PRIMORDIAL SOUNDS: BIG BANG ACOUSTICS

Mark Whittle
University of Virginia

Press Release: AAS meeting, Denver, 6/1/04
(see also http://www.astro.virginia.edu/∼dmw8f)

Have you ever wondered what the “Big Bang” actually sounded like? Surely, you may be thinking, this is a trick question — didn’t it just sound like, well, a really big bang? Surprisingly, perhaps, the answer is “no, not really”. As is often the case with Nature, things are not so simple, and a more accurate description would be something like this: a descending scream, building into a deep rasping roar, and ending in a deafening hiss. As if this were not impressive enough, the entire acoustic show is itself the prelude to a wonderful transformation: the highest pitch sounds ultimately spawn the first generation of stars, while the deep bass notes slowly dissolve to become the tapestry of galaxies which now fills all of space. The birth of the Universe, it turns out, had its own primal scream.

Before describing how we know what these primordial sounds were like, let’s answer some immediate questions and concerns. First, weren’t we all taught in high school that sound cannot exist in the vacuum of space? Well yes, but space wasn’t so empty when the Universe was young. Remember, the Universe is expanding so it was smaller in the past, and all the matter we now see in stars and galaxies was spread out uniformly to make a hot thin gas, a kind of cosmic “atmosphere”. It is within this atmosphere that sound waves could form, grow and move.

If we were alive back then, could we hear these sounds? Regrettably, the answer is a very definite no, for several reasons. Even ignoring the fact that we would instantly suffocate and roast in the searing heat of that early fireball, the sound **pitch** is way too low for us to hear, by about 50 octaves. Imagine a sequence of progressively deeper “bass pianos” extending below the lowest notes of a regular piano: the cosmic concerto is played on the **seventh** piano in this sequence. Compared to concert pitch A, for which 440 sound waves pass us each second, a typical cosmic wave takes more like 50,000 years to pass by. The musical term “basso profundo” hardly seems adequate. On the other hand, the sound **volume** does remain within the human range, at least for the first million years or so. For example, about 400,000 years after the big bang (an important time, as we shall see), we observe pressure variations of about one ten thousandth of the average pressure, and this ratio corresponds to a sound volume of about 110 decibels — about as loud as a rock concert. At that time, the cosmic sound was neither pathetically quiet, nor fatally cacophonous. It was simply powerfully loud.

If we could shift the sound up by 50 octaves so that our ears could hear it, what would it sound like — a single note, a chord, a roar, a crackle, or what? Remarkably, the sound is somewhere between a chord and a roar. There are specific notes present in a sequence just like a normal chord, but whereas the notes in a normal (musical) chord are “pure”, the notes in the cosmic chord each contain a significant spread of pitch, and this renders the sound somewhere between a roar and a chord. Pushing this slightly further, by looking at the ratio between the pitch of the notes, we can figure out what the actual chord is. For the two deepest (and loudest) notes in the cosmic chord, we find a slow change across the first million years from a major third

(4 semitones) to a minor third (3 semitones). Stated more poetically, the Universe's symphony opens, appropriately, with a positive majestic major chord, but as time passes the mood shifts to a sadder one as the minor chord builds.

Unfortunately, it is difficult to recognize this subtle transition because of two other effects which dominate the sound. The first is a continuous drop in pitch spanning many octaves, as deeper notes are added to the mix and the waves are all stretched by cosmic expansion. This, combined with the steady increase in volume, makes a sound quite similar to a high speed jet plane passing close by, although in that case the sudden drop in pitch is caused by a different process, the Doppler effect. For the Universe, the drop in pitch is caused partly by its growth in size — smaller wine glasses vibrate with a higher note than bigger wine glasses — but also because there has been more time for deeper notes to form, the wavelength of a note being set by the length of time it takes for gas to fall into and bounce out of the slight gravitational variations that were present at that time. Notice that while the sound is dropping in pitch, it is also growing in volume. Ironically, the Big Bang started out silent! Only with the passage of time did sound begin to grow.

