Product Review: Smyth Research Inc. Realiser A8

Jan Didden and Stuart Yaniger

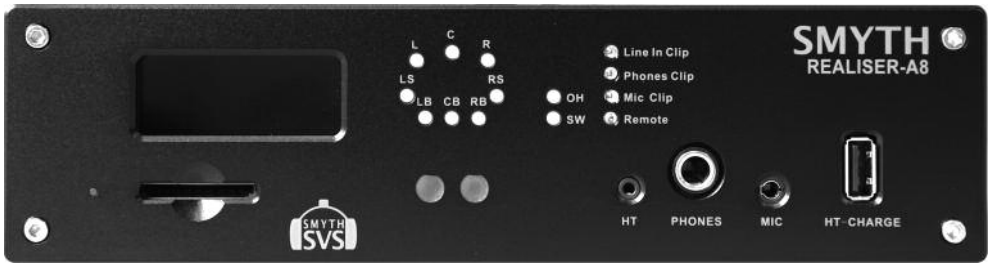

Introduction

One of us (JD) heard a demo of the A8 Realiser in Germany several years ago. The demo made quite an impression but due to other ongoing projects, follow-up was filed under 'when I next have some time to spare'. While preparing for the 2013 October series of audio events (Burning Amp, Rocky Mountain Audio Fest, Audio Engineering Society convention) the Realiser came back to the fore as an interesting piece of equipment to demonstrate and discuss at these shows. Those demos would also provide user opinions that could be integrated into a review. Smyth Research gracefully agreed to provide a unit with matching headphones several weeks ahead of the first event. This was important – the Realiser is a very unusual piece of equipment so getting familiar with it and its operational possibilities requires some time.

Background

Headphones are a convenient way to listen to music, but they have several disadvantages. One major problem is that the sound seems to be localized inside the listener's head, rather than spaced away from the listener (usually in front) as in live music or conventional stereo or multichannel. The problem seems to be exacerbated with multichannel sources. Over the years, this issue has been understood by considering the crosstalk that real acoustic sources have to the ears, including amplitude and timing, as well as frequency response due to diffraction and absorption by the head and pinnae. To solve this problem, engineers have developed a dazzling variety of signal-processing methods that are intended to provide a more natural sense of localization.

Interaural Timing Difference (ITD) and Interaural Level Difference (ILD) are known phenomena that help explain how we know something is to the left or right, but they cannot convey front/back or up/down, or range. The ITD and ILD will be identical at a particular azimuth angle whether in front or behind, above or below, near or far. It is the head-related transfer function (HRTF) that supplies the full angular localization information (azimuth and elevation). Range (distance to source) is determined by the brain from changes in the HRTF with head movement, and also to some extend using the ratio of direct to reverberant energy.

The earliest applications of these models to the issue of localization in headphones of which we are aware is that of Benjamin Bauer in the 1960s, who offered some analog methods of ameliorating the problem- Bauer's approach used some clever equalization and delay of a signal cross-fed from one channel to the other in an attempt to simulate the effect of ears in two different locations.

These analog implementations can improve the headphone "performers in your skull" problem, but in our experience, the improvement is rather limited. With the recent availability of cheap and fast signal processing in the digital domain, designers have been able to experiment with delay, EQ, and other methods with a flexibility that they have never had in the past when limited to opamps and wire. Thus, it isn't terribly surprising that software and firmware approaches have pretty much displaced analog approaches- the amazing power of computation and the ease of modification/upgrade are far too attractive to pass up.

What it does and how it does it

And this brings us to the Smyth Research Inc. Realiser A8, which incorporates the signal processing into firmware. The easiest way to describe the Realiser is probably as follows: 'a listening venue virtualiser'. The A8 provides, through its headphone replay, the perception that you are listening to the sound in your listening room, using your own audio system including speakers, with your own ears. Of course, to be able to do this, the A8 must 'know' about your sound system, your room and your ears. This information is gathered through a set of calibration and measurement runs. The A8 comes with a pair of miniature in-ear microphones (**photo 1**) which are fitted in the ear canal, and receive the sound at almost the same way your ears would receive it. It is not perfect, as the mics themselves are blocking the last fraction of an inch of the canal, but it is close. More on this later.

The calibration starts with the A8 playing some test sounds through the speakers to help the listener calibrate the replay level so that the signal to noise ratio at the in-ear mics is sufficient for the required measurement accuracy.
Next, sweeps are played through each speaker in turn. In a stereo system, the wearer of the mics is instructed to 'look left', 'look right' and 'look center'. This process allows the A8 to record the combination of speaker response, room acoustics and the user's Head Related Transfer Function (HRTF). The set of recorded data is saved as the Personal Room Impulse Response (PRIR). Finally, the user is instructed to place the headphone over the ears with the mics still in place, and test sweeps are

Photo 1 The miniature measurement mike sits on the outward side of the ear plug.

played through the headphones to record the transfer function from headphones to the ear. This latter set of data is saved as the Head Phone EQualization set (HPEQ). PRIR and HPEQ can be labeled to identify the user and/or the venue where they were recorded.

