Virtual Acoustic Prototypes: listening to machines that don't exist

Andy Moorhouse (1)

(1) Acoustics Research Centre, University of Salford, Manchester, UK

ABSTRACT

The concept of Virtual Acoustic Prototypes (VAPs) is explained and is illustrated with examples. A VAP is a computer representation of a machine (e.g. a lawnmower), such that it can be heard without it necessarily having to exist as a physical assembly. It is argued that, whereas visualisation tools are well developed in the field of visual design, equivalent tools for auralisation, such as VAPs, are still in their infancy. Examples of VAPs for a fridge, a telecommunications base station and a washing machine are presented, through which it becomes clear that considerable sophistication is required to include all the various excitation and transmission mechanisms found in real machines. It is explained that VAPs cannot be purely 'virtual' and that some measured data will be needed for the foreseeable future, particularly to characterise active components. Some of the advantages of working with VAPs are outlined.

INTRODUCTION

One of the most difficult and interesting aspects of acoustics is that it spans two realms, the physical and the psychological. In acoustics, unlike most other technical disciplines, we deal with 'causes' that are physical, but 'effects' that occur literally between the ears of the listener, and are primarily psychological. Whereas the field variables in the physical domain are precisely measurable quantities, the 'field variables' describing human reaction are of a quite different nature. The feelings evoked by listening to a particular sound are not easily measurable or even repeatable. Furthermore, they are not necessarily related to the physical variables in a straightforward way. For example, most indicators for environmental noise now use the A weighted Leq, effectively quantifying the average energy of the sound after filtering by the ear. These work well where the character of the sound is fairly constant as for example in the case of traffic noise. However, it is well known that the $L_{A,eq}$ is not a reliable method to compare sounds with a different character. For example, an impulsive sound could be judged as being louder than a steady sound even when it contains less acoustic energy. For these reasons, in the automotive sector, and increasingly in domestic products, designers are coming to the conclusion that the only reliable way to judge the effect of a particular sound is to listen to it.

Let us look at this problem from the point of view of a designer of say domestic or outdoor equipment. The designer wants the sound to evoke positive feelings in customers and other listeners. The design targets are therefore in the psychological domain (essentially they are 'feelings'), whereas the parameters under the control of the designer are physical quantities of quite a different nature (thickness of plates, type of material etc). How then do we assess the effect of a design change (physical) on the target quantity (psychological)?

Architects and visual designers face a similar problem: the feelings evoked by a particular shape can best be judged by looking at the shape, and a wide range of visualisation tools has been developed ranging from hand sketches, drawings, physical scale models, computer models, prototypes, and, increasingly, virtual tools like virtual prototypes and virtual environments. These tools all aim to convey an accurate impression of the looks of the product to an appropriate level of accuracy and detail depending on the stage to which the design has progressed. The level of detail needed in the visual representation may vary from a sketch taking under a minute, to a full 3D model taking weeks or months to assemble.

What tools of a similar nature are available for acoustic design? The answer at current state of the art is very few. In concert hall design auralisation techniques have been available for some time whereby the reverberation in a computer-modelled hall can be added to music recorded under anechoic conditions. In the automotive sector, simulations of the sound and vibration of vehicles are nowadays produced prior to prototypes becoming available to give an impression of what the driving experience will be like. Looking at these two examples it is fairly clear that the level of technology required to achieve the acoustic equivalent of an architect's model is fairly sophisticated. This is surely a main reason why acoustic design tools are years behind those used for visual design.

It can be argued that, if acoustic design is to succeed in shaping sounds then designers need design tools, i.e. techniques to convey an accurate impression of the sound of a product whilst still on the drawing board. Furthermore, we can argue that an <u>array</u> of tools is needed to cater for different stages in the design process, ranging from simple methods to more 'high level' techniques. In this paper we will explore the use of 'virtual acoustic prototype' techniques mainly in the domestic and outpoor products sector. The examples of VAPs given were developed during the recent Nabucco project, funded by the EU.

WHAT IS A VIRTUAL ACOUSTIC PROTOTYPE?

For the purposes of this paper a Virtual Acoustic Prototype (VAP) will be considered as a computer representation of a machine, e.g. a washing machine, fridge, lawnmower etc, such that its sound can be heard without it necessarily having to exist as a physical machine. Like a real machine, a VAP is constructed from 'components', although the VAP components do not necessarily correspond to the physical components of the real machine. Each constituent part of a VAP is a *representation* of some vibro-acoustic mechanism taking place within the machine. We will attempt to explain what is meant by this by using the example of concert hall auralisation, which is more familiar than VAPs.

Substructuring into active and passive parts

Figure 1 shows a schematic of how acoustic designers can listen to a concert hall while still on the drawing board. The main steps in this process are:

- A recording is made of a musician playing in an anechoic chamber
- The impulse response function of the hall is measured, or calculated by numerical methods such as ray tracing
- The anechoically recorded music and the impulse response function are combined in the computer by convolution and the result is auralised, for example by playing over headphones.

Thus, the original music and the reverberance etc. of the room are both heard on the auralisation. This approach is now fairly well known and fairly widely used.

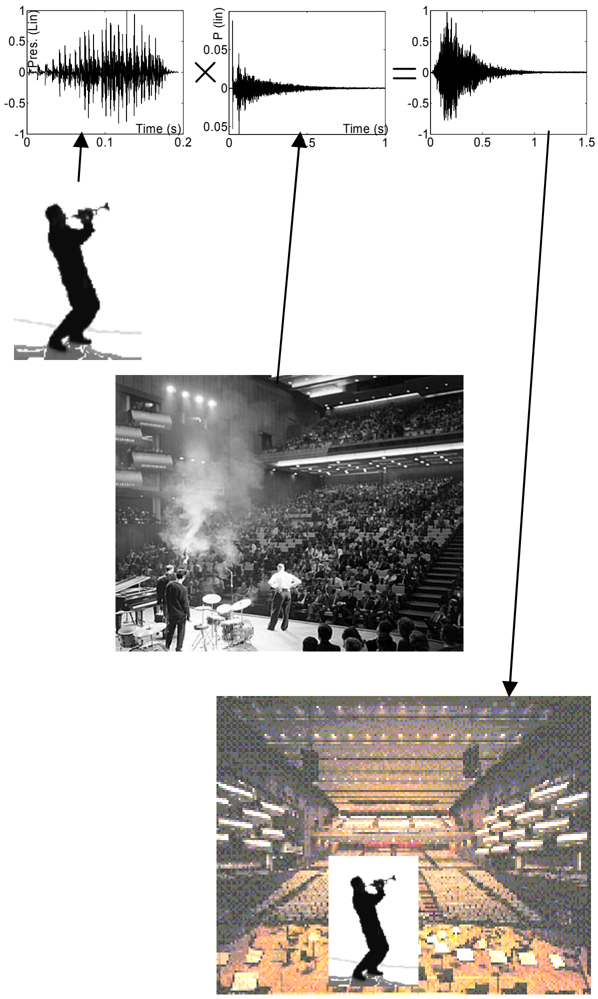

Figure 1 Auralisation techniques are now widely used for concert halls: music is recorded in anechoic condictions; the impulse response of the hall is measured or calculated (centre photo shows the Royal Festival Hall, London being tested with a pistol shot); the music and room response are combined in the computer (courtesy Prof Trevor Cox).

One of the features of this technique, which is relevant for VAPs, is that the data representing the source (the music) and the room are independent of one another. Thus, source and response data can readily be interchanged so that for example a musician can be heard 'playing' in a hall they have never visited.

