

## CONFINEMENT OF VIBRATION BY ONE-DIMENSIONAL DISORDER, II: A NUMERICAL EXPERIMENT ON DIFFERENT ENSEMBLE AVERAGES

C. H. HODGES<sup>†</sup> AND J. WOODHOUSE

*University Engineering Department, Trumpington Street, Cambridge CB2 1PZ, England*

*(Received 28 March 1988, and in revised form 26 September 1988)*

When producing theoretical estimates of the effect of irregularity on the vibrational behaviour of a structure, as has been described previously [1–3], one aims to predict the behaviour of a typical irregular structure within the set considered. It might be thought that this would correspond to the ensemble-averaged behaviour over the whole set, but this need not be the case. In fact there are several different possible ways of taking “the” ensemble average, which can produce remarkably different results. Some of these do indeed correspond to the typical member of the ensemble, but some of them certainly do not, since rare configurations giving very large contributions to the average distort the answer. This problem is investigated here via a numerical experiment on a very simple system: the object is to complement the discussion in a companion paper [4], which gives theoretical analysis. The aims here are threefold: to know which average is the most appropriate in a given context; to understand why the other averages are different; and to identify the anomalous configurations making them different so that these can be avoided in designing real structures. A by-product is a better understanding of the averaging procedures used in the past—it has not always been recognized that these various averages can give radically different answers, and one can find instances where what was used was perhaps the easiest approach rather than the most appropriate one. In addition, the numerical experiment described here gives a good illustration of the basic phenomenon of Anderson localization in an acoustic context.

### 1. INTRODUCTION

In three previous papers [1–3] we have introduced and demonstrated experimentally the phenomenon of vibration confinement by irregularity in a structure. In particular, it was shown there that for a one-dimensional structure with slight random irregularity in some of its properties, each normal mode of vibration is localized around some particular region of the structure, on average decaying exponentially as one goes away from that region in either direction. This is in marked contrast to a regular structure, in which all normal modes are extended throughout the structure.

The most important consequence of this difference emerges when we ask whether vibrational energy injected at a given point can penetrate to distant parts of the structure. In a regular structure the vibration can propagate unattenuated arbitrarily far from the driving point for frequencies within a pass band of the structure. In an irregular structure, on the other hand, the energy is confined near the driving point at all frequencies, to a greater or lesser extent. This extent is determined by the degree of randomness in the structure compared to the strength of coupling between adjacent bays of the structure, as discussed in detail previously. The greater the degree of randomness, the more rapid

<sup>†</sup> This paper describes work on which Dr Hodges was engaged up to the time of his tragic death in 1986.

is the attenuation of the response as one moves along the structure away from the driving point, for frequencies within a pass band of the corresponding regular structure.

The calculations of the localization effect given previously (following the extensive literature on the subject in solid state physics) predicted that the response to driving a typical disordered structure would decay exponentially away from the driving point. Quantitative predictions of the rate of decay were given, and these were checked against an experiment which used slightly irregularly placed masses on a vibrating string. The Introduction to that article contains a much fuller account in qualitative terms of the localization phenomenon, and an understanding of it to the level given there is assumed here. While the approach used in those previous papers was that which is standard and well accepted in the solid state applications of Anderson localization, readers from the acoustics field have questioned some of the results, and in particular have pointed out that our "typical" result disagrees substantially with the ensemble-averaged response of the same model.

In this paper and a companion one [4] we give a fuller discussion of this issue, not only to show that our original approach is indeed the most appropriate to the problem in hand but also because investigation of some of the issues raised proves quite illuminating. The fact that different averages (in this case arithmetic or geometric means) can give significantly different answers immediately tells us that there must be a wide spread of behaviour among the configurations included in the ensemble, since the different averaging procedures only represent a difference of weighting between these same quantities. Thus it must be the case that one procedure gives much more weight to the extreme configurations so that, although rare, they contribute significantly to the average. This fact is of practical importance, since among other things it tells us that if we are trying to design a structure to take advantage of the localization effect, we must take great care to avoid those extreme configurations.

We approach the analysis of the various averaging procedures by a combination of theory, given in the companion paper [4], and a numerical experiment which is described here. A very simple model of a chain of ten coupled oscillators has been simulated, for a large number of different realizations drawn from an ensemble whose statistics allow a particular degree of randomness to the natural frequencies of the individual oscillators. The vibrational behaviour of each realization was calculated, and the various kinds of average under investigation were accumulated. The results give an idea of the behaviour of the various averages, including some measure of the variance with respect to the mean. Thus they give some data against which the theory can be tested, and also enable the exceptional configurations which distort certain of the averages to be identified and studied.

