

Introduction to saxophone acoustics

This introduction to the physics of the saxophone requires no mathematics beyond multiplication and division, nor any technical knowledge of acoustics. For some preliminary information about sounds and vibration, read the introduction to [How do woodwind instruments work?](#). For background on topics in acoustics (waves, frequencies, resonances, decibels etc) click on "Basics" in the navigation bar at left.

- [Overview](#)
- [The reed controls the air flow](#)
- [Playing softly and loudly](#)
- [The saxophone as a 'closed' conical pipe](#)
- [Harmonics of a closed conical pipe](#)
- [How the reed and pipe work together](#)
- [Spectrum of the saxophone](#)
- [Tone holes](#)
- [Register holes](#)
- [Cross fingering](#)
- [Other effects of the reed](#)
- [Cut-off frequencies](#)
- [Frequency response and acoustic impedance](#)
- [More information and links](#)

To set the mood, [listen to a movement of a quartet](#) for saxophone, flute, bassoon and cello.

Overview

The saxophone player provides a flow of air at a pressure above that of the atmosphere (technically, a few kPa or a few percent of an atmosphere: applied to a water manometer, this pressure would support about a 30 cm height difference). This is the source of power input to the instrument, but it is a source of continuous rather

than vibratory power. In a useful analogy with electricity, it is like DC electrical power. Sound is produced by an oscillating motion or air flow (like AC electricity). In the saxophone, the reed acts like an oscillating valve (technically, a control oscillator). The reed, in cooperation with the resonances in the air in the instrument, produces an oscillating component of both flow and pressure. Once the air in the saxophone is vibrating, some of the energy is radiated as sound out of the bell and any open holes. A much greater amount of energy is lost as a sort of friction (viscous loss) with the wall. In a sustained note, this energy is replaced by energy put in by the player. The column of air in the saxophone vibrates much more easily at some frequencies than at others (i.e. it resonates at certain frequencies). These resonances largely determine the playing frequency and thus the pitch, and the player in effect chooses the desired resonances by suitable combinations of keys. Let us now look at these components in turn and in detail.

The reed controls the air flow

The reed is springy and can bend. In fact it can oscillate like a spring on its own---for a saxophonist this is bad news: it's called a squeak. Normally, the reed's vibration is controlled by resonances of the air in the saxophone, as we shall see. But it's also true that the reed vibration controls the air flow into the saxophone: the two are interconnected.

Let's imagine steady flow with no vibration, and how it depends on the difference in pressure between the player's mouth and the mouthpiece. If you increase this pressure difference, more air should flow through the narrow gap left between the tip of the reed and the tip of the mouthpiece. So a graph of flow vs pressure difference starts off almost proportionally. However, as the pressure gets large enough to bend the reed, it acts on the thin end of the reed and tends to push it upwards so as to close the aperture through which the air is entering (the arrow in the sketch at left). Indeed, if you blow hard enough, it closes completely, and the flow goes to zero. So the flow-pressure diagram looks like that in the graph sketched at right.

The reed (as any saxophonist will tell you) is the key to making a sound. The player does work to provide a flow of air at pressure above atmospheric: this is the source of energy, but it is (more or less) steady. What converts steady power (DC) into acoustic power (AC) is the reed. The first part of the graph (and the dashed line) represents a resistance: flow proportional to pressure difference. Just like its electrical analogue, an acoustic resistor loses power. So in this regime, the saxophone will not play, though there is some breathy noise as air flows turbulently through gap between reed and mouthpiece. The operating regime is the downward sloping part of the curve. This is why there is both a minimum and maximum pressure (for any given reed) that will play a note. Blow too softly and you get air noise (left side of the graph), blow too hard and it closes up (where the graph meets the axis on the right).

Readers with a background in electricity, seeing the region of the curve in which flow *decreases* with increasing pressure, will recognise this as a negative (AC) resistance. Whereas a positive resistance takes energy out of a circuit, a negative resistance puts energy into the circuit (as happens in eg. a tunnel diode oscillator). In the saxophone, it is indeed this negative AC resistance that provides the energy lost in the rest of the instrument. Most of the energy is lost inside the bore, in viscous and thermal losses to the walls, and a relatively small fraction is emitted as radiated sound.

