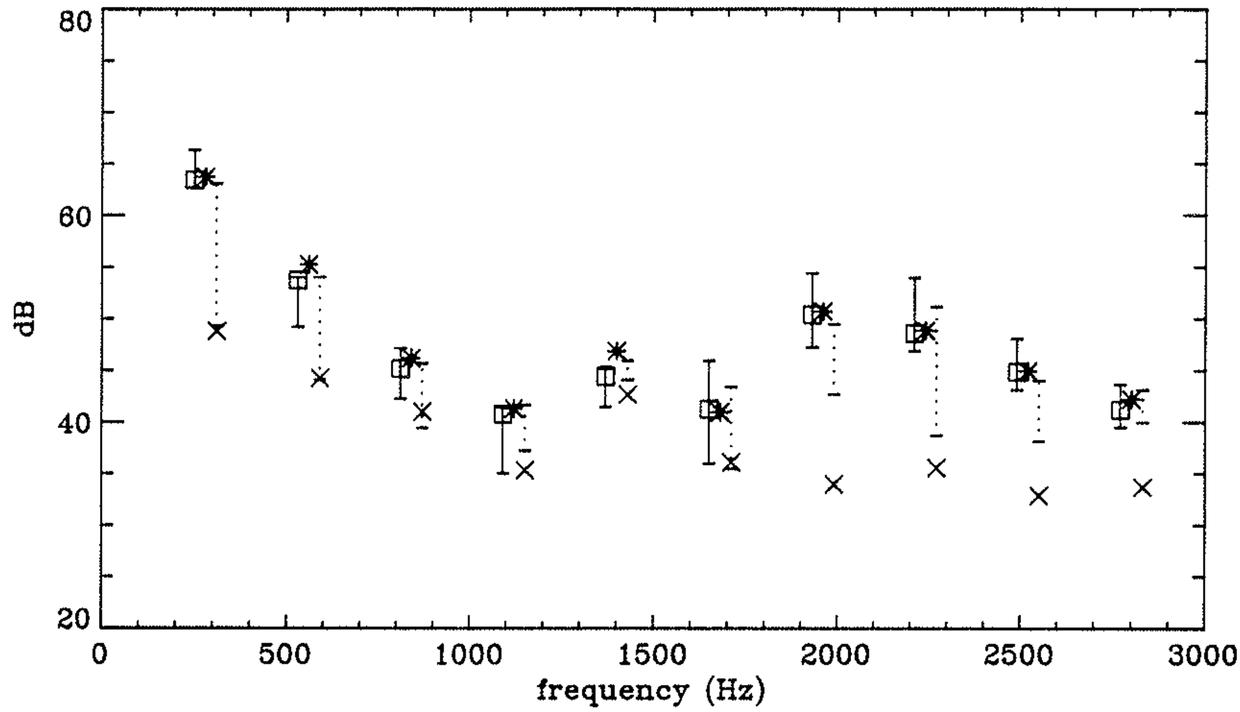

Figure 5a 90% RMS predictions for case F1N1

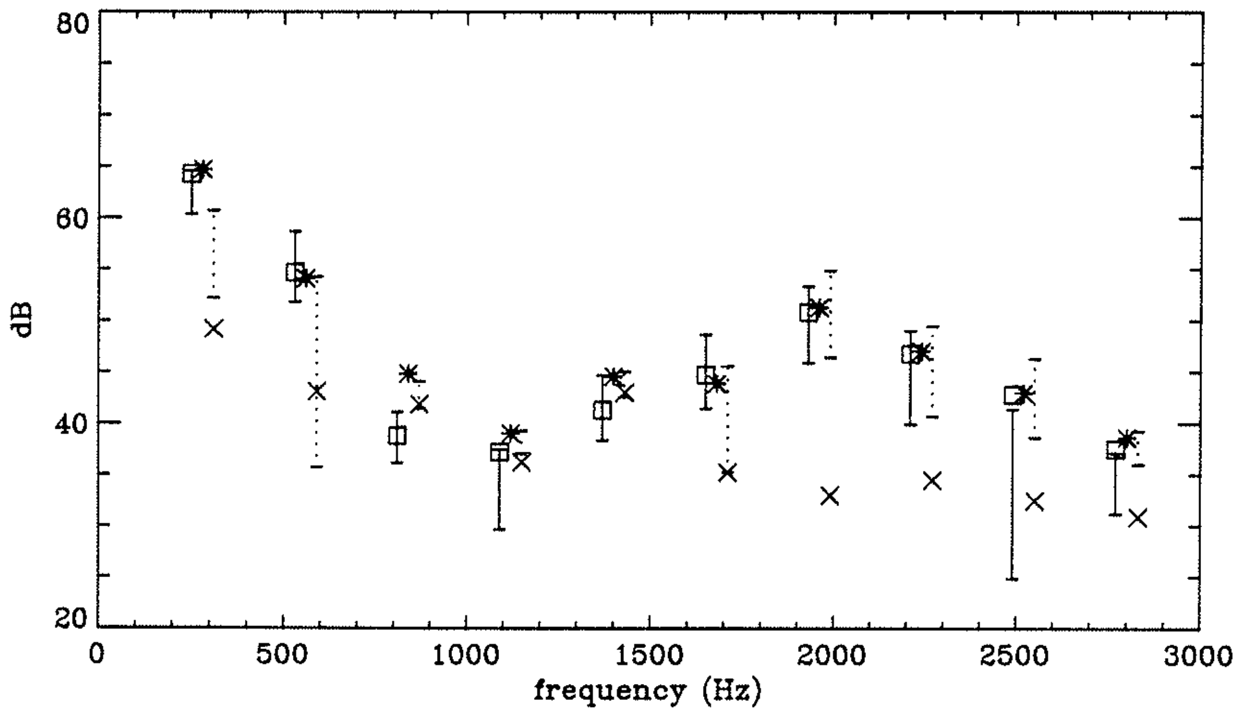