The investigation is of significance for several reasons. Firstly, it enables us to find out which average we should be using in theoretical work—our criterion for correctness is that we wish to predict the behaviour of a *typical* disordered member of the ensemble, so that an average unduly affected by large contributions from rare anomalous configurations is to be avoided. Secondly, by giving a deeper understanding of why the various averages give different results we gain insight into the existing work on the subject, in which different averages have been used at different times. Thirdly, the identification of the anomalous configurations is of practical importance since we would wish to avoid them in designing a real structure to take advantage of the localization phenomenon.

One interesting result of the study is that the most "obvious" ensemble average of the response along the structure does not correspond to the response pattern of *any* individual members of the ensemble. The average represents a balance between a large number of cases giving more-or-less exponential decay away from the drive point with a known

exponent, and a small number giving extremely large responses at points distant from the drive. The ensemble average shows an exponential decay with a very different exponent! This result serves as a useful warning against the indiscriminate use of ensemble averages to describe the “typical” behaviour of a system under any circumstances: its implications go beyond the specific area of Anderson localization.

The results discussed here have a direct and very significant application to the approach to vibration analysis known as Statistical Energy Analysis (SEA). We do not discuss that application here, but we have given a short discussion of some of the implications elsewhere [3]. It may well be that the ideas discussed here will lead to a significant broadening of the scope of SEA, and an increased understanding of the meaning of a SEA vibration prediction.

## 2. THE FINITE CHAIN OF OSCILLATORS: A NUMERICAL EXPERIMENT

We present in this section some results of a numerical experiment, together with a qualitative interpretation of them. We consider a model consisting of ten oscillators identically coupled together to form a one-dimensional chain as illustrated in Figure 1. The squared natural frequencies of the separate oscillators are chosen with the aid of a random number generator, from a population whose probability distribution is uniform in a certain range and zero outside that range. For a given realization of such a chain we can compute a variety of measures of the response of the chain to driving, and then by repeating this for many different realizations we can build up averages of these measures. Details of the model and computer program are given in Appendix A.




Figure 1. Sketch of the chain of ten coupled oscillators which constitutes our model.

Ten oscillators may seem a very small system to consider when examining the essentially statistical procedure of ensemble averaging. However, it is not the system itself which is of interest here: our aim is to examine a range of qualitative phenomena concerning the averaging process, and to validate the theoretical modelling of the companion paper. Both of those aims turn out to be well served by this very small system, while keeping the cost of computations within reasonable bounds.

For our first illustration we consider the simplest case. We calculate the squared-amplitude response of each of the ten oscillators to driving by a sinusoidal force of unit amplitude applied at the first oscillator. We calculate both arithmetic and geometric means of these responses over many realizations of the chain (the geometric mean being calculated by averaging the logarithm of the response). The result is shown in Figure 2, for a case in which the driving frequency was at the centre of the pass band of the equivalent regular chain, and the level of distributed damping was such as to produce a  $Q$  factor of  $10^6$  for each mode. The oscillator squared frequencies were distributed over a range whose half-width was 10% of the centre frequency. The coupling between oscillators was such that the pass-band width for the corresponding regular chain was 8% of the centre




Figure 2. Averages of the single-frequency squared-amplitude response to driving the first oscillator in the chain: the upper curve is the linear average and the lower curve is the logarithmic average. Model parameters are given in the text. The averages are taken over 13 000 individual configurations of the chain.

frequency. The graph has position along the chain plotted linearly on the  $x$ -axis and response plotted logarithmically on the  $y$ -axis.

The two curves are strikingly different. The logarithmic average shows a rather accurate exponential decay along the chain. On the other hand, the linear average shows a much less accurate exponential decay of response along the chain, with a lower decay rate and a large offset. Since the two averages represent merely a difference of weighting on the same data, this divergence of the results indicates clearly that there must be occasional configurations with very large responses. These can affect the linear average much more than the logarithmic average, of course, and this would imply that the linear average is controlled by a special subset of the configurations which do not represent typical behaviour of members of the ensemble. The logarithmic average, on the other hand, has some claim to represent typical behaviour—that claim is rigorously justified, at least for certain cases, in the companion paper [4]. In our previous work we have used logarithmic averaging to estimate typical behaviour, for this reason.

As a first test of this idea, Figure 3 shows, on the same scale as Figure 2, the individual responses of 50 configurations out of the 13 000 which made up the average. It is clear that the bulk of the results cluster around the logarithmic average, while the linear average does not represent the behaviour of many, if any, individual members of the ensemble. Indeed, for this particular selection of 50 cases only one individual case exceeds the linear average response, and then not at all points.