Playing softly and loudly

This diagram allows us to explain something about how the timbre changes when we go from playing softly to loudly. For small variation in pressure and small acoustic flow, the relation between the two is approximately linear, as shown in the diagram below at left. A nearly linear relation gives rise to nearly sinusoidal vibration, which means that the fundamental in the sound spectrum is strong, but that the higher harmonics are weak. This gives rise to a mellow timbre.

As we play more loudly, we increase the pressure (which moves the operating point to the right) and we also increase the range of pressure. This means that the (larger) section of the curve we use is no longer approximately linear. This produces an asymmetric oscillation, whose spectrum has more higher harmonics. (Centre diagram.)

When we blow even harder, the valve closes for part of the part of the cycle when the pressure in the mouthpiece is low due to the standing wave inside the instrument. So the flow is zero for part of the cycle. The resultant waveform is 'clipped' on one side (diagram at right), and contains even more high harmonics. As well as making the timbre brighter, add more harmonics makes the sound louder as well, because the higher harmonics fall in the frequency range where our hearing is most sensitive (See [What is a decibel?](#) for details).

While talking about decibels, we should mention that spectra are usually shown on a decibel scale. This means that one notices easily on the spectrum a harmonic that is say 20 dB weaker than the fundamental, even though it has 10 times less pressure and 100 times less power. What is important is that your ear notices it too, because of the frequency dependence referred to above. However, it is much more difficult to notice the presence of harmonics if you look at the waveform.

Soon after this page went up, saxophonists wanted to know a lot more about playing more loudly. Usually, blowing harder makes it louder: more pressure times more flow gives more power. Further, blowing harder eventually takes you into the non-linear and clipping ranges that produce stronger high harmonics, and therefore a sound that is both louder and brighter. However, if you blow too hard, you close the reed completely and it stays closed. A hard reed requires greater pressure differences to bend, and so allows you to blow harder without entering the non-linear or clipping ranges. On

the other hand, the smaller nonlinearity makes the sound more mellow (and therefore less loud, all else equal). This is complicated by the embouchure: because of the curvature of the lay, the position where you place your lower lip and how hard you bite makes a difference to the effective length and therefore stiffness of the reed. It's like a lever whose fulcrum is moved. Mouthpiece design also influences the amplitude at which non-linear or clipping ranges begin: a mouthpiece with an obstruction near the reed tip becomes less linear at lower pressures and so produces the brighter (harsher?) and louder sound beloved of some saxophonists. Alright, I admit it, I have such a mouthpiece in the case too.

The saxophone is a 'closed' pipe

The saxophone is open at the far end or bell. But it is (almost) closed at the other end. For a sound wave, the tiny aperture between reed and mouthpiece---a much smaller cross section than the bore of the instrument---is enough to cause a reflection almost like that from a completely closed end. The rest of the saxophone is approximately conical. Of course most saxophones are bent, and these sections are not conical. There's also the bell, some of whose effects we shall discuss below.

For the purposes of this simple introduction to saxophone acoustics, we shall now make some serious approximations. First, we shall pretend that it is a simple conical pipe---in other words we shall assume that all holes are closed (down to a certain point, at least), that the bore is conical, and that the mouthpiece end is completely closed. This is a crude approximation, but it preserves much of the essential physics, and it is easier to discuss. Of course the saxophone doesn't come to a sharp point: it has a mouthpiece. The mouthpiece is shorter and fatter than the cone it replaces, and it has approximately the same volume.

At left is a schematic of a soprano saxophone, an idealised conical bore and a truncated cone. At right is a photo of an alto saxophone. The larger saxophones are bent to bring the keys within comfortable reach of the hands. The bends make only modest differences to the sound, so we shall picture straight saxophones in the

diagrams here. This [photo](#) shows three experienced saxophones (sop, alto and tenor).

The natural vibrations of the air in the saxophone, the ones that cause it to play notes, are due to standing waves. (If you need an introduction to this important concept, see [standing waves](#).) What are the standing waves that are possible in such a tube?

