

Understanding the Acoustics of

The Native American-Style

Flute

By Mike Prairie

Preface

This paper represents a compilation of my understanding of the physics of the flute so far, and is currently an evolving document. It is written from the perspective of the Native American style flute with a fixed sound hole configuration and a limited second-octave range. The treatment is appropriate for the penny whistle as well. Most of the treatment is also consistent with mouth-blown flutes like the simple-system Irish flute or the shakuhachi, but no effort is made to account for embouchure variations.

The original works I read on the physics of flutes were written by Lew Paxton Price in a pair of books titled *Secrets of the Flute: The Physics, Math, and Design of Non-Mechanical Folk Flutes* and *More Secrets of the Flute: More of the Physics, Math, and Design of Non-Mechanical Folk Flutes*, and I owe a great deal to the author for providing me with that jump start. His work is probably the most referenced explanation in the NAF community, so I have made an effort to be consistent with his notation and descriptions whenever I can. The balance of the notation style is intended to follow that of Nederveen.

Price's books do well in the first octave, but I was looking for some quantitative explanations of the physics in the second register. So I turned to Arthur Benade. His writings gave excellent qualitative explanations of such things as perturbation weight functions and tapered bores, but I still couldn't find the mathematical construct for which I was searching. I found it in Cornelis Nederveen's book *Acoustical Aspects of Woodwind Instruments*. The book is a thesis with a relatively detailed mathematical treatment of acoustic impedance that models well the effects of variables describing the various elements of flutes. I began the process of digesting his work with the backdrop of Price and Benade looming ever present.

Nederveen's book does not have all the answers, but it does provide the tools I needed to find the ones I have found so far. I had developed my own translations of how certain effects can be modeled, and noted that there was an unorganized pile of drawings, writings, analyses and computer code collecting on my hard drive. So I began to organize those thoughts before I forgot the concepts, or where I put them. This document is the result, and represents more a notebook for my own purposes than an attempt to publish a book. It is not nearly finished, but I will continue to add to this as I can find the time.

I'll add a note regarding my expectations of the theories. The manufacturing process of a wooden flute has limits in accuracy, and the wood tends to change dimensions over time. Furthermore, the number of interrelated variables in a flute is large, and any inaccuracies caused by approximating or omitting any physical effect are compounded. As such, I DO NOT expect to be able to precisely model a complete flute. I DO expect, however, to model *trends* in changes to the intonation of a flute caused by changes in physical variables. With an understanding of these trends,

changes in design can be guided with confidence, and the number of trial-and-error steps can be minimized.

I am releasing it in its current form so that other interested folks can give my explanations a try if they are having trouble with other references. I am attempting to make it as accessible as I can, so someone with an extensive background may be put off by some of my repetition of simpler concepts or long-winded explanations of concepts that may take some extra thought. And I also expect that some of my explanations will be confusing to some (and some explanations may have erroneous information I have yet to catch!). It is what it is, and I appreciate constructive suggestions. This work has not been peer-reviewed and some ideas have yet to be fully tested, so I offer no “warranty.” I currently claim no copyright, so I encourage folks to use the material, but I do ask for the professional courtesy of a reference if you do use this work outside of your own flute making and designs.

Enjoy the journey!

Mike Prairie
Maryland, USA
January 2006

Table of Contents


Preface	2
The Air Column	5
A Vibrating Column of Air	5
Resonance	9
Higher Harmonics	9
Making “Cents” Out of Frequency	12
Temperature effects	13
A Note on End Corrections	14
Acoustic Length	15
k_1 —the Acoustic End Correction (or “False Length”) at the Foot of the Bore.....	15
k_2 —the Acoustic End Correction (or “False Length”) at the TSH	17
The Playing Holes	21
A Note on Playing Hole End Corrections	24
Acoustic Impedance.....	26
Using the Impedance to Describe a Vibrating Air Column	26
Finger Holes and Shortening the Acoustic Air Column.....	35
End Corrections at the Sound Hole.....	40
Backset Effects	46
Bore Shapes and Perturbations	48
A Segmented Bore	48
Comparison to Nederveen.....	54
Perturbation Weight Functions	55
Diversion from the Simple Cylinder.....	58
Estimating the Change in Pitch.....	59
Bibliography	62

The Air Column

A Vibrating Column of Air

The flute can be thought of as a container that holds a column of air. In its simplest form, it is a tube that is open at both ends. The air is “compressible,” which means it has some springiness to it. Another word for it is “elastic.” And although it is as light as air (since it IS air...), it does weigh something so it will have some “inertia,” or resistance to change in motion.

So if someone pushes on the air column at one end of the tube, what happens? If the pushing is done slowly, the column will just slide out the other end. If the push is sudden, the end of the column that was pushed will compress like a spring at first (the elastic property) and the column beyond the compressed part will not move right away (the inertial property). But the coiled spring in the compressed part will push on the column beyond, and eventually the compressed part will work its way to the other end of the tube. The same sort of thing will happen if the end of the column is pulled; the stretched part will move down the tube until it hits the end.



The air itself does not flow all the way down the tube, but a pressure wave does. At any point in the tube, a chunk of air is pushed or pulled by its neighbor behind it, and then it pushes or pulls the next chunk in the line in order to go back to its original position. The wave is the disturbance to the air molecules that moved down the tube. It's the same with “The Wave” football fans do at a stadium—the people don't move, but the wave does.

If the end is pushed and pulled in a regular periodic cycle, the alternating compressions and elongations of the spring will travel down the tube and eventually cause the opposite end of the air column to move in and out of the end of the tube. These alternating compressions and elongations of the air column are sound waves, and the speed at which the waves move down the tube is the speed of sound. The speed of sound is

$$v_s = \sqrt{\frac{B}{\rho}}$$

where **B** is called the “bulk modulus of elasticity,” which is a measure of the springiness of the air, and **ρ** is the density, which is a relative measure of the heaviness of the air. Except for temperature, all the variables that define **B** and **ρ** for air are fairly constant. Putting those numbers in to the equation leaves

$$v_s = 331.39 \sqrt{\frac{T(K)}{273.15}} \left[\frac{m}{s} \right]$$

where **T(K)** is the temperature in Kelvins, and the velocity is in meters per second. Converting to inches per second and using **T** in Fahrenheit, this becomes

$$v_s = 12600.5 \sqrt{\frac{459.4 + T}{459.4}} \left[\frac{in}{s} \right].^1$$


At 68° F, that’s 13,500 inches per second. Let’s say the tube is 13.5 inches long. If the column is pushed in at one end, how long will it take for the wave to reach the other end? If the speed of sound is 13,500 inches per second (...how convenient!), then a pressure wave will take 13.5/13500, or 1/1000 of a second to reach the other end.

If the pushing and pulling was timed so that the “pull” happened just as the previous “push” hit the end of the tube, the ends of the air column would be moving out of the tube at both ends at the same time. In our example, that would be if the column is pulled back 1/1000 of a second after it was pushed. The air in the center of the tube would be pulled equally in both directions in that case, so it wouldn’t move at all. Then when the “push” came along, it would coincide with the last “pull” reaching the other end of the tube and pulling air into the end. So then the ends of air column would be moving into the tube at both ends at the same time. The air in the center of

¹ This is the equation given by Lew Paxton Price in his “Secrets of the Flute” book. If variables and physical constants from different sources are used, an equation with slightly different numbers results. In the final analysis, however, the differences in pitch are within about 0.5 Hz. This translates to less than 2 cents, which is below perceivable differences.

the tube would be squeezed equally in both directions in that case, so once again it wouldn't move at all. If this is repeated continuously (and the pushing and pulling is smooth and periodic), then an oscillation is set up that has a period of $2/1000=1/500$ of a second, or a frequency of 500 cycles per second. This would produce a sound at 500 Hz.

The figure below shows how the air column stretches and contracts during the vibration cycle. The medium blue represents zero acoustic pressure, the lighter color occurs when the pressure drops because the column is stretched, and the dark color represents a compressed part of the column where the pressure is higher. The things to note are that the ends of the column remain at atmospheric pressure (zero acoustic pressure) and move in and out of the end of the tube, and the center of the column experiences large pressure changes and remains stationary.




As this cycle continues, the center of the air column would stay still at what is called a “flow node,” and the ends of the column would flow in and out in a mirror image. The flow is at a maximum at the ends, and that is called a “flow antinode.”

An interesting thing to point out is what is happening to the pressure in parts of the air column. This will be important later on, but for now it is just “interesting.” The ends of the air column are connected to the air around the tube, so the pressure at the ends is


fixed to the same pressure as the atmosphere. So the acoustic pressure, which is the changes compared to the atmosphere², is about zero here; this is called a “pressure node.” In the middle of the tube, the air is getting squeezed and stretched at a maximum. This is a “pressure antinode.”

Looking at the same cycle from a slightly different perspective, if the column is thought of as a long spring, we can look at how the spring moves, or doesn’t move, throughout the cycle. The picture below shows the spring as it contracts (turning red in the drawing) and expands (turning blue). The rest position of the spring is shown with thinner black lines.



Combining the contracted, rest, and compressed states of the spring in one drawing highlights the important concept about the motion of the air column—that the center stays fixed at one point, and the motion gradually increases toward the open ends of the tube where the ends of the column spring move freely. This is shown below in the second drawing (the first is the column spring at rest). The third drawing what happens during the second mode of oscillation, which will be discussed shortly, and the fourth picture shows mode 3.

² The term “acoustic pressure” refers to the local changes in pressure that will propagate through the air as sound. So whatever the total pressure is at an instant in time, we subtract the atmospheric pressure to get the acoustic pressure. Even though the sound we hear from a flute is a result of radiating pressure waves, the positive and negative acoustic pressures at the ends are miniscule compared to what is going on at the center of the flute. So the ends are effectively at zero acoustic pressure compared to the rest of the air column. Unless otherwise stated, the use of the term “pressure” will refer to “acoustic pressure” for convenience.



Resonance

One thing not mentioned yet is the fact that the air column coming out of the tube at the far end has some outward momentum, so it will not just stop when it reaches the original length. It will continue moving a short distance and cause the column to stretch inside the tube. That stretched air is like a stretched spring, and it will pull the end of the column back in. On its way back into the tube the end of the air column will have inward momentum, and it will crash into and compress the air in the tube. This will form a new compression wave that will be sent as a reflected wave back into the bore in the opposite direction.

The reflected wave will reach the end where the pulling and pushing is happening and try to push the end of the air column out of the tube. If a push is happening at that time, then the push and the pressure wave will collide and the total energy left will be the difference between the push and the reflected wave. This is called “destructive interference.” If the wave arrives at the end just as a pull is happening, then the two are moving in the same direction, and they will add together and be even stronger. This is called “constructive interference.” That combined wave will be reflected again, and will start down the tube again just as the next push comes along. You could think of this like a child on a swing. Think of the child as the pressure wave that keeps moving back and forth, and at one end mom or dad gives the child a little push just as the swing starts forward again. This effect is called resonance, and it preferentially increases the energy of oscillations at certain frequencies, and destroys the oscillations at others. This is the case with the air column in the tube. The push-pull action at one end is the “generator,” and the tube is the “resonator,” and when the generator frequency matches the natural frequency of the resonator, they are in “resonance.”

Higher Harmonics

Resonance can happen at higher frequencies as well. These are called “harmonics” because they are multiples of the lowest, fundamental frequency. Here’s how it works with the pressure waves. In the example above, our generator was working at 500 Hz,

and that resonated with the 13.5-inch tube because the reflected waves constructively interfered with the generated waves. If we run the generator at 1000 Hz, we can launch a compression wave down the tube and launch an elongation wave just as the first compression wave gets half way down the tube. As the first compression wave hits the end of the tube, we launch another one. As the first one reflects and starts back up the tube, the generator launches another elongation wave. The first compression wave makes it to the half-way point on its way back to the generator as another compression wave is generated. Finally, the initial reflected wave hits the end of the tube as the generator is “pulling” another elongation wave. Since the pull and the wave are moving in the same direction, they constructively interfere.