Figure 3 enables us to understand one of the features of Figure 2: the smoother curve of the logarithmic average reflects the better-behaved statistics of that average compared with the linear average. The fact that the logarithmic average is not only smooth but is rather accurately an exponential decay lends strong support to the theory of Anderson localization as applied to this simple model in reference [2], and we may note also that the exponent deduced from the graph agrees quite well with the theoretical estimate given there, which was indeed a logarithmic average. This check is described in more detail in Appendix B of this paper.

What are the anomalous configurations which are having such a strong effect on the linear average? A plausible first guess is that they are configurations such that a normal mode frequency happens to fall on or near our driving frequency. In that case, the typical




Figure 3. Some of the individual responses making up the averages shown in Figure 2. The scales are the same as in that figure, and the two averages are superimposed as heavy curves for ease of comparison.

response magnitude should be of the order of the  $Q$ -factor of the mode in question. In other words, we might expect the linear average to depend strongly on the damping assumed in the model. The logarithmic average, on the other hand, should be insensitive to damping provided this is sufficiently small, according to the theory [1-3]. So our next test is to run the program with more damping. Figures 4 and 5 show the corresponding results with  $Q$ -factors of  $10^4$  and  $10^2$  respectively. It is clear that the behaviour is broadly as we have just anticipated, so lending support to the interpretation of the anomalous configurations.

Thus we are led to a picture of the individual response patterns making up the average of Figure 2. Most of the configurations have a response with no modal resonance frequencies near to our driving frequency, and these tend to produce a response pattern exponentially decaying along the chain. These dominate the logarithmic average. Occasionally we obtain a configuration with a mode frequency sufficiently close to our driving frequency




Figure 4. Averages as in Figure 2, but with the modal  $Q$ -factors reduced from  $10^6$  to  $10^4$ . Scales are the same as in Figure 2. Only 100 configurations were included in the average in this case.

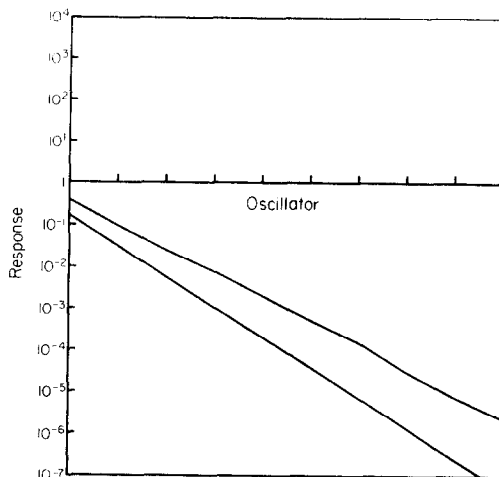


Figure 5. Averages as in Figure 2, but with the modal  $Q$ -factors reduced from  $10^6$  to  $10^2$ . Scales are the same as in Figure 2. Only 1000 configurations were included in the average in this case.

that there is a very large response of the whole chain, with a pattern naturally reflecting the corresponding mode shape. These rare configurations respond sufficiently strongly to dominate the linear average. Now, the typical mode shape is localized about some region of the chain, as discussed at length in the previous papers [1-3]. Thus the individual resonant responses which determine the linear average do not tend to look anything like the average: rather than decaying smoothly along the chain, they grow from the drive point to a peak somewhere, then decay from there to the end of the chain. Sketches of the two types of behaviour are given in Figure 6. The ten mode shapes computed for a single configuration of our model are shown in Figure 7. This qualitative picture of what determines the forms of the two averages justifies, at least heuristically, our earlier remark that the logarithmic average represents the "typical" behaviour of the chain, while the linear average (the usual "ensemble average") does not describe the response pattern of *any* individual members of the ensemble. Something of this behaviour can be seen in Figure 3. If we perform one experiment, we are most likely to see something like the logarithmic average.

We now have a qualitative explanation of why the two averages are different, and this is enough to explain the offset of the linear average relative to the logarithmic one. This

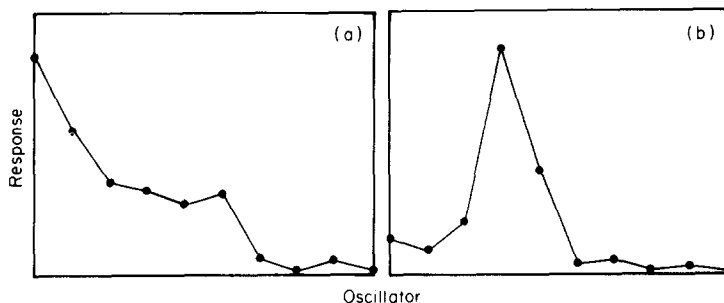


Figure 6. Sketches of (a) the typical response behaviour of the chain when the driving frequency does not correspond to a mode frequency of the chain; and (b) a typical response pattern when the driving frequency corresponds to that of a mode strongly localized in some region of the chain. The averages discussed above represent a balance between a large number of responses like case (a) and a small number of responses like case (b) with much greater amplitude.




