

Death Metal/Throat Vocal Analysis

Chuck Stelzner

PHYS498POM

What is growling?

Wolves do it. Dogs and bears do too. Should humans? “On a regular basis” is often the answer, and one need only tune in the right radio station to hear it. Popularized partially as an outgrowth of the semi-screaming style of punk music vocals and the hoarse shouting of thrash metal, death metal growls, screams, and grunts are guttural and, some say, invoke a pseudo-evolutionary response – a “Pay attention!”, if you will. Bands like Possessed, Massacre, Morbid Angel, and Death sowed the seeds, along with their numerous European contemporaries (Celtic Frost, Venom), for a seemingly atonal range of vocals, produced essentially by blowing all the air out of one’s lungs as quickly as possible through tortured vocal cords, with the help of a raised Adam’s apple (contrary to traditional, operatic vocals). The followers and colleagues of these bands, and of course many more, birthed a new age, where atonal singing is seen as the new way to express extreme emotion – goodbye unamplified vocal projection, goodbye true vocal volume to signify emotional vocalizing! In much of popular music, especially in Europe, gone is the ‘clean’ vocal except as a side effect of a healthy career; to wit, many of death metal’s frontrunners only seek out classical voice training after their band has achieved some manner of fame, so to expand their fan base. And the rest, who stick with seemingly atonal vocal stylings, run the frequency spectrum: from screams so high they put squeaky car brakes to shame to growls so guttural that one is transformed into a rusty, human garbage disposal, slogging through the last of the stringy celery leftovers.

Again, what is growling? There are certainly those who think they have never heard such a thing, however it is almost impossible to avoid at least a reasonable facsimile of growling (leastwise because of its similarity to gargling and the real guttural sounds made everywhere in the animal kingdom). For a demonstration, a tutorial on reproducing the basic growling vocal style is as follows:

- First, drink some water. Milk, if possible. A simple demonstration will not give you throat lesions, but it might make you cough otherwise.
- Next, attempt to hack up some spittle. Make a sound like a cat with a hairball, while arcing your jaw shut several times. This is the sound like that social misfit in church, or right before your father spits from the car window. Yes, like that.
- Swallow. Ensure your throat is lubricated. Drink more water or milk if necessary.
- For the first try, pretend that you are gargling. This simulates the sensation of growling. Using real mouthwash is entirely acceptable.
- Once you are comfortable with this, jut your jaw forward and put your mouth nearly closed, into an exaggerated O.
- While using as much air as necessary, attempt to cough the ‘word’ “ooooooh”.
- Alternately, force your mouth open as wide as possible while trying to cough the

word “lie”, opening wide after the sound of the letter L.

Many online discussion forums touch on this (metaltabs.com, kerrazy-torrents.com), and every major metal band, whether they use a growling/screaming vocal styling, has a thread in their online forum where some person new to the scene asks how to reproduce the vocals of X band or Y singer. Many questioners are only harassed, as it has been realized that traditional instructions for death vocals are not in the least bit universal. Some people are only able to reproduce one style, some can do anything that they can hear with minimal practice, and some simply run home to their Chloraseptic spray and Zinc lozenges after trying for thirty seconds. Among the suggestions for ‘coating the throat’ are bananas, milk, hot ‘n spicy teriyaki beef jerky, hot tea, warm water, room-temperature water, and even baking soda dissolved in water. The issue with these is that many see them as crutches along the way – would you really keep a glass of warm milk onstage with you at a show, to gulp furiously between verses? These solutions are widely panned as simply ways to simulate growling or screaming and not true ways to perfect one’s vocal form, a poor and transitory substitute akin to temporary tattoos. Furthermore, baking soda mixed with water, when drunk after sessions of death vocals, actually thickens the vocals cords on a semi-permanent basis, eventually changing the timbre of one’s voice; this side-effect is not included along with postings of this ‘golden tip’ as often as it should be.

Furthermore, there are other potential risks to screaming continuously at the top of one’s lungs, shouting atonally into a microphone for two plus hours a day, or cough-growling your way to success in the metal underground. When done at all for long periods of time (though food-based lubrication can help in the short term), the inner surface of the throat is laid open, raw, and is therefore more susceptible to infection and sickness. Lymph nodes in the neck/throat area swell, attempting to protect from potential disease, and at the very least one is usually left with a sore throat. It is the sort of thing which much be eased into, which is only truly satisfying to do for a relatively small portion of the population, and which seemingly inevitably causes lesions in the throat and on the vocal cords, making speech and simply breathing a somewhat painful thing later in life. Being that the vocal styling is largely treated like homosexuality used to be – with a general outsider consensus opinion of “that’s fine for you, but we don’t want to know about it” – there has been relatively no research done on death vocals as they relate to the prevalence of things that most people worry about – throat cancer, true vocal cord damage, and the like. In first wave of people who have made their livings gargling and growling and screaming onstage, most of them have yet to reach age 40, let alone 60 or 70, when the human warranty runs out and every questionable habit makes its consequences known.

Why growling?

With all of the pessimism, the unavoidable Neo-satanic stereotyping, and the pain, why would anyone persist in abusing themselves like this? One might as well ask why people procrastinate, or why people diet, why they fight, or even why they explore space. It seems primal, unavoidable, a testing of human boundaries. And why do a project on it?

