MEASURED TONES

The interplay of physics and music

Ian Johnston

1989

Institute of Physics Publishing
Bristol and Philadelphia

	American system	German system
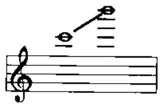	C_6 to B_6	c‴ to b‴
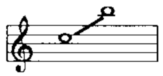	C_5 to B_5	c″ to b″
	C_4 to B_4	c′ to b′
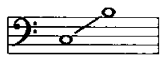	C_3 to B_3	c to b
	C_2 to B_2	C to B
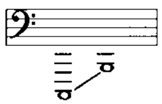	C_1 to B_1	C′ to B′

I think there is no question but that the American system is the more sensible, and that is the one I have used consistently throughout this book.

Appendix 2

Logarithms

A logarithm is a mathematical device which may be understood by thinking about the convention of writing large numbers in terms of powers of 10. For example,

the number 100 is equal to 10×10, and can be written 10^2;
the number 1000 is equal to $10 \times 10 \times 10$, and can be written 10^3.

Consideration of the logic involved in this convention will convince you of the validity of this table:

1 000 000	10^6
100 000	10^5
10 000	10^4
1000	10^3
100	10^2
10	10^1
1	10^0

The table is by no means complete. It is possible, for example, to have *negative* powers if you want—you just get them by continuing to divide by 10.

The main usefulness of this system, from a computational point of view, is that, when you multiply two numbers together, you *add* the corresponding powers. For example,

$$100 \times 10\,000 = 1\,000\,000$$

or:

$$10^2 \times 10^4 = 10^6$$

It is simple enough to check others for yourself.

In about the year 1600, the Scottish mathematician, John Napier, introduced the idea of concentrating attention on the number written in the superscript position, and he invented a new word for it—**logarithm**. So instead of writing, for example,

$$10^3 = 1000,$$

the same information was conveyed by writing

$$3 = \log 1000.$$

Then the previous table could be rewritten as

NUMBER	LOGARITHM
1 000 000	6
100 000	5
10 000	4
1000	3
100	2
10	1
1	0

Next came the idea of logarithms (or powers) which were not whole numbers. One simple way to appreciate this concept is with square roots. We know that $\sqrt{10}$ is equal to 3.1623, or in other words,

$$3.1623 \times 3.1623 = 10.$$

Now because it makes sense to write

$$10^{0.5} \times 10^{0.5} = 10$$

then it must follow that,

$$\log 3.1623 = 0.5.$$

Here's a sample table. Don't worry about how they're calculated: that's what Napier spent half his life doing, and even today it's rather messy.

NUMBER	LOGARITHM
1	0
2	.3010
3	.4771
4	.6051
5	.6990
6	.7781
7	.8451
8	.9031
9	.9542
10	1.000

It isn't so long since school students (and others) regularly used logarithms to perform complicated long multiplications and divisions; and books consisting of nothing but columns of numbers and their logarithms—called **log tables**—were laboriously calculated and published. You see, they still have the important property that, when you multiply two numbers together, you add their logarithms. So, for example, you can check for yourself that

$$\log 2 + \log 3 = \log 6,$$

or,

$$.3010 + .4771 = .7781.$$

A mechanical device which did exactly the same job was the **slide rule**. This consisted of two scales which could be lined up along side one another (which performs the function of adding two lengths together). What was special about these scales was this. In an ordinary ruler, for example, the numbers 1, 2, 3, 4 ... are arranged with equal spacing between them, like this:

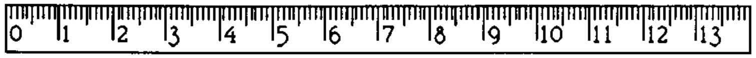

This is what is called a **linear scale**

But it is also possible to arrange the numbers so that their spacing is proportional to their *logarithms*. This kind of scale is called a **logarithmic scale**.

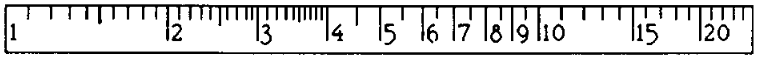

This is what was used in the slide rule, and when a length from one scale was 'added' to the length from the other, this effectively *multiplied* the two numbers on the scales.

Nowadays pocket calculators have made log tables and slide rules obsolete, but the logarithmic scale still has not disappeared; and this is because it is useful for measuring quantities which typically increase multiplicatively rather than additively (particularly psychological ones). The most obvious of course, which I have been talking about throughout this book, is musical pitch.