Figure 7. A complete set of mode shapes for a sample configuration of the chain. In this case, of course, we plot oscillator displacement linearly on the  $y$ -axis to give true pictures of the mode shapes. It is apparent that we have one mode localized in the vicinity of each oscillator, as we would expect.

arises simply from the difference between resonant and non-resonant behaviour at the driving frequency. It remains to ask why the linear average also has a different rate of decay along the chain. This difference does not depend on damping, as Figures 2, 4 and 5 show. We can obtain a strong clue to what does cause it by evaluating a different measure of the response of the chain to driving. Instead of calculating the squared response of each oscillator to single frequency driving at the first oscillator, we calculate the integral of this squared response over all frequencies. This is the same as the total squared response to driving by a random force with a white spectrum at the first site, and is a measure of how the response of the chain over the whole pass band is modified by the presence of disorder. Details of how it is calculated are given in Appendix A.

Results for the same conditions as Figure 2 are shown in Figure 8. We see that both the linear and logarithmic averages yield smooth curves, neither being very precisely exponential. The slope difference between the two averages broadly follows that of Figure 2. The linear average curve is in fact quite close to the corresponding curve in Figure 2, as we might expect since the process of integrating over frequency has a somewhat similar effect to including many more configurations in the single-frequency average. In particular, the magnitude of the frequency integrated response is of the order of the  $Q$  factor of the modes since the resonant peaks are taken fully into account by the integration process.




Figure 8. Averages over the same set of configurations as Figure 2, of the integral over all frequencies of the squared response amplitude. The scales are the same as in Figure 2. Again, the upper curve is the linear average and the lower curve is the logarithmic average.

The fact that the slope difference between linear and logarithmic averages is present in the frequency integrated curves as well as in the single frequency ones shows that this difference arises from anomalous behaviour of certain configurations as a whole, rather than from mere anomalous *frequencies* as was the case with the offset between the single frequency curves. The first question we should ask about these special configurations is whether they are very rare with very large responses (as was the case for the single-frequency resonant configurations), or fairly common with fairly large responses. Figure 9 shows 50 of the individual frequency-integrated response patterns in a similar way to Figure 3. We see from this that the linear and logarithmic averages both lie well within




Figure 9. Some of the individual frequency-integrated responses making up the averages shown in Figure 8. The scales are the same as in that figure, and the averages are superimposed as heavy curves for ease of comparison.

the band of the individual cases, so that the anomalous configurations in this case are rather common, and do not have especially large responses.

We are not able to specify exactly what these special configurations are, but it has become clear from a detailed inspection of the results for many individual cases that they are in some sense “quasi-regular” configurations. In our average over all configurations of ten oscillators with squared frequencies distributed within given limits, we will from time to time encounter configurations in which all oscillators have frequencies very close together, or with the pattern of oscillator frequencies symmetric about the centre of the chain, or with some more complicated regularity in the oscillator frequencies. All such configurations give anomalously large responses at points distant from the drive. However, since there is a wide variety of such quasi-regular configurations, it is not easy to give any simple argument which gives an idea of their cumulative effect on the linear average. This point is mentioned briefly in Appendix B, and is considered further in the companion paper [4].

We now remark on yet another possible measure of the response of the chain to driving, whose average we might sometimes be interested in: the transmission coefficient from the driving point to other points on the chain for driving at a single frequency. This is simply the squared response at each oscillator divided by the driving point response. In other words, it is the squared response to single frequency driving on the first oscillator with a specified *amplitude* rather than a specified force. The results of linear and logarithmic averages of this quantity for the same conditions as Figure 2 are shown in Figure 10.




Figure 10. Averages over the same set of configurations as Figure 2, of the single-frequency response to driving the first oscillator with unit displacement rather than unit force. The scales are the same as in Figure 2. Again, the upper curve is the linear average and the lower curve is the logarithmic average.

The shape of the logarithmic average curve is hardly surprising: it is simply the same as the logarithmic average of Figure 2, shifted upwards so that the driving-point response is unity. The linear average is perhaps more surprising. From a value of unity at the driving point it jumps up at the second oscillator to a value approaching that of the linear average response of Figure 2, and thereafter decays in a way which follows that response, albeit in a more irregular way. The explanation for this lies in the equations of motion of the chain of oscillators: inspection of these (given in Appendix A) reveals that the

transmission from the first oscillator to a point further down the chain is identical to the response to driving the second oscillator with unit force while holding the first oscillator fixed (given coupling springs between the oscillators of unit strength). Thus the linear average of Figure 10 shows the response to unit-force driving of a chain of length 9, plus a unit response at the first oscillator. The same fact is evident in the logarithmic average: note that the value at the second oscillator in Figure 10 is the same as the driving-point value in Figure 2.