My interest lies in the masking of tones. With a classical music background, I found myself at a crossroads, bombarded by bands that I found myself unable to traditionally

‘sing’ along with. In the case of Slayer (post-Reign in Blood), the shouts follow the tones of the lead guitar fairly well. I found myself wondering exactly what there was to such a style, and for some people it is easier to simulate a growl than to simply shout in a semi-tonal way. Possessed of nearly no technical musical background, the wonderment ended there, when I failed to duplicate Tom Araya’s vocals (the lead singer of Slayer).

Flash forward 7 years or so, to an almost unbelievable hybrid class called The Physics of Music and Musical Instruments. We were (re-)introduced to MatLab, and introduced to Wave Analysis Toolbox, a series of analysis tools written by Joe Yasi (a former Physics of Music student) for MatLab, for determining the harmonic content, phase information, and myriad other aspects of a sound sample, as well as extracting those same harmonics for individual playback. I had long wondered, given the way in which it is produced, whether growling could be married to any one fundamental frequency. Certainly there were higher vocal stylings, the ‘death scream’ common to black metal (an outgrowth of death metal that deals mostly in medieval themes, Satanism and the occult, and grisly death, often characterized by a banshee shriek type of vocals), so it could be easily said that metal vocals run the gamut and touch at least on selected parts of the frequency spectrum. But were they, especially in the garbage-disposal depths of grindcore and extreme death metal, actually hiding a note? Could they be stripped down to something that sounded like a tuning fork, or a normal operatic clean note? Furthermore, as somewhat of a whimsical, fantastical dream, I wondered if it would be possible (and simple) to reverse-engineer the process, to be able to add a ‘growl effect’ to normal vocals? I decided to find out, with the help of MatLab, SourceForge (a sound playback and manipulation program) and the Wave Analysis Toolbox (WAT).

What are you talking about?

Here are a few terms, some of which I’ve already used above. Once defined, they will greatly help to understand the pages to follow, which will have them couched in somewhat less explanatory syntax.


- death vocals – any of the family of semi-tonal vocals, on a continuum from *growling* to *shouting and yelling* to the high shrieks and screams of *hardcore* and *black metal*;
- clean vocals – classically-trained singing, such as one would hear at the opera;
- growling – the lower register of death vocals; similar to gargling while coughing;
- shouting and yelling – exactly what one would guess: loud and angry vocals, like in an argument;
- hardcore – a type of music that serves as a hybrid of emo (overly-emotional, maudlin music, with *clean vocals*) and death metal (complete with *growling* and/or *screaming*), switching back and forth between the two at will; can also sound like *grindcore*;
- black metal – to quote from above: “an outgrowth of death metal that deals mostly in medieval themes, Satanism and the occult, and grisly death, often characterized by a banshee shriek type of vocals”;
- screaming – vocal styling prevalent in *hardcore*, *black metal*, and *thrash metal*; similar to the squawks of birds and to the high-pitched nasal voice of fairy tale witches;
- grindcore – subset of *death metal*, characterized by playing ‘as hard and as fast as we can’; unintelligible, deep *growling* prevalent, with occasional *screaming*

-death metal – originally, metal centered entirely around death; very loud, often unintelligible vocals, anywhere on the spectrum of *death vocals*; lyrics often discuss death, anger, dominance;


-thrash metal – type of metal only one step removed from mainstream American metal; characterized by *shouting and yelling* that border on higher *growling*; faster tempo than typical metal and *death metal*;

What is harmonic analysis?

To begin, picture a wave; perhaps from your favorite movie or crime TV show, where they analyze someone's voice to find important clues to a murder. It might look like this:



Jean Baptiste Joseph Fourier, a scientist of the 18th and 19th centuries, developed a method of representing any periodic function (such as the wave above) with a series of component waves built from sines and cosines. His work enabled the math and science communities to represent otherwise difficult functions with easy-to-analyze periodic functions. A sine wave looks like this:



Smooth and very predictable. A cosine wave is simply one of these waves, shifted to the left a little; Fourier said that many of them can be added to approximate (and with an infinite number of them, exactly duplicate) any periodic waveform, even the messy one above. A string of sine and cosine terms that can be added to produce an approximation of a given periodic function is called – the whole string of terms, that is – a Fourier series. Harmonic analysis is simply the “computation and study of Fourier series” (MathWorld).

For any given waveform, a harmonic is one of the sine or cosine terms that makes up the Fourier series of that waveform; it's one of the pieces that you can put together to get the original waveform again. Imagine, if you will, adding the sine wave above to mutations and shifts of itself and then finding that you have the first, messy waveform above...

The reverse can also be done: it's entirely possible to find the harmonics of a waveform; that's exactly what WAT does. But before we call that definition complete, here's some further clarification: all harmonics are terms in the Fourier series of a waveform, but only “integer multiple[s] of the fundamental frequency” are harmonics (Harmony Central).

That means that not all terms in the Fourier series can be called harmonics. And what is the fundamental frequency? The lowest prominent frequency in a waveform – usually the one you hear and would identify it as (‘That’s a C#!’).

Notice that I didn’t say it’s the absolute lowest frequency in the Fourier series – most times, with sounds like the one used to produce the first picture above, there is very low frequency noise below, above, and between the harmonics, but no one hears a sound and can pick out the frequency of the noise. Thusly, the lowest prominent frequency, usually the loudest one, is the fundamental frequency. All of the harmonics are integer multiples of it, making it fairly easy to predict where they’ll be found during harmonic analysis.

