Side Ports

Over the past few years there has been something of a trend for cutting holes of various sizes and shapes in the sides of guitars. While a hole in the side of the guitar might or might not directly effect the way the top and back, the main sound producing surfaces on most guitars, vibrate, it could certainly be expected to effect the vibrations of the air inside the box. I really became interested in this when a customer asked if I could build a guitar with a port. Since I'm the sort of guy who likes to have a map when I'm going into new territory, I looked up what information I could find. There was a certain amount on air resonant modes, but not much treating ports, so I decided to do an experiment.

'Ideal' Air Modes

Although the air inside a box can be made to vibrate at any frequency there are certain pitches where the vibration will be better organized, if you will. At these pitches any given input of power will result in a much higher amplitude of sound. These organized ways of vibrating are called 'resonant modes', or simply 'modes'. Most of the studies done on guitar air modes have used either thick walled guitar shaped boxes (1) or actual guitar bodies that were buried in sand to deaden the top, back, and sides (2) (3). This makes a lot of sense. The air modes are complicated enough themselves, and factoring in the effects of vibrating walls just adds confusion at first.

All air resonant modes involve a confined volume of air that trades energy back and forth between inertia (flow) and pressure. Envision a puff of air with a certain mass which is moving toward a pipe that has one end open and the other closed. As the air rushes into the open end it compresses the air inside, and the rising pressure slows down the flow. Eventually the flow stops, and the excess pressure in the pipe forces the air back out. The inertia of the out flowing mass actually carries out more air than went in in the first place, and this causes the internal pressure to drop below that outside. More air is drawn back into the pipe, and the cycle starts over.

If the opening of the pipe is smaller than the cross section and the pressure rises or falls at the same time in all parts of the pipe we usually call the result a ’Helmholtz’ type mode. It’s also possible to produce resonant modes in pipes that are open at both ends; think of air masses moving toward both openings at the same time and ‘bouncing off‘ each other. The ’pipe’, of course, can be any shape, as well. A pipe with both ends closed can be thought of as two open-ended pipes joined together, in which case there can be vibrations without any opening at all, although there will be no way for the sound to get out. There are many different possible resonant modes for any given pipe, at different pitches. Allen’s article an AL#1 gives a very good description of all of this.

Here are most of the lower frequency air mode shapes I found in the ‘corker’ guitar I built for the experiment. I’m using a heavy line to indicate places where the pressure changes a lot, and gives a high Sound Pressure Level (SPL). The dotted line indicates a ‘null’; a place where the energy is in flow, and the SPL is zero, or close to it. When you cross a dotted line the phase changes: positive pressure at any one instant becomes negative. The dotted line is like the fulcrum of a seesaw, in a sense. The A-0 mode is the ‘main air’ or ‘Helmholtz’ mode, and there is no phase change anywhere inside the box. The nomenclature is largely that used by Fletcher and Rossing (op. cit.).

‘Real’ Air Modes

It’s obvious that pressure changes within the guitar body are going to push or pull on the walls, and since these are not rigid on real guitars, they will move. They’re more likely to move at frequencies near their own resonant pitches, but powerful air modes, such as the ’Helmholtz’ mode, will cause the top, back, and sides to vibrate even if they’re fairly heavy and far off-resonance, as in the ‘bass reflex action‘. One indicator that this is happening is to see the same body or air mode shape showing up at more than one pitch.

On this guitar there was strong ‘ring mode’ motion in both the top and the back at three different frequencies, indicating that both plates are coupling strongly with the ’Helmholtz’ air mode, and with each other.

One of the air modes that gets missed a lot is a second Helmholtz-type mode, involving air moving through the sound hole and in and out of the lower bout only. It seems to depend both on the geometry of the guitar, with a pronounced waist and a sound hole near it, and also on motion of the wood in the plates. I could not find this mode in a rigid ‘pipe‘ with a waist and hole. Elejabarrieta et al (4) did see this, at 242 Hz, in a sophisticated computer model of a guitar that took the wood and air coupling into account. It is also often missing on ‘Dreadnought’ shapes guitars, which lack a pronounced waist. This mode seems to have an effect on the way the guitar sounds. Since the air is moving out through the hole when the top is moving inward, the top and air tend to cancel each other, and reduce the sound output in front of the guitar near that frequency.