It may not be immediately clear why we should be interested in transmission coefficients rather than responses to unit-force driving, for finite structures such as our chain of oscillators. The answer is that when considering driving of any structure, we have to allow for the mechanical nature of the device which is doing the driving. If the driver has low impedance relative to the driving point impedance of the driven structure, the response will be similar to that for constant-force driving. However, if the driver has high impedance relative to the driven structure, the response will be more like that for constant-amplitude driving. Between these two extremes, unless the driver exhibits resonant behaviour of its own, the response changes from one to the other in a well-behaved manner. Thus we need to know both constant-force response and constant-amplitude response if we are to predict the actual response to driving by a range of realistic drivers.

### 3. RESULTS FOR OTHER DEGREES OF DISORDER

We continue our presentation of the results of the numerical experiment on a chain of coupled oscillators by giving some results for different degrees of disorder. First, Figure 11 shows the six averages we have already seen for the case of 10% disorder, all on one set of axes for easy comparison. Figures 12 and 13 then show the same six averages for the cases of 5% disorder and 20% disorder respectively. The graphs are all shown on the same vertical scale, which is different from that of the previous figures. We can see a number of interesting features in this set of three figures, which go some way to confirm both the correctness of the program and our qualitative understanding of the phenomena at work. All the numbers from which these graphs are plotted are given in Appendix B.




Figure 11. Summary of the six averages plotted in Figures 2, 8 and 10. The vertical scale is now different from those figures, for comparability with Figures 12 and 13.




Figure 12. A set of six averages as in Figure 11, but with half the disorder strength of that figure. The exponential decay of the responses along the chain is thus smaller here. The averages were taken over 13 000 configurations.

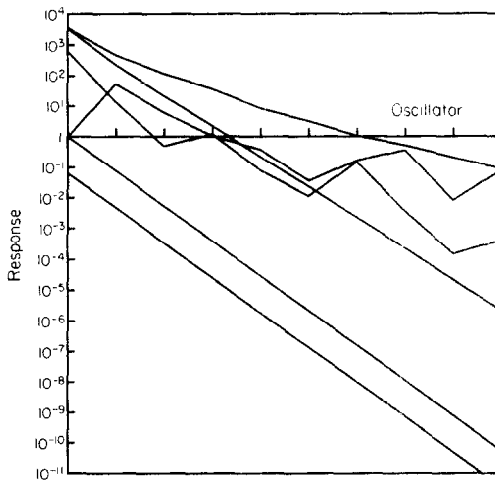


Figure 13. A set of six averages as in Figure 11, but with twice the disorder strength of that figure. The exponential decay of the responses along the chain is thus greater here. The averages were taken over only 1100 configurations in this case.

The behaviour of the logarithmic averages follows our previous discussion. In each case, we see that the driving-point value of the single frequency response average is quite accurately equal to the value of the transmission coefficient at the second oscillator. This obviously results in the driving-point response at a single frequency going down as the level of disorder goes up. This trend is also apparent in the linear averages of the single-frequency response. It is worth considering in a little more detail why this should be, since it goes against the simple intuition that more disorder means more confinement of the energy near the driving point, giving a *larger* driving point response as disorder rises.

Note first that the simple intuition is correct for the case of the frequency integrated curves: the driving point response increases as the disorder increases. It will approach a

limiting value as disorder becomes sufficiently strong. The intuitive argument relies on the fact that the rate of energy input to the structure is independent of the configuration, so that it does not vary in a way which weights the averaging process. This is true for broadband forcing, as explained in reference [5] at the beginning of section 4. The rate of energy input in that case depends only on the mass at the driving point, and on no other aspect of the structure (this remark being true for general structures, not just for the model of coupled oscillators in a one-dimensional chain which we are discussing here). However, it is not true for single-frequency driving, since the driving-point impedance varies greatly among configurations, and so the energy input rate varies similarly. Thus the frequency-integrated response curves follow the simple argument, while the single-frequency response curves do not.

It is not difficult to see roughly why the single-frequency linear average goes down relative to the frequency-integrated response as the disorder increases. We have already seen that the linear average is dominated by the occasional configurations which happen to have a resonant frequency close to our driving frequency. Now, as the level of disorder increases the resonant frequencies of each configuration are spread out over a wider range, so that the probability of hitting one at our driving frequency goes down, provided our driving frequency lies within the pass band of the corresponding regular structure. Thus the total effect on the linear average of these resonant configurations is decreased, and in particular the value of the average response at the driving point goes down for such frequencies. (Of course, for frequencies just outside the original pass band the effect may be very different since these frequencies can now fall within the broadened band.)