The second noticable feature of the sound is a dramatic transition at about 400,000 years. After this time, a hiss becomes apparent, at first almost imperceptible, but as time passes it builds to cacophonous levels and completely drowns out the earlier notes of the deep chord. This transition marks an important moment in cosmic history: the Universe suddently turns transparent. Before that time, the Universe was filled with a hot but opaque gas – a glowing fog. In this fog, light traps the gas and plays the dominant role, preventing the sound waves from growing too much. However, as the expanding Universe cools it turns transparent, quite quickly when the temperature dips below about 5000 F. As the fog clears, the gas is no longer trapped by the light, and for the first time it is free to move. Where does the gas go? Immediately, it starts responding to the gravitational pull of a hidden giant, dark matter — that illusive but ubiquitous substance which outweighs atomic matter six to one. All the while, dark matter had been gathering in clumps, drawn in by its own gravity. When the atomic gas starts falling into these clumps the sound volume increases enormously, and because many of the clumps are small, the dominant pitches are high and the sound become a loud hiss.

At this point, the sound transforms once again, and as gravity begins to overwhelm pressure, the gas falls into the smallest dark matter clumps, ultimately condensing to become the first generation of newborn stars. In a sense, then, these first stars were born from primordial sound. In a final poetic image, it is the light from these stars which re-illuminates the Universe, bringing to an end the first cold dark night with a new dawn. That first night had itself begun shortly after the Universe turned transparent, as the expanding fireball cooled and faded and darkness came, bringing to a close the most brilliant day of all, when the Universe burst forth in a fanfare of light and sound, announcing Nature's first dawn.

Having described these remote and wild times, let's now come back to Earth and look at how it is possible to know, with any certainty, what creation sounded like. The framework for our understanding has been in place for almost a century now, starting with the discovery of cosmic expansion in the 1920s by Edwin Hubble. The expansion implied the Universe was born in an initial "explosive" event which was later called, slightly mockingly by Fred Hoyle in 1950, the "Big Bang" — a name which, for good or ill, has stuck. For the present acoustic story, the two most important subsequent breakthroughs were the discovery of a faint microwave glow across the whole sky in 1963 by Arno Penzias and Robert Wilson, and then the further discovery in 1992 of slight patchiness in this glow by NASA's COBE satellite and science team. This patchiness is extremely slight, comparable to the height of a bacterium on a bowling ball, and its detection and measurement pose a very difficult experimental challange. Because of its importance, however, during the 1990s a number of groups worked hard to make ever more detailed maps of the microwave sky, culminating most recently in 2003 with the all-sky microwave map produced by NASA's WMAP satellite and science team.

It is this Cosmic Microwave Background (CMB), and in particular its patchiness, which holds the key to Big Bang acoustics. Let's briefly review what the microwave background actually is. Rather unbelievably, it comes from an ancient piece of the infant Universe which was only 380,000 years old. The fact that we can directly observe ancient history is an old astronomy trick: the light arriving from a distant object left it long ago, and so we see it as it was then, not as it is now. Look far enough away and you can see back almost to the Big Bang itself. But not quite. Remember that the early Universe contained a bright glowing fog. Hence, we can look out through the transparency of space as far as, but no further than, that glowing fog. Here's the spooky part. Because all directions look back in time, we see the fog in all directions — the whole sky should be glowing with the light from the Big Bang. And it is! We just don't see it with our eyes. Cosmic expansion shifts the light to become microwaves — as many microwave photons fall to Earth from the sky as do light photons from the full moon. If we had microwave sensitive eyes, even at night we could find our way by the light of the creation of the Universe! This extraordinary ability to witness, directly and in full panorama, the light from the Big Bang must be one of Nature's most remarkable gifts.

So far so good, but why have studies of the microwave sky spurred the development of Big Bang acoustics? Because most of the finest scale patchiness visible in the microwave maps shows, more or less directly, the peaks and troughs of sound waves moving through the hot gas of the young Universe. One can actually see the primordial sound waves, not moving of course, but frozen in place as they crossed the wall of fog, caught just as the Universe turned transparent. The situation is not unlike looking down over the ocean and taking a photograph: a whole collection of water waves is visible, little ones on top of bigger ones on top of even bigger ones, all superposed. Analysing the complex pattern of patches, using a computer, can yield the relative number and strength of waves of different sizes — in other words the relative loudness of high and low pitch notes. A graph of this is called the "power spectrum" and is a precise way to characterize the collection of waves, and hence the quality and loudness of the sound. Our knowledge of the CMB power spectrum has slowly improved over the last decade, and now spans about 10 octaves, of which the highest 5 correspond to acoustic waves (the lower octaves are another story, and paint a picture from an even earlier time, less than a nanosecond after the Big Bang). Remarkably, the upper power spectrum shows many of the features of a vibrating object: a strong fundamental at a wavelength of about 220,000 light years, and a sequence of higher harmonic peaks with shorter wavelength. It is this power spectrum which provides the starting point for the recreation of the primordial sound for human ears.