Constructing a VAP is similar in that the first step is to separate the sources from the remaining passive parts of the machine. The components in a VAP, as in the concert hall example, are therefore of two types: active and passive. Active components are associated with physical components that initially generate the excitation, e.g. fans, pumps, compressors, electric motors etc. All remaining parts of the machine are categorized as passive and are collectively termed the 'frame'. The frame does not generate excitation, but modifies that from the active components on its way to the receiver (listener) position. The frame can be thought of as a filter, attenuating or amplifying certain frequencies by the action of resonances, diffraction etc. Thus, like the concert hall, it plays an important role in what is heard by the listener without generating any initial disturbance itself.

An important requirement for the VAP is that the data chosen to represent the source must be independent of the frame and vice versa. If this requirement is met then each source can be combined with any frame and vice versa. In the concert hall analogy we saw that this allows musicians to 'play' in any hall, existing or not. In the case of the VAP, it allows a source, say a fan, to be installed (virtually) into any frame and the sound of the assembly to be heard.

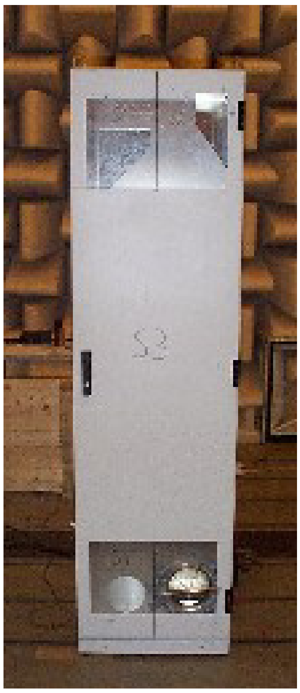

Figure 2 A telecommunications base station with cooling fans in the door (from Moorhouse et. Al. 2003)

Figure 3 Inside the base station door, showing two fans (from Moorhouse et. al. 2003)

A simple example

A fairly simple example of a VAP is shown in Figure 2 and Figure 3. The machine to be represented is the cabinet of a telecommunications base station, in the door of which are mounted two nominally identical fans as part of the thermal control system. This example is almost exactly analogous to the concert hall case: the fan (the source) corresponds to the musical instrument, and the ductwork and cavities to the hall; the only major difference was that the calculations were carried out in the frequency domain. To model the base station, a fan was removed and mounted in an acoustically transparent box, using the methodology of ISO10302 (1996) so that the sound power could be measured in the absence of the frame (see Figure 4). From the measured sound power, the 'source strength' was back-calculated. What is meant by source strength? In this case it was taken to be the net fluctuating force acting on the air by the impeller, which was assumed equivalent to the force exerted by a single dipole acting at the centre of the impeller along the axis of the fan (as illustrated in Figure 5). The frame was then represented by the transfer function from this excitation point to the external receiver position (Figure 5). This transfer function included the combined effects of cavity/ duct resonances, transmission from the end of the duct etc. and was measured using a reciprocal technique (Moorhouse et. al. 2003, see also Pavic et. al. 2003).

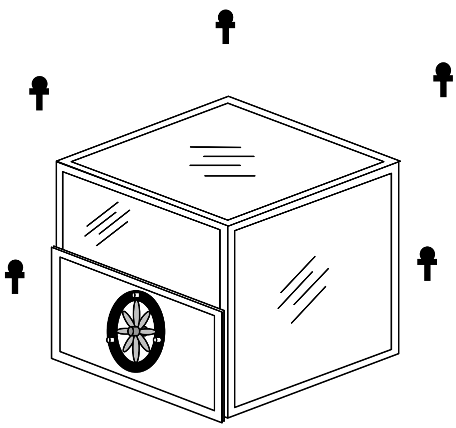

Figure 4 The fan is removed from the frame and tested in an acoustically transparent box according to ISO10302 (1996)

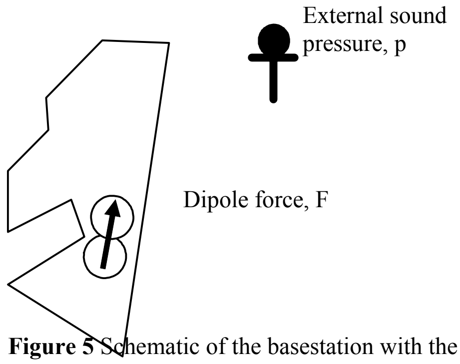

Figure 5 Schematic of the basestation with the fan represented by an equivalent acoustic dipole at the fan centre. The source strength is the dipole force and the transfer function the external sound pressure per unit dipole force.

The spectrum of sound pressure at the receiver location, calculated by combining the source strength and transfer function, is shown in Figure 6 and agrees well with the directly measured sound pressure spectrum. Furthermore, the auralised sound at this position was similar to the real sound (it will be described later how the auralisation was achieved from frequency domain data).

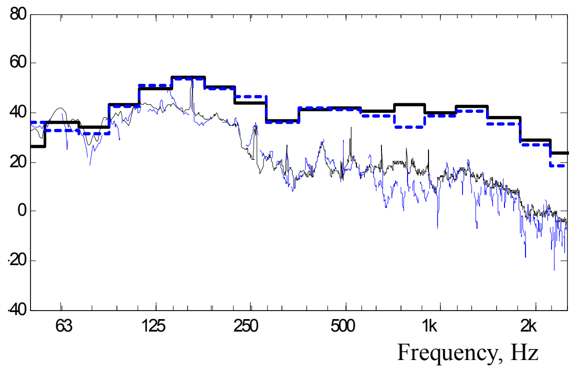

Figure 6 Sound pressure level measured on the real prototype (solid lines) compared with that predicted from the VAP (dotted lines). Narrow band in dBlin, third octave band values in dB(A).

In section 3 a more complicated example of a domestic fridge will be described. First however, we will examine some of the practical limitations on VAPs.

Practical limitations on VAPs

The 'holy grail' for designers is to carry out all their design in the virtual domain, ultimately working purely from electronic data and avoiding physical prototypes completely until the design is finalised. The advantages in terms of cost and time-to-market of this approach can be huge. However, it may have been noticed in the above example that the two sets of data, representing the active and passive components, were obtained by measurement. Clearly, this implies that the machine, or at least the parts of it, must already exist. What then are the advantages of the approach, and furthermore, how can we even claim that our model is a *virtual* prototype at all?

Before answering this question we must consider the realities of the situation: auralisation of machines 'from scratch', i.e. starting from drawing board data is a very ambitious aim, and is still many years off.

One reason is that auralisation places particularly high demands on the numerical prediction of sound transmission and radiation since neither finite element methods nor statistical energy analysis can generally cope with the entire audible frequency range. Hybrid methods (Shorter and Langley 2005) show some promise of overcoming this hurdle, but even so, many machines are likely to prove too complicated for some time. Take the example of a washing machine frame: to calculate the transfer function from inside to outside requires an acoustic model of the cavity, coupled to a structural model of the cabinet (with its own difficulties of modelling), coupled to another acoustic model of the external air. Hybrid methods could simplify this considerably, but for the foreseeable future it is unlikely that reliable numerical methods will be available to designers in 'medium technology industries like domestic and outdoor products.

Despite the difficulties, it is even now possible to model some machine frames numerically. An example is the base station shown above, where it proved possible to calculate the transfer function from a boundary element model of the ducts and cavities. This purely calculated data was then combined with the source strength data, which was based on the measurement. The predicted sound pressure spectrum at the receiver was similar to that shown in Figure 6.

The second reason that modelling from scratch is ambitious is that understanding of source mechanisms is not sufficiently advanced. It is beginning to become possible to predict the sound generated by some sound sources, for example fans, although this is likely to remain a highly specialised job for some time at least. However, other mechanisms, like stick slip friction are still not sufficiently understood for reliable models to have been developed. Therefore, even when transfer functions can be modelled numerically, representations of the active components will almost invariably be based on measurement for some time to come.