The HRTF contained in the PRIR allows the A8 to 'externalize' the perceived sound and to provide the illusion that one is not listening to headphones but rather to speakers at some specific distance and angle from the listener. HRTF refers to the way sounds are modified as they strike the upper torso, head, and especially the external ear (pinna) before the sounds enter the ear canal. These modifications are changes to the spectrum, but the brain does not hear the filter effect consciously, rather, it translates this into direction.
Because heads and pinnae vary widely, only a measurement through one's own pinnae can provide an emulation which is very nearly perfect, though one could listen to someone else's measurement for casual listening and there would be no ambiguity as to basic virtual speaker position. It just would not be precisely right, and the EQ would also be off since various heads sound brighter or duller to a foreign listener

Besides HRTF processing, there is a secondary mechanism by which the A8 helps the externalization, namely Head Tracking. A small IR transmitter is placed on the headphone headband on top of the listeners' head (**photo 2**), and a matching receiver is placed at a convenient place on the wall where the virtual speakers should be perceived. As the listener rotates his/her head, either consciously or subconsciously, the receiver detects this and causes the A8 to move the sound field in the opposite direction, thereby giving the illusion that the sound field stays in the same lateral position. This reinforces the externalization process, and the combined effect is a very convincing illusion that one is listening to his own speakers in his own room, complete with all the room influences that play a role.

Hardware Implementation

The Realiser is packaged in a cute little black box with an external "wall wart" supply. It has RCA analog inputs (unbalanced) and both RCA outputs and a ¼" jack for headphones and to provide a bypass to send the signal to the usual power-amp and speakers. The digital I/O is HDMI- no coaxial or XLR-type S/PDIF nor AES/EBU, which we would consider a major oversight- this required some major reconfiguring of SY's system, since his usual digital sources are run to the coaxial S/PDIF input of an

Photo 2 The head angle transmitter on top of the headphone band.

electronic crossover. If you're a multi-channel user or have HDMI-compatible digital sources, you're good to go all digital, but if not, you'll have to do all the I/O in the analog domain. Curiously, there is a Toslink fiber optic output- does anyone still use these?

If you use the recommended and excellent Stax headphones (which we did), the analog outputs of the A8 drive the headphone amp directly. Dynamic headphones may be plugged into the ¼" phone jack provided on the front panel.

The fun begins with the head tracking setup, which comprises a head-mounted IR transmitter and an optical angle-sensitive receiver. The receiver is mounted front and center of the listener, using a plastic clip which is big enough to clip on top of most relatively thin TV sets. The transmitter clips onto a bracket sized for the Stax headphones- other makes, be prepared to improvise. When mounted on the head, it gives the user a bit of a Teletubbies appearance. To avoid terminal embarrassment, we suggest the head tracker be used only in private or in the presence of consenting adults who promise not to tell.

Besides the Dork Factor, the main disadvantage of this implementation is that the optical receiver needs to be connected to the A8 black box via a long wire. That's always nice to have running down the middle of the living room floor. It's less of a problem for viewers of desktop or laptop video screens, but it also means that the head tracking function can only be used in a room where you can mount the receiver. This seems a curious implementation, given that tilt and angle sensing is now a commodity, being contained in most smartphones and game controllers. One could built an entire unit, including a wireless connection between the sensors and the A8 black box, to fit into the current head-mounted transmitter's form factor, and do it inexpensively. This would also allow "Realised" files to be listened to in portable situations. (On the other hand, the choice for this particular setup may have made perfect sense when the A8 was designed, five or six years ago, when integrated sensors were not generally available and rather expensive. Technology moves very fast these days).

Another key hardware piece is the previously mentioned pair of electret microphones, which are inserted into the user's ears, flush with the opening of the ear canals. These are used in two steps- the

calibration of the room, performed with the headphones off, and the calibration of the headphones, performed (obviously) with the headphones on. Yes, the process looks as dorky as the head tracker, especially for the room calibration, where the head tracker needs to be mounted on a hair band (the cosmetic kind, not the rock kind!). But fortunately, this only needs to be done once per user.

Use and documentation

The documentation seems to have been written by an engineer rather than by a skilled technical writer. It is difficult to follow, does not emphasize the important points, presents information out of order, no "quick start" guide for most common uses… think of trying to learn Microsoft Word by reading the manual! Like with Word, there are so many features, settings, and options that one can easily get lost in the thicket, and there's no guide to extricating oneself.