Note that the single-frequency logarithmic average also goes down at the driving point as the disorder increases. This is for a related reason: we have seen that the logarithmic averages are dominated by the behaviour between the resonant peaks rather than on them. With increased disorder so that the peaks are spread more widely apart, the valleys between them are deeper so that the typical response at a given frequency in such a valley is decreased.

The next feature of Figures 11, 12 and 13 worthy of note is the shape of the single-frequency linear average curves. With 20% disorder, it is plausible that underneath the statistical fluctuations, this curve is roughly straight in our log-linear plot. However, by the time we have come down to 5% disorder, the curve is by no means straight. It appears to be modulated by an oscillatory shape having some two complete cycles in the length of the chain. This is not hard to explain in general terms. Our driving frequency is at the centre of the pass band of the equivalent periodic chain. This would correspond to a wavelength of four bay-spacings, since the Bloch wavenumber varies across the pass band from zero (infinite wavelength) to Nyquist (wavelength of two bay-spacings), and midband corresponds to the mean of these extreme Bloch wavenumbers. Now with disorder, the mode frequencies are moved around, and our linear average is dominated by those modes whose frequencies occasionally pass our driving frequency. With small disorder, the only modes which can be moved far enough to do this are ones which live close to that frequency in any case: in other words ones with a wavenumber close to the periodic chain midband wavenumber. Conversely with sufficiently strong disorder, the modes are shifted by an amount large compared with the original pass-band width, and by then all modes and hence all wavenumbers are more or less equally probable at our driving frequency. Thus with very small disorder the spatial form of the linear average is dominated by the midband wavenumber, and as disorder is increased this is progressively modified until with sufficiently high disorder the average becomes smooth. The smallest degree of disorder shown here, 5%, is not so small as to make the midband wavenumber alone evident. However it is sufficiently small that only a part of the original pass-band is within

reach of our driving frequency, so we see an intermediate regime in which the curve contains a spread of wavenumbers but is not entirely smooth.

The final comment on details of the figures for various levels of disorder relates to the shapes of the logarithmically averaged curves. With strong disorder, all three logarithmic averages are remarkably straight in the log-linear plots, with hardly any trace of boundary effects near the two ends of the chain. As the disorder goes down this changes slightly. Particularly in the frequency integrated curves, boundary effects become increasingly evident extending towards the centre of the chain from both ends. This is exactly what we would expect. The length scale of any boundary effects will inevitably be of the same order as the modal localization length scale, which is very short for 20% disorder, but occupies a significant proportion of the length of the chain by the time we come down to 5% disorder.

#### 4. CONCLUSIONS

In summary, we have examined the results of a numerical experiment on the averaged response properties of a one-dimensional chain of coupled oscillators. We have seen very clearly that the usual linear ensemble average can be a very poor guide to the behaviour of a typical member of the ensemble, and we have shown that a geometric average over the ensemble gives a better measure of this typical behaviour. By simple arguments we have been able to understand the results qualitatively in some detail, and in the companion paper [4] some of the results are investigated quantitatively. (The companion paper does not deal only with the model of a finite chain of oscillators discussed here, but also with the related problems of a finite number of mass or spring constraints on a vibrating string or beam.) This serves to settle some questions from previous work on the subject, and also to give useful insight into the behaviour of the different averages which extends beyond the particular model we have used in the numerical work. As a by-product, the results of the computer program give useful confirmation of the general phenomenon of Anderson localization for the model studied, which may serve to convince sceptics of the potential relevance of this phenomenon to acoustical applications.

#### ACKNOWLEDGMENTS

The authors are extremely grateful to Professor D. G. Crighton and Drs I. Roebuck and J. F. Scott for invaluable discussions. The work was carried out while both authors were in the employ of Topexpress Ltd, and it was supported by the Procurement Executive, Ministry of Defence.

#### REFERENCES

1. C. H. HODGES 1982 *Journal of Sound and Vibration* **82**, 411-424. Confinement of vibration by structural irregularity.
2. C. H. HODGES and J. WOODHOUSE 1983 *Journal of the Acoustical Society of America* **74**, 894-905. Vibration isolation from irregularity in a nearly periodic structure: theory and measurements.
3. C. H. HODGES and J. WOODHOUSE 1986 *Reports on Progress in Physics* **49**, 107-170. Theories of noise and vibration transmission in complex structures.
4. C. H. HODGES and J. WOODHOUSE 1988 *Journal of Sound and Vibration* **130**, 237-251. Confinement of vibration by one-dimensional disorder, I: theory of ensemble averaging.
5. J. WOODHOUSE 1981 *Journal of the Acoustical Society of America* **69**, 1695-1709. An approach to the theoretical background of statistical energy analysis applied to structural vibration.