Advantages of VAPs

What are advantages of a VAP if measurements are still needed? Not much if one only has a single component and a single frame, one might as well assemble them and listen to the real machine. However, once one starts to assemble a database of several active components and several frames the possibilities start to become interesting. Take for example a washing machine. This is typical of many products where complete redesign from scratch is rarely needed. Rather, a new design is usually a gradual evolution from a similar existing design. Therefore, many of the components in the new design will already exist and their vibro-acoustic properties can be measured.

Consider the apparently simple choice of how to select the optimum washing machine motor from the many possibilities on the market to suit a particular frame? The potential advantages of a virtual approach in terms of cost and time are clear. It is less obvious that the virtual approach could actually be more reliable than real prototypes. This is because to make a comparison one needs to keep everything constant except the variable of interest, i.e. in this case one should change the motor and keep the frame constant. This is easier said than done: one cannot guarantee that the frame will be identical before and after disassembly. Indeed there is growing evidence that small structural changes can bring about significant differences in vibro-acoustic behaviour (see for example Kompella and Bernhard 1993). For example, a slight change in the angle of connected plates can alter sound radiation significantly (Rebillard and Guyader 2000). On the other hand, using a virtual approach one can guarantee that the frame is *identical* by using the same transfer function data. Then one can listen to different motors in exactly the same frame. Naturally, the sound of the VAP cannot be a better likeness than the real thing in an absolute sense, but here we are interested in differences before and after modifications, i.e. *relative* effects, and here the VAP has some advantages over a conventional approach.

Perhaps the biggest advantage of the VAP approach is the insight that comes from breaking down the machine in a systematic way. For industrial participants in the Nabucco project this was considered the most useful aspect of the approach. The systematic approach can perhaps best be illustrated by an example, which is given in the next section.

AN EXAMPLE VAP: A DOMESTIC REFRIGERATOR

The previous example of the base station was a relatively simple one since only airborne excitation of the frame was involved. A refrigerator is considerably more complicated since the excitation from the compressor is a combination of airborne, fluid-borne and structure-borne.

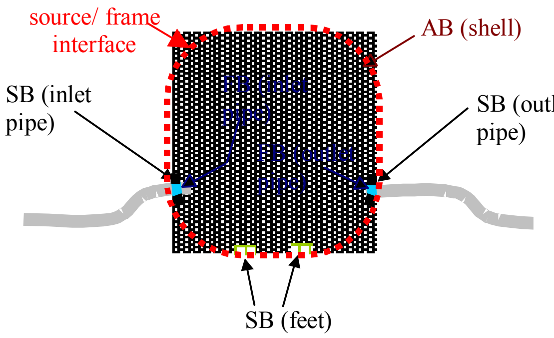

Figure 7 Source/ frame interface around the compressor showing structure-borne (SB), airbonre (AB) and fluid-borne (FB) excitation of the frame

General scheme for the refrigerator

These excitation types can be seen when we sub-structure the complete machine into source and frame. Where the source stops and the frame starts is sometimes arbitrary, for example the refrigerant pipes could be considered to be part of either. To define the substructures, a boundary is drawn around the source (Figure 7) so that all sound generating mechanisms are inside and what is outside is purely passive. The excitation of the passive frame is then considered to be purely through excitation over the interface. Where the interface cuts through solid structures, the excitation is considered to be structure-borne (SB) in origin, where it intersects fluid in a pipe it is fluid-borne (FB), and where there is air on the frame side of the interface, it is considered airborne (AB). Applying this logic the excitation of the frame is a combination of:

- direct sound radiation from the hermetic shell (AB)
- vibration in discharge and suction pipes (SB)
- gas pulsation in discharge and suction pipes (FB)
- vibration of the compressor feet (SB).

The scheme for the VAP is then as laid out in Figure 8. Although in the physical prototype there is only a single active component, in the VAP there are four sources, corresponding to the above excitations, which must be treated separately.

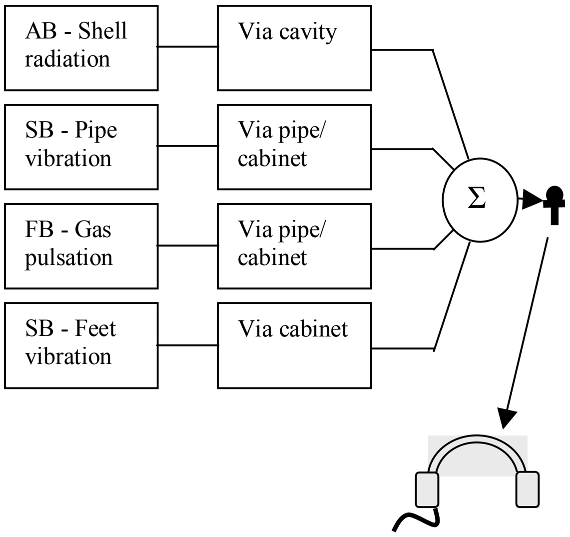

Figure 8 Vibro-acoustic scheme for refrigerator

Note that energy can be converted between different forms on its way to the receiver location. For example the FB excitation of the fluid in the pipe must be converted to pipe wall vibration and then into sound waves in the surrounding air in order to be heard at the listener location. All these effects are grouped together as FB sound because this is the nature of the initial excitation. Similarly, SB sound starts as vibration, but must be radiated into the surrounding air to be heard. In the next subsections we will consider how each of the excitation types is dealt with in constructing the VAP.

Airborne sound

As described above, in order to construct the VAP we need to find a way to *represent* the excitation and transmission mechanisms. In this subsection we consider how to find the data to fill the top line of boxes in Figure 8. For AB sound the real compressor shell is idealised as a vibrating and breathing sphere. In other words we represent the true shell by a combination of a monopole and three perpendicular dipoles as shown in Figure 9.

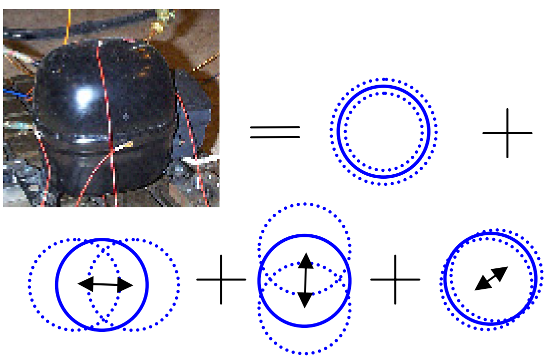

Figure 9 The real compressor shell is idealised as a combination of a monopole (breathing sphere) and three perpendicular dipoles (oscillating sphere) for AB sound.

The strengths of these elementary sources are obtained using vibration measurements. It turns out that the above four types of motions can be extracted from the readings from six accelerometers positioned on the poles of the principal axes as shown in Figure 10. The strength of the monopole is given by the average in-phase outward acceleration:

$$Q_{monopole} = 4\pi r^2 \left(A_1 + A_2 + A_3 + A_4 + A_5 + A_6\right)/6j\omega \quad (1)$$

where $Q_{monopole}$ is the source strength of the monopole, A_1 etc. are the accelerations (Figure 10), r is the radius of the sphere and ω the radian frequency. The strength of the x direction dipole is proportional to the rigid body acceleration in the x direction, i.e.

$$D_x \propto \left(A_1 - A_2\right) \quad (2)$$

where D_x is the strength of the x direction dipole (similarly for y and z). We end up with four frequency spectra representing the strength of the monopole and the three dipoles.

