replying to, so it is more specific than "Re: Contents of Sirsound-list digest..." so that it matches the post you are replying to.

Also, please EDIT the quoted post so that it is not the entire digest, but just the post you are replying to.

Today's Topics:

1. Re: Eigenmike (Fons Adriaensen)

Message: 1
Date: Tue, 11 Jun 2013 22:02:12 +0000
From: Fons Adriaensen <***@linuxaudio.org>
Subject: Re: [Sursound] Eigenmike
To: ***@music.vt.edu
Message-ID: <***@linuxaudio.org>
Content-Type: text/plain; charset=us-ascii

> Hi Jorn,
> Why do you say that you would not want to use the Eigenmike for musical recordings?

(pre scriptum: please do not copy the entire digest when replying...)

Joern will probably answer your question and provide some good arguments for his opinion. But I do share his view on this. Of all the musical recordings made using the Eigenmike I've heard so far there were none that I really liked. It's all too easy to get carried away by the initial quite spectacular surround effect, but once that wears off, in all cases I want to reach for some EQ. And using EQ usually improves things, but never to the point that the result is really OK.

When you use a recording technique such as Blumlein, XY, MS, or ORTF everything depends on how well the mics preserve their polar pattern in function of frequency - or their frequency response in function of direction - for the simple reason that almost all sources will be off-axis. That's an entirely different situation compared to multi-mic techniques, where such things matter much less, and it's why people tend to use high quality expensive mics in such a role.

When using the Eigenmic (or a first order tetrahedral), you're in the same situation. You will use either virtual directional mics or Ambisonic components which are polar patterns as well. Simple fact is that with the current state of the art these synthesised patterns are not really what they should be - they are only correct over a rather limited frequency range which gets narrower as either Ambisonic order or directivity goes up. Which in the end means you can either get the direct sound right, or the diffuse (reverb) part, but not the two at the same time. Which is why EQ fails to correct things.

In some applications that is not a problem, for example when you need some virtual spot mics to follow one or a few actors on a film set or stage. But when the result depends on capturing sound from all directions in a balanced way - as it does for music recordings which include e.g. the concert hall acoustics as well, this is a real problem.

There are currently two SW apps for processing the Eigenmic signals into either a set of directional 'spot' mics or into third order Ambisonics: the original SW coming with the Eigenmic, and the one developed by Farina's team here in Parma. While I'm not personally involved in the latter, I'm following that project closely. But so far neither of them have things as right as they could be. Part of the problem may be that even while the Eigenmic capsules are of very high quality, the remaining random differences in sensitivity and frequency response are still too large to allow accurate higher order processing without individual calibration. And until someone finds a better method than renting an anechoic room for half a day, such calibration is an expensive exercise.

Earlier today I've been listening to some organ recordings made using the Eigenmic and processed into 3rd order Ambisonics. This is one of the applications that work reasonably well - frequency response errors on organ sounds are much more forgiving than on e.g. a string quartet, or voices. But I found the difference between a 1st order decode and a 3rd order one rather marginal while it is normally quite easy to hear that difference. Which probably means that higher order components are not as accurate as they should be.

Ciao,

--

FA

A world of exhaustive, reliable metadata would be an utopia.
It's also a pipe-dream, founded on self-delusion, nerd hubris
and hysterically inflated market opportunities. (Cory Doctorow)

Sursound mailing list
***@music.vt.edu
https://mail.music.vt.edu/mailman/listinfo/sursound

End of Sursound Digest, Vol 59, Issue 12
****

-------------- next part --------------
An HTML attachment was scrubbed...
URL: <https://mail.music.vt.edu/mailman/private/sursound/attachments/20130615/07349173/attachment.html>

Fons Adriaensen 2 years ago

> *Post by Paul Power*
> so it seems that the microphone is still not a perfect solution for HOA recording and requires some highly technical tinkering before it will give half decent results.

It depends very much on the application. Some of those recordings made by Farina's team are quite good, for example the collection of 'ambient sounds' of the city they made earlier this year. But the musical ones so far failed to convince me. It's the processing SW that isn't yet what it could be - I've no doubts about the quality of the mic itself.

I've been present at some of the music recordings, so I can compare my own impression of the event and the reproduction. One of these was at the opera house in Parma. The mic was placed just behind the conductor, and maybe half a meter above his head. Not and ideal position, but it provides a reference point. I was sitting in the first row, less than a meter from the mic stand. The recording was 'decoded' using a number of spot mics looking down on the orchestra and four or so covering the stage (an angle of 120 degrees or so as seem from the mic). The spot mics were then panned into 3rd order AMB, and also used as primary sources in the WFS system of the Sala Bianca.