## APPENDIX A: SOME DETAILS OF THE COMPUTER PROGRAM

The equations of motion in the absence of dissipation of our chain of ten oscillators are

$$(\Omega_n^2 + 2 - \omega^2)x_n - x_{n-1} - x_{n+1} = \delta_{1n}, \quad n = 1, \dots, 10,$$

where  $x_n$  is the displacement of the  $n$ th oscillator,  $\Omega_n$  is its natural frequency,  $\omega$  is the driving frequency, and the strength of coupling between adjacent oscillators is taken to be unity.  $\delta_{1n}$  is the usual Kronecker delta, denoting driving force applied to the first oscillator. We use boundary conditions at the ends of the chain corresponding to zero motion of oscillators  $n = 0$  and  $n = 11$ . The  $\Omega_n^2$  are selected with the aid of a random number generator, so that they have a uniform probability distribution in the range  $50(1 \pm \alpha)$ , where the values of  $\alpha$  are given in Appendix B. Having calculated a set of ten of these oscillator frequencies, the program then calculates the normal modes and corresponding eigenfrequencies of the chain, using a standard routine for eigenvectors and eigenvalues of a real, symmetric, tridiagonal matrix. For the  $k$ th normal mode, we write  $w_{kn}$  for the motion of the  $n$ th oscillator, and  $\omega_k$  for the eigenfrequency. The response of the whole chain to driving at the first site is then calculated by means of the usual formula

$$Y_n = \sum_{k=1}^{10} \frac{w_{kn} w_{k1}}{\omega_k^2 + 2i\omega\Delta_k - \omega^2}, \quad (\text{A1})$$

where damping factors  $\Delta_k$  are chosen to give all modes the same  $Q$  factor, by  $\Delta_k = \omega_k/2Q$ . The squared modulus of  $Y_n$  gives us our first quantity of interest, as plotted in Figure 2.

The frequency-integrated intensity  $I_n = \int_0^\infty |Y_n(\omega)|^2 d\omega$  can be calculated from equation (A1) by straightforward contour integration, yielding

$$I_n = 2\pi\varepsilon \sum_{j,k} \frac{w_{j1} w_{jn} w_{k1} w_{kn}}{(\omega_j + \omega_k)(1 + \varepsilon^2)[(\omega_j - \omega_k)^2 + \varepsilon^2(\omega_j + \omega_k)^2]},$$

where  $\varepsilon = 1/2Q$ . This can now be computed readily, and is the quantity plotted in Figure 8.

## APPENDIX B: THE RESULTS OF THE PROGRAM AND SOME REMARKS

The numbers produced by the computer program, which were plotted in Figures 11, 12 and 13, are given in Tables 1, 2 and 3. In each table the six columns contain the following, respectively: (i) the linear-averaged, frequency-integrated response; (ii) the logarithmically averaged, frequency-integrated response; (iii) the linear-averaged, single-frequency transmission; (iv) the logarithmically averaged, single-frequency transmission;

TABLE 1  
Results for 5% disorder ( $\alpha = 0.05$ )

2.09E+03	1.94E+03	1.00E+00	1.00E+00	1.01E+04	4.96E-01
7.48E+02	6.65E+02	8.39E+02	6.49E-01	4.36E+03	3.22E-01
4.37E+02	3.27E+02	8.46E+02	3.89E-01	2.82E+02	1.93E-01
2.73E+02	1.86E+02	2.90E+02	2.49E-01	2.23E+03	1.24E-01
1.86E+02	1.11E+02	5.58E+02	1.51E-01	2.61E+03	7.50E-02
1.31E+02	6.94E+01	1.96E+03	9.35E-02	1.05E+03	4.64E-02
9.64E+01	4.42E+01	8.53E+02	5.81E-02	2.74E+02	2.88E-02
7.38E+01	2.96E+01	5.58E+02	3.52E-02	3.91E+01	1.75E-02
5.97E+01	1.96E+01	1.45E+02	1.96E-02	2.75E+02	9.75E-03
5.43E+01	1.41E+01	1.08E+02	2.00E-02	1.54E+02	9.94E-03

TABLE 2  
Results for 10% disorder ( $\alpha = 0.1$ )