Thus, the AB source strength is characterised without the need to perform any acoustic measurements. This indirect measurement has several advantages over a conventional sound power measurement, not least that the measurements are almost completely immune to background noise. The use of four equivalent elementary sources is also more accurate than using the compressor sound power, because the sound power is not an invariant quantity when there are reflecting surfaces and cavity modes in the near field of the source. For example anti-symmetric modes couple well with a dipole with the same alignment, thereby amplifying its contribution, an effect which is not accounted for using a sound power approach.

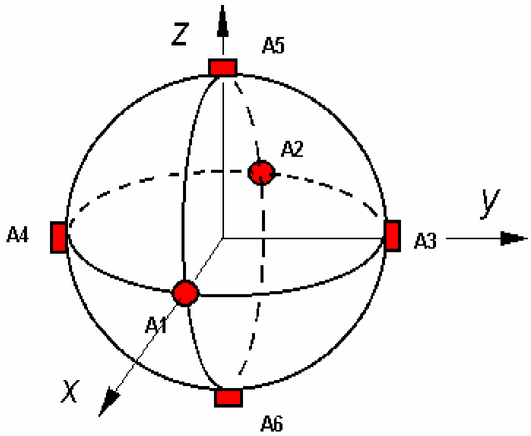

Figure 10 Six accelerometer arrangement used to measure AB and SB source strength of the compressor shell

Next we consider the transmission from the compressor shell to the external receiver location. This is represented by transfer functions, one for each of the four equivalent sources of Figure 9. The external sound pressure (due to AB excitation) is then given by:

$$p_{AB} = Q_{monopole} H_{monopole} + D_x H_x + D_y H_y + D_z H_z \quad (3)$$

where $H_{monopole}$ is the transfer function for the monopole and H_x etc. for the dipoles. This equation is shown diagramatically in Figure 11.

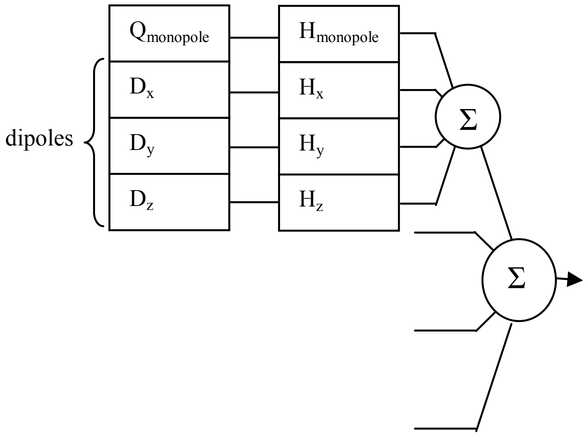

Figure 11 Development of Figure 8 showing contributions to the AB sound from the compressor shell

The transfer function $H_{monopole}$ corresponds to breathing of the compressor shell. It quantifies the external sound pressure per unit volume velocity of the shell, (the volume velocity is the volume of air displaced by the breathing per cycle, given in m^3/s). To measure this transfer function conventionally we would need to replace the compressor shell with an idealised sphere of the same size, pulsating with a known volume velocity, and to measure the sound pressure at the external receiver positions. The arrangement is shown schematically in the left of Figure 12. However, it is difficult or impossible to perform tests in this way because of limited space within the cavity where the compressor is housed. There is also the difficulty of obtaining a suitable source. Instead, we make use of the principle of acoustic reciprocity (Fahy 1995) whereby the source and receiver positions are interchanged. Therefore, we place a monopole at the external receiver position and measure the average, in-phase sound pressure over the surface of the shell. This is done by placing microphones close to the surface at the same positions as were used for the accelerometers in Figure 10. The measurement setup is illustrated on the right of Figure 12.

The same microphone arrangement can be used to extract the x, y and z dipole transfer functions, but rather than the in-phase pressure we use the difference in pressure of the microphones on the x axis in a similar way to equation (2), (similarly for y and z).

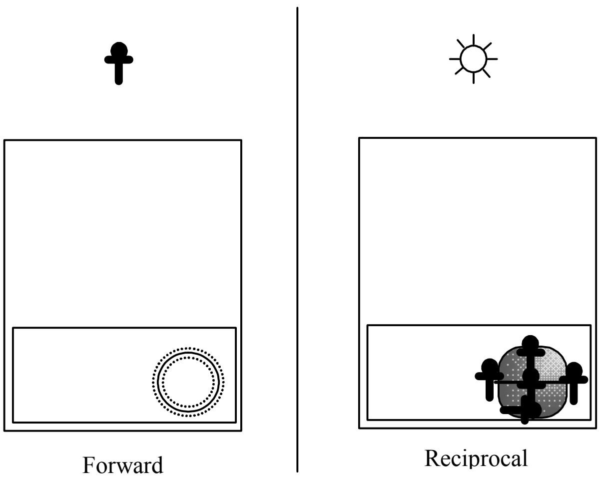

Figure 12 Forward measurement of the transfer function (left) and reciprocal measurement (right)

Fluid-borne sound

The fridge was one of two products studied in the Nabucco project where fluidborne noise needed to be taken into account, the other being a washing machine with a pump. For pumps, it is usual to characterise the source by a ‘source strength’ and a ‘source impedance’ as defined in ISO 10767 (1996). However, this same arrangement cannot be easily adapted to the fridge compressor. One of the reasons is that the refrigerant fluid undergoes changes from liquid to gas phases as it passes around the pipe circuit. Consequantly its properties, wave speed in particular, at any point in the circuit are unknown. The measurement of source impedance is also problematic, even for specially equipped laboratories, and even more so for industries with limited time and resources.

Fortunately, a simplification was possible in that the suction pipe contribution was found to be negligible. The question of how to characterise the compressor as a source independently of the frame still needed careful consideration. Since the pipe is narrow compared with a wavelength, we can assume plane wave propagation only in the pipe. The pressure pulsations in the discharge pipe are then made up of the superposition of an outgoing and a reflected wave as illustrated in Figure 13 (the terms ‘pressure pulsation’ or ‘pressure ripple’ are usually used by refrigerant engineers where acousticians would use the term ‘sound pressure’).

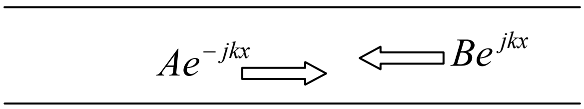

Figure 13 The sound pressure in the discharge pipe is a superposition of outward and reflected waves.

The strength and phase of the reflected wave are determined by downstream discontinuities and are therefore affected by the frame. On the other hand, we can reasonably make the assumption that the outgoing wave is determined only by the source and is independent of the frame. The independence criterion needed for the VAP can then be met if we could find the amplitude of the outgoing wave.

The reflected wave cannot be physically removed. However, it is possible to remove it using appropriate signal processing techniques. It is well known that using two pressure transducers in a waveguide one can obtain the amplitudes of the outgoing and reflected waves, provided the wave speed is known (similar algorithms are used in two microphone impedance tube measurements to obtain reflection coefficient). It turns out that if a third transducer is added then the wave speed (which remember is not known for the refrigerant fluid in this case) can also be deduced. Therefore, the measurement rig adopted consisted of three equally spaced pressure transducers in the discharge pipe as shown in Figure 14. The source strength of the compressor was then characterised by the amplitude of the outgoing wave.

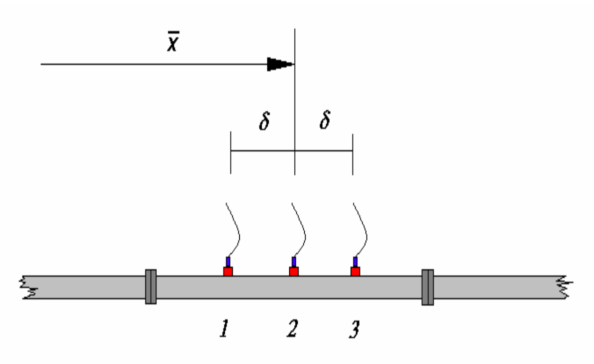

Figure 14 Rig for measurement of outgoing wave amplitude and wave speed in the refrigerant fluid.