Another difference between ‘real’ and ‘ideal’ air modes is that on real guitars the ‘A-1’ or ‘lengthwise pipe’ mode often shows up at two frequencies, with the nulls displaced. This seems to depend somewhat on both geometry and wood motion. A rigid box with no sound hole and an above-center waist will show the pressure null displaced toward the waist, while a box with a sound hole just above center on the length will show the null moved away from the hole. A rigid box with both a waist and a sound hole will show a single A-1 type mode, with the null at some compromise position, but when the walls are not rigid it can show modes at two frequencies, with the nulls displaced. Since one of the modes has the null closer to the sound hole it does not radiate as effectively. However, the other has the null displaced away from the hole, and can be a strong radiator of sound.

Air modes also can couple quite strongly with the wood of the top and back, modifying the way they vibrate and thus effecting the sound output even when the air resonance itself cannot radiate directly. This is probably one major way the ’higher’ air modes (above the pitch of the A-0 mode) effect the tone of most guitars. Air modes often have relatively high losses, and will thus tend to ’waste’ energy if they are not strong direct radiators of sound.

On the guitar in the experiment, the air modes and their frequencies (in Hertz) were:

	Ports closed	#2 port open
A-0-1	97	103
A-0-2	251	251
A-1-1	360	374
A-1-2	429	440
A-2-1	577	577
A-2-2	620	623
A-3	701	714
A-4	765	775

The Effect of a Port on an Air Mode

In general, the way a port will effect an air mode depends on where it is relative to the pressure and flow maxima of the mode, as we already saw with the sound hole and the A-1 modes, and how big the port is. If the port is in an area where there is a lot of pressure change it will radiate energy from the mode. This usually cuts down the internal SPL, but it generally increases the energy radiated from the guitar. It will also usually raise the pitch, as it did with the A-1 modes in the experimental guitar. If the port is at or near a null of the mode, where there is no pressure change, it will tend to drop the pitch a little, and may make the internal SPL of the mode stronger. The bigger the port, the greater the effect, in general.

The Experiment

As you can see, even though the qualitative effects of a port are fairly simple to predict on a single air mode in isolation, the complexity of a guitar makes the overall effect of putting a hole in some part of the side more problematic. Thus I built a guitar to test things out. The initial concept was based on Carleen Hutchins' 'La Grueyer' violin, which has more than sixty small holes bored in the ribs, that can be plugged with small corks. Opening them all virtually eliminates all of the air modes of the instrument. (5) I did not need to be that thorough, so I only used ten pairs of 5/8" holes, evenly spaced along the right (upper) side from the neck to the tail block. The guitar itself is a 'standard' classical size, with a reasonably nice European spruce top, Indian rosewood sides, and a Padauk back. It has been tried by several first-rate players and builders, and is a 'pretty good' (usually expressed with surprise, as it's also pretty ugly) guitar. Here are the 'wood' resonant modes of the assembled experimental guitar, as found using the Chladni pattern technique. (6) The first set of numbers are the frequencies of the modes with the guitar in a 'normal' configuration, and the numbers in parenthesis were the frequencies seen when the #2 ports, near the wide part of the upper bout, were opened.

Some resonant modes showed up clearly at more than one pitch, and this is an indication that there is 'coupling' going on with some other mode, which is either close in pitch or interacts strongly for some other reason. The T-1 and B-1 modes, for example, clearly interact with each other, and that makes sense since they are similar in shape and frequency. However, they both also couple strongly with the A-0-1 mode, which is much lower in pitch, simply because that mode can push strongly on the top and back through internal air pressure changes. Although I did not list it in the table of air mode pitches, there is strong output at the sound hole at 215 Hz, near the T-1 and B-1 pitches. When modes couple it is common to find that the different parts reach their maximum amplitude at somewhat different frequencies.