2.96E+03	2.82E+03	1.00E+00	1.00E+00	2.04E+03	2.36E-01
6.72E+02	4.78E+02	9.83E+03	2.51E-01	2.31E+02	5.92E-02
2.74E+02	1.20E+02	5.29E+02	6.09E-02	1.44E+02	1.44E-02
1.28E+02	3.58E+01	2.92E+02	1.50E-02	6.54E+01	3.55E-03
6.10E+01	1.07E+01	2.14E+02	3.64E-03	2.00E+01	8.58E-04
3.01E+01	3.23E+00	1.05E+03	9.12E-04	9.69E+00	2.15E-04
1.71E+01	9.84E-01	5.57E+02	2.20E-04	3.39E+00	5.18E-05
8.77E+00	3.02E-01	9.43E+01	5.36E-05	1.76E+00	1.27E-05
5.41E+00	9.11E-02	1.77E+01	1.29E-05	3.21E-01	3.05E-06
3.18E+00	2.80E-02	1.08E+01	3.63E-06	1.89E-01	8.57E-07

TABLE 3  
Results for 20% disorder ( $\alpha = 0.2$ )

3.53E+03	3.44E+03	1.00E+00	1.00E+00	6.25E+02	6.30E-02
4.67E+02	2.26E+02	5.40E+01	7.58E-02	1.35E+01	4.77E-03
1.17E+02	2.15E+01	6.06E+00	5.25E-03	4.84E-01	3.30E-04
3.70E+01	2.24E+00	1.04E+00	3.97E-04	1.20E+00	2.50E-05
8.44E+00	2.16E-01	3.70E-01	2.88E-05	7.85E-02	1.81E-06
3.15E+00	2.29E-02	3.55E-02	2.16E-06	1.11E-02	1.36E-07
1.04E+00	2.25E-03	1.65E-01	1.63E-07	1.48E-01	1.02E-08
5.06E-01	2.25E-04	3.52E-01	1.15E-08	3.56E-03	7.24E-10
1.98E-01	2.12E-05	8.08E-03	7.74E-10	1.47E-04	4.87E-11
8.61E-02	2.13E-06	7.90E-02	5.39E-11	3.94E-04	3.40E-12

(v) the linear-averaged single-frequency response; (vi) the logarithmically averaged single-frequency response. The numbers are presented in computer "E-format": the number following the letter E in each number is the power of ten by which that number is to be multiplied. The response of the driven oscillator (1) is in the first row of each table, and the other oscillator responses are in order in succeeding rows.

The parameters used in these runs of the program were as described in Appendix A, with the addition of a value  $Q = 10^6$  and a squared observation frequency  $\omega^2$  of 52, which is the centre of the pass band of the corresponding regular structure.

An interesting way to interpret some of the information in the above tables is to investigate the dependence on disorder strength of the rates of exponential decay along the chain of the various averages. Theory suggests that the attenuation of response at the  $n$ th site on the chain relative to the first site will be of the form  $(AV/W)^{\beta n}$ , where  $A$  is some constant,  $V$  is a measure of coupling strength ( $V = 1$  in our calculations),  $W$  measures the disorder strength, and the value of  $\beta$  depends on which average is being taken. For the logarithmic averages, we expect  $A = e$  and  $\beta = 2$  when the disorder is strong (see reference [4], equation 12), while for the linear averages  $\beta$  is expected to be no more than unity. The easiest way to estimate  $\beta$  from the tables is to plot the average factor by which the response is attenuated per site against  $W$  on a suitably logarithmic scale: a scaling law for attenuation as a function of disorder in the form given above will then produce a straight line for this plot, the slope of which will be  $\beta$ .

Such a plot is shown in Figure 14, with the logarithmic average of the single-frequency response and the linear average of the frequency-integrated response taken as the best




Figure 14. Plot of the attenuation factor per site against the log of the disorder strength, as described in the text. The points are taken from the computer runs summarized in Figures 11, 12 and 13, whose detailed results are given in Tables 1, 2 and 3. The two lines joining plotted symbols represent the single-frequency, logarithmically averaged response (squares) and the frequency-integrated, linearly averaged response (circles). Also shown on the graph as a dashed line is the theoretical prediction for the logarithmic average based on a strong-disorder limit. This agrees well with the calculated values at the higher levels of disorder. Identical (logarithmic) scales are used for both axes, so that slope magnitude may be measured directly.

representatives of the two types of average for this purpose. It will be seen that the results for the logarithmic average agree well with theory, as  $W$  becomes large (the theory is a large- $W$  limit). For the linear average the three points lie in a plausible straight line, the slope of which is less than unity. We have no way at present of predicting its exact value, but a simple argument given in reference [4] shows that there are at least enough anomalous configurations to explain why  $\beta$  is no greater than unity.