In order to measure the transfer function, the sound pressure due to FB excitation was measured directly. In order to do this it was necessary to eliminate the contributions of AB and SB sound. This was achieved by running the fridge from a compressor in an adjacent room (to eliminate AB contributions), connected via long pipes with flexible sections (to remove SB contributions). The test setup is illustrated in Figure 15. Note that it was also necessary to measure simultaneously the source strength. This is because the transfer function is defined as the output (external pressure) per unit input (FB source strength in this case), so the strength of excitation must be known. The source strength was measured using the three pressure transducers as described earlier.

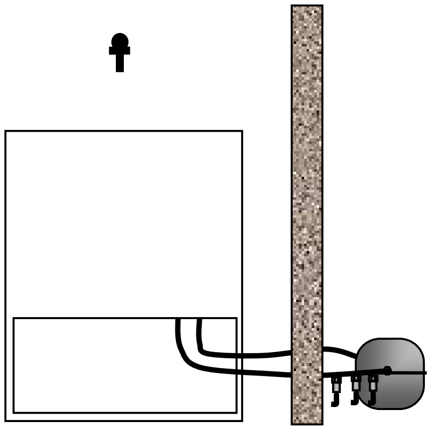

Figure 15 Test setup for the FB transfer functions. The sound pressure at the receiver position is measured and is normalised according to the measured FB source strength.

Structure-borne sound

Having looked at AB and FB, the final contribution to consider was the effect of SB sound from the pipe and the feet. The source strength parameter for the pipe was defined as the sum of squared vibration amplitudes in three orthogonal directions. The feet were treated in a similar way.

These source strengths could be obtained by a transformation of the six velocity measurements used to characterise the AB sound source strength (Figure 10). Thus, no additional measurements were needed to obtain source strength data.

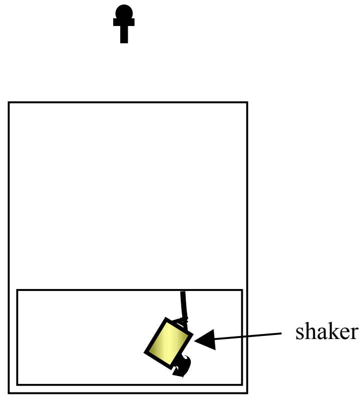

Figure 16 Test setup for measurement of SB transfer functions: a known force is applied to the end of the pipe from a shaker and the sound pressure is measured.

Transfer functions were measured using a forward technique in which a force was applied to the end of the pipes from a shaker through a force transducer. The resulting sound pressure was measured and then divided by the force input to obtain the sound pressure per unit force. The force was applied in several different directions and an average calculated.

It should be noted that the representation of the SB excitation and transmission mechanisms was significantly simpler than for some other cases because there was a large impedance mismatch between the pipe and the shell. It was then possible to assume that the vibration of the pipe was the same as that of the shell, i.e. that the compressor was a velocity source with respect to the pipe. In general however, it is necessary to include an impedance or mobility matching step to calculate the contact forces (see for example Moorhouse 2003).

What the customer hears: combining AB, SB and FB contributions

Having obtained source strength and transfer function data for AB, SB and FB excitation types, the sound pressure spectrum due to each could be calculated and summed together to predict the overall external sound at the receiver position according to the scheme of Figure 8. The sound pressure spectrum for each case is shown in Figure 17. (The SB contribution through the feet was relatively small, so only that through the pipe is shown).

In Figure 17, for each excitation type a harmonic series is evident, with strong peaks at multiples of the compressor speed. However, the spectrum shapes are significantly different. The AB component, which is dominant in terms of dB(A), contains a wide range of frequencies. The SB component is predominantly low frequency, consisting of a rapidly falling set of harmonics. The FB component contains strong peaks up to about 600Hz, and the envelope of the peaks has a pronounced bell shape between the second and sixth harmonics. This shape would be caused by a resonance, probably somewhere in the pipe circuit, and is significant in terms of sound quality as will be discussed later.

Also shown Figure 17 is the combined spectrum (p_{total}), which was calculated assuming the AB, SB and FB sound pressure contributions (p_{AB} etc.) to be incoherent, i.e.

$$|p_{total}|^2 = |p_{AB}|^2 + |p_{SB}|^2 + |p_{FB}|^2 .$$

In theory, the separate AB, SB and FB contributions are coherent because they ultimately come from the same source, the compressor. However, in practice the transfer paths for each are rather different, and the relative phase with which they arrive at a receiver point will vary depending on the precise location. The preceise spectrum shape will therefore vary with position. In practice, this effect can be heard by in a room with a fridge by moving the head around slightly: the sound can be heard to vary perceptibly from one position to another due to variations in phase. However, these differences are not significant in terms of sound quality, so the incoherent averaging represents the combined effect at an 'average' position and is reliable for auralisation as will be discussed later.

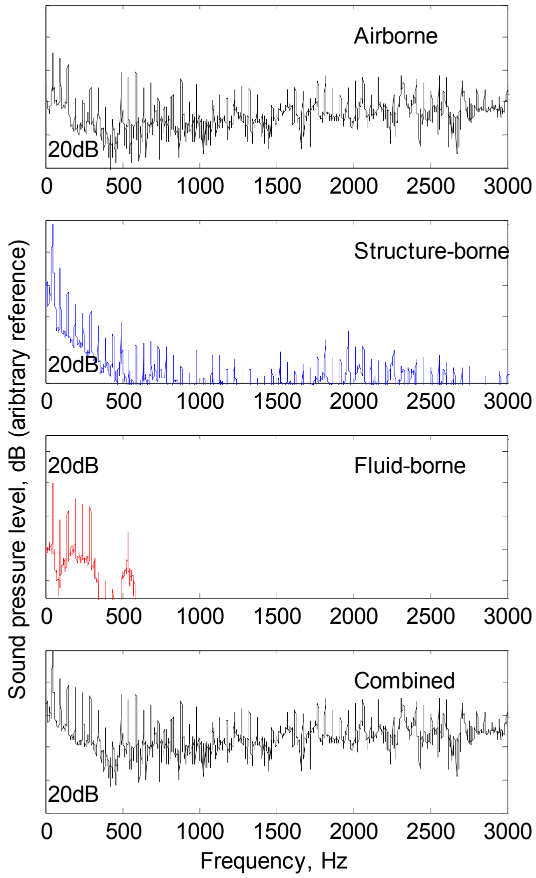

Figure 17 Contributions to the external sound pressure from AB, SB and FB excitation for the fridge. (y axis referencesare suppressed for confidentiality reasons)

The breakdown given in Figure 17 allows the designer to perform a conventional rank ordering by comparing for example the dB(A) levels associated with each excitation type. The contributions in order of importance to the dB(A) level in this case were AB, FB, SB (pipe) and SB (feet). As all noise control engineers know, a rank ordering is the first essential step to designing effective noise control, so the ability to carry out rank ordering is an important function of a design tool.

As well as being used for rank ordering, the value of the above spectra can be extended considerably by auralising the results. The designer is then able to listen, not just to the combined effect as with the real prototype, but also to individual contributions, perhaps to identify the source of an unpleasant feature of the sound. The techniques to achieve the auralisation from frequency domain data are discussed in the next section.

AURALISATION: LISTENING TO THE VIRTUAL MACHINE

So far, all the calculations shown have been in the frequency domain, in the form of narrow band spectra. This is a different situation to that taking place in the concert hall analogy described in the introduction. There, the calculations were done in the time domain, the time history of the source being convolved with the impulse response function of the hall to obtain the modified sound.