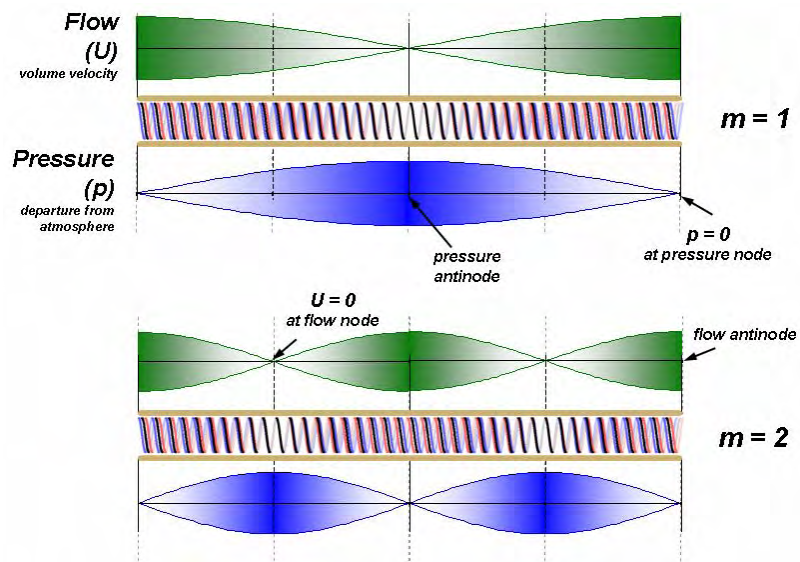
The distance over which the sound wave travels during one full cycle of pushing and pulling by the generator is called the “wavelength.” In the 500-Hz example, the wave traveled down and back up the tube before the next cycle started, so the wavelength was twice the length of the tube, or $\lambda = 2L$. In this case, the second cycle of the 1000-Hz wave starts as the wave finishes the one-way trip down the bore, or $\lambda = L$. The general relationship between frequency and wavelength is

$$f = \frac{v_s}{\lambda} .$$

The relationship between the frequency and the length of the tube is

$$f = m \frac{v_s}{2L} ,$$

where the factor “**m**” is introduced as the mode number, which is the number associated with the harmonic of the fundamental. It is always an integer, and the fundamental pitch occurs for **m**=1. The figure below shows how much the flow (green) and pressure (blue) vary along the length of the simple air column for mode 1 in the upper drawing and mode 2 in the lower drawing. The shape of the curves will be revisited in a later section.



Making “Cents” Out of Frequency

When tuning a flute, the pitch is usually referenced to the Western chromatic scale with the A above middle C defined as 440 Hz. The scale has 12 notes, or “semitones” distributed across each octave. A chart of these notes is below.

MIDI	Note	Pitch (Hz)	MIDI Fraction (cent)										
			-50	-40	-30	-20	-10	0	+10	+20	+30	+40	+50
57	A3	220	214	215	216	217	219	220	221	223	224	225	226
58	A#3	233	226	228	229	230	232	233	234	236	237	239	240
59	B3	247	240	241	243	244	246	247	248	250	251	253	254
60	C4	262	254	256	257	259	260	262	263	265	266	268	269
61	C#4	277	269	271	272	274	276	277	279	280	282	284	285
62	D4	294	285	287	289	290	292	294	295	297	299	301	302
63	D#4	311	302	304	306	308	309	311	313	315	317	318	320
64	E4	330	320	322	324	326	328	330	332	333	335	337	339
65	F4	349	339	341	343	345	347	349	351	353	355	357	359
66	F#4	370	359	362	364	366	368	370	372	374	376	379	381
67	G4	392	381	383	385	387	390	392	394	397	399	401	403
68	G#4	415	403	406	408	411	413	415	418	420	423	425	427
69	A4	440	427	430	432	435	437	440	443	445	448	450	453
70	A#4	466	453	456	458	461	463	466	469	472	474	477	480
71	B4	494	480	483	485	488	491	494	497	500	503	505	508
72	C5	523	508	511	514	517	520	523	526	529	532	535	539
73	C#5	554	539	542	545	548	551	554	558	561	564	567	571
74	D5	587	571	574	577	581	584	587	591	594	598	601	605
75	D#5	622	605	608	612	615	619	622	626	629	633	637	640
76	E5	659	640	644	648	652	655	659	663	667	671	675	679
77	F5	698	679	683	686	690	694	698	703	707	711	715	719
78	F#5	740	719	723	727	731	736	740	744	749	753	757	762
79	G5	784	762	766	771	775	779	784	789	793	798	802	807
80	G#5	831	807	812	816	821	826	831	835	840	845	850	855
81	A5	880	855	860	865	870	875	880	885	890	895	901	906
82	A#5	932	906	911	916	922	927	932	938	943	949	954	960
83	B5	988	960	965	971	976	982	988	993	999	1005	1011	1017
84	C6	1047	1017	1023	1029	1034	1040	1047	1053	1059	1065	1071	1077
85	C#6	1109	1077	1083	1090	1096	1102	1109	1115	1122	1128	1135	1141
86	D6	1175	1141	1148	1154	1161	1168	1175	1181	1188	1195	1202	1209
87	D#6	1245	1209	1216	1223	1230	1237	1245	1252	1259	1266	1274	1281
88	E6	1319	1281	1288	1296	1303	1311	1319	1326	1334	1342	1349	1357
89	F6	1397	1357	1365	1373	1381	1389	1397	1405	1413	1421	1430	1438
90	F#6	1480	1438	1446	1455	1463	1471	1480	1489	1497	1506	1515	1523
91	G6	1568	1523	1532	1541	1550	1559	1568	1577	1586	1595	1605	1614
92	G#6	1661	1614	1623	1633	1642	1652	1661	1671	1681	1690	1700	1710
93	A6	1760	1710	1720	1730	1740	1750	1760	1770	1780	1791	1801	1812

Since the pitch is often given as referenced to the closest note plus or minus a fraction of the difference between that note and the next higher or lower note, respectively, a convenient way to describe that fraction is to use “cents.” Each semitone step is divided into 100 parts, or cents. The MIDI scale is divided this way as well, so using its mathematical definition is convenient in converting the frequency in Hz to a MIDI number that refers to a pitch and how flat or sharp it is. The equation for the MIDI number given the frequency is

$$MIDI = 12 \frac{\ln(f/6.875)}{\ln(2)} - 3 ,$$

And to convert MIDI to frequency, use

$$f = 6.875 \times 2^{\left(\frac{3+MIDI}{12}\right)}.$$

In the example above, the MIDI number for 500 Hz is 71.21. That corresponds to B₄ plus 21 percent of the way from B to C, or B₄ + 21 cents. The 1000-Hz frequency is an octave higher, so it is B₅ + 21 cents.


Temperature effects

One of the properties of the air column is that its properties change with temperature. There are also some minor effects due to humidity and altitude, but the temperature effect is far more dominant. As the temperature rises, the density of the air drops, and this causes the velocity of sound to increase as was described earlier. So when the elements of a flute are designed, the design must be done for a specific playing temperature. If the flute is played or tuned by the flute maker at a different temperature, the difference can be accounted for by a simple shift in pitch. The equation for frequency is given by

$$\begin{aligned} f &= \frac{m}{2L} v_s \\ &= \frac{m}{2L} \left(12600.5 \sqrt{\frac{459.4 + T}{459.4}} \right) \end{aligned}$$

Plotting this as a function of temperature, one can see how much of an offset in pitch is required if the playing temperature is different from the design temperature. The chart below shows the offset in pitch one would use in a cooler or warmer room when tuning a flute designed to be played in tune at 72° F.

Tuning Offset for Temperature Variations



In the case of our 13.5-inch air column resonating at $B_4 + 21$ cents (500 Hz) at 68°F , if we were to increase the temperature to 80°F , the pitch would be $B_4 + 41$ cents (506 Hz).


If we wished to make the tube resonate at a perfect B (494 Hz), we would have to cool it to 54°F . Or we could use a longer tube. At 72°F , for example, our tube would resonate at 494 Hz if we increased the length to 13.72 inches.

A Note on End Corrections

As a final note in this section, we must acknowledge that the length of the physical tube does not accurately describe the acoustic length of the air column. The air at the ends of the column is actually interacting with the air outside the tube. The pressure node described above is a point at the end of the acoustic length of the air column. The end of the column is defined where the mathematical projection of the acoustic pressure wave outside the tube would reach zero. The wave is never really pinned to zero pressure because it is actually radiating the pressure waves we hear as sound, but we can define a virtual length of tubing that extends beyond the physical end that defines where the pressure would be pinned to zero if it did indeed work that way. This virtual length is called an “end correction” or “false length,” and it allows us to model the air column as if it had an ideal length that had precisely-defined pressure nodes at its end point. This is discussed in some detail in the next section.

Acoustic Length

The acoustic length of a flute sounding the fundamental note is half the wavelength of that note. However, the bore length (measured from the plug to the foot) is significantly shorter than that length. To account for the rest of the length, the acoustic corrections at the ends of the flute must be considered. The total acoustic length, then, is the sum of the bore length plus the corrections at the foot (k_1) and the sound hole (k_2).




k_1 —the Acoustic End Correction (or “False Length” at the Foot of the Bore

The vibrating air column in the bore extends past the openings of the tube. This can be approximated by a bubble that is a half sphere with a volume given below. The bubble can be converted to a cylinder with the same volume, and the length of that cylinder is defined as k_1 . k_1 is acoustically how far the air column extends past the end of the foot.

$$V_{half\ sphere} = \frac{1}{2} \left(\frac{\pi D^3}{6} \right) = \frac{\pi D^3}{12}$$

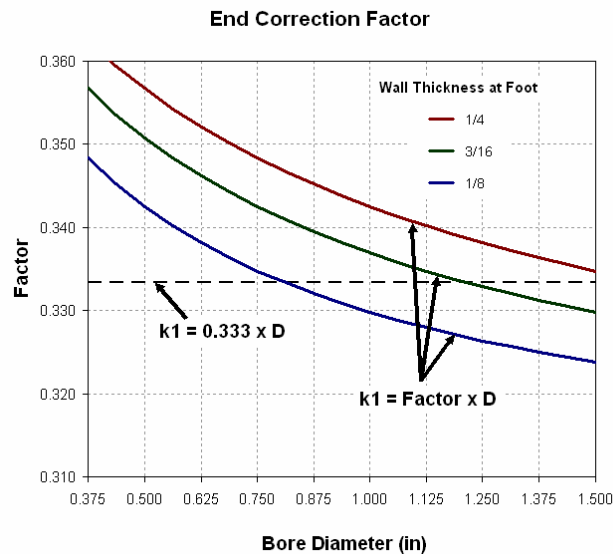
$$V_{cylinder} = \frac{\pi D^2 L}{4}$$




Before continuing, note that this is a general approximation that works fairly well within the precision of crafting a flute. If a more precise estimate of the end correction is desired, the end correction worked out by Benade can be used. That correction is

$$Factor = 0.411 - \frac{0.065}{(0.42 + 2T / D)^{0.54}}$$


where **T** is the thickness of the flute wall at the foot. This factor is plotted below for various bore diameters and wall thicknesses along with the constant factor of 1/3 for reference.



k_2 —the Acoustic End Correction (or “False Length”) at the TSH



On page 17 of “Secrets of the Flute” by Lew Paxton Price, k_2 is described for a round blow hole, or embouchure. An adjustment for the false length above the hole is given when the air column experiences partial containment by the player’s upper lip. This is the same effect as what I describe as “bird encroachment.” In this case, the embouchure is cylindrical, so the development of the idea is fairly straightforward. Later, geometric conversions of the air column contained within the hole are shown which are needed (but not necessarily the whole story...) to account for the rectangular TSH and the expansion of the hole due to the cutting edge ramp.




The new height of the false-length cylinder on top of the blow hole is for this specific case. It includes an adjustment factor that is more easily measured than calculated, but I will go over my own concept that attempts to explain it in a moment.

Now the corrected cylinder needs to be converted to an acoustically-equivalent cylinder with the same cross section as the main bore. That is done with what Lew Paxton Price describes as the "magic ratio," which is given by the cross-sectional area (labeled **S**) of a tube is divided by the length (**L**) of that tube.

$$\text{magic ratio} = \frac{S}{L}$$

This ratio is used to convert a finger hole, for example, to an effective length of tube with the same diameter as the bore and the same magic ratio. The ratio is used as follows:

$$\frac{S_1}{L_1} = \frac{S_2}{L_2}$$



In this case, both the embouchure and bore cross-sectional areas are proportional to their diameters, so the constant factors used to calculate the area cancel, and the ratio becomes

$$\frac{E^2}{H} = \frac{D^2}{k_2}.$$

Solving for k_2 yields the acoustic length of the embouchure.