The A8's manual can be a valuable reference once the basic setup process is understood, but it's tough going if you want to use it to learn how to do the basics of setup. The user should take advantage of the experience of Smyth's dealers and insist on a training session and an initial setup. The default setup provided with the A8 is, to be honest, lame in the extreme. We could hear that things were different processed versus unprocessed, but neither of us got an illusion of "out of the head" imaging. Once we struggled our way through the process, we were rewarded with a stunningly real illusion- but the learning curve was unnecessarily steep.
A few hundred dollars to a good technical writer could fix this issue completely.

How does it sound?

Before going 'on tour' as it were, with the A8, we did several measurement and listening sessions at SY's residence. After several attempts, mainly due to the unfamiliar terminology and structure of the user guide, we produced PRIR and HPEQ measurements for both our ears in his listening environment. Once that was done, the A8 produced extremely convincing and accurate reproduction of the sound stage on the headphones. The remote control allows switching between speakers and headphones and we found it very difficult to decide whether we were listening to the speakers or the headphones. This may sound, on the face of it, as an exaggerated statement, but it has to be experienced to be fully appreciated. At one point at an early morning testing session, while SY's wife was still asleep, with the headphones on and the A8 switched in, a musical selection started and JD reflexively ripped off the headphones and dived for the volume control, thinking that the loud sounds were coming from the speakers! This sort of complete fool-the-ear reaction is a rare and amazing thing – except with the A8!

Still, we are both very well aware of the usual bias and other traps involved in subjective listening and wanted to do some more testing with less involved individuals. This chance came when discussing the A8 with Siegfried Linkwitz, a known expert on speakers, room acoustics and all the factors involved in providing convincing music reproduction. Mr. Linkwitz kindly invited us to bring the A8 to his residence and compare its performance with his personal, highly optimized speaker system.

After performing the PRIR and HPEQ calibration runs we proceeded to listen to Linkwitz' favorite music. We used a remote control to switch the speakers on and off while Linkwitz took the headphones on and off, to allow an instantaneous comparison[1]. Linkwitz commented that, if we timed the switchover right, he could not hear any difference in sound stage between the speakers and the headphones. He was impressed.
There was however a slight timbral difference between the two situations. According to Smyth Research, this may be attributable to the fact that the mics measure at the exit of the ear canal ("blocked meatus") while ideally one would want to measure at the eardrum. There are obvious practical problems with doing that, and a kind of 'universal correction' is also not possible because people's ear canals are often very different.

However, the timbral difference can be compensated using the "Manual EQ" method described in the A8 manual. In this mode, the Realiser generates bands of pink noise, and for each band the listener compares headphones on / headphones off. For most bands the two will be the same, but when they are not, an adjustment can be made. In this way the timbral difference can be completely removed. It is not a difficult procedure, but it is tedious. One should ask whether any slight difference is relevant, that is, is it on the order of the same model speaker made this year instead of last year, or a slightly different position in the room, etc. According to Smyth Research, most pro users find the difference, if any, irrelevant. Mastering engineers sometimes use Manual EQ and they seem satisfied with the result. Manual EQ is an adjustment to the HPEQ file, not the PRIR.

It should be noted that the above described ear canal effect is subtle and not necessarily perceptible. A large difference suggests an error in measurement, perhaps poorly located mics; we found this out the hard way when we got a bizarrely unbalanced result from the first PRIR measurement performed using SY's ears. Following advice from the manual, we retried the measurement with the microphones correctly reinserted in his ear canals, which fixed that immediately.

Casual demos

As planned, we took the A8 to various audio events that took place in October 2013. Burning Amp is a gathering of audio diy aficionados in San Francisco; the Rocky Mountain Audio Fest is a three-day event featuring manufacturer listening rooms, a parts, equipment and records/CDs market place and technical presentations, and takes place in Denver, Colorado. Finally the Audio Engineering Society held its 135th convention in New York. At each event, we demonstrated the A8 to interested participants.

At these casual demos, it was not possible to record the PRIR and HPEQ data for each listener as there was no speaker system and listening room available. The system was played with the PRIR and HPEQ recorded at an earlier session. Therefore, it was not possible to compare the headphone rendering

[1] There is an automatic function in the A8 that switches between headphones and speakers when the headphone is tilted downwards, but we did not use this at the time.

with speaker rendering, but it was still possible to judge the believability of the externalization. The head tracking system was used in each case, to help the process. The remote has the capability of switching the processing on or off, so one can instantaneously compare the A8 externalization processing with the situation of just L and R headphone stereo with its familiar 'in the head' sound field. In all cases, listeners were very impressed with the results of externalization. Generally, when we switched to 'just headphones' in the middle of a listening session, listeners reacted strongly to the collapsing sound field and were often literally speechless. Many people asked where they could buy it and what it cost[2]. While such impulsive reactions do not necessarily lead to a purchase, they do illustrate how convincing the experience is.