How then might we be able to perform the auralisation starting from frequency domain data? Naturally, any spectrum can be transformed to the time domain using Fourier analysis, but there are two problems which must first be overcome:

- the number of data points is sufficient only for a very short time sample
- phase data is missing.

Considering the phase issue first, we recall that sound pressure signals recorded in the real world are real functions. However, when transformed to the frequency domain a complex Fourier spectrum results, i.e. the spectrum associated with a given time history has both magnitude and phase. If both magnitude and phase are known then the time history can be reconstructed exactlyby inverse Fourier transformation. However, the phase will be lost if any kind of averaging is carried out on the spectral data. In the example of the fridge discussed earlier, the most effective source strength for the SB component of excitation proved to be sum of the squares of the vibration levels in three orthogonal directions. During the summing operation, which effectively is a sum of the energy of the vibration in each direction, all phase information is lost. The resulting spectrum for SB sound therefore includes magnitude information only. In some cases it would be possible to retain phase data, for example the AB source parameters and transfer functions for the fridge could in theory all be measured as complex spectra. However, in practice retaining the phase information creates measurement and data handling difficulties, particularly for medium technology industries that are the intended users of such techniques.

Therefore, we have to accept that we will not have explicit phase data for the auralisation. Fortunately, and somewhat surprisingly, it turns out that for steady state sounds this is not a major drawback. Such sounds can be auralised by assuming the phase spectrum to be random. The phase spectrum is therefore generated as a sequence of random numbers with magnitudes in the range 0 to 2π, (except for the zero frequency and maximum frequency values which are set to zero). The resulting phase spectrum is then added to the measured or calculated magnitude spectrum before inverse Fourier transformation and playback of the time history.

The question of the length of the time record obtained is now considered. Narrow band spectra do not usually contain more than 3200 frequency lines, and sometimes fewer. When transformed to the time domain the sampling rate is determined by the bandwidth of the spectrum. For example, for the fridge, the calculated spectrum extended to 6kHz (only the lowest 3kHz of which is shown in Figure 17). Because of aliasing, the sample rate must be at least twice the maximum frequency, i.e. at least 12kHz. The 3200 samples will therefore produce less than a quarter of a second of audible sound, which is insufficient for auralisation. A short sample produced in this way could be lengthened by looping, but this usually creates an audible 'looping' artefact that gives a misleading impression of the sound.

The solution is the increase the number of data points without increasing the frequency range, which is done by interpolating the magnitude spectrum as will be illustrated in Figure 18.

A final problem to overcome is that the spectrum should be doublesided for inverse transformation, whereas so far we have only a singleside spectrum. Therefore, one more step is required which is to add to the singlesided spectrum an alias spectrum, or a reflection of the spectrum with reversed phase.

Procedure for auralisation

Combining the above operations we arrive at a procedure for auralisation of steady state sounds, as summarised in Figure 18. The sequence is as follows:

i. The magnitude spectrum is interpolated to give as many data points as required. (It improves processing speed if the number of spectral lines is set to one more than a power of 2, but this is not essential.)

ii. For each spectral line of the magnitude spectrum a phase angle is generated. For the zero and maximum frequencies this phase is set to zero, and in between it is set to a random value between 0 and 2π.

iii. The spectrum is reflected about the maximum frequency (the Nyquist frequency) to give a double sided spectrum, i.e. if the there are *n* samples in the single sided spectrum then the magnitude spectrum is extended to $2(n-1)$ samples with magnitude M given by:

$$M(n+p) = M(n-p)\,,\ \ p = 1,2,...n-2$$

iv. The phase spectrum, Φ, is reflected in a similar way except that the values above the Nyquist frequency are negative so as to form an antisymmetric spectrum:

$$\Phi(n+p) = -\Phi(n-p)\,,\ \ p = 1,2,...n-2$$

v. The doublesided antisymmetric spectrum is then transformed to the time domain.

For example: for the fridge, a double sided spectrum with 36k data lines gives a three second audio clip at 12kHz sampling frequency. Therefore, 18k samples are required in the interpolated single sided spectrum. In practice we would probably choose 16385 points for the singlesided spectrum so that the doublesided spectrum contains 32768 lines, which is a multiple of 2 and can be transformed more efficiently.

This procedure is not restricted to VAPs, but can be used (with caution) to auralise any narrow band spectrum.

Interpretation of auralised data

The results of auralisation on the fridge example from the previous section are quite revealing. The AB component sounds as one would expect a fridge compressor to sound. The SB component sounds like a ‘low hum’, although the higher harmonics are also audible giving slight ‘rattle’ to the sound. The FB component gives a pronounced ‘boxy’ tone to the sound, which, on its own, sounds fairly unnatural. This feature is not particularly noticeable in the combined sound, being masked by the more dominant AB component. However, were the AB component to be reduced then the ‘boxiness’ would become audible. This scenario can be simulated fairly easily once the data is available by simply adjusting the levels before combining the different components. This tells the designer to be wary: a design modification that reduces the relative contribution of the AB component would make the ‘boxiness’ more pronounced, which may or may not be a desirable feature. It is even possible that a reduction in dB(A) could result in a sound that is subjectively louder.

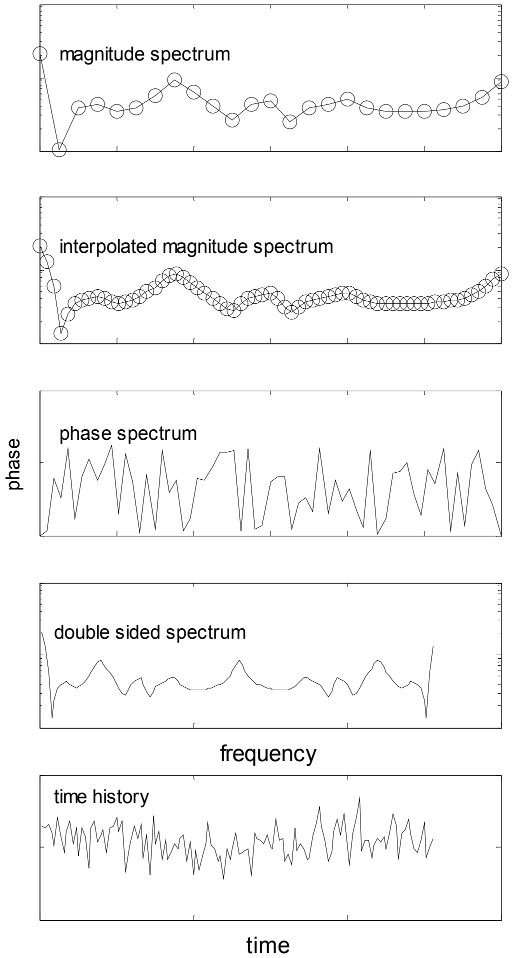

Figure 18 The sequence of operations for auralisation from a magnitude spectrum. The spectrum is interpolated, reflected, random phase is added and the complex spectrum is transformed to a time signal

SOME APPLICATIONS OF VAPS

It has already been mentioned that one of the main uses of VAPs is for comparison of different active components in the same frame. An example of comparing washing machine motors is given in Moorhouse and Pavic (2005). In this section some other types of comparison are considered.