To answer the question, we must take into account the fact that the saxophone is approximately conical. This means that sound waves 'spread out' as they travel down the bell. This means that the amplitude of the waves gets smaller as we go from mouthpiece to bell. The fact that the saxophone is open to the air at the far end means that the total pressure at that end of the pipe must be approximately atmospheric pressure. In other words, the acoustic pressure (the variation in pressure due to sound waves) is zero. The mouthpiece end, on the other hand, can have a maximum variation in pressure: it is an antinode in pressure. *If we were dealing with a cylindrical pipe (such as a flute or clarinet), where the standing waves are sinusoidal, we would expect the maximum and the zero of a wave to be one quarter wavelength apart. But the variation in the amplitude of the wave due to the variation in cross sectional area complicates the story.*

So we have devoted a whole page to comparing [cylindrical and conical pipes](#) and, if you want the details, you should read that page now. However, the result is this: the standing waves in a cone of length L have wavelengths of $2L$, L , $2L/3$, $L/2$, $2L/5$... in other words $2L/n$, where n is a whole number. The wave with wavelength $2L$ is called the fundamental, that with $2L/2$ is called the second harmonic, and that with $2L/n$ the n th harmonic.

The frequency equals the wave speed divided by the wavelength, so this longest wave corresponds to the lowest note on the instrument: Ab on a Bb saxophone, Db on an Eb saxophone. (See [standard note names](#), and remember that saxophones are transposing instruments, so that the written low Bb3 is actually Ab2 for a tenor saxophone in Bb, Db3 for an alto in Eb, and Ab3 for a soprano in Bb. Hereafter we refer only to the written pitch.) You

might want to measure the length of your instrument, take $v = 350$ m/s for sound in warm, moist air, and calculate the expected frequency. Do you get a better answer if you use the real length, or if you use the length of the cone made by extrapolating back to a point? Then check the answer in the [note table](#).

So, with all of the holes closed, you can play the lowest note: (written) Bb₃, with a wavelength roughly twice the length of the instrument. With this fingering, however, you can also play other notes by overblowing---by changing your embouchure and changing the blowing pressure. The harmonic series on Bb₃ is shown below.

How the reed and pipe work together

To sum up the preceding sections: the bore of the saxophone has several resonances, which are approximately in the ratios of the simple harmonics, 1:2:3, but successively more approximate with increasing frequency--we'll see why below under [frequency response](#). The reed has its own resonance--which is approximately what you hear when you produce a squeak. One good way to produce a squeak is to put your teeth on the reed. In normal playing, with your lower lip touching the reed, you damp (ie reduce the strength of) the reed's resonances considerably. This allows the resonances of the bore to 'take control'. To oversimplify somewhat, the saxophone usually plays at a strong bore resonance whose frequency is lower than that of the reed. (We shall see below how [register holes](#) are used to weaken the lower resonance or resonances and thus make one of the higher resonances the strongest.)

When the saxophone is playing, the reed is vibrating at one particular frequency. But, especially if the vibration is large, as it is when playing loudly, it generates harmonics (see [What is a sound spectrum?](#)). These set up, and are in turn reinforced by, standing waves. Consequently, the sound spectrum has strong components at harmonics of the fundamental being played.

Register holes are discussed in more detail on the page about clarinet acoustics. See [more about register holes on the clarinet page](#).

Spectrum of the saxophone

In the section on [standing waves](#), each wave is considered as a pure sine wave. The sound of the saxophone is a little like a sine wave when played softly, but successively less like it as it is played louder. To make a repeated or periodic wave that is not a simple sine wave, one can add sine waves from the harmonic series. So C4 on the saxophone contains some vibration at C4 (f_0), some at C5 ($2f_0$), some at G5 ($3f_0$) etc. The 'recipe' of the sound in terms of its component frequencies is called its [sound spectrum](#).

Opening tone holes

If you open the tone holes, starting from the far end, you make the pressure node move up the pipe, closer to the mouthpiece---it is very much like making the pipe shorter. Starting near the bell, each opened tone hole raises the pitch by a semitone, which requires a pipe that is about 6% shorter. After you open all of the right hand finger holes, as shown below, you have the fingering for G4, which is shown below. (Black circles mean hole closed, white circles mean hole open.)

For the moment, we can say the an open tone hole is almost like a 'short circuit' to the outside air, so the first open tone hole acts approximately as though the saxophone were 'sawn off' near the location of the tone hole. This approximation is crude, and in practice the wave extends somewhat beyond the first open tone hole: an end effect.