I took the instrument to the Healdsburg festival in 2003, right after completing it, so that people could try it out. A number of trials were made, assessing the change in tone both to players and listeners when a single pair of corks was removed. Because of the noise level, and the fact that the tests were not 'blind', it was difficult to come to firm conclusions.

When I returned from the show I was able to make other measurements over a period of about a month, and some further testing has been done since. My shop is decidedly not an ‘anechoic chamber’, nor is it ’perfectly reverberant’, so there is a considerable room effect that cannot be eliminated. However, by carefully keeping the conditions constant in making pairs of tests it was possible to achieve a good level of reproducibility, and this allowed for valid comparisons between those tests.

The most useful tests involved driving the bridge with a ’stinger’; a modified loudspeaker having a small clip on a post attached to the coil, and with most of the cone removed. This converted the freely hanging guitar into a speaker, which could be driven with a sweep signal from my computer’s sound card. Being careful not to disturb anything I could successively remove and replace each pair of corks for a sweep. A dB meter could be used as a microphone to pick up the output, which was recorded on my computer. The recorded sound was put through a Fast Fourier Transform program, which isolated things like the 60 cycle hum from florescent lights. The ‘real’ and ‘imaginary’ output of the FFT was used to generate charts using a standard spreadsheet program.

Other tests included the use of a ‘microphone on a stick’ to measure internal air modes, Chladni patterns, and the use of a small accelerometer to measure top and back motion at given points during a stinger sweep. All of these results could be fitted together to give a reasonably coherent picture of the effects of opening a hole at any location on the sides. Most tests were run for most port locations (resulting in two large file folders of data!), but, as listening tests indicated that the most useful port placement would probably be near the wide point of the upper bout (my #2 set of holes), most of the results I will present relate to that port location.

To begin with, we can look at the output spectrum of the guitar as picked up with the microphone as close as possible to the sound hole. This will present only a partial picture, missing much of the complexity. In particular, it is difficult to find a microphone position on the top that will pick up all of the modes. On the other hand, close micing like this does minimize the effects of the room, and will be a useful first step. Here is a chart of the sound picked up about one diameter away from the normal sound hole when the bridge was driven with a sine wave sweep from 50-1000 Hz. I will note that these ‘close mic’ tests were done at a much later date, and in drier conditions, than the original testing, and the frequencies of many of the modes have shifted somewhat from those obtained earlier. Although there are changes in the spectrum when the port is opened, they are hard to see at this resolution.

By expanding the low range we can see the shift in the ‘main air’ resonant frequency, caused by opening up the port. The output level has also risen some, although a probe microphone found that the internal SPL at the lower block had decreased about 30%.

More sound is being radiated, and consequently less is stored as pressure. There is some increase in output around the ’main top’ resonant pitch, around 185 Hz, and more at the ’main back’ frequency, at about 225 Hz.

At higher frequencies the main changes were a large increase in output between 425-500 Hz, and a lesser increase from 550-650. The lower frequency shift is most likely associated with the A-1-2 mode and the T-4 ‘cross tripole’ mode. There are several top, back and air modes that could be associated with the upper band shifts. At a higher frequency the A-4 mode, which shows as a rather deep dip around 760 Hz in the ‘closed’ spectrum is a much less deep dip in the ‘open’ spectrum. This mode does not radiate well from the sound hole, but takes power from the top that might have gone to moving air through the hole. It is likely that the open #2 port shifts the null a bit relative to the sound hole, which would account for the rise in level, even as there is still a dip at the A-4 pitch. I will note that the A-4 mode was one that had the internal SPL rise, by 119%, when the #2 port was opened.

How does the sound output of the top change? With the #2 port open there are increases in output in the ‘main air’, ‘main top’ and ‘main back’ frequency bands. The microphone was close to the node line for the T-3 ‘long dipole’ mode, and did not pick it up. The complicated peak structure between 150-300 Hz shows the interaction of two top modes that the microphone could pick up, and two back modes that are taking energy out of the top. The shape of the spectrum changes between 400-500 Hz., due to alterations in the pitches of the two A-1 modes, but the total energy output remains much the same. At higher frequencies it is difficult to speak with any confidence about the small changes that were seen. At this point the resonant pitches are about as close together as their bandwidths, and in this 'resonance continuum' it is difficult to sort things out. The gains overall seem to about balance the losses.

The largest change, as one would expect, takes place at the port location itself. Here is the low range chart. I will note that more microphone gain had to be used for this test with the closed port than for the close mic tests so far. The output was so low with the corks in that a higher gain setting was required to get the signal up to a useful level. The large increases in output from opening the port were based on small beginnings. Still, the increases are impressive in relative terms: more than 4dB averaged over this range. Higher up in pitch the output from the A-3 and A-4 modes can be clearly seen.

At this point, with gains in output at the top, the sound hole, and the port location, it looks like adding a port would be a clear winner. Let's see what happens when we back off a bit. Here is the low range chart produced with the microphone a meter out in front of the guitar, on the centerline of the top, and pointing between the sound hole and the bridge.

What gives? We would not expect to see the sorts of huge gains that we did at the port itself, but comparing this with figures 6 and 7 is still a bit confusing. Both of the 'close' charts, for example, showed gains of around 2dB at the 'main air' pitch, while here, although the frequency has shifted, the output has remained the same. The reason for this it that the air motion at the hole and the top motion are out of phase with each other. As the top and back both move outward at 102 Hz air is sucked into the sound hole, but is pushed away from the top. The extra output of the top and the sound hole simply tend to cancel each other out at this microphone location. There is an actual drop in output around 225 Hz in this chart, which is harder to explain, but may have to do with the mic position as well. At higher frequencies the shift in the A-1-2 mode, just below 450 Hz, is quite clearly seen as an increase in power. Otherwise there are only small gains over very narrow frequency bands.

In another series of tests, an accelerometer was used to monitor the back motion while the top was being driven by the ’stinger’ at the bridge. It showed more movement with the #2 port open up to the pitch of the ’main top’ mode, and less above that. This is consistent with the idea that a lower internal SPL allows the back to move more freely at low frequencies, but that there is also less internal sound to drive the higher back modes.

Finally, to get an idea of how it all sorts out for the player, the mic was placed about where the players ear would be, and the sound of the sweep driven at the bridge recorded from there. The results were not as remarkable as the close mic’d ones from the port itself, but did show clear gains in levels right around the ’main air’ pitch, and at specific frequencies above about 450 Hz, with small gains between 200-250 Hz. Most of the rest of the ’open’ curve was either just above or just below the level of the ’closed’ one.

Another Study

I was able to find one published study that included some data on a guitar with a sound port (7). Chaigne and Rosen did two tests to assess the effects of opening the port, one using string plucks, and the other a mechanical driver. Although the results of their mechanical driver test seems quite different from those we’ve seen so far, I think it is possible to explain both sets of data without doing major violence to either. Sadly, it has not been possible to duplicate their study closely, which does reduce the level of certainty.

Their first test involved placing the guitar in an anechoic chamber and plucking the open strings several times each by hand. Increases in output on the order of 2 dB were seen for the lower three strings, with smaller gains as one went up and no change for the high E string. Since I did see gains, at least in the ‘close miced’ tests, up to about the pitch of the open G string, and less certain changes above that, this result fits reasonably well. It is possible that the larger port that they used helped avoid some of the phase cancellation I noted, or it may be simply the outcome of a different mic position or test protocol.

For their other test Chaigne and Rosen mounted the guitar vertically on a rotating table, and drove it with a mechanical driver (B&K 4807) at the bridge. The input signal was white noise, band pass filtered between 50-350 Hz, so that it had energy covering the range of the fundamentals of the open strings. The output of the guitar was measured with a microphone at some unspecified distance, and the guitar was rotated to find the output in all directions. The output was quite uniform in all directions, and was 4 dB greater (more than twice the power) with the port open.