A yellow cylinder is shown with a vertical curly bracket to its right. To the right of the bracket is the text "total height 'H'". Below this, a dashed rectangular box contains the equation $H = 0.61 E + 0.33 E + T$. Below the dashed box is the equation $k_2 = (D_B / E)^2 H$. Below that is the word "or". At the bottom, a yellow rectangular box contains the equation $k_2 = (D_B / E)^2 ((0.61 + 0.33) E + T)$.


$$H = 0.61 E + 0.33 E + T$$

$$k_2 = (D_B / E)^2 H$$

or


$$k_2 = (D_B / E)^2 ((0.61 + 0.33) E + T)$$

For the NAF, some geometric gymnastics are required, as mentioned before. I will show the pictures, and leave the detailed calculations for the reader (!). Note that I have not tested this (yet), but it is a consistent extension of the approach used by Price, and it shows a conceptual explanation of the bird effects. So let's begin...



The amount of bird encroachment depends on the bird configuration. A bird with no chimney will have a small effect, while a bird with a deep chimney will have a much larger effect and reduce the cross section of the end correction cylinder substantially. I have not worked out a sensible way to quantify the encroachment yet, but will likely end up measuring it. To do that, I will assume I have everything else accounted for, calculate the k_2 with no encroachment, then back the difference up to the point where the cross section of the top end-correction cylinder is adjusted to account for the encroachment. I hope to eventually be able to generate some curves for different bird geometries so that the final k_2 prediction will be more accurate. Once that is done, frequency-dependent effects could be added to the theory. This also extends to the encroachment of the plug on the lower bubble, which seems to be important at higher frequencies. But I digress...

Now, we convert the cylinders to new cylinders with the same cross section as the main bore with diameter D_B .



I mentioned earlier that I would “leave the detailed calculations for the reader.” This isn’t so much a cop-out than it is a practicality. I can do the math—and probably will eventually to see how good the approach is and generate those curves I mentioned above—but such a long series of approximations and outright guesses for specific values will likely be fraught with errors. In addition to that, there are some other effects that cause k_2 to shorten as a function of frequency; those effects will be discussed later. For now, the determination of k_2 for the fundamental pitch is sufficient for determining the acoustic length of the overall flute. For these reasons k_2 is best measured.

To measure k_2 , simply play the fundamental note and figure out the half wavelength—that defines the acoustic length of the flute. Next, calculate k_1 which is simply $D/3$. Finally, subtract the bore length and k_1 from the acoustic length. The answer is k_2 .

The Playing Holes


The playing holes, or finger holes, have the effect of shortening the acoustic length of the flute. When a playing hole is opened, a short cylinder with a small cross-sectional

area is created that provides another path for the vibrating air column to reach the ambient atmosphere around the flute. With the open hole, there are two paths to the outside (or more if additional holes are open). These two paths act in parallel just as two resistors in parallel act in an electric circuit; the total resistance is made smaller than the individual resistances. Acoustically, the effect is to shorten the length of the flute and raise the pitch.

The analogy with the electric circuit holds if the effective length of the finger hole is adjusted so the diameter equals that of the bore. First the corrected length of the playing hole is found in a similar fashion as the embouchure hole above, except there is no encroachment to be considered. Then the magic ratio is once again invoked, and the effective length **F** with an equivalent diameter **D** of the playing hole with diameter **P** and real height **T** is

$$F = \left(\frac{D}{P} \right)^2 \left(\frac{2}{3} P + T \right) .$$

This is depicted in the picture below.



Now we can continue with the analogy of the parallel resistors by replacing the resistances by the acoustic lengths of the sections. The finger hole is open at a distance **L_R** (which includes the end correction at the foot), so the acoustic length **X** of the combined finger hole and bore section at the foot is

$$\frac{1}{X} = \frac{1}{L_R} + \frac{1}{F}$$

or


$$X = \frac{L_R F}{L_R + F} .$$

When X is added to the remaining upper segment on the left (L_L , which includes the acoustic end correction length k_2), a new substitution length L_S is the acoustic length of the air column. The new wavelength is $\lambda = 2L_S$, which corresponds to the new, higher frequency $f = v_s / (2L_S)$, where v_s is the velocity of sound.

Since we need to consider multiple holes, let's call F , X and L_S for the first hole F_1 , X_1 and L_{S1} . To determine the frequency of the next open hole, the process above is repeated by replacing L_B with L_S , L_R with $X_1 + L_{12}$ and F with F_2 . Here, L_{12} is the distance between hole number 1 and hole number 2, and F_2 is the acoustic length of the second playing hole. A new X_2 is found along with a new L_{S2} that defines the next higher pitch. X_2 is

$$X_2 = \frac{(X_1 + L_{12})F_2}{(X_1 + L_{12}) + F_2} .$$

This is depicted below in the figure.



A pattern emerges as

$$X_3 = \frac{(X_2 + L_{23})F_3}{(X_2 + L_{23}) + F_3} ; \dots ; X_6 = \frac{(X_5 + L_{56})F_6}{(X_5 + L_{56}) + F_6} .$$

If a six-hole flute is tuned to the pentatonic modes 1 and 4 using the popular fingering, some adjustments need to be made to the simple pattern above. Namely,


$$X_4 = \frac{(X_2 + L_{24})F_4}{(X_2 + L_{24}) + F_4} \quad \text{and} \quad X_5 = \frac{(X_3 + L_{35})F_5}{(X_3 + L_{35}) + F_5}$$

By adjusting the playing hole parameters (location, diameter and wall thickness), the desired pitch for each hole can be reached.

A Note on Playing Hole End Corrections

The end corrections for the playing holes above were treated assuming they were simple cylinders. In reality, the ends of the mini cylinders of the playing holes are curved because the bore is round, and the diameter of the playing hole is large enough compared to the bore diameter to require an adjustment. Benade worked out this relationship, and the adjustment to the correction can be upwards of 5 to 10 percent higher depending on the size of the hole and bore. The simple calculations above, however, tend to yield values that are quite close to the actual dimensions after tuning the first octave! This may be due to the fact that the calculations above are based on approximations that neglect impedance effects that tend to shorten the corrections. Another factor is that the length of the hole is defined by something less than the thickness of the wall at the hole, although that may not account for the entire difference. The impedance approach will be discussed later, but for the simple case in the first octave, the approximations essentially seem mask those higher-order effects.

When we use the impedance calculations, particularly for the second octave, the more precise playing hole end corrections may be important, and are given here. The inner factor is given by E_i , and the outer factor is E_o .



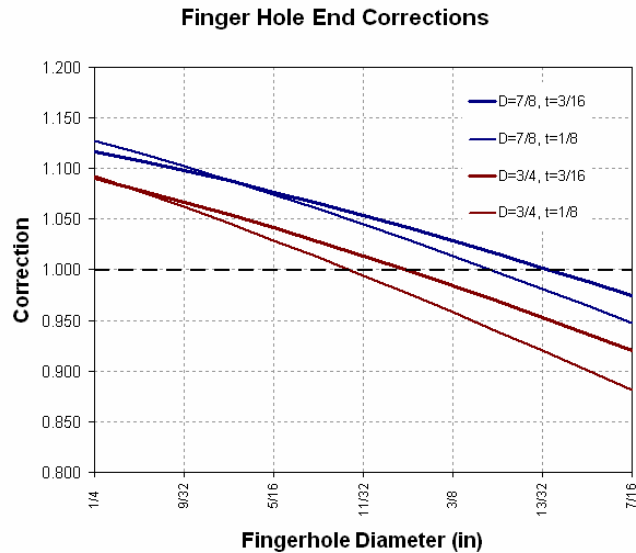
$$E_i = 0.65 - 0.45 \left(\frac{P}{D} \right)$$

$$E_o = 0.32 + 0.103 \ln \left(0.3 \frac{D + 2T}{P} \right)$$

The length of the hole is slightly less than the thickness of the wall, and is given by

$$L_H = \frac{1}{2} \left(\sqrt{(D + 2T)^2 - P^2} - D \right)$$

When the total corrected length is computed and compared to the corrected length from the previous simple calculation, one finds that for typical hole sizes compared to the bore sizes, the corrected length is generally higher by about 5 percent or more using Benade's equations. The results for 3/4- and 7/8-inch bores with either 1/8- or 3/16-inch wall thickness are shown below.



At this point, all the basic calculations for designing a flute have been discussed. These equations are good for estimating the dimensions of a cylindrical flute to play the first octave in tune. Once laid out, the tuning should proceed by the usual method of cutting the holes small and gradually increasing the diameters until the notes are in tune. The second octave, however, presents several challenges. A simple cylindrical flute will not play overblown notes in tune with those of the first register, so some modifications must be made. The following sections go deeper into the various effects that different variables have on the intonation, and provides options for adjustments that will improve the intonation of the upper-register notes.

Acoustic Impedance


Using the Impedance to Describe a Vibrating Air Column

In “Secrets of the Flute,” Lew Paxton Price talks about the *magic ratio* which is given by the cross-sectional area (labeled **S**) of a tube is divided by the length (**L**) of that tube.

$$\text{magic ratio} = \frac{S}{L}$$

This ratio is used to convert a finger hole, for example, to an effective length of tube with the same diameter as the bore and the same magic ratio. The ratio is used as follows:

$$\frac{S_1}{L_1} = \frac{S_2}{L_2}$$



This converted length is the acoustic length of the finger hole, with which calculations of combined lengths of bore sections and finger holes can be performed.

Let's look at an example. If we have a bore with a diameter **D** and a finger hole with a diameter **P** and a thickness **T**, we wish to find a length **F_P** that is the acoustic length of the finger hole (in other words, the equivalent length of the finger hole if the diameter was **D**). First, we can find the cross-sectional area of both the bore and the finger hole as

$$S_D = \pi R^2 = \pi \left(\frac{D}{2} \right)^2$$

and

$$S_P = \pi \left(\frac{P}{2} \right)^2 .$$

For the length of the finger hole with its associated end corrections, consider that it is simply a very short tube that is open at both ends. The end corrections for each end are **P/3**, so the length **L_P** is

$$L_P = 2 \left(\frac{P}{3} \right) + T .$$

The magic ratio has the same value whether it uses the real dimensions or the converted effective dimensions; that's why it is considered "magic." In this example, the equivalent ratios are

$$\frac{S_p}{L_p} = \frac{S_D}{F_p} .$$

F_p is the only thing we don't know in that equation, so we can solve for that to find the answer:

$$F_p = \frac{S_D}{S_p} L_p$$

$$F_p = \frac{\pi(D/2)^2}{\pi(P/2)^2} \left\{ \frac{2}{3} P + T \right\}$$

$$F_p = \left(\frac{D}{P} \right)^2 \left\{ \frac{2}{3} P + T \right\}$$

which is the same equation used for the finger hole acoustic length found on page 36 of "Secrets of the Flute."

We can think about the physical meaning of the magic ratio in terms of the resistance to air flow. The idea can be understood if we think about blowing through a narrow straw, like one of those tiny coffee stirrers, or a cocktail straw. Now think about blowing through a soda straw, which has a much larger cross-sectional area. If both straws are the same length, the soda straw will be much easier to blow because the opening is much bigger. However, what we are trying to do is find a straw that is big in diameter like the soda straw, but takes the same effort to blow through as the cocktail straw. This can be done if we make the soda straw much longer. The reason a long straw is harder to blow through is that we are pushing against the inertia of the entire column of air the length of the straw, and that inertia makes up for the bigger opening. The inertia resists the blowing.

In a like manner, a cocktail straw could be cut smaller so that the effort of blowing through it would be the same as through a normal soda straw. But that raises a question about flutes; why not convert the bore to an equivalent diameter equal to the finger hole instead? This clearly would not work because the lengths would be much too small to represent a vibrating air column that must be a half wavelength of the pitch in the first octave! So what is going on?

By looking closer at the dynamics of the air in the vibrating air column, we find that the magic ratio works well in areas of high flow and low pressure fluctuations, such as at the ends of the air column. (This is called a *pressure node*, or a *flow antinode*.) But as we look at parts of the column that have higher pressure changes and less air motion, the magic ratio becomes increasingly less accurate. In fact, in the center of the

bore where the pressure amplitude change is highest and the flow velocity is lowest, the equivalent acoustic lengths flip! (This area is a *flow node*, or a *pressure antinode*.)