(For the technically minded, we could continue the electrical analogy by saying that the air in the open tone hole has inertia and

is therefore actually more like a low value inductance. The impedance of an inductor in electricity, or an inertance in acoustics, is proportional to frequency. So the tone hole behaves more like a short circuit at low frequencies than at high. This leads to the possibility of cross fingering, which we have studied in more detail in classical and baroque flutes.)

The frequency dependence of this end effect means that the low note played with a particular fingering has a larger end effect than does the corresponding note in the next register. If the saxophone really were a perfect cone, then the registers would be out of tune: the intervals would be too wide. This effect is removed primarily by the the shape of the mouthpiece.

Register holes

Holes can also serve as register holes. For instance, if you play E4 (call this frequency f_0) and then push the octave key, you are opening a hole part way down the (closed part of the) instrument. This hole disrupts the fundamental, but has little effect on the higher harmonics, so the saxophone 'jumps up' to E5 ($2f_0$). One can imagine a saxophone that had a separate register hole for each note, but that would be a lot of keys. In fact, only two register holes are used for the second register from D5 to F#6. This is of course a compromise: the register hole is never at the pressure node of the standing wave of the required upper note. This is not a big problem in practice. The register hole is small, so it is not really a 'short circuit', except at low frequencies. So it does not too much affect the higher frequency standing waves. It does however disrupt the fundamental, and that is its purpose: to stop the instrument dropping down to its bottom register.

Just in case you haven't noticed on your own saxophone, the octave key is automated: one key opens one or other of two register holes, according to whether or not the third finger of the left hand is depressed. So the upper register hole (right at the top near the mouthpiece) opens for notes above G#5, whereas from D5 to G#5

the lower hole is used. This is an example of a mechanical logic gate, which is well worth examining closely. The octave keys on oboes are only partly automated and on bassoons---trust me, you don't want to know.

The octave key register holes are used from D5 to F#6. In higher registers, other register holes are used. One is the key for F6. This hole is designed primarily as a tone hole, so it is bigger than it need or should be for an ideal register hole. Many players adjust the mechanism for the alternate F6 fingering so that it opens that hole only a little way. This improves its performance as a register hole, but compromises the intonation of the alternate high F fingering.

Cross fingering

On the saxophone, successive semitones are usually played by opening a tone hole dedicated to that purpose. Being a closed, cylindrical pipe the saxophone overblows an octave, and so one would need eleven tone holes to cover the keys in one octave before repeating fingerings using the octave key for the register holes. Because players don't have this many fingers, the keys and clutch mechanisms are supplied, so that one finger can close or open two or more holes. This fingering system is very similar developed by Theobald Boehm for the flute. However, cross fingering is also used. For instance, one of the fingerings for the note F#4 or F#5 is a simple fingering: the right hand index finger closes its hole, and the right ring finger opens a tone hole with a key (upper figure below). The other is a cross fingering: all three fingers on the left hand, plus the middle finger on the right close their holes. Why do these both give the same note?

An open tone hole connects the bore to the air outside, whose acoustic pressure is approximately zero. But the connection is not a 'short circuit': the air in and near the tone hole has mass and requires a force to be moved. So the pressure inside the bore under a tone hole is not at zero acoustic pressure, and so the standing wave in the instrument extends a little way past the first open tone

hole. (There's more about this effect under [Cut-off frequencies](#).) Closing a downstream hole extends the standing wave even further and so increases the effective length of the instrument for that fingering, which makes the resonant frequencies lower and the pitch flatter.

The effect of cross fingerings is frequency dependent. The extent of the standing wave beyond an open hole increases with the frequency, especially for small holes, because it takes more force to move the air in the tone hole at high frequencies. This has the effect of making the effective length of the bore increase with increasing frequency. As a result, the resonances at higher frequencies tend to become flatter than strict harmonic ratios. Because of this, often one cannot use the same cross fingerings in two different registers.

A further effect of the disturbed harmonic ratios of the maxima in impedance is that the harmonics that sound when a low note is played will not 'receive much help' from resonances in the instrument. (Technically, the bore does not provide feedback for the reed at that frequency, and nor does it provide impedance matching, so less of the high harmonics are present in the reed motion and they are also less efficiently radiated as sound. See [Frequency response and acoustic impedance](#). To be technical, there is also less of the mode locking that occurs due to the non-linear vibration of the reed.) As a result, cross fingerings in general are less loud and have darker or more mellow timbre than do the notes on either side. You will also see that the impedance spectrum is more complicated for cross fingerings than for simple fingerings.