Before discussing that, however, lets clarify why the CMB power spectrum has been a holy grail for cosmologists. It certainly is not just to find out what creation sounded like, though that is an added bonus. There are two nice metaphors which illustrate why it is so important. First, if you strike a wineglass and then a teacup you instantly know that these two objects differ in many ways. The sound an object makes is like a fingerprint: it is unique to the object, and reveals much about its nature. The same is true for the Universe: its primordial sound carries a great deal of information about its structure and properties. Indeed, the CMB power spectrum has played a crucial role, when combined with several other astronomical datasets, in establishing more than a dozen basic properties of the Universe. There is another, perhaps more appealing, human metaphor which also illustrates the importance of the microwave background. Just 380,000 years into the life of the Universe is equivalent to just 12 hours into the life of a human. Now, 12 hours after conception, a human is tiny and formless, and all that is present is its DNA. Yet within that DNA, hidden and encoded, is information which determines much of what the developing child and adult will become. So too with the microwave background. It depicts a compact Universe which is virtually formless, and yet hidden within its delicate patchiness is encoded a huge amount of information, much of which determines how the Universe will subsequently evolve and grow. In a sense, then, studying the microwave background is to astronomy what the human genome project is to the life sciences. It is astronomy's "Cosmic Genome Project". Both genome projects, human and cosmic, present huge scientific challenges, but also promise huge scientific rewards.

With all this background now in place, let's return to the more playful nature of this current work — reproducing the primordial sounds suitable for human ears. The obvious first step is to use the observed CMB power spectrum to create a sound, remembering to shift the pitch up by about 50 octaves. The actual shift is arbitrary, but a good choice is to place the fundamental peak at about 220 Hz, corresponding to the A below concert A (a technical reason to choose this shift is to match the acoustic frequency in Hz to the measured "ℓ" harmonic number, which corresponds to a wavelength on the sky of roughly $180/\ell$ degrees). With this done, the sound is like a deep rasping roar which, if played at the correct volume of 110 decibels, is really quite powerful. Because the microwave background is a static image, this roar doesn't change in time — it is an acoustic "snapshot".

Now, as far as using real data is concerned, that is about all one can do — there is afterall only one measured CMB power spectrum. Fortunately, however, it is possible to go much further by using computer models of the early Universe. These are highly sophisticated programs which have been developed by a number of workers over the past decade, mainly to help interpret the observed CMB power spectra and extract cosmic properties. These computer programs are publicly available and relatively easy to use (the two I have used are CMBFAST written by Uros Seljak and Matias Zaldarriaga, and DASh, which is built around CMBFAST, and is written by Manlo Kaplinghat, Lloyd Knox and Constantinos Skordis). It is these programs which really allow one to access a host of aspects of the primordial sounds. For example, it is possible to generate CMB power spectra for different kinds of Universe, and then turn these into sounds. An overdense Universe with closed geometry, for example, has a deeper pitch than an underdense Universe with an open geometry. More subtle differences occur for Universes with more or less atomic matter content. (These differences have, of course, been known for a long time; this is simply an acoustic rendering of the differences).

The programs offer two extremely important options. The first is to correct for several non-acoustic distortions which significantly affect the patchiness seen on the microwave sky. In this sense, the Universe is not a perfect concert hall, and the primordial sounds suffer a number of distortions, both in situ and en route to us. For example, the foggy wall is not perfectly flat, it has significant depth, and so any sound waves shorter than this depth get blurred out and become invisible. This has the effect of supressing the highest pitches, and is analogous to a thick carpet or drapes in a concert hall which absorb high frequencies and deaden the sound. Other effects have no simple analog. For example, some patchiness arises due to Doppler shifts from moving gas, or to gravitational shifts due to light crossing density variations. Fortunately, the computer simulations work principally with the undistorted sounds, and only at the end of the calculation do they fold in all these other effects to yield the "observed" microwave sky. By catching the calculation before this final stage, one can extract a cleaner, more accurate acoustic power spectrum, relatively free of distortions (in cosmological parlance, these two power spectra are called $C(\ell)$ and P(k), "C of L" and "P of K", and one can think of them as the "observed" and "true" power spectra). As one might expect, these cleaned power spectra have narrower harmonic peaks which are better separated from each other, giving a somewhat more chord-like sound, though still resting upon a broad roar.