In short, harmonic analysis is computation and/or experimentation to determine how many harmonics a given waveform actually contains, and how relatively loud they are compared to the fundamental. In the case of a perfect sine wave, like the second picture, there is only the fundamental frequency and no harmonics. If one were to compute the Fourier series of a sine function, one would get... a sine function. A single term, meaning that there are no higher harmonics anywhere in the waveform. The same holds true for a perfect cosine wave, but every other periodic waveform in existence has a multiple-term Fourier series.

A triangle wave, or a wave that looks like the teeth on a saw, has a fair amount of harmonics, as you might guess – it takes the addition of quite a few cosine and sine functions to approximate a sawtooth wave very well. For a square wave, it takes even more, and this tells you that the square wave has more harmonics – if the range of frequencies of the terms of the Fourier series is bigger, there will be more harmonics hiding there.

What exactly did you set out to do?

My initial questions were:

1. Can death metal vocals be reduced to a clean, fundamental note?

This is explained above, just above the “What are you talking about?” heading.

2. Does the frequency that a growl appears to be in any way match the analyzed frequency?

To wit, does the apparent frequency of someone growling match the frequency that WAT gives you?

3. If the answer to #1 is yes, then can the harmonics and the original sounds be combined in any way to capture only the noise, or to add the noise to a clean note to make it a growl?

Can SoundForge or any other sound analysis program simply add, subtract, or mix these raw sounds and analyzed sounds to produce only the noise inherent to growling, and can this be added to a sound sample of the correct frequency to produce a growl?

4. Does the size of the mouth opening effect death vocals in the same way as it does for clean vocals?

Opening the mouth wide and relaxing the tongue and throat produce better-sounding operatic vocals – what is the effect on growling?

What did you do, and how did you do it?

For the ‘how’, included with this report is a document detailing the instructions on how to use WAT and SoundForge. They are not included here because they are very precise and beyond the interest of the casual reader. My procedural walkthrough (third section below) highlights some of the exact steps I took to harmonically analyze a given sound and record the results, but only to the extent that it helps to understand the process. Again – for instructions on operating MatLab and WAT, please see the included (separate) instructions document.

Research

Before I began, it was the strong advice of Professor Errede that each of us draw from the wisdom of others. Thusly, I did google searches for such word strings as “growling death metal vocals”, “growling analysis”, and “harmonic growling analysis”. I came up with a great deal of really fascinating things, such as scientific analyses of a growling vocal style used in ethnic music styles and a paper written about growling while playing a saxophone, but nothing about harmonic analysis as I envisioned it. I even stumbled upon a band called Hatebeak, who use a trained parrot as their source of death vocals, but nothing to draw upon. Thusly, I started fresh, in to my knowledge largely uncharted territory.

Appealing to the Death Metal Community

Is there such a thing? A death metal community? Not in the sense that they all get together for musical heritage reunions every summer (though the Wacken Open Air festival in Wacken, Germany comes close), but if you search hard enough you can usually find an official webpage for any given band, let alone an email for the webmaster or even the individual band members. On a suggestion from Professor Errede, I took the initiative to compile two massive lists of death/thrash/black metal vocalists email addresses and to email each them individually. For my own sanity, I made a master email and changed the content as appropriate for each vocalist, adding comments or removing the fact that I live in the USA so as not to patronize the American musicians. This file is included as *core_email.doc* along with this report.

I also started a new thread in the vocal zone of the metaltabs.com forum (link at the end of this report), as part of my surprise that there was no evidence of anyone having investigated the musicality of death vocals. I included my email address, and besides getting seriously flamed on the board, I got one non-professional applicant, Will from Catharsis, a band in Wisconsin.

Out of the 42 professional vocalists I contacted, only 6 replied, only 2 of these contributed directly. Furthermore, as will be detailed in the **How Did It Come Out?** section below, only one of these was in true death vocal style, being that of Bobby Bray of The Locust. The other was a slightly off-topic contribution from John Whoolilurie, who instead submitted a fascinating sample of dual-note throat singing, whose harmonic analysis in turn inspired further analysis of Bobby’s vocals...

I attribute the poor turnout for professional vocalist contributions to the fact that, while they do vocals in English (all but a choice few), not many are fluent to the level that I wrote the request email. Many of them are from Sweden, Germany, Finland, etc., and while I shouldn't expect them to know English, it became somewhat disappointing, thinking that some of them might have contributed had I known anyone who could translate the request into their native tongue(s)...

Also, more than a few of them were on tour during part of the semester, recording a new album, or – shock of shocks – working non-metal-music-related jobs, living their lives.

Recording

It was the suggestion of Matt Winkler, TA for the class and an excellent guy to have around for anything music-, MIDI-, and physics-related, who suggested that I simply go ahead and record my own vocals while I waited for the professional vocal submissions to 'pour' in (my own sarcastic wording). In the darkroom near the lab, we recorded two different sessions of my clean and death vocals, employing the use of a keyboard during the second session for pitch matching.


Matt's part in things was largely the same during both sessions – he ran the recorder, checked the volume level (which gave constant problems), and hummed back his guess of what my growling pitch was. I growled, he hummed while I ran out of air, and then I would try singing and growling that same pitch so that we could match them. After we settled on a pitch, we recorded first the clean note and then the growl for several different pitches. In this first session, no real heed was given to mouth positioning/aperture width.

In the second session, I wore a single headphone attached to the keyboard, which allowed us to augment both the growls and the clean vocals to match a key on the keyboard. I took much better notes, including where on the keyboard each note was supposed to be, and we recorded two of each vocal (death and clean) – first open-mouthed, then almost closed-mouthed.