In the example of the straws, consider that we are blowing a puff of air from both ends now, and we're trying to push air into the tube with the same effort. In this case, let's assume the ends are made from the soda straw, and the middle section is the diameter of the cocktail straw and we are trying to find the equivalent soda-straw length for the cocktail-straw-sized section. The cross sectional area is smaller in the middle of the tube now, and when the air blown in from both ends meets in the middle, the local pressure will be higher since we're trying to push the same amount of air through a narrower tube. This higher pressure generates a larger opposition to the flow, which is another way of saying the springiness of the air is increased. This, in turn, results in a higher sound velocity, which means that a sound wave would travel across the thin tube as fast as it would travel across a fat tube that was much *shorter*.

What this is saying is a section of tube near the middle of the bore will have an equivalent length that behaves in an opposite manner to the magic ratio—a long tube with a small cross-sectional area is acoustically equivalent to a short tube with a large cross section *in the region of low flow velocity and high pressure amplitude*.

The magic ratio says that a *short* tube with a small cross-sectional area is acoustically equivalent to a *long* tube with a large cross section; but this is true only *in the region of high flow velocity and low pressure amplitude*. When the two rules are combined, it explains why a half wavelength-long tube (end corrections included) will have the same resonance frequency regardless of the diameter.

You may have noticed that I changed the rules of the game in the second straw example and said we were blowing a puff of air in from both sides instead of just trying to blow through the tube. That's because we have to be talking about a vibrating air column in order for the idea to make sense and be applicable to the flute. Recall that the vibrating column is like a spring. The spring stores energy as it is compressed, releases the energy as it expands back out, stores more energy as it is stretched, and releases the energy again as it contracts. The spring reacts to being pushed or pulled by pushing or pulling back.

At the ends of the tube, the situation is different because the ends of the air column are not constrained. In this case the ends can be thought of as slugs of mass that are pushed in and out of the tube. They do not store energy and give it back—they use it up. Instead of pushing or pulling back like springs, they are just inertial masses that resist being pushed or pulled.

In the above two paragraphs, I have chosen two words carefully: the springs *react* and the inertial masses *resist*. These are key concepts in the description of electrical circuits in terms of *impedance*. Electrical impedance is a measure of the opposition to a sinusoidal or alternating-current (AC) signal, and has a reactive component as well as a resistive part. The reactive components in an AC circuit are the capacitors and

inductors, and they are electrical energy storage devices just as our spring is a mechanical energy storage device. The electrical impedance is defined as the AC voltage divided by the AC current. If the voltage and current are in phase (their peaks and valleys line up perfectly), then the impedance is purely resistive and the power in the circuit dissipates. If a capacitor is added to the circuit, the change in voltage will be slowed or delayed compared to the current. If the voltage and current are out of phase (the peaks and valleys of one lag behind those of the other), then part of the energy is alternately stored and released, and the impedance has a reactive component.

In an acoustic “circuit,” the sound pressure is analogous to voltage, and the flow velocity is analogous to current. For example, in the contracting phase of the column, as the pressure in the middle of the bore builds, the air particles continue to flow toward the middle until the force of the stored energy in the spring can overcome the momentum of the air particles. At that point in time, the pressure is at the maximum and the flow is zero. If the pressure is described with a sine curve, then the flow would have to be described by a similar curve, except that it would have to be 90 degrees out of phase to be zero when the pressure was at a peak. In other words, the flow would have to be described by a cosine curve.

The acoustic impedance is defined as the sound pressure divided by the flow velocity, or

$$Z = \frac{p}{U}$$

where p is the acoustic pressure wave in a cylindrical tube, which is the amount the pressure differs from the surrounding atmospheric pressure. The acoustic pressure is given by

$$p = \hat{p} \sin(kx + \psi) e^{j\omega t}$$

and U is the volume flow velocity given by

$$U = \frac{-\hat{p}S}{j\rho c} \cos(kx + \psi) e^{j\omega t}$$

These equations are “complex” (i.e., mathematically they have “real” and “imaginary” parts), but for almost everything on which we will use these equations, the complex math can be ignored. For completeness here, I will mention that j is $\sqrt{-1}$, \hat{p} is a complex constant that will soon be cancelled out, c is the velocity of sound, and $e^{j\omega t}$ is a phase term that causes the pressure and flow to chase each other around in complex space, but that will cancel out as well. The other terms will be defined in a moment. The impedance, then, is

$$Z = \frac{-jpc}{S} \tan(kx + \psi)$$

where S is the cross-sectional area of the tube, x is the distance along the acoustic length of the vibrating column, and ψ is a phase term defined by the conditions at the ends of the section of tube under consideration. The wave number k is given by

$$k = \frac{\omega}{c} = \frac{m2\pi}{\lambda}$$

where λ is the wavelength of the sounded frequency that is related to the angular frequency ω . The value of π is 3.14159..., and m is the mode number such that $m=1$ for the fundamental mode and $m=2$ for the first overtone (over-blown notes).

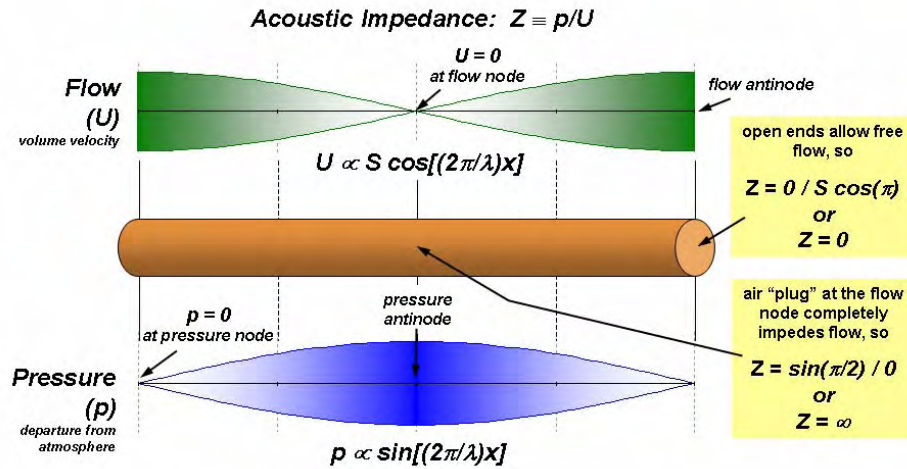
For a simple pipe open at both ends, the pressure at the ends is fixed at atmospheric pressure, so the sound pressure is fixed at zero, which means the impedance is also zero. If $x=0$ at one end, then $\tan(\psi) \equiv 0$, or $\psi = 0$. The other end is at $x=L$, so

$$\tan(kL) \equiv 0$$


$$\Rightarrow k = 0, \frac{\pi}{L}, \frac{2\pi}{L}, \dots$$

The useful solutions are for $k = m\pi/L$, with $L = \lambda/2$. So the impedance of a simple cylindrical tube is

$$Z = \frac{-jpc}{S} \tan\left(m\pi \frac{x}{L}\right).$$



If the impedance for the first mode is plotted (ignoring the constant terms), the value goes from zero at the ends to infinitely positive or negative in the center of the air column. (Note that the negative values for impedance just mean that the spring is pushing in the opposite direction than in the positive region). The infinite impedance at the center of the column means that no air actually flows past this point, which is what is actually happening in the flute. Another way to say it is the *admittance* is zero in the middle of the column, and it is infinite at the ends. Admittance is the inverse of impedance, and will be discussed in the next section.



Notice that near the ends of the vibrating column, the tangent curve is fairly straight for almost 20 % of the length of the tube. In this area, the impedance can be approximated by a straight line given by $Z = kx/S$ when $\psi = 0$. Let's say we want to replace 10 percent of the end of the tube with a narrow tube, and we want to find out how long the tube would have to be. An example of this might be a choke at the foot of the flute to shorten the flute while keeping the same fundamental frequency. A choke is basically a plug with a hole drilled in it, and the hole is smaller than the diameter of the bore.

EXAMPLE: Let's set up the problem. Let's say the half wavelength-long air column is 16 inches, and the 10-percent mark is 1.6 inches, which we will label X_1 . The diameter of the bore is 0.75 inches with a cross-sectional area we'll label S_1 , and we wish to make a choke with a 0.375 inch-diameter hole with a cross-sectional area of S_2 . The value of the impedance at X_1 in this linear part of the column is

$$Z(X_1) = \frac{kX_1}{S_1}$$

where k is fixed by the wavelength of the fundamental note, or $k = 2\pi/\lambda$. We now wish to know how long the choke needs to be, which is X_2 . In order for this to work, the impedance at X_2 for the narrow tube must be equal to the impedance at X_1 for the wide tube. So,

$$Z(X_2) = \frac{kX_2}{S_2} = \frac{kX_1}{S_1} = Z(X_1)$$

If we simplify and invert this, we get

$$\frac{S_2}{X_2} = \frac{S_1}{X_1}.$$


You may have noticed that this is nothing more than the magic ratio! Now we can solve for X_2 to get

$$X_2 = \frac{S_2}{S_1} X_1 = \left(\frac{D_2}{D_1}\right)^2 X_1 = \left(\frac{0.375}{0.75}\right)^2 1.6 = 0.4.$$

If we were to use the impedance functions, we would have

$$X_2 = \frac{1}{k} \arctan \left[\frac{S_2}{S_1} \tan(kX_1) \right] = 0.413.$$

This is the exact solution, which is about 3 percent higher than the solution given by the magic ratio. If we wanted to replace a quarter of the air column, or 4 inches, then the error would be about 25 percent, and if half the column were replaced, the error would be 400 percent! A graphical representation of the 4-inch case is shown below to illustrate the process for finding the solution with the impedance function. Both tube diameter sizes are represented in the graph on the left with their respective impedance curves. At the 4-inch mark (X_1) the value of the curve for the large diameter tube is marked with a blue cross. The same value for the small diameter tube is marked with a red cross. That point lies at the location X_2 , and is the length of the narrow tube required to replace the 4-inch section of the larger diameter tube. The final choked tube is shown in the second picture along with the impedance curves shifted so the interface between the two tubes is in the same place. Note that the impedance curve for the new tube is continuous, and is a combination of the two curves.




One point to be made here is that simplifying approximations are very useful if they are used where intended, but can lead to erroneous results if misused. In this example we can recall the earlier comment about the magic ratio being good in the region of high flow velocity and low pressure amplitude, which characterizes a pressure node. The end of the tube is such a place, but as we go further into the bore, the character of the vibrating air column changes from that of a pressure node to that of a flow node where there is low flow velocity (and high pressure amplitude). The transition from pressure node to flow node is gradual and continuous, so while there are regions in which certain approximations can be used, the impedance function can be used along the entire column.

Finger Holes and Shortening the Acoustic Air Column

The acoustic admittance \mathbf{Y} is another very useful concept, and it is simply the inverse of the impedance, or,

$$Y = \frac{1}{Z} = \frac{-S}{jpc} \cot(kx + \psi).$$

As with electronic circuits, the admittance is useful in evaluating parallel circuits. This happens acoustically in the flute when a finger hole is opened and the vibrating column branches out through the hole. From a point below the finger hole, both the branch section and the tube between the hole and the foot start with the same impedance value and both end outside the flute with an impedance of zero. If we use admittance instead of impedance, these two branches can be simply added together as with two parallel circuit elements.



Consider the case where a finger hole at location \mathbf{X}_H is opened in a tube with a total length \mathbf{L}_T . Opening the hole will cause the flute to play a higher frequency that has an acoustic length that is less than \mathbf{L}_T . The new length can be represented by a substitution tube of length \mathbf{L}_S and the new wave number is defined by $k = m\pi/L_s$.

The finger hole has a length $\mathbf{L_H}$ and a cross-sectional area $\mathbf{S_H}$, whereas the cylindrical tube sections all have a cross-sectional area of \mathbf{S} . Note that all the lengths described here include adjustments for end corrections.