It is an interesting exercise to draw lines on a sheet of logarithmic graph paper (paper in which the lines are spaced logarithmically rather than linearly), corresponding to the notes which define the lines of the ordinary musical stave. What you can see clearly is that the ordinary stave is in fact a good approximation to the 'natural' arithmetical scale against which you should measure pitch.

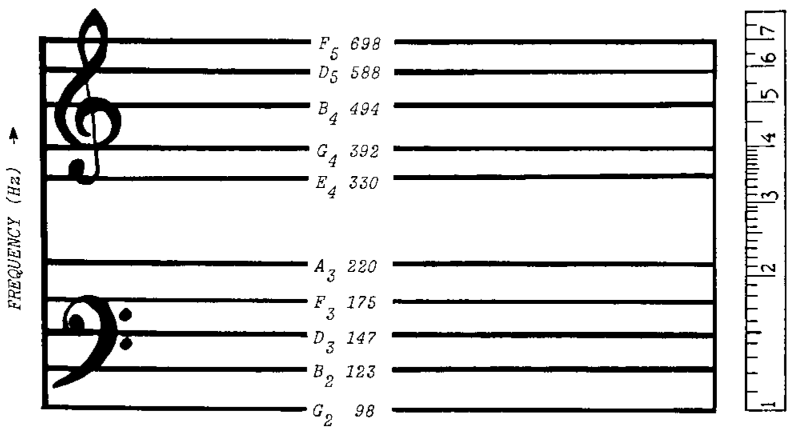

Appendix 3

Measurement of pitch intervals

Any system which attempts to quantify the perceptual entity, pitch—as opposed to the physical quantity, frequency—must be based on the fundamental observation (discussed at length in Chapters 1 and 2) that intervals of pitch which are judged to be equal to one another, have frequencies which are in the same numerical ratios. The ratios corresponding to the most important pitch intervals in music may be represented on a bar graph like this:

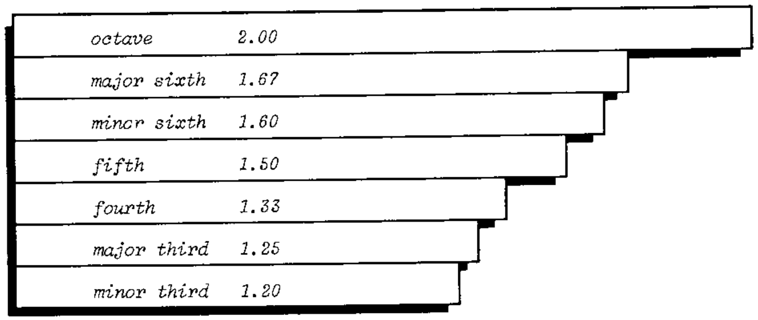

In this diagram a *linear scale* is used to represent the ratios; but we know it would be more sensible to use a *logarithmic* one (see Appendix 2). In which case, exactly the same information would be contained here;

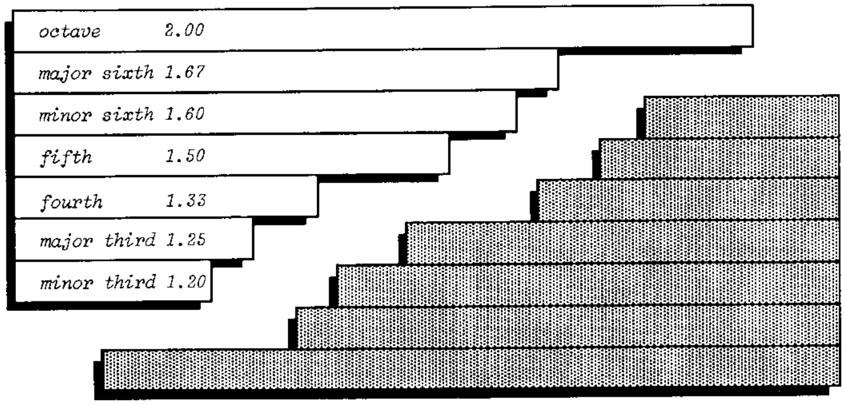

The convenient thing about this diagram is that the 'lengths' associated with the different intervals show clearly the known results of 'adding' certain pitch intervals together:

fourth	+	fifth	=	octave
major third	+	minor sixth	=	octave
minor third	+	major sixth	=	octave
minor third	+	minor third	=	fifth
...etc.				