The second wonderful aspect of these programs is that one can follow the time evolution of the sounds, from just moments after the Big Bang, through the great transition of fog clearing, right up to the time when the first stars begin to form. It is these evolving sounds which have the characteristic downward scream descending into the deep roar which becomes drowned out by the loud hiss. A new concern enters with these evolving sounds — in truth the durations should be hundreds of thousands of years, but we dont have that kind of time available! What duration to choose? Currently, I've chosen a duration of a few seconds. This allows one to notice most of the interesting changes, and also evokes a suddeness which one naturally associates with an explosive expansion. It can also be useful to break the sounds into different epochs. For example, allocating 5 seconds to each of the three major stages: $0 - 50,000$ years for the descending scream, $50,000 - 500,000$ years for the deepening roar, and $0 - 100$ million years the growing hiss. By stretching time one can include all three: first second for $100 - 1000$ years; the next second for $100 - 1000$ years; the third for $1000 - 10,000$ years; and so on. Although one loses the authenticity of the sound's evolution, you gain in being able to nicely follow the major developments. A second important choice is whether or not to include the change in sound volume. The volume changes so much in the later epoch that if you make sure not to be deafened by the hiss, the descending scream and roar become almost inaudible. Again, one can make different versions, one which keeps the volume fixed throughout the progression and one which lets the volume change as it should. A number of these sounds are available on my website: http://www.astro.virginia.edu/∼dmw8f together with movies of the evolving power spectra. Hopefully, the web material will continue to improve as I have time to explore more aspects of this delightful subject.

Lets end this press release with two rather different but important statements. The first is to clarify the status of this work, what it is and what it is not. It is principally pedagogical. There is no new science here, at all. The existing knowledge base of the cosmology research community is vast and deep and subtle, and I could in no way contribute further to it (my own research area is in active galaxies, a quite different subject). What this work has done, perhaps, is introduce a new way of presenting some aspects of this exciting and rapidly developing field. When communicating scientific topics, especially remote or abstract ones, it is always important to use diagrams or images or, in this case, sounds. All of them are representations, none of them are perfect, but each brings the listener closer to grasping the concepts and ideas, and can even help establish an emotional connection with the subject. I think these sounds add to that repertoire in a novel and potentially powerful way.

A final point concerns the authenticity, or accuracy, of the sounds — did creation really sound like that? I have tried to point out the various ways in which the primordial sounds have to be modified for human hearing (their hugely sub-audible pitch; their lethally high temperature; the presence of distortions) but there are always more caveats lurking in the details. To illustrate, here's one of the more obvious ones: whereas normal acoustic power spectra represent the variation of pressure with time, the CMB power spectra represent variation of pressure with location. Depending on the situation, such distinctions can introduce uncertainties in the frequency content of the sound. In this particular case, such details might at worst alter the sounds slightly, so I decided to simply treat the power spectra in the standard way, as variations in time. Further concerns and details can be found on my website. Overall, however, I have tried hard to keep the sounds as close to what our current scientific understanding suggests, and in this sense they are authentic. When some aspect of the sound has been modified (such as maintaining fixed volume, or using a compressed time scale) this is noted, and the motivation is usually to maintain clarity. In other words, the sounds have not been modifed for "artistic" reasons.

Acknowledgements.

I have benefitted greatly from converstions and assistance from many people: Joe Wolfe and Alex Tarnopolsky in the UNSW Acoustic Physics group helped with generating sounds from power spectra; Charlie Lineweaver and Louise Ord in the UNSW Astronomy group, and Pedro Ferreira at Oxford, helped me to understand many aspects of CMB physics; Constantinos Skordis at Oxford, taught me how to use DASh and CMBFAST; and Mike Tuite at UVa helped make the movies with attached sound files. Finally, as a non-specialist, I have had to read quite widely, and have found extremely useful review articles by, among others, Wayne Hu, Max Tegmark, Joe Silk, and Charlie Lineweaver.

For more information, including a powerpoint talk and downloadable movies and sound files, go to http://www.astro.virginia.edu/∼dmw8f