There are several ways to analyze this problem, and they essentially amount to choosing a view of the problem and carefully keeping track of the variables and signs. I choose to focus on the elements to the right of the point \mathbf{X}_H . We know that \mathbf{k} is fixed, even if we don't yet know its value. We also know that the open ends of the real tube, the finger hole and the virtual substitution tube all have an admittance of infinity (impedance of zero) because they are in contact with the outside atmosphere, and are therefore at a pressure node. We also know that the admittance looking to the right of \mathbf{X}_H for the combined finger hole and right-side tube section is the same as the admittance of the virtual tube section, or

$$Y_H + Y_R = Y_V$$

where Y_V is the virtual section of length $L_S - L_L$. The admittance at the open end of the substitution tube on the right at the pressure node is infinite, so we can choose $x=0$ there. By doing so, the boundary condition forces $\psi = 0$, and the admittance is

$$Y_V(X_H) = -\frac{S}{jpc} \cot(k[L_S - L_L]) .$$

Acoustically, we can argue that the ends of the finger hole and right-side tube sections are open at $x=0$ as well, so in like manner,

$$Y_H(X_H) = -\frac{S_H}{jpc} \cot(kL_H)$$

$$Y_R(X_H) = -\frac{S}{jpc} \cot(kL_R)$$

Putting these together, we find

$$-S_H \cot(kL_H) - S \cot(kL_R) = -S \cot(k[L_S - L_L]) ,$$

If we rearrange this a little, we get

$$S \cot(k[L_L - L_S]) + S_H \cot(kL_H) + S \cot(kL_R) = 0 .$$

If we know all the dimensions and wish to find the pitch, we can recognize that this is a function of k and use a computer program to solve for the zeros. This can be done graphically by plotting the function using a range of values of $k = m\pi/L_S$ that are sure to contain the answer, and choosing the k value where the function crosses the x axis. This is particularly interesting if all the variables are known in the first octave and we wish to compute the pitch of the second octave note. This will be discussed further in a later section.

If we know what we want k to be to yield the desired pitch, and wish to find the finger hole placement L_R , we can rewrite the equation as

$$S \cot(k[(L_T - L_R) - L_S]) + S_H \cot(kL_H) + S \cot(kL_R) = 0$$

and solve for L_R .

EXAMPLE: Let's say we have an F# flute and want to find the distance from the foot to place the first hole, which will play an A. We can figure out the acoustic length (half wavelength) for the F# by

$$f = 370 \text{ Hz}$$

$$Temp = 72^\circ \text{ F}$$

$$v_{sound} = 12600.5 \times \sqrt{\frac{459.4 + Temp}{459.4}} \text{ inches/sec}$$

$$L_T = \frac{1}{2} \times \frac{v_{sound}}{f} = 18.31 \text{ inches}$$

By opening the hole, we want to shorten the acoustic length to that of an A. Using the same approach above, this turns out to be 15.4 inches. We'll call this the substitution length, L_S . Now let's find the length of the finger hole. If the desired hole diameter is $P = 5/16$ inches and the wall thickness is $T = 3/16$ inches, the hole length with end corrections will be

$$L_H = 2 \frac{P}{3} + T = 0.397 \text{ inches}$$

using the simple end corrections. If Benade's corrections are used, namely

$$E_i = 0.65 - 0.45 \left(\frac{P}{D} \right)$$

$$E_o = 0.32 + 0.103 \ln \left(0.3 \frac{D + 2T}{P} \right)$$

$$T_H = \frac{1}{2} \left(\sqrt{(D + 2T)^2 - P^2} - D \right)$$

$$L_H = (E_o + E_i)P + T_H$$

then $L_H = 0.426$ inches.

The cross-sectional areas of the finger hole and bore are S_H and S , respectively. But since the area appears in all terms, we can drop the common terms that will cancel anyway. In other words, $S \propto D^2$ (S is proportional to the diameter squared), so our equation becomes

$$F(L_R) = D^2 \cot(k[(L_T - L_R) - L_S]) + P^2 \cot(kL_H) + D^2 \cot(kL_R) \equiv 0.$$

We know all the terms in the equation except $\mathbf{L_R}$, which is the variable we are trying to find, so we can say that the equation is a “function of $\mathbf{L_R}$ ” and give it the label $F(L_R)$. The correct value of $\mathbf{L_R}$ is the one that makes $F(L_R) = 0$. One way to solve this is to first recognize that

$$\begin{aligned} \cot(k[(L_T - L_R) - L_S]) &= \cot(kL_T - kL_R - kL_S) \\ &= \cot(kL_T - kL_R - m\pi) \\ &= \cot(kL_T - kL_R) \end{aligned}$$

Since the function is periodic, a shift of $m\pi$ will give an equal value, so we can remove that variable from the equation. Then we can use a trigonometric identity that can be found in a math book that says

$$\cot(kL_T - kL_R) = \frac{\cot(kL_T)\cot(kL_R) + 1}{\cot(kL_R) - \cot(kL_T)}$$

Taking that on faith (or deriving it for the fun...) we can now write the equation as

$$D^2 \frac{\cot(kL_T)\cot(kL_R) + 1}{\cot(kL_R) - \cot(kL_T)} + P^2 \cot(kL_H) + D^2 \cot(kL_R) \equiv 0$$

This can be rearranged using straight algebra to get

$$(\cot(kL_R))^2 + \left[\left(\frac{P}{D} \right)^2 \cot(kL_H) \right] \cot(kL_R) + \left[1 - \left(\frac{P}{D} \right)^2 \cot(kL_H) \cot(kL_T) \right] \equiv 0$$

This is a quadratic equation that can be solved for $\cot(kL_R)$ using the famous “quadratic formula” to get


$$\cot(kL_R) = \frac{1}{2} \left\{ - \left(\frac{P}{D} \right)^2 \cot(kL_H) \pm \sqrt{\left[\left(\frac{P}{D} \right)^2 \cot(kL_H) \right]^2 - 4 \left[1 - \left(\frac{P}{D} \right)^2 \cot(kL_H) \cot(kL_T) \right]} \right\}$$

Using a bore diameter of 7/8 inches and the other values defined above, we use the positive solution (the one that makes sense...) that turns out to be $\cot(kL_R) = 0.572$. Then we can solve for L_R using

$$L_R = \frac{1}{k} \cot^{-1}(0.572) = 5.15 .$$

This is the hole location from the end of the acoustic column of air. To find the location from the foot of the bore, we subtract the end correction, $D/3$, to get $L_R = 4.86$ inches.

An alternate solution is to plot the function for a range of L_R values, and the solution is where the function crosses the x axis.



Note that the solution using the approximation of $\tan(kL) \cong kL$ is shown, and is lower in placement by about 1/3 inch, and almost 1/4 inch if a simple finger hole correction is used. If the approximate solution were used and the finger hole ended up slightly larger than predicted, or shallower or undercut when tuned, this could be a reasonable explanation. The alternative that the “standard” model gave good results may indicate that there are other effects at play that we have not yet addressed for which the approximate solution serendipitously compensated.


Subsequent finger hole placements can be found by replacing L_S by L_T , defining a new L_S for the next note, and repeating the process until all the finger holes are

complete. Note that the end correction will not have to be subtracted from the subsequent results as long as the locations of the previous L_S lengths are carefully tracked.

Before proceeding further, we should look at the effects of the impedance at the other end of the flute. Since the sound hole behaves more like a finger hole than an open end of the bore, the impedance and associated end correction will have a dependence on frequency. Furthermore, the extent of a backset between the plug and the sound hole will induce additional frequency-dependent corrections. Since the end corrections at the sound hole are dependent on frequency, the acoustic tube length L_L to the left of X_H will be slightly different for each finger hole.

End Corrections at the Sound Hole

The shape of the tangent curve should also affect the length of the end correction at the sound hole. As stated earlier, the length of this correction, k_2 , is most easily determined by measuring fundamental pitch and translating it into a total acoustic length, then subtract the bore length and the foot end correction. But at higher pitches, the sound hole interacts with a larger fraction of the air column, and that means a larger part of the tangent-shaped impedance curve is in play. If the measured k_2 is taken as a baseline, the relative shift of higher pitches can be estimated. Such a shift is plotted below. The two curves represent notional fundamental k_2 values that would depend on the configurations of the birds, and perhaps the proximity of the plug to the sound hole. (note: I have yet to fully test this concept with further experiments...)



Mathematically, this is based on the equation

$$\frac{1}{S_{bore}} \tan(kk_2) = \frac{1}{S_{TSH}} \tan(kT_{TSH})$$

where T_{TSH} is the thickness of the flute at the sound hole, and D_{TSH} is the average equivalent diameter of the sound hole. Solving for k_2 yields

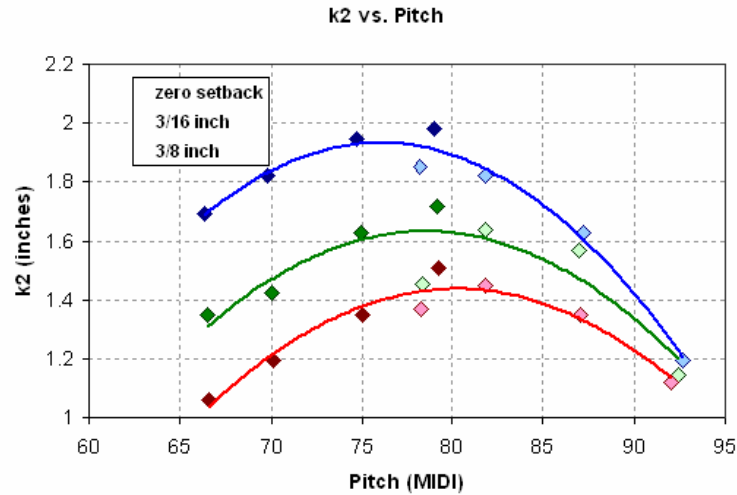
$$k_2 = \frac{1}{k} \arctan \left[\frac{D_{bore}^2}{D_{TSH}^2} \tan(kT_{TSH}) \right] .$$

Recall that k is the wave number

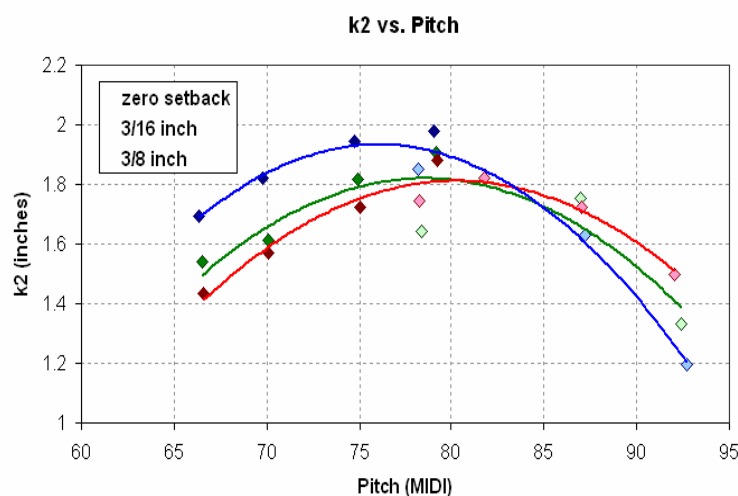
$$k = \frac{2\pi n}{\lambda} ,$$

so as the pitch goes up, the wavelength gets shorter, and the wave number k gets larger. After crunching the numbers, one finds that since D_{TSH} , D_{bore} and T_{TSH} are constants, the value of k_2 drops as the pitch goes up. Note that if the *magic ratio* was used, k_2 would not change with pitch.

To test this, I performed an experiment in which k_2 was found in a flute with no finger holes, but of various lengths. This allowed the effect of opening finger holes to sound higher frequencies, but the physical length of the bore was known, and the value of k_1 was assumed to be a constant $D/3$. The following shows the result of the experiment.



In this case, the bore was measured from the plug, which was placed at three different locations relative to the northern edge of the TSH. The solid points represent the fundamental pitch of the tube for each length, and the lightly filled points are the overblown note. There is a distinct bow in the curves that shows that k_2 gets *longer* at first, then gets shorter after peaking at some pitch that grows as the backset increases. The graph gets more interesting when the bore is measured from the back of the TSH, as shown below. Using three different backset values and using the distance to the TSH (standard practice in classical flute theory) yielded some curious insights into what may be happening.



In this case, the curves cross in a region of the graph. It turns out that the crossover occurs approximately at the point where k_2 is $1/8$ the acoustic length of the pitch. This is where the air column is half way between a pressure node and a flow node. Later when bore perturbations are discussed, we will show that an increase in the bore diameter in the vicinity of a pressure node will have a sharpening effect, and near a flow node it will have a flattening effect. In this case, if the backset is considered to act as a bore enlargement just inside the TSH, then for longer bore lengths (lower pitch) the enlargement should have a sharpening effect proportional to the depth of the backset since the perturbation is nearer the pressure node outside the TSH than it is to the pressure antinode, or flow node, inside the bore.

In the bore perturbation approach, enlargements in high-flow areas of the vibrating air column (near a pressure node with little pressure variation) result in a sharpening of the pitch. Likewise, enlargements in lower-flow regions (high pressure fluctuations) result in a flattening of the pitch. The region half way between the pressure node just outside the TSH (at the end of the column section described by k_2) and the first flow node (pressure antinode) inside the bore has equal flattening and sharpening influences, so there is no change in pitch for an enlargement in this region.