We have studied cross fingerings more extensively on flutes than on saxophones, by comparing baroque, classical and modern instruments. (There is of course no baroque saxophone.) See [cross fingering on flutes](#) or [download a scientific paper](#) about cross fingering.

Other effects of the reed

As well as controlling the flow of air, the reed has a passive role in saxophone acoustics. When the pressure inside the mouthpiece rises, the reed is pushed outwards. Conversely, suction draws the reed in towards the bore. Thus the reed increases and decreases the mouthpiece volume with high or low pressure. (Technically, we say it is a *mechanical compliance* in parallel with the bore.) Indeed, it behaves a bit like an extra volume of air, which could also be compressed and expanded by changing pressure in the mouthpiece. It has the effect of lowering the frequency of each resonance a little. However, soft reeds move more than hard reeds, so soft reeds lower the frequency more than do hard reeds. Further, this effect is greater on high notes than on low, so soft reeds make intervals narrower and hard reeds make them wider. This is useful to know if you have intonation problems. (See also [tuning](#).)

The effect of reeds and their hardness is discussed in more detail under the "effect of hardness" section on the [clarinet page](#).

Cut-off frequencies

When we first discussed [tone holes](#), we said that, because a tone hole opens the bore up to the outside air, it shortened the effective length of the tube. For low frequencies, this is true: the wave is reflected at or near this point because the hole provides a low impedance 'short circuit' to the outside air. For high frequencies, however, it is more complicated. The air in and near the tone hole has mass. For a sound wave to pass through the tone hole it has to accelerate this mass, and the required acceleration (all else equal) increases as the square of the frequency: for a high frequency wave there is little time in half a cycle to get it moving.

So high frequency waves are impeded by the air in the tone hole: it doesn't 'look so open' to them as it does to the waves of low frequency. Low frequency waves are reflected at the first open tone hole, higher frequency waves travel further (which can allow [cross](#)

fingering) and sufficiently high frequency waves travel down the tube past the open holes. Thus an array of open tone holes acts as a high pass filter: something that lets high frequencies pass but rejects low frequencies. (See filter examples.) This is one of the things that limits the ability to play high notes on the saxophone. The stiffness of the reed is another: a saxophone will only play notes with frequencies lower than the natural frequency of the reed.

Frequency response and acoustic impedance of the saxophone

The way in which the reed opens and closes to control the air flow into the instrument depends upon the acoustic impedance at the position of the reed, which is why we measure this quantity. The acoustic impedance is simply the ratio of the sound pressure at the measurement point divided by the acoustic volume flow (which is just the area multiplied by the particle velocity). If the impedance is high, the pressure variation is large and so it can control the reed. In fact, the resonances, which are the frequencies for which the acoustic impedance is high, are so important that they 'control' the vibration of the reed, and the instrument will play only at a frequency close to a resonance. (There is further explanation on What is acoustic impedance and why is it important?). The section below shows how the major features of the saxophone's shape give rise to its acoustic impedance spectrum, and thus to how it operates.

This figure shows in black the calculated impedance spectrum for a simple cone (the impedance is given in decibels: $20 \log_{10}(Z/\text{Pa}\cdot\text{s}\cdot\text{m}^{-3})$). A suitable reed attached to the input of this tube would play near the frequencies of the peaks, which are in harmonic ratios 1:2:3:4 etc.

It is interesting to compare this pair of curves with the comparable curves for the clarinet. The soprano saxophone is actually a little

longer than a clarinet (about 700 vs 670 mm). However, the first resonance of the clarinet (call it f_0) occurs at a lower frequency than that of the saxophone (call it g_0). Further, the peaks of the saxophone curve occur at all the harmonics ($g_0, 2g_0, 3g_0$ etc), whereas those of the clarinet curve occur only at the odd harmonics ($f_0, 3f_0, 5f_0$ etc). The differences among open cylindrical pipes (flutes), open cylindrical pipes (clarinets) and closed conical pipes (saxophones, oboes, bassoons) are explained in [Pipes and harmonics](#).