Figure 5b 90% RMS predictions for case F2N1

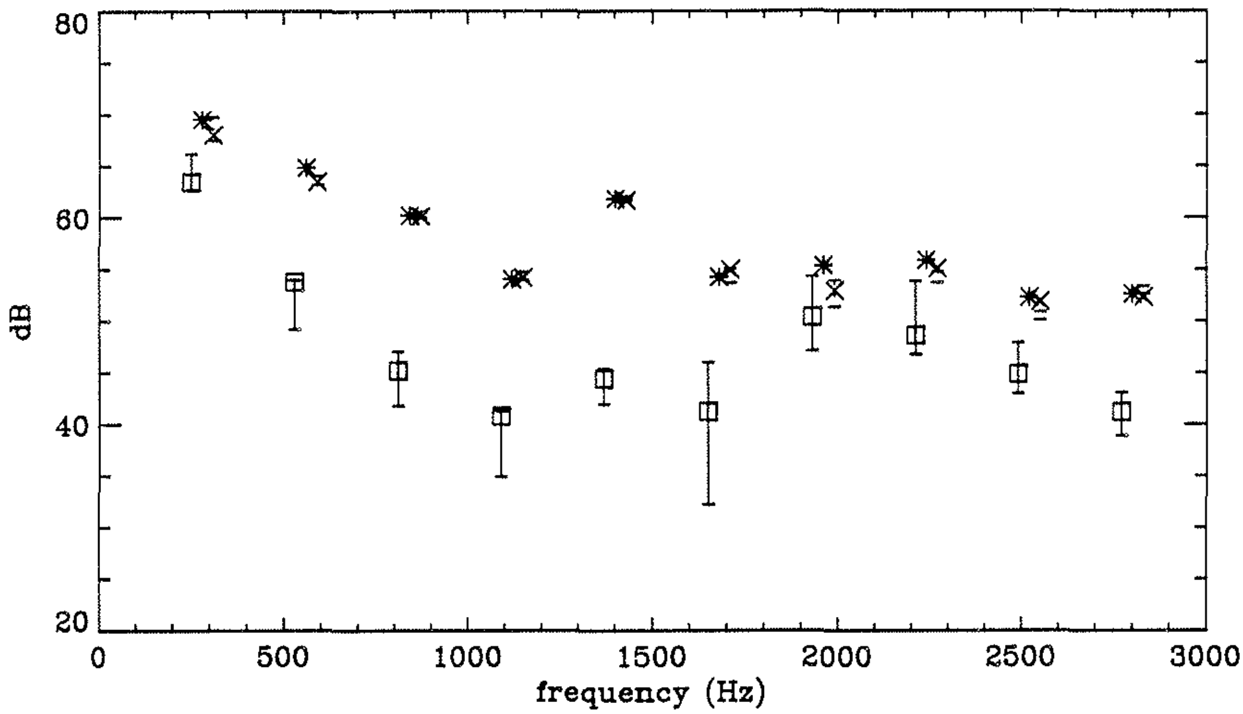