Note that for even the zero-backset case, an enlargement perturbation still exists because of the geometry of the flat plug. A concave plug shape may exist that would result in no perturbation, but that is yet to be determined. So, for longer physical bores, the enlargement perturbation will occur proportionally closer to the pressure node at the end of the acoustic air column, and a sharpening of the pitch will be more pronounced in longer bores. The sharpening is due to the effective shortening of the acoustic length of the bore when the perturbation is in a high-flow region. This means the physical dimension of the bore is longer than the acoustic dimension. When the k_2 is calculated as

$$k_2 = L_A - (L_B + k_1)$$

the measured L_B is really too long, which means the calculated value of k_2 will be too small. A more precise equation might be

$$k_2 = L_A - (L_B + \Delta_p + k_1)$$

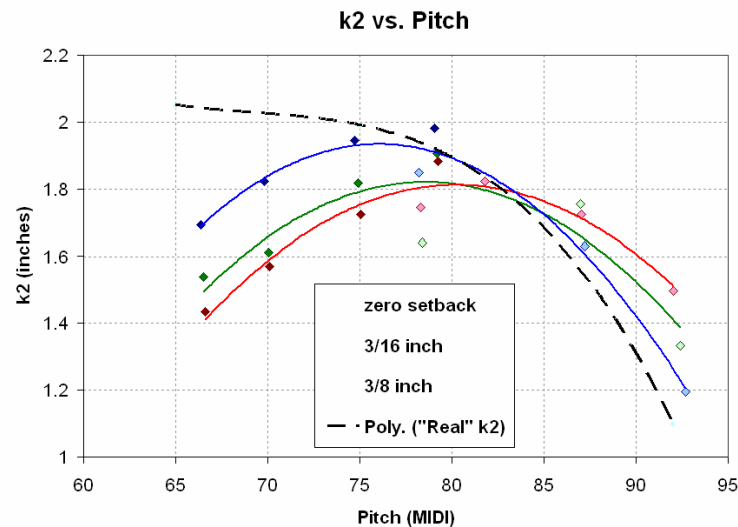
where Δ_p is the change in acoustic length of the physical bore due to the perturbation.

In the case of the perturbation causing a sharpening of the pitch, the value of Δ_p would be negative because the physical bore is effectively shortened. Likewise, for the higher pitches of the shorter physical bore lengths or overblown notes, the perturbation region resides closer to the first internal flow node, and the sharpening effect will diminish as the perturbation gets closer to the midpoint between the pressure and flow nodes. For higher pitches for which the perturbation resides beyond that “nodal midpoint” on the side closest to the flow node, the effect will be to flatten the pitch, or to increase the effective acoustic bore length. In that case, Δ_p will be positive, which means the actual value of k_2 will be smaller than what was observed in the graph. Again, the effect will be enhanced by a deeper backset that will produce a larger perturbation in the bore.

This implies that a reversal in the observed (or incorrectly-measured) change in k_2 length will occur when the perturbation crosses the nodal midpoint. Such a reversal is observed if the experimental k_2 vs. frequency plot is derived by using the distance to the TSH instead of the plug is used for the measured bore length, as shown in the second chart above with the bore length measured to the TSH. If the measuring point is chosen carefully (north edge, center, or some other point along the TSH), the zone where the lines cross will occur at the pitch for which the backset perturbation occurs close to the nodal midpoint. In other words, the k_2 at the crossing zone will represent the true value of k_2 —at least for a particular frequency.

If the true value of k_2 was constant (i.e., not a function of frequency), then the perturbation effects would be nearly symmetric (odd symmetry) about the nodal

midpoint value. If it was a function of frequency as described by an impedance calculation, the symmetry would have a downward-sloping bias to it. So, the experimental curves appear to be described by the added effects of the impedance (shorter k_2 for higher frequencies) and the backset perturbation (shorter for lower frequencies and higher for higher frequencies). This is shown below where the second graph has been redrawn with a notional line (dashed heavy black line) representing what the true value of k_2 might be for this case.

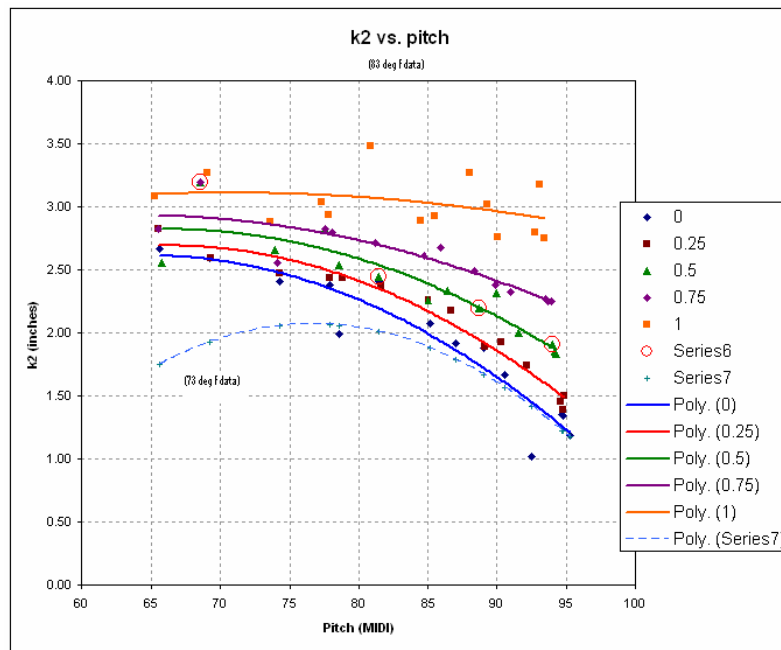


So why do the octave notes usually play flat in a cylindrical flute? From the first graph in this section, we can see that k_2 should get shorter, but the proposed concept of a perturbation effect near the plug can account for the apparent lengthening of k_2 if the perturbation (or some other effect that explains it) is neglected. With finger holes involved, the situation is even more complex, and some of the flatness occurring in the second octave can be attributed to the impedance associated with the open playing holes. Another contributor may be some nonlinear effects occurring from the interaction of the air jet on the cutting edge. These are ideas that I hope to develop later...

...Ok, it's a little later, but I still have more to do.

The following is the experiment above repeated using the flute as a filter for a white noise source. The measurements are the frequencies for which the impedance is a minimum, which includes the higher frequencies in the spectrum for each length and backset. The dashed line is taken from the previous experiment (note that the temperatures for the two experiments were different by about 10 degrees F).

The trend suggests that the impedance approach at the beginning of this section has merit, but also suggests that there is an impedance effect induced by the air jet. Fletcher & Rossing have an explanation of such a phenomenon that may explain the difference between the passive (white noise) and active (sound generation using the air jet). If that explanation works out, it means that the air jet induces a phase shift in the impedance that is more pronounced in the high-flow part of the vibrating air column, and effectively makes the flute acoustically shorter than what the passive theory would predict. This sort of makes sense because such a phase shift usually is manifested in the imaginary part of a complex function. For such a “gain” term, the shift would be positive, whereas a negative shift would be the result of a “loss” mechanism (rough walls, etc) that may make the flute appear longer (rough walls slow the propagation). Keep in mind that I have not yet thought this “explanation” through, so I could have a couple things backward or otherwise off. But the data shows what has been observed, which is what it is.



...more to come later...from here to the end of the section is essentially a draft outline with a few partially-developed concepts

Backset Effects

The volume of air between the sound hole and the plug has a significant effect on the relative tuning between the registers. The perturbation treatment mentioned above may be valid, but it may also only be a partial explanation of the plug position's influence on the frequency-dependent tuning.

Outline:

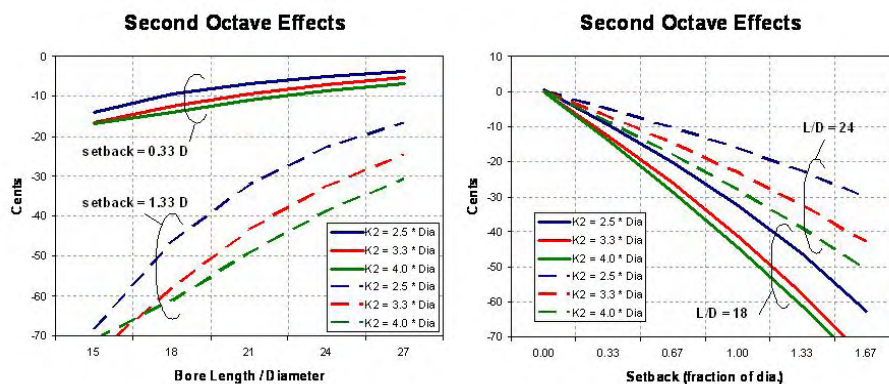
- Qualitative description
- Helmholtz resonator description
- Admittance model from Nederveen

$$\frac{1}{\tan(kL_E)} + \frac{1}{\tan(kL_B)} - \tan(kL_S) = 0$$

$$k = \frac{2m\pi}{\lambda} ; m = 1, 2$$


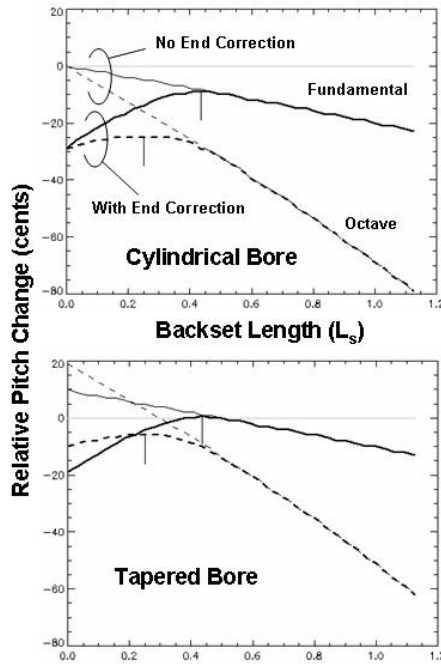
- Model results & discussion

The effect of different bore L/D ratios and TSH designs (as described by K2 which gets larger as the TSH). Note that the model is for a perfect cylindrical flute, and it says that the overblown fundamental is in tune in all cases if there is no setback. I have yet to look at what happens when finger holes are in the mix. The distances in the graphs are all referenced to the bore diameter since the ratios are what matters and are valid for any size flute.



The model shows that the setback has a dramatic flattening effect in the second octave when combined with fatter flutes (i.e., lower L/D ratios). The K2 also has a significant effect when a TSH / Bird design results in a large K2 (i.e., narrow TSH or a deep chimney in the bird) and will amplify the setback's flattening effect.

So, a narrow bore, large TSH and no chimney in the bird are design elements that should help mitigate a flat second octave. Minimal setback should help, but there seems to be at least one other yet-to- be-determined variable that flattens the second octave since many zero-setback flutes are still flat in the second octave. This could be due to a plug-encroachment effect that works in a similar fashion to that of a chimney or overhang added to a bird at the outer side of the TSH. The end correction in the figure below refers to the contribution from the plug's proximity to the sound hole. The form of the encroachment effect is notional at this point.



$$\frac{1}{\tan(kL_E)} + \frac{1}{\tan(kL_B)} - \tan(kL_S) = 0$$

$$k = \frac{2m\pi}{\lambda} ; m = 1, 2$$

L_S is the physical backset length from plug to the edge of the sound hole. L_E is the end correction at the sound hole, and is a function of L_S while the plug is very close to the edge of the sound hole. The zeros of the equation are computed for the various backset lengths.

Bore Shapes and Perturbations

A Segmented Bore

Recall that the acoustic length of a column of air in the first octave is $L_A = \lambda/2$, where the wavelength of the sound is λ . The wave number is $k = m2\pi/\lambda$ where m is the mode number if higher registers are involved, and rearranging this gives $kL_A = m\pi$. This is a result of the boundary conditions placed on the standing wave in the bore containing the vibrating air column. The impedance of the bore is given by

$$Z = \frac{-j\rho c}{S} \tan(kx + \psi)$$

where S is the cross-sectional area of the bore, which is a function of the location x along the column, and ψ is the shift in the phase of the tangent function. In the following, the phase shift is used to move the tangent function associated with the impedance of each section of the bore along the x direction so that the tangent curves line up at the interface between sections. There will be more discussion on that shortly. Finally, $-j\rho c$ includes the air density and speed of sound, which was described earlier. It will be neglected in the following section because it is essentially constant throughout the air column and algebraically cancels in the equations.