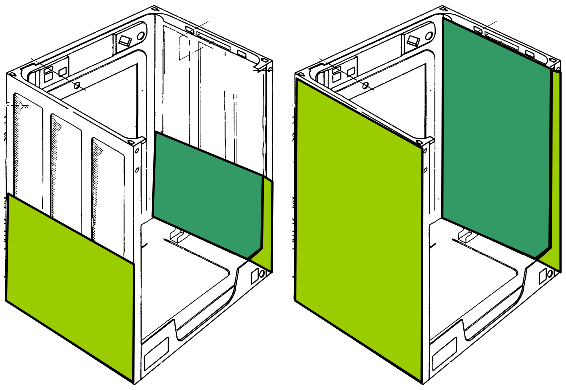

Figure 19 Washing machine cabinet before and after removal of some absorbent lining (courtesy Electrolux)

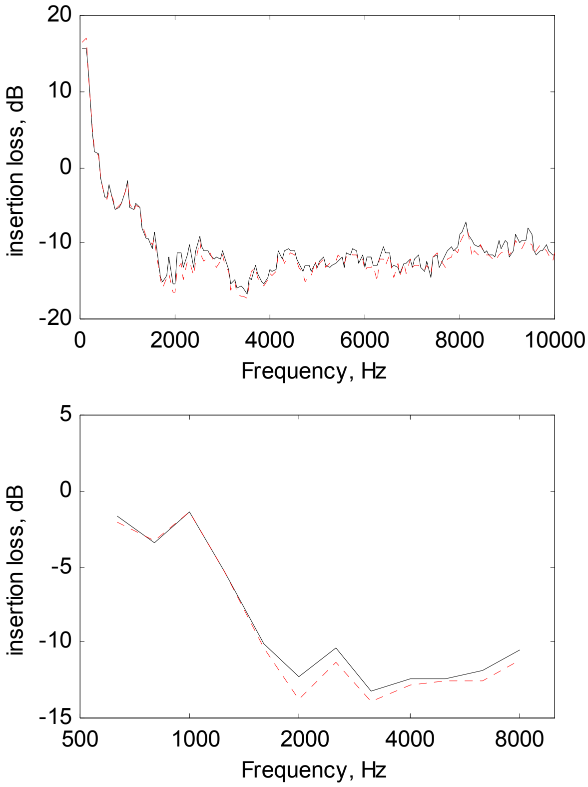

Figure 20 Insertion loss of the washing machine cabinet measured before and after modification. Upper plot=narrow band; lower=third octave bands.

Modifications to the frame

Another example from a washing machine is shown in Figure 19. Here, rather than replacing the motor, it was the frame that was modified by removing a small amount the acoustically absorptive lining from inside of the cabinet. The conventional method to quantify the change would be to run the machine before and after modifications and to compare the measured sound power or pressure. However, even with a sophisticated motor control system it is impossible to guarantee identical operation of the motor for both tests. This introduces an uncertainty into the results such that differences of the order of 1-2dB could not be attributed with certainty to the frame modifications, but could equally be due variations in the motor. On the other hand, using the VAP approach the frame performance is quantified using transfer functions, which are measured using a controlled sound source, and the reproducibility is extremely good. Figure 20 shows the insertion loss of the washing machine cabinet obtained from the measured transfer functions before and after modification. Differences of up to 2dB were observed at some frequencies. Importantly, a check on the reproducibility of the results produced two curves that could not be distinguished, even on a narrow band plot. This means that the 2dB differences can clearly be attributed to the frame modifications. The VAP techniques therefore prove to be rather a 'sharp instrument' capable of reliably measuring small differences in performance. As designs become more refined, differences of one or two dB are becoming more significant, so the capability of being able to measure small changes is becoming more and more important.

Listening to an installed motor under load

The second application also involves a washing machine, but in this case the objective was to auralise the sound of the motor inside the frame when operating under realistic loading conditions. This sound cannot be recorded directly using a physical prototype, because the load itself (the washing machine drum) generates noise, which would contaminate the recording. In keeping with the philosophy of VAPs, the data obtained for the motor should be independent of the frame and other sources (like the drum) such that the data is transferable. Therefore, the approach was taken of characterising the motor by measurements on a separate test rig. The source data obtained was then combined with measured transfer functions for the frame to predict the external sound. Both AB and SB sound from the motor were significant and had to be included in the VAP.

Initially a validation test was conducted on an unloaded motor. For the AB component, the motor was represented by an equivalent monopole whose volume velocity was back-calculated from the sound power of the motor running in free field conditions. The transfer functions were measured by a reciprocal measurement procedure, similar to that shown in Figure 12, in which a calibrated monopole source was placed outside the frame and microphones were placed inside at the position of the motor. These microphones were precisely positioned using a 'dummy' wire-frame motor as shown in Figure 21.

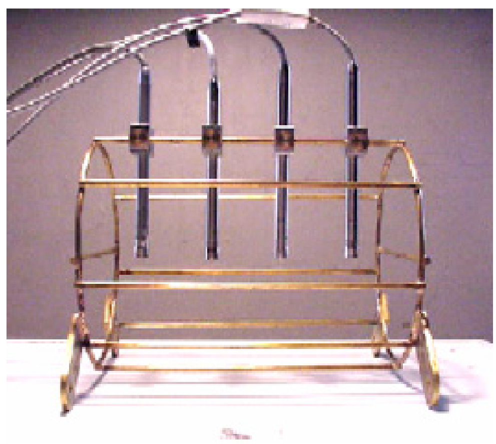

Figure 21 Wire frame motor for positioning microphones at the position of the motor during reciprocal transfer function measurements.

Handling the SB component was considerably more complicated. In the fridge example given above it was found that the vibration of the compressor shell was not affected by the attached pipework. However, in this case no such simplifying assumption could be made since the motor velocity was significantly altered when attached to the frame. It was therefore necessary to include an additional step in the calculation to account for the 'mobility matching' between source and frame. This step is analogous to the impedance matching of imperfect voltage sources with a load impedance, but is considerably more involved because four feet were involved, each with vibration in three perpendicular axes. A description of the procedures adopted can be found in Moorhouse (2003) which are not discussed here.

The SB source strength was initially taken as the free velocity (i.e. the velocity at the feet with the motor suspended on very soft springs). The SB transfer functions were measured using a reciprocal technique similar to that described above, but with accelerometers attached to the frame at the motor contact positions (illustrated in Figure 22). The approach was validated for an unloaded motor by comparing measured and predicted sound pressure levels, the results from which are shown in Figure 23 (validation for the loaded case is not possible because of noise generated by the load as described above). The agreement is good over most of the frequency range, the only significant discrepancies occurring at low frequencies due to errors in the prediction of the SB sound contribution. Such differences are common for this type of calculation and probably cannot be eliminated completely.

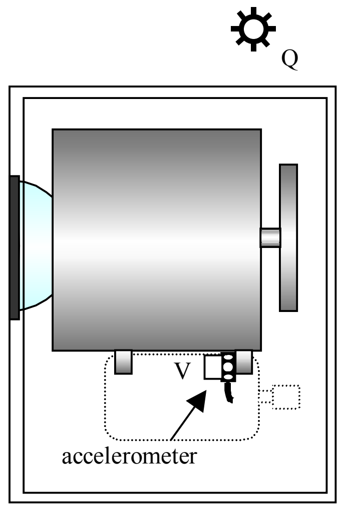

Figure 22 Reciprocal measurement of SB transfer function

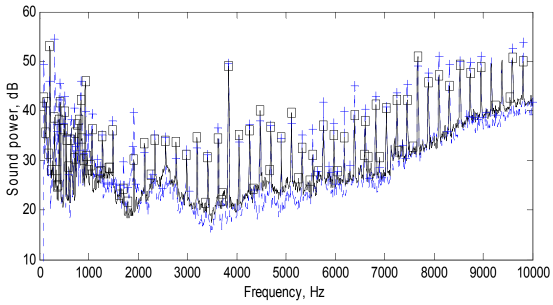

Figure 23 Comparison of measured (solid lines) and predicted (dotted lines) sound power level for the washing machine with an unloaded motor (from Moorhouse and Pavic, 2004)..