Surround sound use

In all above cases, the A8 has been used in 2-channel stereo mode, as we both predominantly listen to stereo. However, the A8 is clearly targeted to surround users as well. One of us did some experiments in this mode, using a very clever function in the A8 to set it up for full surround using only a single speaker! Remember that to calibrate the system for the stereo setup in a room, the user has to alternatively 'look left', 'look right' and then the sweeps are played and the response in that look angle recorded as part of the overall PRIR.

This concept can be extended to 7.1 surround in several ways. One method is to use just a single speaker, and instruct the A8 to record the PRIR successively as follows. Just look forward (speaker at 0^0 angle) for the first sweep, then move the speaker to the $+30^0$ listening angle position and repeat the sweep, continuing in 30^0 increments until a full circle has been recorded. Thus the A8 records the responses obtained as if there were a full system of surround speakers in each of those positions. And you can even do it simpler: the 'ONE' procedure allows you to just leave *one* speaker in *one* fixed position, say front center, and instead move your head to each of the called out listening angles. There is a difference between the two procedures of course: when moving the one speaker to each position, the A8 does actually record the room response in each position. When leaving the speaker stationary but changing listening angle, you can record a kind of 'artificial' room response using the preferred position in your room. Which one gives the preferred rendering depends, among other factors, on the symmetry of the room, and can even provide better replay acoustics than would be possible in the actual room.
This method to synthesize a full surround system from a single speaker has one great advantage: you can use one very good, very expensive speaker instead of having to buy 7 or 8 of them! If you own a very high quality stereo speaker system, you can now synthesize it into an equally high quality surround system!

After setting up the system with one stationary speaker as described, we proceeded to listen to the surround version of the Avatar movie DVD. Again, as experienced with the stereo tests, the surround

[2] The retail price for the A8 Realiser was $ 2910 at the time of writing. While the A8 can be used with any headphone, a Stax SR-207 with matching SRM-2525 amplifier is recommended and can be provided by Smyth Research, and would add $ 760 to the total.

rendering was totally convincing. When watching the movie, the sounds were placed in an external 360^0 sound field and matched very precisely the way the characters in the movie looked towards the sound sources.

Wrap-Up

We generally find the physical pressure and sometimes 'sweatiness' of headphone ear pieces objectionable. But when listening to the Realiser, you are unconsciously more focused on an imaginary set of speakers. The result is that you tend to 'forget' that there is a set of headphones on your head, and the objectionable physical presence recedes to the back of your conscious focus.

We also find the effect of virtualization of speakers and sound field more 'natural' if you can see an actual set of speakers, even if they are not playing. We'd like to test that also with just speaker mockups visible, or maybe just pictures of speakers on the wall. Don't laugh – perception is all about make-believe! At any rate, this effect tends to diminish when you listen more often to the A8 and get used to hearing a speaker system without one being visible.

A common meme in fashion audio circles (especially in response to skepticism about things like "sound" of amplification or wires or demagnetization of CDs) is that, "Scientists and engineers know very little about how humans hear or the details of psychoacoustic processing in the brain." Although one could suggest in vain that people who say this spend some time reviewing the hundreds of books and thousands of papers written about this very subject, perhaps it would be faster to have them spend ten minutes with the remarkable Smyth Realiser A8. Designed using solid principles of engineering and well-established understanding of human hearing, with no magic or voodoo, its function is startlingly good.
Our only reservations are with some of the ergonomic features and documentation, not with the actual concept and execution of this remarkably convincing signal processing device. If you do a lot of headphone listening, or you want to 'take your system with you', as it were, consider an audition of the Smyth Realiser Inc. A8.

Further Reading:

Benjamin Bauer, "Stereophonic Earphones and Binaural Loudspeakers," JAES, April 1961.
Benjamin Bauer, "Improving Headphone Listening Comfort," JAES, October 1965.
Louis A. Abbagnaro, , Benjamin Bauer, and Emil Torik, "Measurements of Diffraction and Interaural Delay of a Progressive Sound Wave Caused by the Human Head," JASA September 1975.
Frederic Wightman and Doris Kistler, "A Model of Head-Related Transfer Functions Based on Principal Components Analysis and Minimum-Phase Reconstruction," JASA, March 1992.
Frederic Wightman and Doris Kistler, "Headphone Simulation of Free-Field Listening" JASA, February 1989.
Stephen Smyth, Personalized headphone virtualization, US Patent 7,936,887.
JM Loomis , C Hebert , and JG Cicinelli, "Active localization of virtual sounds," JASA, October 1990.