For the column of air, the impedance at the ends has to equal to zero because the ends are connected to the air outside the bore. This means that $kx + \psi = 0$, which implies that $\psi = 0$ at $x = 0$, and $\psi = -kx = -m\pi$ at the end of the column where $x = L_A$. This is the “boundary condition” that makes $kL_A = m\pi$ when the air column is vibrating in resonance with the bore. Note that the physical length of the air column, including end corrections, is not necessarily the same as the acoustic length. That will become apparent later. Another note about the air column is the wave number k is the same everywhere in the bore, no matter what the shape is; k is defined by the wavelength of the standing wave, or the pitch of the note. This also means that the *period length* of the tangent function describing the impedance of any given section of the bore will be the same as for every other part of the bore. The tangent function will vary in magnitude because the cross-sectional area S of that section of the bore defines the factor that scales the tangent function up or down. Another thing that can change about the tangent function is the phase—whether the function is shifted to the right or left—which is further explained below.


If the bore is divided into sections, additional boundary conditions will exist at the interfaces between the sections. Physically, we know that the impedance between one

section and the next must be equal for both sections *at that point*. The boundary condition at the interface is not simply zero, so we first have to calculate what the value of the impedance (the tangent function scaled by $1/S$) should be at that point. This is given by

$$Z_1(x) = Z_2(x), \text{ or}$$

$$\frac{1}{S_1} \tan(kx + \psi_1) = \frac{1}{S_2} \tan(kx + \psi_2)$$

A bore with an arbitrary shape can be approximated by a series of cylindrical sections that are mathematically easy to handle. By making this approximation, each section can be solved separately as long as we keep track of how each previous section defines the boundary condition between it and the new section. As the bore is divided into more sections the accuracy of the calculation will get better, and small steps in cross-sectional area between sections should be used. A good way to see if the sections are fine enough is to do the calculations for increasingly finer steps until the answer essentially doesn't change any more.



Here's the method: First we make an educated guess about what the wave number k should be. The value of k will probably be changed a few times until the answer converges to the correct solution, but that will be addressed later. A good starting point is $k = m\pi/L_B$ where L_B is the bore length (including end corrections). We then calculate the location on the impedance curve (a.k.a. the *angle*, or the *angular distance*) that coincides with the end of the first section; this is simply kL_1 where L_1 is the length of the first section (including the end correction...). Plugging this in to the impedance equation defines the first section of the overall impedance curve for the bore, and it gives us the impedance value at that first interface. Now we make sure the *value* of the impedance curve for the start of the second section is the same. Note that

the angular distance of the beginning of the second section is the same as that of the end of the first section. Since the second section may have a tangent function of a different magnitude, the curve may not line up properly, so the function needs to be moved right or left (a phase shift) to line up the curves at that point. The phase shift for that second section is ψ_2 , and we now need to find its value.

For the very first section, we know that the boundary condition at the opening forces $\psi_1 = 0$, so at $x=L_1$ we have

$$\frac{1}{S_1} \tan(kL_1) = \frac{1}{S_2} \tan(kL_1 + \psi_2)$$

Solving for ψ_2 we get

$$\psi_2 = \arctan\left[\frac{S_2}{S_1} \tan(kL_1)\right] - kL_1$$

The whole impedance curve describing the second section will now be shifted left or right so that the starting point is no longer at $x = 0$. The difference between the shifted starting point and zero is the *phase shift*. Then the section of the curve from kL_1 to kL_2 that corresponds to the second bore section is added to the first segment of the impedance curve. Drawings of this process are presented below to help visualize these steps.

For the next section, we repeat the steps above. To find the impedance value at the next interface, find the angular distance defined by the length of the section (call it L_2), and that angle is kL_2 , add that to the total argument in the tangent function for the current bore section, and find the new impedance value. That is the value that must be matched at the start of the third section.

$$\frac{1}{S_2} \tan(kL_1 + kL_2 + \psi_2) = \frac{1}{S_3} \tan(kL_1 + kL_2 + \psi_3)$$

The argument in the tangent functions will continue to grow as sections are added. This could get messy, so a convenient approach is to lump all the section lengths into a single variable. Let's introduce it as a new substitute length X so that

$$kL_1 + kL_2 + \psi_2 = kX_2 + \psi_2$$

$$kL_1 + kL_2 + \psi_3 = kX_2 + \psi_3$$

then

$$\frac{1}{S_2} \tan(kX_2 + \psi_2) = \frac{1}{S_3} \tan(kX_2 + \psi_3)$$

Now, solving for ψ_3 gives us

$$\psi_3 = \arctan\left[\frac{S_3}{S_2} \tan(kX_2 + \psi_2)\right] - kX_2 .$$


Note that a pattern has emerged. For each section, which we can label with “ i ” the phase ψ_i is

$$\psi_i = \arctan\left[\frac{S_i}{S_{i-1}} \tan(kX_{i-1} + \psi_{i-1})\right] - kX_{i-1}$$

where

$$\begin{aligned} X_{i-1} &= L_1 + L_2 + L_3 + \cdots + L_{i-2} + L_{i-1} \\ &= \sum_{j=1}^{i-1} L_j \end{aligned}$$

This process continues until the end of the bore is reached. The drawing below shows this process in which a new section of bore is getting added to the impedance curve. Call this section B, which lies on the blue dashed curve. The previous section (A) lies on the red dashed curve, which already has a phase shift of ψ_A relative to the starting point of the bore. In the left picture, the B curve is plotted before a phase shift is applied, so it has the same phase as the A curve. Plotting the impedance of the new section on this curve leaves an impedance mismatch at the interface. This is because the height of the curve depends on the bore cross section, and the two sections are different in this case. Since the impedance must match, the B curve is moved to the left until the impedance curves line up. The amount the curve must be moved is the phase shift ($\Delta\psi$) relative to the previous section, and the total phase shift for the B curve is ψ_B .



As we go from section to section along the bore, that phase shift will move the impedance curves for the bore sections back and forth, but by the time we get to the end, the shift must be such that the impedance is zero. If the impedance isn't zero, then the wave number k needs to be adjusted and the process started over. This is an iterative process that is best done on a computer!

At the end of the air column, the length is L_B (the bore length, including the end corrections). The boundary condition requires that the impedance equal zero at the end, but the phase shift is not necessarily zero. Instead, the total of the argument of the tangent function must be a multiple of π , or

$$kL_B + \psi_n = m\pi$$

where ψ_n is the final phase shift. Recall that $k = m\pi/L_A$, so

$$m\pi \frac{L_B}{L_A} + \psi_n = m\pi$$

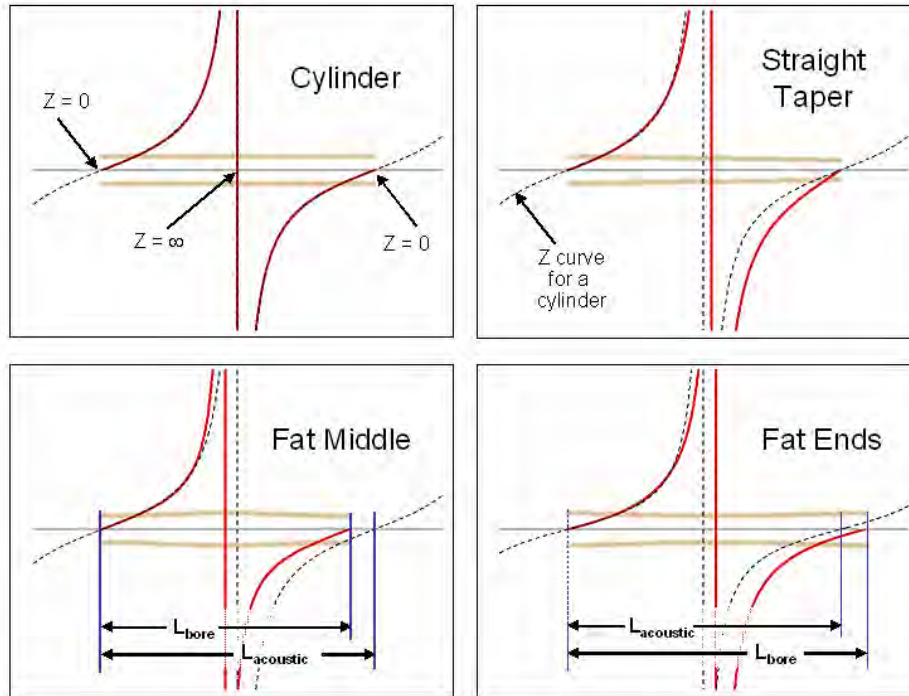
or

$$L_B = L_A \left(1 - \frac{\psi_n}{m\pi} \right) \quad \text{or} \quad L_A = L_B \left(\frac{m\pi}{m\pi - \psi_n} \right).$$

In either case, ψ_n is small, and as ψ_n approaches zero, L_A and L_B become equal. For positive phase shifts, the overall bore length will be less than the acoustic length, and vice-versa.³ The picture below demonstrates this concept. There are four bore shapes

³ Note that the phase angle is the distance from the starting point of the tangent function to the $kx=0$ point, so a positive phase shift means that the tangent curve is shifted to the *left*. This can be a bit confusion for people who have not used phase angles before!

shown: a straight cylinder, a straight taper, a tube with dilation in the middle, and a tube with flared ends. The shapes are depicted in each graph.

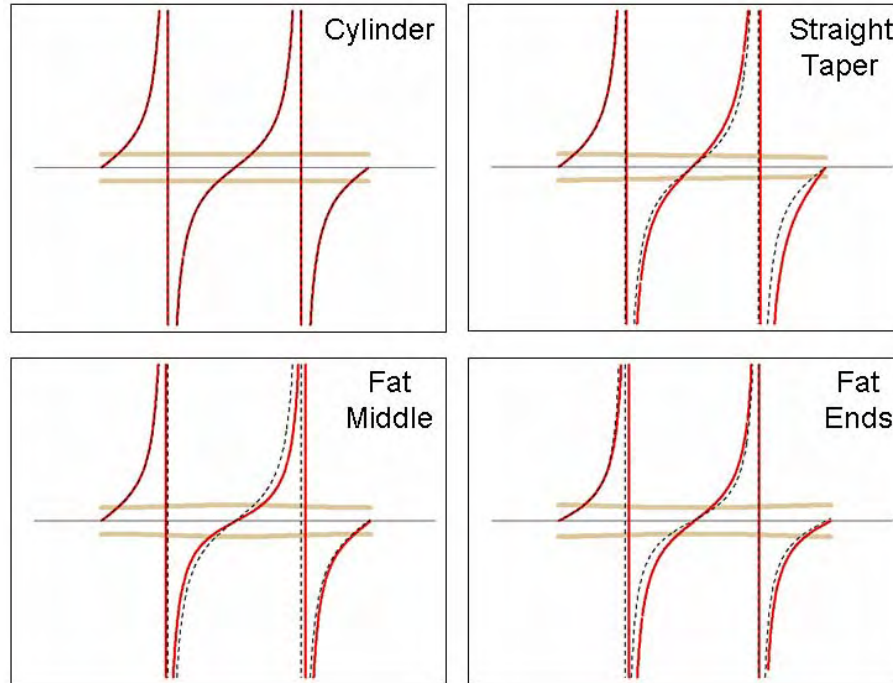


Note that the shape of the curve for the cylinder follows that of a tangent curve with a constant scale factor (since the cross section is constant throughout the bore). The dashed curve is the tangent curve for this case, and the same curve is drawn in each case for comparison. The straight taper shows a distortion in the curve that shifts the center of the curve toward the narrower foot of the bore. Notice that the bore length and the acoustic length are equal in this case.

The lower two pictures show the impedance curves for the other two types of perturbations. The bore with the fat middle can also be thought of as a bore with constrictions on the end. In an earlier discussion, the case of a constriction at the foot was discussed, and the result was a shorter bore length for the same sound wavelength. This is consistent with that discussion. The last bore with the fat ends turns out to have the opposite effect; the bore is longer than the acoustic length. Recall the discussion on the validity of the *magic ratio* when the section of a tube in question was far from the ends where the impedance was zero. This case demonstrates mathematically why a constriction at a flow node will actually decrease the acoustic length. This effect of perturbing the bore in the vicinity of flow and pressure nodes and antinodes is the basis for a qualitative approach to predicting their effects on the

intonation of a flute. This is what Benade referred to as a *perturbation weight function*, and will be discussed later.

The location of the nodes and antinodes are particularly important when the first and second octaves are to be in tune. The plots below show the same cases above for the second octave.