This means that it is possible to use the 'length' of the ratios along the logarithmic scale (i.e. the logarithm of the numerical ratio) as a direct *measure* of the size of the intervals. But to make such a measurement useful, we need a *unit* of measurement that everyone agrees on. Such a unit was suggested in 1895, by the English musicologist (and translator of the works of Helmholtz), Alexander Ellis. He chose the hundredth part of an equal tempered semitone, and called it the **cent**.

Recall, by referring back to page 77, that the ratio corresponding to the equal tempered semitone is 1.0595 (actually 1.059463 is a more accurate value); and the logarithm of this is .025086, (which you can check with a pocket calculator). Then the cent is a hundredth of this,

$$1 \text{ cent} = .00025086 .$$

Therefore, if you have two frequencies, f_1 and f_2 (where f_2 is the higher), you can calculate the size of the pitch interval between them by this important formula:

$$\text{number of cents} = 3986.3 \times \log(f_2/f_1)$$

The number 3986.3 appearing here is simply the reciprocal of 0.00025086. Some sample calculations based on this formula, showing the number of cents in various intervals are:

PYTHAGOREAN TUNING

semitone	tone	fourth	fifth	octave
90	204	498	702	1200

JUST TUNING

semitone	minor tone	major tone	minor third	major third
112	182	204	316	386
fourth	fifth	minor sixth	major sixth	octave
498	702	814	884	1200

EQUAL TEMPERED TUNING

semitone	second	minor third	major third	fourth
100	200	300	400	500
fifth	minor sixth	major sixth	seventh	octave
700	800	900	1100	1200

Appendix 4

Measurement of loudness

The loudness of a sound is the psychological perception that our brain ascribes to that property of an acoustic wave, for which the most obvious corresponding physical quantity is the rate at which energy is carried by the wave into our ears.

Any freely vibrating body loses energy as its vibration decays, and the rate at which this energy escapes can be measured in so many joules/second (the **joule** is the unit of energy in the S.I. system). Technically this quantity is called **power** and its unit of measurement is given the name of the **watt** (after guess who). To give you some feel for numbers, here are a few generally accepted measured values for the acoustic power output of some musical instruments.

INSTRUMENT	ACOUSTIC POWER (W)
Clarinet	0.05
Double bass	0.16
Trumpet	0.31
Piano	0.44
Cymbals	9.5
Bass drum	25
Orchestra	67

By way of comparison, light bulbs which give out radiant energy (heat as well as light) at a rate of 100 W are common. But in the case of most acoustic systems, the processes by which this escaping energy is transformed into sound are, in general, remarkably inefficient. A good average figure is that, of the power lost by a freely vibrating mechanical system, at most 1% will end up as acoustic energy.

Acoustic energy travelling through ordinary space spreads out, so that at some distance away, the amplitude of the wave depends not only on the power output of the source, but also the *area* over which it has spread. Any 'measuring' instrument, like an ear for example, will respond only to that energy which crosses the area of its detecting surface. The physical quantity being 'measured' is called the **intensity** of the wave, and its units are watts/square metre. Again typical values, at various distances away from some of the same instruments, are given in this table.

INSTRUMENT	INTENSITY (watts/m^2) at 1 m	at 10 m
Clarinet	0.004	.00004
Trumpet	0.012	.00012
Cymbals	0.75	.0075
Orchestra	5.3	.053

These figures are only average values of course. Sound from any source will not necessarily spread out equally in *all* directions. Indeed a careful series of measurements around these instruments shows surprising variations of intensity in different directions for various frequencies. Such work is obviously important in determining the appropriate seating arrangements for players in an orchestra; but for the present we can ignore this kind of fine detail.

In general, the softest sound that most ears will respond to has an intensity of about 10^{-12} watts/m^2—an almost unbelievably small figure (which just goes to show how sensitive the ear is). As a comparison, a 100 W light bulb emits an intensity of radiant energy at a distance of 10 m of just under 0.1 watts/m^2. The acoustic intensity at the threshold of hearing is therefore a factor of 10^{11} times less than this.