Only the A-0-1 mode will radiate well in all directions, since it is the only one with a wave length much longer than the size of the box. In my own tests I noted an increase in output both from the ‘main’ soundhole and the port at the A-0-1 pitch, although that increase seemed to be cancelled out at a distance by the increase in amplitude of the out of phase motion of the top with the port open. It is possible that a larger port, which is what they seem to have used, would give a greater gain in air mode radiation without much of an increase in top amplitude. This would also help account for the results of their string pluck experiment.

Conclusions

At this point the reader could be excused for feeling somewhat confused. There’s a lot going on here, and the interactions are complicated. Close miced readings, and the Chaigne and Rosen data from an anechoic chamber test, show appreciable increases in the output of a guitar with a side port, seemingly particularly in the lower part of the ‘bass reflex’ frequency range from about 100-250 Hz, and at certain pitches in the range above, say, 400 Hz, depending on the location and size of the port. On the other hand, tests in a more ‘realistic’ environment and at some distance in front of the guitar suggest that some or most of the increased output simply cancels itself out, particularly in the lowest frequency range. Summing over my ‘sound in front’ spectra yielded increases overall of about 1-2% with an open port, with small drops in output over a broad band canceling the narrow band increases. Interestingly enough, this seemed to hold true no matter what the port location or how many corks I removed. On the other hand, the ’ports’ in my experiment were small: each location being two holes 5/8” in diameter. A larger port might well be more efficient, if only because the ratio of edge to area is smaller, and there is thus less drag to air flowing through the hole.

Opening any port alters the timbre of the guitar, both through greater output from internal air modes that are not normally strong radiators, and because of changes in the way the wood resonant modes act. The changes that seem to have the greatest effect at a distance are in the range from 350-1000 Hz, where normal hearing is becoming increasingly sensitive. Since our senses are set up to detect small changes like that it seems likely to me that this accounts for much of the strong reaction people have when a port is opened. In the playing tests I have done it is common for people to remark on the ‘stronger’ sound when a port is opened, only to express reservations a short time later when they have gotten used to the new timbre. More tests, particularly of the ‘blind’ type which have not been done yet, will be needed to confirm this.

The one undeniable benefit of a side port in a location facing the player is as a ‘monitor’, feeding the player a little more sound than they would normally get directly from the guitar, particularly at higher frequencies. This can be a benefit when one is playing in a well sound proofed room that is large and has a certain level of background noise (think; ’restaurant gig’). This is how I’ve been using it. Some care must be taken with the size and location of the port if the normal timbre of the guitar is to be preserved, but fortunately it does not take a large port to make a useful monitor. Making the normal soundhole smaller on these guitars compensates for the expected shift in the A-0-1 mode pitch. I will note that it is a little hard to predict how much to reduce the main soundhole size when using a port, since the larger the port and the further it is from the hole the greater the pitch shift, in general. There is always the possibility that certain port locations and sizes do, in fact, yield a useful increase in output and efficiency, and if that is the case, they will not be a fad, but rather a standard element of guitar design in the future.

~~~~~~

(1) W.D.Allen "Basics of Air Resonances" American Lutherie #1 March 1985 pp 16-20

(2) Fletcher and Rossing 'The Physics of Musical Instruments" , Springer Verlag 1991 fig 9.9 p. 215

(3) Jansson, 'Acoustics for Violin and Guitar Makers' Fourth edition, 2002, fig. 6.12 available in free download at: http://www.speech.kth.se/music/acviguit4

(4) Elejabarrietta, Ezcurra and Santamaria, "Vibrational Dynamics of the Resonance Box of the Guitar: Finite element Method and Modal Analysis" Catgut Acoustical Society 'Journal', Vol. 4, #4 (Series II) Nov. 2001, pp. 37-41

(5) Rogers, 'Another Look at the Swiss Cheese Violin' Catgut Acoustical society 'Journal', Vol. 3, No. 4, (Series II) Nov. 1997, pp. 17-23

(6) my plate tuning articles in AL #28-30

(7) Chaigne and Rosen, 'Analysis of Guitar Tones for Various Structural Configurations of the Instrument' Catgut Acoustical Society 'Journal', Vol. 3 #8 (Series II) November '99 pp 24-31
~~~~~~