Harmonic Analysis

1. Before I took the analysis very far, I created an MS Excel spreadsheet (included as *master listing.xls*) to house the results.
2. Just prior to analysis, I ensured that the spreadsheet was on the correct sheet, page, etc. I then booted up MatLab, ran WAT, and loaded a sound file.
3. The next available step in WAT is to select part of the waveform for analysis. Anyone doing their own harmonic analysis with WAT beware: analyzing a two second clip on a 1.5GHz machine takes a minute or two. Thusly, I selected, at maximum, a two-second clip, from the meat of the sound. By this I mean that I disregarded any lead-in time and any die-out time, where the amplitude of the sound was not constant and where the harmonic content was bound to be unreliable.

4. I then chose the harmonics for analysis. A picture of an example of this (a clean vocal):



5. This has been vertically compressed to save space; it is a logarithmic plot of frequency vs. amplitude for the entirety of a two-second sample of a clean vocal (average amplitude for the duration of the sample). If the picture were larger (full-sized harmonic display graphs to follow), one could see that, to the left of the first peak, there are much smaller peaks right above 0Hz. These are below the threshold of concern when in the presence of such large signals such as the peak directly to their right. This leftmost large peak is the fundamental frequency, and all of the other evenly-spaced peaks to its right are harmonics, corresponding to integer multiples of the fundamental. For each analysis I chose roughly as many harmonics as possible (selecting more harmonics also increases the analysis time significantly), up to and including nine total harmonics. Before analysis, I also took a snapshot of the get_harms screen (such as the above picture) and saved it.
6. After that, the computer did its job, separating out each harmonic specified. If there was a system beep during the analysis, I simply closed and opened MatLab again and tried to pick more specific harmonics (thinner frequency ranges) or left out some of the more 'iffy' harmonics. The harmonics for clean vocals are very straightforward, but growls look like trash, and it's understandably difficult to pick out analyzable frequency ranges for the individual harmonics.
7. When the analysis was complete, I saved all of the associated files, to include sound files of the harmonics all by themselves and many graphs depicting statistics about the harmonics.
8. After each analysis, I then closed MatLab completely. Failing to do this means that WAT might fail during the next round of analysis or might produce erroneous data or files.
9. During the course of analysis, MatLab produces separate sound files for the analyzed harmonics as well as for the unanalyzed clip of the original waveform (the two second clip). After closing MatLab, I loaded both the clip and the fundamental frequency waveforms in SoundForge. Using the Statistics option, I took the Zero Crossings result to be the overall frequency of the clip. It is important to note that, since this number is computed by measuring the number of times that the waveform crosses the zero point and dividing by the length of the clip, it is not an accurate gauge for the frequency of any unanalyzed clips. They are still composed of many separate terms in a Fourier series, all of which have different frequencies, thereby throwing off the zero crossing

calculation. For the fundamental frequency and harmonic clips, however, it is extremely accurate.

10. After recording the frequencies in the spreadsheet, I played the two sounds and compared them. I asked myself, “Does the fundamental frequency clip sound like the prominent frequency in the unanalyzed clip?” I then recorded the quality of their frequency match in the spreadsheet.
11. I then furiously backed up every file I had open (which I highly recommend), even the spreadsheet. This did not help very often, as weeks worth of spreadsheet data disappeared at the drop of a hat, from one lab session to the next.
12. If further analysis of the same waveform was needed, I loaded MatLab and WAT again and performed it. Sometimes I wanted to investigate large peaks that occurred below the expected fundamental frequency (for my own growling/clean vocal comparisons), and sometimes the waveform seemed to contain several fundamental frequencies – one can only analyze harmonics with respect to a single fundamental frequency at a time.

How did it come out?

Besides three things, the results were much better than expected:

1. The fact that the newest copy of the master spreadsheet seemed to disappear from one lab session to the next, necessitating backlogging the recorded numbers.
2. The fact that each lab session was 6 hours long and that MatLab seemed to need a complete system reboot (the computer itself) after 4.5 hours of running WAT (in order to complete any harmonic analyses).
3. The fact that sometimes I would wait 3 minutes, just to find out that WAT somehow choked on the last harmonic out of 5 that I chose. And then doing it all over again in the next 3 minutes...

But these were minor compared to the thrill of analyzing death vocals for college credit, and compared to what I found. Below are the relevant analyses of first my own vocals and then of the contributions I received, complete with graphs of harmonic content.


Chuck (Me) – session 1

Session 1 had three clean vocals and three growling vocals, which were essentially, frequency-wise, medium, high, and low. They were recorded to match the pitch of their associated growl, and the pitch-matching was only somewhat successful when the numbers rolled out of SoundForge. Overall, there was an average of 12.21Hz between the clean and growled vocals, giving an error of 9.6% overall (when compared to the growls, which were the supposed baselines for the pitches).


Clean 1 vs. Growl 1

They certainly sound close to one another – a generally good match – but the growl came out (after analysis) to be 14.02Hz above the clean note. As with any of the results, this could be due to the excellent adaptability of the human ear, averaging frequencies to reduce sonic dissonance during the dead reckoning of pitch that Matt and I did, or it could be due to imperfect filtering of the frequencies. Perhaps I even chose the

boundaries for the harmonics poorly within WAT. Percentage error was 10.8% for this set of vocals.