This difference gives the clarinets a big advantage: a clarinet of a given length can play lower notes than can conical instruments of the same length. (Baritone sax players envy the small case that bass clarinetists lug around.) Clarinets pay for this advantage: the odd harmonic series means that they overblow a twelfth instead of an octave, which makes the fingering more awkward, especially around the 'break' between registers. They are also less loud than a comparable conical instrument. (The baritone sax player can blow the bass clarinetist away on that score.)

Still looking at the black curve for the simple cone, we see that, for frequencies above about several hundred hertz, the resonances become weaker with increasing frequency. This is due to the 'friction' of the moving air against the inside of the instrument (technically, the viscous and to a smaller extent thermal losses in the boundary layer). This affects higher frequencies more than low.

However, we also see a problem faced by the conical instruments but not by cylindrical instruments: the first resonance is weak. We saw above under [cut-off frequency](#) that, at low frequencies, it is easier to move the air backwards and forwards, because less acceleration is involved. This means that the impedance of the 'air at the downstream end of the instrument' (technically the radiation impedance at the end) is low for low frequencies. Further, the cone is good at matching the low impedance of the large end to the high impedance at the small end. Together, these effects make the lowest resonance of conical instruments weaker than those of their

cylindrical cousins. That is why the whole curve turns down to lower values at low frequency. (Compare with [curves for the clarinet.](#))

Effect of the bell. Let us now return to the graph of the impedance spectrum, which we reproduce below for convenience. The curve in red is for the same cone, with a simple bell at the end. Note that the bell makes the pipe longer, so each peak and trough has been moved to lower frequencies, as expected. Note however the change in the overall shape: all of the resonances are now weaker (extrema are smaller). This is because the bell 'helps' the sound waves in the bore to radiate out into the air. (Incidentally, the presence of a large, effective bell is what makes brass instruments loud: try playing a trombone with the tuning slide taken off.) More sound radiated means less sound reflected, so the [standing waves](#) are weaker.

Beyond the third peak, this effect increases with frequency: as the frequency increases over this range, the resonances are more weakened by the bell at high frequency than at low. This is because the bell is much smaller than the wavelengths of the low frequency waves, and so is less effective at radiating these waves. (Incidentally, this frequency dependent effect of the bell is what makes brass instruments brassy: the bell-less trombone has a sound that is darker, as well as softer.)

The effect of the bell is that of a high pass filter. We could say that this is rather like the [cut-off frequency](#) effect of a series of open tone holes. In fact, the purpose of the saxophone's bell chiefly is primarily to provide a high pass filter for the lowest few notes, so that they have a cut-off frequency and so behave more similarly to the notes produced with several tone holes open.

Note, however, that the bell also reduces the first resonance, by further improving its impedance matching to the radiation field. The bell is really only important to a few of the lowest notes (and to highly cross-fingered notes notes in the fourth octave). Saxophone

and oboe players will know the acoustic effect of this weak first resonance: it makes it difficult to play notes at the very bottom of the range. It is especially difficult to play these notes softly. We saw in [How the reed and pipe work together](#) that the higher harmonics of the reed motion could be 'helped' by resonances of the bore. When you play loudly, you generate proportionally more power in the higher harmonics (this is due to increasingly non-linear motion of the reed). So the lowest notes are not so difficult to play softly, where the strong, high harmonics of the reed are supported by resonances of the bore, compared with soft notes, where you are relying on the the weak fundamental resonance.

We have not mentioned the effect of the mouthpiece. First, we note that thje bore is not a simple cone: it is a cone truncated at a comfortable diameter to take a mouthpiece. This truncation affects the tuning: informally, we can think of it as making the pipe slightly like a cylinder, which stretches the frequency gap between resonances. This means that, unless compensated, it would stretch the interval between registers to over an octave. Now the geometry of the mouthpiece is a little complicated, but its main contribution to the acoustic response is that it compensates for the 'missing volume' of the cone. Indeed, its volume is comparable with that of the missing cone.

To the effect of the mouthpiece volume, we may add the [compliance of the reed](#), discussed above. This acts in parallel with the bore, and its impedance decreases at high frequency, so its effect is to reduce the rise in impedance with frequency: softer reeds give lower overall impedance at high frequency. Further, the very high resonances are weaker and occur at lower frequency when you use a soft reed. This is discussed in more detail under the "effect of hardness" section on the [clarinet page](#).