Figure 6a 90% RMS predictions for case F1N2

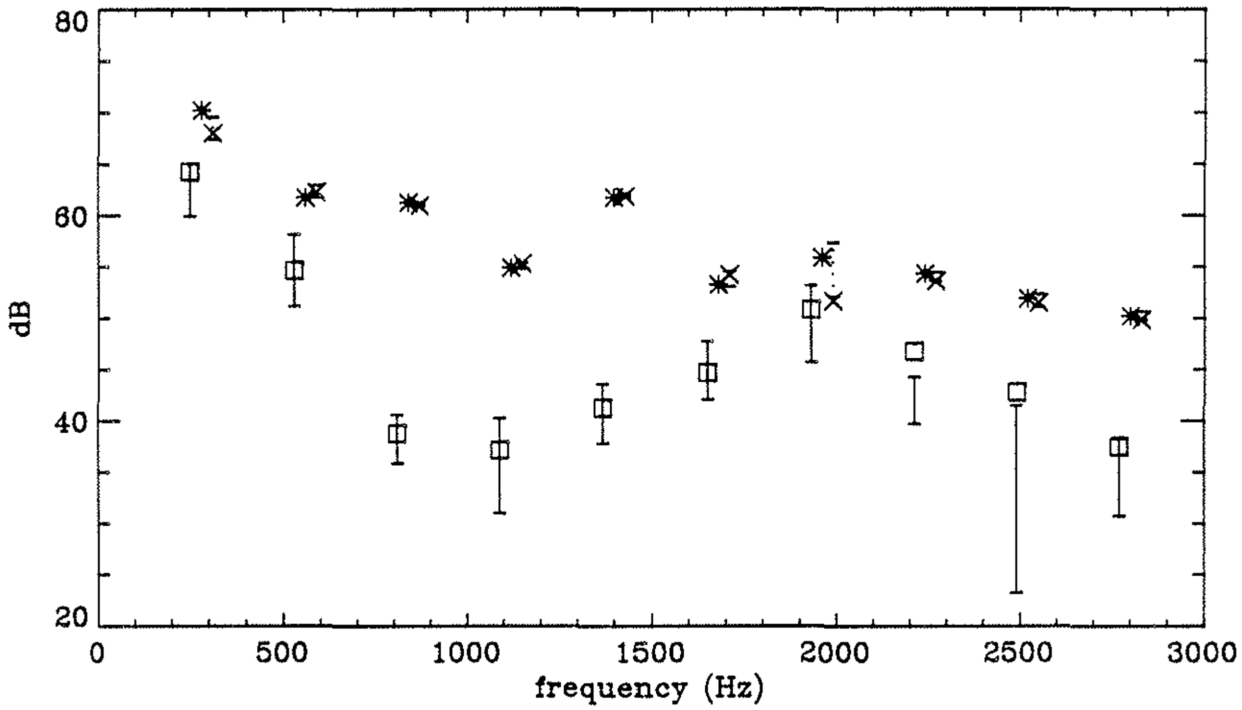

Figure 6b 90% RMS predictions for case F2N2