The next step was to obtain source data for the loaded motor. In order to achieve this a special test rig was constructed as illustrated in Figure 24 in which the motor was attached to a rigid steel block. Lateral loading of the motor shaft was achieved by loading with a belt at the same tension as in the real machine. Torque loading was achieved with a non-contact eddy current brake, the discs of which could be moved further apart or closer together to decrease or increase the loading.

The AB source strength under load was obtained with a sound power measurement from which the volume velocity of the equivalent monopole source was calculated as before.

For the SB source strength, free velocity could no longer be used because the feet were effectively blocked by the mounting block. Instead, the blocked force was used, i.e. the force exerted by the motor on the block. These forces were measured wth force transducers set into the mounting block at each foot.

Interestingly, the torque proved to make little difference either to the AB or the SB source strength. On the other hand, the lateral loading by the belt significantly increased the SB source strength. This could be due to the extra static load on bearings, and/ or distortion of the motor casing due to the lateral load.

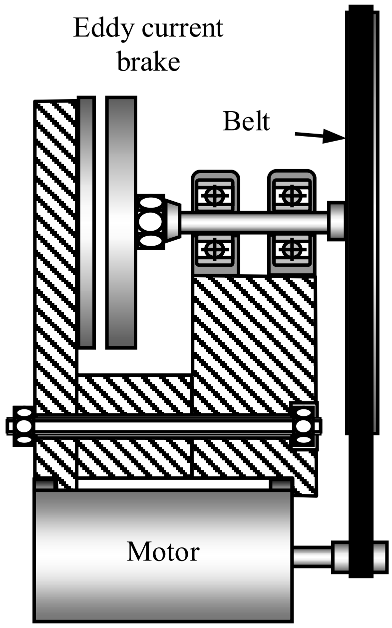

Figure 24 Mounting block with belt and eddy current brake for loading the motor.

Having obtained the source strength data the sound pressure level around the virtual washing machine, as well as its sound power could be calculated in the VAP. The results are shown in Figure 25. The upper plot shows the sound power level over the 10kHz range, and the lower plot gives the same data but with a zoom on the lower frequencies where the SB component is more significant. There are some differences in the spectrum, with some peaks changing in level. An interesting change is the frequency shift of the peak at around 650Hz, which is thought to be due to modifications in stiffness of the motor casing due to distortion by the load.

The effect of the loading on the auralised sound was not strong; in fact the differences between the loaded and unloaded cases were hardly detectable by ear, particularly at high running speeds. Having carried out the exercise one would be fairly confident about using source data from an unloaded motor, which naturally simplifies things significantly. However, this result was a surprise; a much larger effect had been anticipated, and one would not have been so confident at the outset. It should be emphasised that a comparison of this sort could only be made using a VAP, or some similar virtual approach because the sound of the installed motor running under load, but free from other sounds, cannot be realised physically.

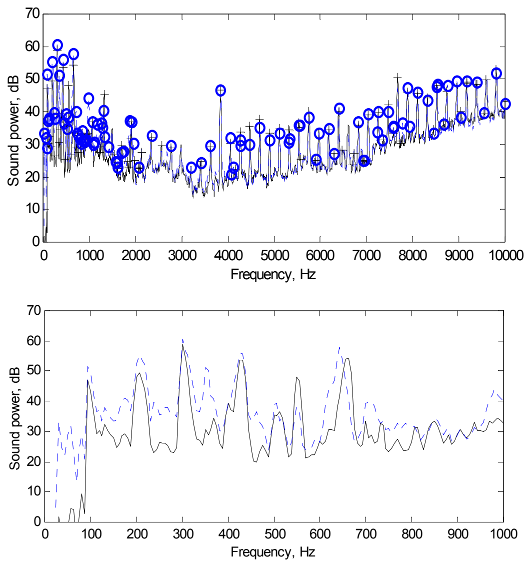

Figure 25 Comparison of (calculated) sound pressure level for the unloaded (solid lines) and loaded (dotted lines) motor installed in the washing machine frame. Upper plot: 10kHz range; lower plot: zoom on 1kHz range (from Moorhouse and Pavic, 2004).

CONCLUDING REMARKS

The above examples illustrate that considerable sophistication is required to auralise the sound of even a fairly simple machine. It is also clear that, for the time being at least, a VAP cannot be purely 'virtual', but will necessarily include some measured data. However, the advantages of the virtual approach are many. Obvious advantages include reducing the cost and time associated with constructing physical prototypes. Less obvious advantages include the observation that the components of a VAP can be exactly reproduced so that more precise comparisons can be made of the effect of varying one element, such as a motor. The above example of the loaded motor also makes clear that some situations can be constructed in a VAP which cannot be realised physically.

There are two main advantages to the designer. Firstly, the process of constructing a VAP provides a systematic framework for understanding how and why the machine makes the sound that it does. Secondly, auralisation is a powerful tool for communicating ideas about acoustics. In the Nabucco project it became evident just how much difference it made to non-specialists like managers and accountants to be able to hear the sound of the machine as opposed to having it described to them. This observation makes clear that there is a strong potential to improve understanding of acoustics issues through auralisation techniques, but at the same time it also illustrates how big is the divide between specialists and non-specialists at current state of the art.

Acousticians in various sectors, buildings, automotive, transport, consultancy etc. are increasingly becoming aware of the potential for virtual techniques and auralisation in acoustics, and VAPs can be thought of as just one part of an increasingly wide movement. In the future we might see design drawings accompanied by sounds, and consultants reports routinely including computer auralisations.

REFERENCES

Fahy FJ 1995, The vibro-acoustic reciprocity principle and applications to noise control. ACUSTICA 81 (6): 544-558.

Kompella M S and Bernhard B J, (1993) 'Measurement of the statistical variation of structural acoustic characteristics of automotive vehicle', in Proceedings of the SAE Noise and Vibration Conference, Warrendale, USA.

Moorhouse A, Berglund P-O, Fournier F, Avikainen T, 2003, Fan characterisation techniques, Proc Fan Noise 2003, Senlis, France.

Moorhouse A T & Pavic G, 'Virtual Acoustic Prototypes of White Goods Products' Proc Inter-noise 2004, Prague, Czech Republic, 2004.

Moorhouse A T, 2003 Use of a hybrid measured-calculated mobility matrix for simplified calculation of structure-borne sound from an electric motor, Proc. 10th ICSV, Stockholm.

Pavic G, Gavric L, Tourret J, Sottek R, 2003 Synthesis of noise from a fan type source placed in a complex installation. Proc Fannoise 2003, Senlis, France.

Rebillard E, Guyader JL 2000 Calculation of the radiated sound from coupled plates ACUSTICA 86 (2): 303-312.

Shorter PJ, Langley RS 2005. On the reciprocity relationship between direct field radiation and diffuse reverberant loading. J Acoust Soc Am 117 (1): 85-95.

ACKNOWLEDGEMENTS

Much of the work presented here was carried out the in EU funded project 'Nabucco' no GRD1-1999-10785, support for which is gratefully acknowledged. I would particularly like to acknowledge the contribution of Goran Pavic who was the mastermind behind much of this project. The fridege example was largely the work of Leon Gavric, Michel Darpas and Jean Tourret at CETIM, France. The base station example was mostly due to Per-Olof Berglund, Mats Abom, Anders Nilssen and Leping Feng, at KTH Stockholm, Sweden. The auralisation procedures were obtained from Roland Sottek and Holger Kaempfer at HEAD Acoustics, Germany. Several other people also contributed to the material in this paper, notably Timo Avikainen, Krystof Kryniski, Marco Clara, Georg Eimer, Walter Angelis, Barry Gibbs, Gary Seiffert, Richard Cookson.