Now we know our ears do not respond *linearly* to acoustic intensity. Instead they seems to compare sounds with one another in terms of the ratio of the energy collected (as do most of our perceptual organs). Therefore the sensible scale on which to measure the psychological quantity **loudness** is a logarithmic one; and the minimum discernible intensity provides a convenient reference point from which to start the scale. The earliest agreed system was one in which the loudness, measured in the unit of a **bel**, was simply defined as the logarithm of the ratio between the measured intensity and the minimum discernible level—i.e.

$$\text{loudness (in bels)} = \log\left(\frac{\text{intensity}}{10^{-12}\ \text{W/m}^2}\right)$$

The greatest intensity that a sound can reach before it becomes physically painful, and therefore cannot really be considered a 'sound' any more, is about 1 W/m^2—which corresponds to a loudness of 12 bels. In order to have a scale that wasn't limited to 12 levels only, it was decided that in general it was sensible to use as the basic unit the tenth part of the bel, or the **decibel**. Hence the fundamental formula for working out the loudness of a sound, once you know the intensity in W/m^2, can be rewritten as

$$\boxed{\text{loudness (in dB)} = 120 + 10\times\ \log\text{(intensity)}}$$

So the table on page 251 can be completed by including the absolute intensities, thus:

	INTENSITY (W/m^2)	LOUDNESS (dB)
Pain threshold	1.0	120
fff	.01	100
fortissimo	.001	90
forte	.0001	80
mezzo-forte	.00001	70
piano	.000001	60
pianissimo	.0000001	50
ppp	.00000001	40
Hearing threshold	.000000000001	0

It is worth pointing out, in case it isn't obvious, that the standard musical loudness markings are anything but absolute. Clearly a symphony orchestra playing *pp* is still much louder than, say, a single clavichord playing *pp*. The correspondences on the preceding table can be taken only as suggestive.

There is another way in which this description of loudness is oversimplified. The human ear is not equally sensitive at all frequencies. The graph on page 245 shows clearly that it responds most readily to sounds at frequencies between 2000 and 4000 Hz (largely because of the dimensions of the auditory canal). This means that it will perceive a sound of fixed intensity as being much louder than otherwise, if the sound has a frequency in this range. Therefore the definition of the decibel is only a crude approximation to true subjective loudness, because it assumes an average sort of frequency for whatever is being measured.

There are in existence several systems which attempt to make allowance for this variation with frequency. Series of measurements have been conducted with many listeners to determine what intensities of *pure* tones are judged to be equally loud. The average results of many such experiments are then used to produce a plot of contours of equal loudness against frequency, and the range between the thresholds of hearing and pain can again be divided up into 120 different levels. These levels are measured in a unit called a **phon**, and the scale is calibrated by making the loudness level in phons of any contour equal to the sound level measured in dB at 1000 Hz. Needless to say this unit is rather subjective.

There is yet another unit you may come across, a **sone**, which is defined to be the subjective response equal to that of a 1000 Hz tone at a sound level of 40 dB. Both these systems of measurement are useful to experts working in the field of noise control and such like, but for the purposes of most discussion in this book, the decibel is really quite good enough.

Appendix 5

Acoustic impedance

The concept of acoustic impedance was introduced in Chapter 6, and I don't intend to define it more carefully here. That would require the use of calculus quite beyond the scope and purpose of this book. However there are some points that can only be made by arithmetical calculations, and therefore I want to say something about the important formula I quoted on page 192.

Calculation of energy reflected

Let me start by again using the analogy of the Newton's cradle, and consider a collision between two balls

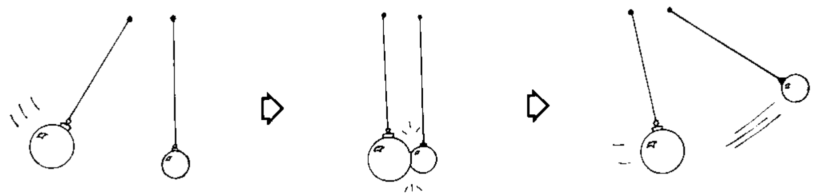

The way you would work out how much energy is given by the first ball to the second in the collision proceeds as follows. First you note that, because of the law of conservation of energy, you can say something about the initial and final kinetic energies (KE) of the two balls.

$$\text{Initial KE of ball (1)} = \text{Final KE of ball (1)} + \text{Final KE of ball (2)}$$

Secondly, because the same force acts (though in opposite directions) on the two balls, for the same length of time (the time they are in contact), the change in the velocity of each will be in inverse proportion to their inertias. So you can say something else about the initial and final velocities of the two balls.

$$\text{mass (1)} \times (\text{initial velocity (1)} - \text{final velocity (1)}) = \text{mass (2)} \times \text{final velocity (2)}$$

(This expression, incidentally, is called the **law of conservation of momentum**. Though I haven't had occasion to mention it before, it is an extremely important law in other contexts, rivalling energy conservation in overall importance to physics.)