In particular the stage mics needed quite some EQ in order to sound more or less natural, but even then the result lacked the presence of the original sound. Parma's opera house is very 'dry' - the result of the layers of plaster added with each restoration without removing the older ones. So sitting in the first row you really get a sound that is quite 'direct', I wouldn't want a seat there normally. But it seemed impossible to reproduce anything like that.

Ciao,

--
FA

A world of exhaustive, reliable metadata would be an utopia.
It's also a pipe-dream, founded on self-delusion, nerd hubris
and hysterically inflated market opportunities. (Cory Doctorow)

Stefan Schreiber 2 years ago

For clarification:

Could the eigenmike also do some 4th order recording (in a "real" sense, not giving some 4th order "output"), or is it a 3rd order microphone?

Thanks,

Stefan

...

Jörn Nettingsmeier 2 years ago

> *Post by Stefan Schreiber*
> Could the eigenmike also do some 4th order recording (in a "real" sense,
> not giving some 4th order "output"), or is it a 3rd order microphone?

with 32 capsules, the eigenmike is capable of measuring the full set of fourth-order harmonics, but only with limited bandwidth.

that's the main problem with the eigenmike: the higher the order, the narrower the bandwidth and the higher the lower boundary. which means that an eigenmike signal set will not usually be downwards-compatible - if you aim for a correct amplitude response for the full 4th order set, it will sound duller and duller the more orders you truncate. it is possible to generate a spectrally correct output for each subset of orders, but that has to be computed from the a-format, with a priori knowledge of the playback order.

--
Jörn Nettingsmeier
Lortzingstr. 11, 45128 Essen, Tel. +49 177 7937487

Meister für Veranstaltungstechnik (Bühne/Studio)
Tonmeister VDT

http://stackingdwarves.net

Stefan Schreiber 2 years ago

...

A short and precise answer, with a lot of stuff inside... :-)

Danke,

Stefan

Sampo Syreeni 2 years ago

> *Post by JÃ¶rn Nettingsmeier*
> with 32 capsules, the eigenmike is capable of measuring the full set of fourth-order harmonics, but only with limited bandwidth.

When you go from the tight confines of a classical SoundField to a general cloud of mics, there's not just one but three separate things you have to bear in mind.

The first one is the obvious change in channel count, which I believe you're thinking about here. In that regard, yes, it is true that as long as you have enough mics and they are reasonably spread out (formally, "in general position"), pretty much any array geometry at all can capture as many harmonics as will fit within the channel count. If four mics, then first order. If nine, then second. And so on.

However, that sort of counting neglects what the particular geometry will capture in addition to the harmonics you're interested in. When you work with a traditional SoundField, first the response pattern of the capsules is very nearly first order, and then the near-coincidence of the capsules makes it relatively easy to reject any extra higher order harmonics via filtering which bled in because of spacing effects. Neither part is true of arrays built from anything other than near-coincident max-first-order capsules, which is to say, of any array above the POA one. They'll always suffer from directional aliasing, so that what order they can capture while rejecting the higher ones almost always becomes a relative business. (Eigenmike is one of the better ones because it's a rigid sphere tessellated with first order capsules, but it too will have to have wildly lower high order rejection than SoundFields do.)

That will cause the prerequisite number of mics to blow up beyond what naïve accounting would suggest for a given order and rejection level. Even in this kind of a sound design, the only way to really be sure is to fully blanket the surface of the sphere with capsules at something like 1/3 wavelength spacing and very even symmetry, which is something nobody even tries because of the cost.

And third, especially with arrays of omnis, or setups where you reach out at the LF from a central Soundfield with extra, spaced capsules, you'll always have to mind the noise amplification characteristic of the A-to-B matrix even if your array geometry is optimal. And once again you'll have to add 2-3 fold more mics in order to get the same performance per mic you got from the basic SoundField. That's because with monopoles you'll be doing differencing between the capsules in order to derive directionality, and even with cardioids or some similar directional pickup, your optimal frequency dependent matrix will be doing similar differences between the cardinal directions of your array, in order to correct for coloring and in order to reject higher order contamination.

So, while I once again haven't calculated this out in full, as a rule of thumb I'd say that extended spherical pickups effectively only have 2/3 of the ideal channel count their capsule count might suggest if mounted on a rigid sphere, optimized right (EigenMike) and only upto a wavelength about thrice the minimum distance between capsules over the spherical surface. If mounted in free space, make that perhaps 2-3 times as many capsules as you would have though when cardioids, and as many as 3-5 if they're monopoles, depending on how well you've optimized the array geometry and what you started up and ended up with, order-wise.