Notice that the acoustic lengths for the lower two cases are very close to the bore lengths. That is because, in these cases, the left and right halves are close to being conical. In other words, they are approximately straight tapers across a half wavelength of the acoustic standing wave, just as the tapered bore in the first octave was. This highlights the effect of introducing bore perturbations to affect the relative tuning between octaves. For instance, the bore with the fat middle will be flatter in the first octave than the second, while the opposite is true in the bore with the flared ends. Again, this is described qualitatively when using perturbation weight functions, but first we'll look at how this approach compares to Nederveen's description of a tapered bore.

Comparison to Nederveen

From Nederveen, the impedance of a bore with a straight taper can be written as

$$Z(x) = \frac{-j\rho c}{S(x)} \left[\frac{1}{\cot(kr(x) - kR_0) - \frac{1}{kr(x)}} \right]$$

where $r(x)$ is the distance from a given point in the bore to a point beyond the end of the bore where the projected diameter would converge to zero, and R_0 is the distance from the end of the air column to that point. Note that the $x=0$ point is now at the convergence point, and all the distances are negative in this description. The resulting equations predict that the length is the same for a tapered bore compared to that of a cylindrical bore, which indeed works out that way. The results are identical to that of the sectional approach described above. For straight tapers, Nederveen's approach is fairly straightforward. But for compound bores, the method does not lend itself to an easy execution.

Perturbation Weight Functions

Based on Benade's description, which is only qualitative in his "Fundamentals of Musical Acoustics" book, there are basically four things that could happen:

1. A contraction of the bore at a large pressure amplitude (antinode) increases the "springiness" of the air at that point, which in turn induces an increase in the vibration frequency, and the pitch rises,
2. An enlargement at a pressure antinode lowers the pitch,
3. A constriction at a point of large flow (flow antinode) increases the local density of the air, which in turn lowers the frequency, and
4. An enlargement at a flow antinode produces increases the pitch.

Recall that the ends of the bore are at a flow antinode, and in the fundamental mode (first octave), the pressure antinode resides near the middle of the bore. Conditions 1 and 4 describe the case of the bore with fat, or flared ends above, in which a constriction in the middle or an enlargement at the ends is in effect the same thing. In that case, the acoustic length was reduced relative to the physical bore length, which equates to an increase in pitch. Likewise, the bore with an enlargement in the middle (fat middle) is the same as one with constrictions at the ends, and the result is a larger acoustic length and lower pitch, as described by conditions 2 and 3.

As playing holes are opened, the wavelength of the air column is shortened and is contained in only a portion of the bore. The pressure and flow antinodes then shift to different locations, which means that bore constrictions or enlargements will be acting

on different parts of the air column. This means that a perturbation in a bore that was used to flatten a particular note, for instance, could have the effect of sharpening a different note. The calculations described above can be used to compute the result, but they give no intuitive sense ahead of time of how a particular design change might affect all the notes at once.

To digress for a moment, consider a scientific idea of “perturbation theory.” Before computers were around, some mathematical problems were difficult or impossible to solve exactly, so approximations were made. The magic ratio is an example of such an approximation. These approximations were usually of simpler problems that behaved in a similar way as the more complex problem. The simple problems didn’t always give the correct answer, but a small change in the simple problem would produce a small change in the answer, and the same small change in the complex problem produced the same small change in its outcome. So this way, an easy problem could be used to find out what would happen when a small design change was made in a complex problem. These small changes are called *perturbations*, and this approach to using easily-solvable equations to predict changes in unsolvable or hard-to-solve problems is called *perturbation theory*. This is very useful in flutes, because we can use a simple model like a cylinder to describe the bore, and introduce small enlargements or constrictions as bore perturbations to predict how they will affect the pitch of the notes.

In a cylindrical-bore flute, the pressure and flow standing waves are described by

$$p \propto \sin\left(m \frac{2\pi}{\lambda} x\right),$$


and

$$u \propto \cos\left(m \frac{2\pi}{\lambda} x\right)$$

where p describes the pressure variation along the bore, u describes the flow, x is the position along the acoustic length of the flute (including end corrections) for a particular note of the given wavelength, and m is 1 for the first octave, 2 for the second, etc. When the absolute value of p is at a maximum, an enlargement will cause a lowering of the pitch, and when u is at an absolute maximum an enlargement will cause the pitch to go up. These equations can be combined in the following way to describe those effects;

$$\begin{aligned} W &= u^2 - p^2 \\ &= \cos^2\left(m \frac{2\pi}{\lambda} x\right) - \sin^2\left(m \frac{2\pi}{\lambda} x\right) \end{aligned}$$

This equation describes what the pitch will do if an enlargement in the bore is made; a positive value means the pitch goes up and a negative value means the pitch goes down. If a constriction is made in the bore, then a negative value yields an increase in pitch, and vice-versa. Note that the value of the W function is between +1 and -1 (friction losses at the side walls has not been considered yet, and I expect that such losses will scale the W function to smaller values in areas of high flow). I chose to use the enlargement as the basis for the pitch change because that's what I currently use to adjust the tuning of flutes. The figure below shows the calculated perturbation weight functions for a cylindrical bore with six nominal finger holes (i.e., this is not for a real tuning—it's just to show the concept). As each finger hole is opened as represented by the tick marks, the wavelength decreases and the end points shift to the left. The extension of the curve to the right beyond the left-most open hole position represents the combined acoustic length of the open hole and the sections below that hole, and the extension to the left of the sound hole represents the end correction at the sound hole. The second overblown mode of three notes in the second octave is shown at the top of the figure.



Using such a set of curves constructed for a particular flute can help identify where certain bore enlargements might help balance tuning. For instance, if the first overblown note shown in the second octave (one open hole) is flat, it could be sharpened by sanding open the bore by a small amount (1/64" or less can have a significant effect) at the highlighted area. In the drawing, the areas under the weight

function curve in the dilated area are colored green if they are positive and red if they are negative. The green areas mean the pitch will be raised, and the red areas correspond to lowered pitch, and if a curve shows equal red and green areas the pitch is not changed. Now notice that that same perturbation will flatten all but the highest two notes in the first register at the same time, and to varying degrees. It will also sharpen the two adjacent notes to a lesser degree, and will have little or no effect on the pitches two notes away.


Diversion from the Simple Cylinder

If the starting bore is not cylindrical, then a weight function curve can be calculated for the general shape of the bore. In a segmented bore, each section has a phase shift ψ associated with it. The phase shift for the bore as a function of x can be obtained by plotting ψ for each section along the bore, then smoothing the curve to get a continuous function (if the curve is not smooth, there will be discontinuities in the W curve). Then the weight function curve is simply

$$W = u^2 - p^2$$

$$= \cos^2\left(m \frac{2\pi}{\lambda} x + \psi(x)\right) - \sin^2\left(m \frac{2\pi}{\lambda} x + \psi(x)\right)$$

The picture below shows W curves for the fundamental and the octave notes for a tapered bore and a bore with a dilated section near the head. The gray dashed curve is the W function for the tapered bore for comparison. Note that the center of the first-octave W function for tapered bore moves closer to the foot. Also, for the bore with the dilated section, the area of the W function near the bulge moves toward the left, and is much more pronounced in the second octave.



Estimating the Change in Pitch

To calculate an estimate of how much the pitch will change based on the perturbation weight functions, we can revisit the magic ratio. Recall the magic ratio was

$$\frac{L_1}{S_1} = \frac{L_2}{S_2} .$$

We'll now see that the weight function can be thought of a scaling factor on the answer the magic ratio produces for a change in the equivalent acoustic length. Recalling the earlier discussion on the validity of the magic ratio near the ends of the bore, notice that the average value of the weight function is close to +1 near the ends of the air column. If the change in acoustic length of a bore dilation calculated from the magic ratio is multiplied by this average weight function value, then the answer will not change very much. If the dilation is made at the center of the bore, the calculated change in acoustic length will be multiplied by an average weight function value of about -1, which reverses the answer that the magic ratio would predict—but this is actually the correct answer! What this means is the magic ratio *is* valid as long as it is used along with the weight function.

So the first step is to calculate the cross-sectional area of the bore segment in which the enlargement is made. If we look at segment i , then the length is L_i , the cross-sectional area before the perturbation is S_i , and after the perturbation it is S_{pi} . Then use

the magic ratio for the perturbed section (L_i/S_{pi}) and the equivalent section with cross section S_i to calculate the un-weighted acoustic length L_{ui} of that section so that

$$L_{ui} = S_i \frac{L_i}{S_{pi}} .$$

The un-weighted *change* in the acoustic length is

$$\begin{aligned} \Delta L_{ui} &= L_{ui} - L_i \\ &= L_i \left(1 - \frac{S_i}{S_{pi}} \right) . \end{aligned}$$

Next, find the average value of the weight function for that section. This will be the area under the W curve for that section, divided by the length L_i of that section, or

$$\overline{W}_i = \frac{1}{L_i} \int_{kx_{i-1}}^{kx_i} W dx .$$

I slipped in a little calculus,⁴ but if the curve is fairly straight between the kx_{i-1} and kx_i points on the curve, then the average weight is

$$\overline{W}_i = \frac{1}{L_i} \frac{W(x_{i-1}) + W(x_i)}{2} .$$

If the curve is not straight, you can break it up into how ever many smaller sections are needed to make the answer sufficiently accurate. Again, the test of “small enough” is when the answer doesn’t change much when the small sections are made smaller.

So then we need to multiply the un-weighted change in the acoustic length times the average weight of that section to get

⁴ If you don’t recognize the symbol in that equation, it is called an *integral*. What the equation does is breaks up the curve in the section of bore between kx_{i-1} and kx_i (which is the angular distance across section i that is kL_i long in angular distance) into infinitely small rectangles that have a height equal to the value of W at that point and a width dx . Then the area $W(dx)$ of all the rectangles are added up, or *integrated* to give the total area under that section of the curve.

$$\begin{aligned}\Delta L_i &= \overline{W}_i \cdot \Delta L_{ui} \\ &= \overline{W}_i L_i \left(1 - \frac{S_i}{S_{pi}} \right) .\end{aligned}$$

This is the weighted change in acoustic length for that section. The change for all perturbed sections must then be computed and added up to get the total change in the acoustic length of the bore ΔL . The new length is then

$$L_{new} = L_{old} + \Delta L ,$$

or

$$\begin{aligned}\lambda_{new} &= \frac{2L_{new}}{m} \\ &= \lambda_{old} + \frac{2\Delta L}{m}\end{aligned}$$

From λ_{new} , the new pitch or frequency can be found as

$$f_{new} = \frac{v_{sound}}{\lambda_{new}} ,$$

and

$$MIDI = 12 \frac{\ln(f_{new}/6.875)}{\ln(2)} - 3 .$$

This whole process needs to be repeated for each of the notes with their respective weight function curves. One thing that quickly becomes clear is that all the notes and all the physical variables that are applied to the bore, finger holes, sound hole design, etc. are interrelated, and getting a good tuning starts becoming more of an art!

Bibliography

Arthur H. Benade, *Fundamentals of Musical Acoustics*, 2nd revised ed., Dover Publications, New York (1990).

Arthur H. Benade and J. W. French, "Analysis of the Flute Head Joint," *J. Acoust. Soc. Am.* (37) (4), pp. 679-91 (1965).

John Coltman, "The Role of the Headjoint in Flute Intonation," *Flute Notes*, London (1985).

Neville H. Fletcher and Thomas D. Rossing, *The Physics of Musical Instruments*, 2nd ed., Springer, New York, (1998).

Cornelis J. Nederveen, *Acoustical Aspects of Woodwind Instruments*, revised ed., Northern Illinois University Press, DeKalb, IL (1998)

Lew Paxton Price, *Secrets of the Flute: The Physics, Math, and Design of Non-Mechanical Folk Flutes*, 2nd ed., Lew Paxton Price, Garden Valley, CA (1994).

Lew Paxton Price, *More Secrets of the Flute: More of the Physics, Math, and Design of Non-Mechanical Folk Flutes*, Lew Paxton Price, Garden Valley, CA (1997).

Joe Wolfe and John Smith, "Cutoff Frequencies and Cross Fingerings in Baroque, Classical, and Modern Flutes," *J. Acoust. Soc. Am.* (114) (4), pp. 2263-72 (2003).

Nelson Zink, "Tuning the Shakuhachi," in *Shakuhachi Design*,
<http://www.navaching.com/shaku/tuning.html>