This is a larger version of the example from above, with the correct aspect ratio. Notice the plethora of harmonics that really just jump out at you, with eight of them below 1000Hz. Again, these clean vocals were amazingly easy to analyze, both for myself with choosing the harmonics and for WAT with filtering out the harmonics.




Growl 1, looking as if it has a bunch of noise to the left of the highest peak (around 500Hz). This is not so. Being that I did 2 analyses for each growl in session 1, I can say that the first, relatively small peak is the actual fundamental frequency, and it matches the unanalyzed growl as the prominent sound that one hears when listening to the unfiltered waveform. The higher peak is due to the formant (“characteristic resonance region”, a range of frequencies where the shape of the instrument – in this case the throat – dictates an amplitude increase for any frequencies that fall within it), whose filtered sound, when viewed in a sound editing program, does a fair job of approximating the overall shape of the unanalyzed waveform.

Clean 2 vs. Growl 2

These also sound close to one another, but they were 7.68Hz apart. This gives a 4.8% error for this set of vocals.




This is Clean 2. I definitely might have done a better job opening up my mouth and aligning my throat for this one, as it's got several dozen harmonics, much more than Clea 1. Note that the formant, as one would expect from the appearance of so many harmonics, is slightly different, flattened in the 2500-3200Hz range.




Here we have Growl 2. It looks amazingly similar to Growl 1 in all ways, to the point of me doing a double-take and redoing the WAT analysis because I thought I'd accidentally analyzed Growl 1 twice. Again we have the fundamental frequency smaller than the highest peak, which is again around 500-600Hz.

Clean 3 vs. Growl 3

Again, these two vocals sound close to one another, but they were 14.93Hz apart. Also, this time the clean vocal was higher than the growl, a first for session 1. This gives a 16.1% error for this set of vocals.



Clean 3: In this case, the formant is again different (outside of the 0-1000Hz range), as was the case with session 2 as well – the growls had remarkably similar formants, but the clean vocals shifted significantly from one note to the next.




Growl 3: This growl was the lowest in session 1. Growls 1 and 2 are closer to one another in pitch than either is to Growl 3, and the formant change shows it, bottoming out from 2000-3000Hz.

Chuck (Me) – session 2


Session 1 had two total pitches, each delivered as an open-mouth growl, a closed-mouth growl, and a closed- and open-mouthed clean note. Note that when I say ‘closed-mouth’, I mean that I attempted to stick out my lips like a fish when delivering the vocal, nearly closing my mouth entirely. Each duet of clean was again done to match the pitch of their associated growls, and the pitch-matching was much more successful for this entire session, being a 8.0466666% overall error (note that this is the % error as calculated without Growl 2b, which could not be analyzed; further note that this would be much better, closer to 4%, if there had not been serious pitch-matching issues with Growl 1b while recording). I attribute this to the use of the keyboard as a reference note. I’ve come to terms with the fact that I don’t have perfect pitch (see elsewhere in the PHYS498 student projects for a class project based on perfect pitch), and having the keyboard keys to repeatedly stab at to remind myself was a huge boon.

Clean 1a vs. Growl 1a

These two vocals sound very close to one another – they are only 3.31Hz apart, with Clean 1a being higher. This gives a 2.8% error for this set of vocals.




Clean 1a: Not much new to see. It's pretty similar in formant shape to Clean 2 from session 1, however they're 30Hz different in fundamental frequency.




Growl 1a: It's interesting to note that, while the fundamental and its associated harmonics are very low, as in all of the growling samples, the formant is a hybrid: from 0Hz to about 2000Hz, it somewhat resembles Growls 1 and 2 from session 1, however from 2000Hz onward, it looks like Growl 3 from session 1.

Clean 1b vs. Growl 1b

For some reason the pitch matching on these was none too great, though there was much effort on my part to stay true to the keyboard. These two vocals sound marginally related – like first cousins – and they are 16.73Hz apart (Growl 1b higher), a gap that one would expect from session 1. This gives a 13.5% error for this set of vocals. The percentage error for these two vocals, and these two vocals only, was calculated with the clean note as a baseline, as it's so obvious upon listening that the growl was too high in pitch.




Clean 1b: Look at those harmonics crumble! Up above the 2500Hz point, the harmonics look terrible when compared to the graph for Clean 1a. I see this as a direct result of the fish-mouthed sort of mouth positioning.




Growl 1b: For the noticeable differences between this and Growl 1a, it appears that the formant above 2000Hz stayed exactly the same. And when you consider that the formant in the 500Hz range was reduced by nearly a factor of 10 from Growl 1a to Growl 1b, you can see the effect of the fish-mouthed positioning.

Clean 2a vs. Growl 2a

These were well-matched in pitch (4.1Hz difference, Growl 2a above), giving a 2.3% error for this set of vocals. These two were right on.




Clean 2a: Interesting formant, given that it has the amazingly abundant harmonic content of Clean 1a but a smoother formant than any of the other clean vocals I recorded.




Growl 2a: A bit of a departure from the previous formants below 2000Hz, but pretty much par for the course above 2k.

Clean 2b vs. Growl 2b

Sadly, I can't tell you much about the difference in pitch between the fundamentals of these two waveforms, as Growl 2b presented so many problems and so many persistent errors that I had to give up on it and leave it unanalyzed. Here are the harmonic graphs, though...



Clean 2b: Interestingly, this looks to be a formant hybrid of Clean 1b and Clean 2a, which makes sense: Clean 1b was also a fish-mouthed vocal, and Clean 2a had the same fundamental frequency (within .53Hz).



Growl 2b: This looks nearly identical to Growl 1b's graph, down to the number (4) of peaks under about 600Hz associated with the fundamental frequency. The fish-mouthed styling worked wonders for making these look like harmonic twins.

John Whoolilurie

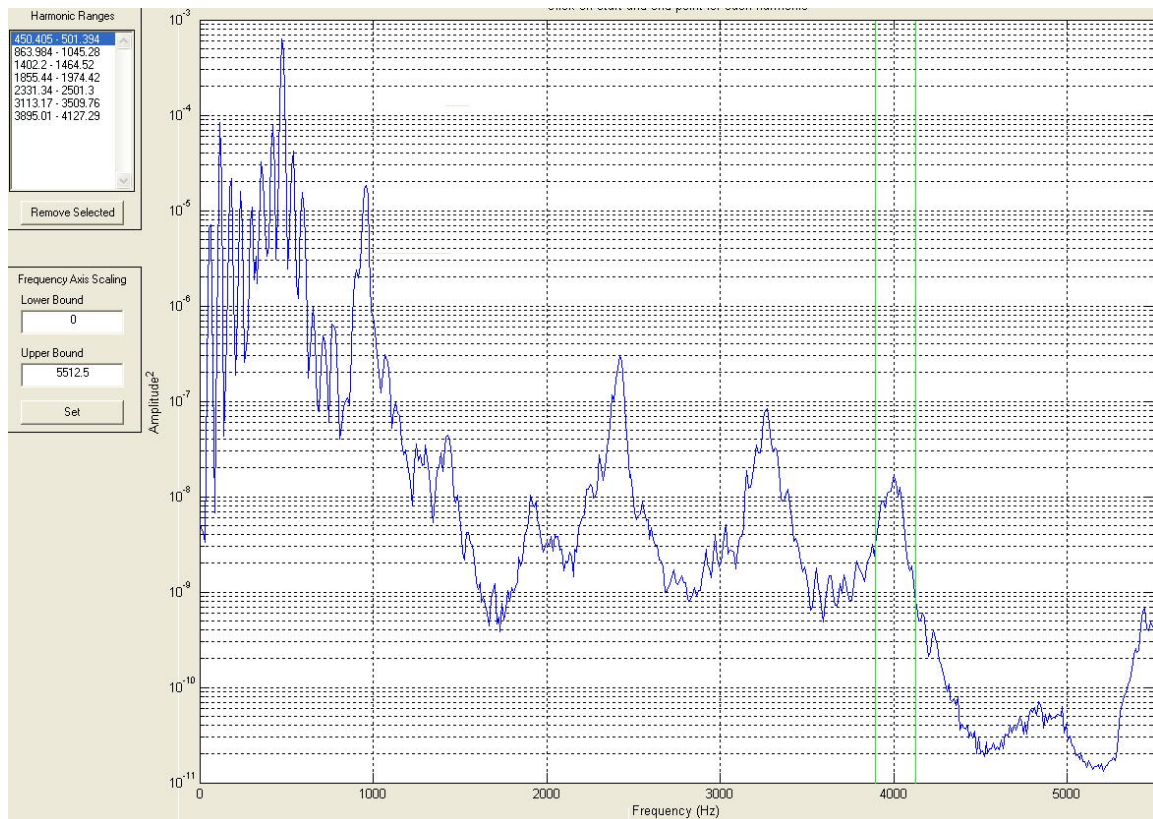
In truth, I analyzed John's vocals last, but it makes sense to put him before Bobby so as to illustrate how I went back and paid more attention to one of Bobby's vocals as a result of John's submission.

John is an amazing musical man, one whom I have had the privilege of seeing perform live, with his former band Estradasphere. To the best of my recollection, at that concert he played bongo drums, saxophone, and tambourine as well as singing cleanly in several different musical styles, not to mention growling. He has released his own CD, with every lick of the musical production performed by him (besides, I believe, some of the mixing and/or mastering), being that he performed on roughly a million separate instruments. With his musical diversity in mind, it therefore came as no surprise that he has advanced his vocals into the realm of dual-tone singing.

Popularized somewhat by Richard Feynman and Tuvan throat singers, dual-tone (and even tri-tone!) singing exists in many different forms. While John's style is not that of the


whistling-type throat singers, who sound like a particularly loud alarm going off in an air duct, his is more mellow.

The bottom fundamental is produced as normal, as one would shape the vocal cords to sing a single note, and the second fundamental is produced by manipulation of the formation of the formant (concentration on and contortion of the throat, etc.). In John's case, the upper fundamental is one octave above the lower.



John 1: It's not exactly obvious where the two tones are, and I used a non-visual approach to determine where the upper fundamental was hiding. I did a WAT analysis based on the obvious bottom fundamental and then listened to each of the filtered harmonic samples to compare it against the unanalyzed waveform. It was clear that the first harmonic of the bottom fundamental was in fact the upper fundamental. It is also good to note that, while there are three separate columns for John in the spreadsheet, the first and third sounds are identical copies of one another and were therefore not included twice graphically in this report.

Also, I did analysis of John's vocals similar to the 'meat 2' analyses of my own vocals – I did a third analysis of John's vocals with the tallest peak as the fundamental frequency. This produced, as I expected, a filtered waveform that matched the amplitude and envelope of the original unanalyzed sound but did not match it in pitch (it sounds like a high whistle, that must be masked entirely in the unanalyzed sound).