If you are good at algebra you might like to try to express kinetic energy in terms of velocity, and then to solve these two equations. If not, just take my word for it that they are equivalent to this single equation for the *ratio* of the final KE of the first ball to its initial KE—i.e. the *fraction* of the energy which stays with the first ball (and can therefore be thought of as having been 'reflected').

$$\text{fraction of energy reflected} = \left(\frac{\text{mass of first ball} - \text{mass of second ball}}{\text{mass of first ball} + \text{mass of second ball}}\right)^2$$

This formula straightforwardly expresses the important point that has carried all through my development of the ideas in Chapter 6. If the two masses are different, irrespective of which is the heavier, the first ball will retain a fraction of its initial energy; and the bigger the mass difference, the greater this fraction. Only when the two masses are equal will *none* of the energy be retained and will all of it be passed on.

The more complicated, but completely analogous, problem of a sound wave travelling down a tube of air is analysed in exactly the same way. You consider the 'collision' of one slice of air in the tube with the next one. You note again that energy has to be conserved, and that the same forces act on the 'slices' for the same length of time (or, if you like, that momentum is also conserved). The only difference is that the resulting equations have to use calculus (they are called **differential equations**). But the *form* of the logical conclusion is still the same, with the role of inertia being taken by **acoustic impedance**.

$$\text{fraction of energy reflected} = \left(\frac{\text{impedance (1)} - \text{impedance (2)}}{\text{impedance (1)} + \text{impedance (2)}}\right)^2$$

If you put numbers into the formula you get a clearer idea of how the fraction of energy which is reflected depends on the *ratio* of the two impedances.

impedance (1) / impedance (2)	% energy reflected	% energy transmitted
0.01	96	4
0.05	82	18
0.1	67	33
0.2	44	56
0.5	11	89
1.0	0	100
2	11	89
5	44	56
10	67	33
20	82	18
100	96	4

From this table you will note the important point that if the two impedances don't match exactly (i.e. if their ratio isn't exactly 1) then some energy will be reflected. And it doesn't matter if the first impedance is greater than or less than the second.

There is however one difference. The quantity inside parenthesis actually represents the *amplitude* of the reflected wave. It is *positive* if the second impedance is smaller than the first, and is *negative* if they are reversed. Of course, when you calculate the energy, you square this factor, so its sign doesn't matter. But there are circumstances when the sign of the amplitude is important, in particular if the impedance mismatch we are thinking about is that which occurs *at the end of a pipe*. An open pipe represents a case when the second impedance is smaller than the first, and a closed pipe a case when the second impedance is greater than the first. We noted that these two cases gave rise to different standing waves because the amplitude of the reflected waves were of opposite signs, and that is exactly what this analysis predicts.

Impedance matching

The fraction of energy *transmitted* from one medium to another can be improved by putting a small length of another medium with intermediate impedance between the two. A simple example will serve to demonstrate this.

Imagine you have two pipes, one half the area of the other, so that their impedances are in the ratio 1:2.

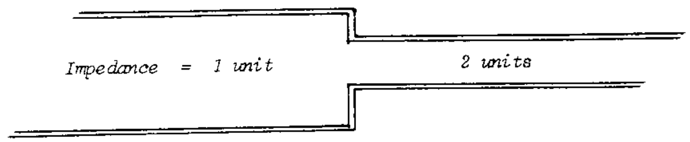

You can calculate easily that 11% of the energy of any wave hitting this junction will be reflected; and therefore only 89% gets through.

But now consider what happens if you insert another piece of pipe between the two, whose area is 2/3 that of the first pipe, so that the three impedances are in the ratios 2:3:4.

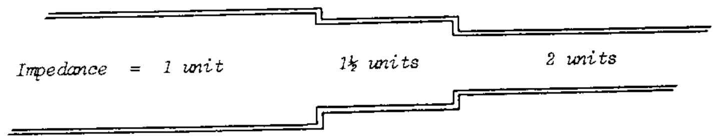

If you do the sums—and please check my calculations for yourself—you will see that 4% of the energy is reflected at the first junction, and 96% let through. Then at the beginning of the next pipe the fraction reflected is only 2%, so that 98% gets past this junction—that is, 98% of whatever falls on it. So the total fraction that gets from the first pipe all the way to the end, is 98% of 96%—which is 94%. And this is a considerably larger fraction than the 89% I calculated earlier.

I hope now you will be able to accept that this result is quite general. Any value of the acoustic impedance of the intermediate pipe, provided it is between the two values on either side of it, will have the effect of allowing a greater fraction of the energy than the original 89% to get through. It relies on the fact (which the formula confirms) that if the impedance mismatch is not too great, then the fraction of energy transmitted is quite close to unity.

The corollary to this should also be plausible without my proving it rigourously. More than one intervening piece of pipe will improve the transmission still further; and, taking the argument to its logical limit, it should be possible to get *all* the energy through by using an infinitely graded set of 'steps'—i.e. a gentle flare.

Brass instruments revisited

This whole subject is obviously very important for any wind instrument whose bore differs greatly from geometric simplicity—in particular brass instruments. When I discussed these in Interlude 1, I hadn't yet established the physical concept of acoustic impedance, so I couldn't say much about the important question of how the makers design the shapes of their instruments to play the notes they want. Let me do that now.

When you look at a trumpet, open at one end, but covered by the player's lips at the other, your first expectation might be that it should behave like a closed organ pipe—that it should play only odd harmonics. Yet you know that it plays a *complete* harmonic series (even and odd). The explanation is that the makers have tampered with its shape so much that they have completely changed its modes of vibration. And they have done this by careful shaping of the bell and the mouthpiece.

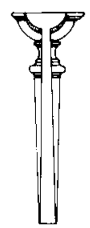

The **mouthpiece** consists of a cup shaped bowl joined on to the narrow end of a short conical tube (which is designed to fit into the main tube of the instrument). There is a severe impedance mismatch at the constriction. Therefore waves will suffer a sharp dislocation at this point—not quite a complete reflection, but close. So this point will have some of the effect of an 'open' end.

The **bell** on the other hand, changes its diameter much more smoothly. It starts off like a gentle cone, but soon curves away from the cone shape more and more.

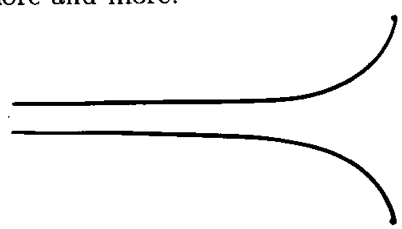

The effect of this is dramatic. Whether or not the flare of an expanding tube seems 'gentle' depends on the wavelength of the standing wave and how it compares with the radius of curvature of the pipe (that is the length over which the bore changes appreciably). At the start of the bell the flare is very mild, the radius of curvature is large. But even so, very long wavelength standing waves (the fundamental and the first overtone) find it too rapidly changing for their taste. So they reflect well down inside the instrument. But the higher overtones (having shorter wavelengths) will reach further out until each one comes to a point where the curvature is much tighter and the variation of the bore is too 'abrupt' for it, Then they reflect—producing a pressure node at that point. So, by the time the bell and the mouthpiece have had their effects, the different standing waves therefore look more or less like this, and you end up with almost a full harmonic series.

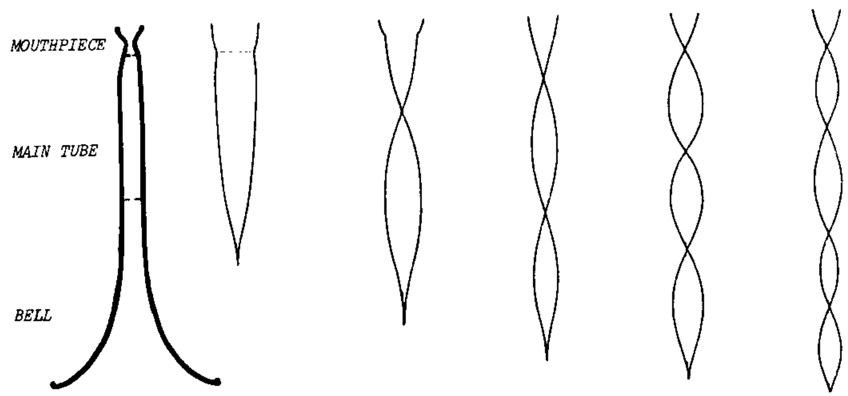

I say 'almost' a full series because the lowest mode is such a distorted standing wave that its frequency is usually quite out of tune with the others and musically unusable.

As one last example of these ideas, I think you could appreciate what happens if you rest your hand lightly inside the bell, as a French horn player does. You constrict the air channel at the point of reflection, increasing its impedance. Therefore you *lessen* the mismatch between the inside of the tube and the outside air, and the pitch is lowered slightly—for exactly the same reason as I outlined when I was explaining the 'end correction' on page 201. As you possibly know, this is exactly what happens.