Judging from data alone, I'd say 3rd order periphony with the EigenMike is already reaching it. Might work in the far field, and even in a reverberant concert hall to a degree, but 4th order would clearly kill it there. In the kinds of cramped spaces Angelo once described, I would be surprised if it could transparently do 2nd order.

...

How does that happen? Okay, it might be that the usable frequency range changes, but once you've done equalizing in the A to B step, don't you then always get just a B-format set which you then decode at your target order?

--
Sampo Syreeni, aka decoy - ***@iki.fi, http://decoy.iki.fi/front
+358-50-5756111, 025E D175 ABE5 027C 9494 EEB0 E090 8BA9 0509 85C2

Jörn Nettingsmeier 2 years ago

...

deriving low-frequency, high-order information from tightly packed capsules (i.e. capsule distance small compared to wavelength) requires ridiculous amounts of gain. at some point, you will only be amplifying random noise, without any meaningful results. so it is important to limit the lower boundary of a high order signal set.

the nice thing about this is that low orders alone are fine for low frequencies (they provide a large sweet spot at LF), and higher orders are only really needed to extend the sweet spot at HF.

so in theory, a setup like the eigenmike has quite useful

characteristics. as it stands today, the big problem is that this bandlimiting effectively kills the biggest advantage of ambisonics, to let the listener/consumer decide the playback order.
sure, for a separated eigenmike signal set, the necessary corrections can be made for arbitrarily low orders, but as soon as you have panned higher order sounds in your b-format master, it gets messy to the point that i don't want to go there...

the same problem arises if you recklessly combine a POA microphone with higher-order spots (like i'm guilty of), but the spectral issues are a lot more manageable. plus it's way easier to keep a four-channel first order signal set separate from the panned HOA set than a 32-ch fourth-order eigenmike set.

--
Jörn Nettingsmeier
Lortzingstr. 11, 45128 Essen, Tel. +49 177 7937487

Meister für Veranstaltungstechnik (Bühne/Studio)
Tonmeister VDT

http://stackingdwarves.net

Sampo Syreeni 2 years ago

> *Post by JÃ¶rn Nettingsmeier*
> deriving low-frequency, high-order information from tightly packed capsules (i.e. capsule distance small compared to wavelength) requires ridiculous amounts of gain.

Yes. I referred to that as the A-to-B noise amplification problem within the matrix.

> *Post by JÃ¶rn Nettingsmeier*
> at some point, you will only be amplifying random noise, without any meaningful results. so it is important to limit the lower boundary of a high order signal set.

Agreed. But how precisely does that affect the definition of the signal set so that it suddenly isn't at your given order?

Suppose that you'll puncture the perceptual noise floor with amplified mic noise at 3rd order at, say, 400Hz. Why can't you just shelf it down to 2nd order on the encoding side and/or even cut out some of the material altogether, to taste?

It's true that on the decoding side you have to worry about the precise order, in order to apply your psychoacoustics right. However, any combination of transfer signals ought to be equally valid, because that signal set is always referenced in fully physical units. It's just a description of yet another physical soundfield, at whatever order it happens to be, perhaps with fields of lower resolution mixed in. In all cases it's just a perfect description of a physical soundfield upto a given order, no matter what lower order, less well-directionally defined stuff happens within it.

So, the decoder might be hosed, true. But evenso, your encoder would be performing just as it should. What you encoded would be just right and it could be archived as-is, to await for a mixed order active decoder which could handle it Dandy. Not to mention it'd prolly sound plenty good over current decoders as well: it isn't as though we mind even second order directivity at the frequencies you are talking about; in a closed space zeroth order often does the job, and even in a fully open space, first order in-phase will typically do the job.

> *Post by JÃ¶rn Nettingsmeier*
> sure, for a separated eigenmike signal set, the necessary corrections can be made for arbitrarily low orders, but as soon as you have panned higher order sounds in your b-format master, it gets messy to the point that i don't want to go there...

I believe it only gets messy if you really need to do perceptual decoding. Otherwise, you can just do a systematic or an in-phase decode and be done with it. Those approaches then work the better the higher the order you're working at. Even at 2nd order and certainly at 4th order an in-phase decode will blow the socks off most pairwise panned formats, with so called "infinite panning accuracy". That's then for literally an infinite number of sources, within the soundscape, if you're so willing to sum them in. ;)

> *Post by JÃ¶rn Nettingsmeier*
> the same problem arises if you recklessly combine a POA microphone with higher-order spots (like i'm guilty of), but the spectral issues are a lot more manageable.