John 2: Very similar in both formant and harmonics. The upper fundamental was again roughly twice the frequency of the lower, and was in fact the first harmonic of the lower fundamental. Also again, harmonic analysis based off the highest peak yielded a waveform that sounded like a whistle and that matched the envelope of the unanalyzed sound.


Bobby Bray

The Locust are a sonic phenomenon. Their songs are about a minute long, if that, and they contain some of the most furious, unintelligible lyrics in all of extreme metal. They dress in grey outfits and fencing masks when playing live (thereby looking like actual locusts), and their song titles often have no bearing on the length, content, or sound of the songs themselves, instead being alliterative whimsies or describing hilarious situations. Bobby Bray is one of two guitarists and one of three vocalists for the band. He was kind enough not only to record vocals specifically for my project but to muse a while on screaming vocals:


“I do think that the whole appeal behind the screaming, high pitch vocal is that it can illicit an evolutionary alertness which can be translated as “PAY ATTENTION”. Perhaps it sounds like someone in extreme agony or someone dying, truly one of the first means of verbal communication we must have had. I kind of see my screaming as an homage to evolution.”

Bobby sent me 12 different screams in all (and one which faded from one pitch to another, prompting me to divide it into two different sounds), and a big regret on this

project is that I didn't have enough time to properly investigate them all. I had made it through perhaps Scream 6 when I returned to the lab yet again to discover that not only had Microsoft warped time, they had warped space as well, swallowing part of my data. Being put behind by that cut the results short. Here are the harmonic graphs for the ones that I was able to analyze:




Scream 1: Being the first sample, there is little to say (once one is at the baseline of thought that one is looking at the harmonic analysis of screaming, that is). It looks almost like there are two different tones around 1000Hz, but the lower is definitely the prominent fundamental when listening to the unanalyzed scream. Furthermore, it could simply be the result of the inherently distorted delivery of death vocals.



Scream 2: This is my favorite sound file for analysis among all of them, even my own vocals, for it has two fundamental frequencies within it... they're not harmonically related to each other at all, certainly not multiples of each other! In a way, it's good that Excel nuked some of my results for Bobby, as I didn't know what I was looking at (hadn't even considered it, really) until I analyzed this vocal for the second time, after having analyzed John's dual-tone singing. I even reconsidered Scream 1, and certainly looked at Scream 3 differently...

Both fundamentals were analyzed separately, and upon listening to them with the unanalyzed vocal, my suspicions were confirmed – they're both present in the original, a dual-tone scream!



Scream 3: The smaller, lesser peak to the right of the fundamental (smaller peak is at about 1300Hz) is a weaker example of the dual-tone screaming vocal, but both fundamentals are present in the unanalyzed vocal.

My only theory on the duality of tones in Bobby's screams is that something might be flapping back and forth very quickly while he vocalizes, akin to quantum states in their inseparability and coexistence. Perhaps his vocals cords are doing something natural, that we all might do if we were to scream that loud, but perhaps part of his range is inherently dual-tone because of something unique to his throat positioning...

Will (Catharsis)

Will's contribution was in response to my posting on the metaltabs.com vocal zone forum, and he kindly sent me a well-recorded mp3 of two different screams.

The sound file began with Will speaking directly to me, telling me where he was from and the name of his band and that he had two screams for me. His speaking voice was somewhat soft-spoken, and so I was nearly blown out of my chair when he shrieked louder than I could have imagined. Truthfully, my next-door neighbor's dog started to bark. And what would you do if you were governed primarily by instincts? The first shriek sounds like a Velociraptor from Jurassic Park, and I'm amazed that Will still has a normal-sounding speaking voice. Kudos to him for going above and beyond.

However, he evidently went above and beyond what WAT could handle. I tried to analyze his vocals each lab session after I got them, and all I got for my trouble was a plethora of error beeps – WAT choked on his vocals and would not analyze them, no matter how few or how many harmonics I selected, and no matter how precisely I selected them. I have found myself wondering if perhaps it is because of the extremely high frequency (over 1000Hz) of the fundamental of his screaming. Then again, Bobby's vocals were up in the 1000-1200Hz range, so perhaps it had something to do with the extremely sharp edge of the sound.

What would you do different next time? What did you learn?

Next time, I would choose much smaller frequency intervals for each harmonic in WAT, and I would back my data up in no less than six different locations, to include burning it to CD after each lab session. I would delve into MatLab coding, to enable customization of the harmonic analysis program, possibly even to the point of preventing the errors that I consistently got when I selected any 'rough' harmonics (jagged but clear peaks). I also would think about recording some manner of my own screams and consider using a talk box to determine the formant of several of my throat positionings, in order to compare them to the harmonic content graphs that WAT produces.

Given the chance to do the project again, I would also endeavor to find a way to translate my submission request email into the native tongue of each vocalist I contact, whether it be through a paid translation service or through a personal contact. This could conceivably increase the participation in future death vocal projects. And, lastly, I would not take the time to save each of the many graphs that WAT outputs, with information about the relative and absolute phase of the individual harmonics, etc. They were very bright and pretty, but as the analysis process and the project evolved and mutated over the course of the semester, I drew essentially zero information from them, instead opting to perform more and more analysis, taking screen captures of the get_harms display screen as the focal point of my analysis.

Lastly, I would finish analyzing Bobby's vocals, to hunt for more dual-tone screams!

I learned many life lessons, not the least of which being that, no matter how paranoid you are about backing up your work, there's always a lightning strike (figurative) to put you out of business. I learned that all is not like the movies, where a doctor clamps a little doo-dad on John Travolta's throat and he magically sounds like Nicholas Cage. Likewise, filtering sound is every bit as difficult as the phenomenal (fake) zooming, sharpening, and general image manipulation that's everywhere in spy movies and crime series on TV. Sounds cannot simply be added and subtracted to determine what breed of dog was barking in the background when Herb Johnson left the message on Aunt Bertha's answering machine – that's basically just fantasy (more on this just below).

That was a lot – please summarize!

The best way I can think to summarize and to conclude is to answer the questions I stated earlier:

1. Can death metal vocals be reduced to a clean, fundamental note?

Simply, yes. It might take quite a lot of teeth-gnashing and wailing, but it can be done in the majority of cases.

2. Does the frequency that a growl appears to be in any way match the analyzed frequency?

Absolutely. It seems to be easier to control and focus the pitch of a growl with your mouth almost closed, an option that is not available in operatic singing. In clean singing, one is encouraged to keep one's Adam's apple as low as possible and to keep the mouth as open as the lyrics will allow. Closing the mouth tends to draw the Adam's apple upward, limiting the good warmth of the tone. In death vocals, this is not a problem, this closing of the mouth, as the Adam's apple is already, by definition, hiked up as far as possible.

3. If the answer to #1 is yes, then can the harmonics and the original sounds be combined in any way to capture only the noise, or to add the noise to a clean note to make it a growl?

They may be combinable and mixable to do amazing things, but I wasn't able to do any of that. The power to do that, I wholeheartedly believe, lies in the magic of code such as makes up WAT, and not in the neophyte SoundForge manipulations of such as myself.

4. Does the size of the mouth opening effect death vocals in the same way as it does for clean vocals?

I essentially already answered this in response to question 2, but let me say again: no. Keeping the mouth opening small made it much easier for me to reign in and control the pitch of my growls.

Further investigation would definitely be needed (and wanted!) to determine how it is that Bobby's screaming is dual-tone. Were this perhaps my actual profession, I would love to do tests similar to what I read about with Tuvan throat singers some weeks ago, where doctors actually put some manner of small tube down their throats and took measurements while they vocalized. It also would be excellent to see if the same holds true for other screamers, even perhaps for Will.

AND: Even though I removed my contact information from the request letter for the death vocals, please feel free to contact me with any questions, comments, or requests (related to this project in any way) at chuckstelzner@gmail.com. Along with this report should be included the two lists of vocalist emails that I used to reach out to the metal community, so if you are inclined to do a similar project, all that I ask is that you keep me informed on anything cool or particularly fascinating that you discover, and of any great submissions you get in response.

Who helped out?

Extreme thanks are due to the following people:

- Bobby Bray**, guitarist/vocalist in The Locust, for his extensive contribution of screaming thrash vocals and for his thoughts on screaming vocals
- John Whoolilurie**, ex-member of Estradasphere and mojow of mojow and the Vibration Army, for his dual-tone throat singing contribution
- Will**, vocalist for Catharsis (WI), for his amazing screaming contribution, though it was just too metal for MatLab to analyze :)
- Toby Driver**, of Kayo Dot (formerly maudlin of the Well), for giving me permission to analyze any of the old maudlin material, though I didn't end up having the time for more analysis...
- Professor Steve Errede**, for running such an incredibly cool course (and for making sure that WAT's amazingly short-lived site license got renewed), and for letting me use his fabulous sound recorder and equipment!
- Matt Winkler**, TA for the course, for a number of things: BSing with me about all manner of music, laughing when I blew everyone's head off with a high-volume playback of screaming vocals, and for invaluable help with the actual recording process; I can't think of anybody else who would debate me whether my growl is closer to F4 or F#4...
- Joe Yasi**, author of WAT, though he likely doesn't yet know this project exists; everyone please note that WAT is my acronym for the software, not his :);

Honorable Mentions:

- Jason Blachowicz** (bass/vocals, Divine Empire) and **Paul Kuhr** (vocals, November's Doom), and **Tommy Rogers** (keyboard/vocals, Between the Buried and Me), for at least replying to my emails...

Where did you get your information?

-Joe Yasi's report; the birth of WAT:

http://online.physics.uiuc.edu/courses/phys199pom/199pom_reu.html

-Fourier series and harmonic analysis: <http://mathworld.wolfram.com/FourierSeries.html>

-Fourier's full name: <http://www-groups.dcs.st-and.ac.uk/~history/Mathematicians/Fourier.html>

-Harmonics: <http://www.harmony-central.com/Guitar/harmonics.html>

-Formants: <http://www.sfu.ca/sonic-studio/handbook/Formant.html>

Do you have any other links for us?

Absolutely!

-mojow and the Vibration Army (John Whoolilurie's band): <http://vibrationarmy.com/>

-The Locust (Bobby Bray's band): <http://thelocust.com/>

-Kayo Dot (Toby Driver's band): <http://www.kayodot.net/>

-the metaltabs.com vocal zone forum:

<http://metaltabs.com/forum/forumdisplay.php?s=e2cebb1ab5255f8dc4eec2dd1e795d5e&f=22>

-kerrazy-torrents.com metal chat forum: <http://www.kerrazy-torrents.com/lofiversion/index.php/f13.html>