If they carry the same signal, that's reckless, because of the time issues. But if you just sink a coherent POA 4-vectr within a 2+ order signal set, I believe that is just fine: the result is a perfectly valid 2nd order signal set, which only decodes suboptimally because we expect so much of the decoder to begin with. Again, there's nothing wrong with the mixed signal set. It's only that the decoder capable of optimizing for directional sources just went from a passive matrix to an active one. No biggie, in this day and age, I'd say.

> *Post by JÃ¶rn Nettingsmeier*
> plus it's way easier to keep a four-channel first order signal set separate from the panned HOA set than a 32-ch fourth-order eigenmike set.

That's why we come down to just the "usual" 25ch 4th order ambisonic set, with no vestige of EigenMike or any other transducer present. We don't care about where the signal came from or where it's going, we only care about its semantics. Which are pretty clear as they stand: a time-variant Fourier-Bessel expansion of a given order of the air pressure field about a sweet spot, about the ambient average. No channels mentioned there. ;)

...

Sampo Syreeni, aka decoy - ***@iki.fi, http://decoy.iki.fi/front
+358-50-5756111, 025E D175 ABE5 027C 9494 EEB0 E090 8BA9 0509 85C2

Jörn Nettingsmeier 2 years ago

> *Post by Sampo Syreeni*
> Agreed. But how precisely does that affect the definition of the
> signal set so that it suddenly isn't at your given order?
> Suppose that you'll puncture the perceptual noise floor with
> amplified mic noise at 3rd order at, say, 400Hz. Why can't you just
> shelf it down to 2nd order on the encoding side and/or even cut out
> some of the material altogether, to taste?

say you have a zeroth order over 10 octaves, fourth order over 4 if you're lucky, and the other orders with a bandwidth somewhere in between. let's also assume we're in the rE domain for all those higher orders, which means the perceptual result is governed by adding the energies, not the amplitudes.
now if the POA subset is spectrally flat, any further order set will add more HF energy, shifting the spectral balance. this can be compensated by using more of the lower orders at LF, with a very messy filter network, to yield a flat response, for any order N.

but then, if you truncate to order N-1, you lose more HF than LF enery, hence the spectral balance will be destroyed.

i suspect this problem is far more fundamental than the comparably minor issues you get from sub-optimal psychoacoustic decoding of normal mixed-order sets.

usual disclaimer: severely math-challenged poster here, who'd love to be disproved and become a happy eigenmike user ever after.

--
Jörn Nettingsmeier
Lortzingstr. 11, 45128 Essen, Tel. +49 177 7937487

Meister für Veranstaltungstechnik (Bühne/Studio)
Tonmeister VDT

http://stackingdwarves.net

Fons Adriaensen 2 years ago

> *Post by Sampo Syreeni*
> How does that happen? Okay, it might be that the usable frequency
> range changes, but once you've done equalizing in the A to B step,
> don't you then always get just a B-format set which you then decode
> at your target order?

As long as the truncation happens in the frequency range where you'd use a 'systematic' decode that is possible, because the decoder coefficients for the remaining lower orders don't change. This is no longer true in the 'max rE' or 'in phase' regimes.

Ciao,

--
FA

A world of exhaustive, reliable metadata would be an utopia.
It's also a pipe-dream, founded on self-delusion, nerd hubris
and hysterically inflated market opportunities. (Cory Doctorow)

Fons Adriaensen 2 years ago

...

For 4th order, the lower limit is around 1kHz, which is well above
the range where you would use a 'max rV' or systematic decoder. And
even that requires extremely accurate calibration. The problem is
that extracting the higher order signals requires amplification of
tiny differences between the capsule outputs, and even the smallest
gain differences will result in the wanted signal being swamped by
amplified errors.

For third order the limit is around 400 Hz, again assuming very
careful calibration (and stability).

Ciao,

--
FA

A world of exhaustive, reliable metadata would be an utopia.
It's also a pipe-dream, founded on self-delusion, nerd hubris
and hysterically inflated market opportunities. (Cory Doctorow)

Eric Benjamin 2 years ago

The eigenmike is supplied with a software application called EigenStudio, which allows the user to output various beams with selectable directivity. Those outputs aren't, or mostly aren't spherical harmonics. Instead, one has available cardioid, hypercardioid and supercardioid outputs up to 3rd order. Of course it is possible to derive the spherical harmonics from those.

Eric Benjamin

----- Original Message -----
From: Stefan Schreiber <***@mail.telepac.pt>
To: Surround Sound discussion group <***@music.vt.edu>
Cc:
Sent: Monday, June 17, 2013 8:52 AM
Subject: Re: [Sursound] Eigenmike (Fons Adriaensen)

For clarification: