H2020 ICT-22-2014
COCOHA
Project No.: 644732

edited by Alain de Cheveigné, 20 Apr 2016

COCOHA report on Acoustic Signal Processing for Hearing Aids – v4

This document provides an overview of acoustic signal processing as usable for a hearing aid device. It reviews the *need* for SNR enhancement, the *constraints* (technical and marketing), the *main approaches* (single channel denoising vs multimicrophone arrays), some relevant *new developments* (MEMS microphones, hand-held devices, wireless, algorithms), and some *solutions* that have emerged, both simple and sophisticated. It attempts to identify the main hurdles, and to determine if, and how, they can be overcome.

Executive summary

1. SNR improvement is the goal, acoustic scene analysis the best way to attain it.
2. The parameter space from which to choose a solution is vast. This is a strength, and also a weakness because any chosen solution is vulnerable to competition from other solutions.
3. A key decision is to associate the hearing aid with external devices (one or more). Despite downsides, this choice alone can provide significant SNR benefit.
4. The greatest benefit may come from ad-hoc distributed networks of nodes communicating by wireless links between each other and with the hearing aid.
5. Acoustic scene analysis involves applying a *multichannel filter* to the microphone streams, implemented in either time or frequency domain. There are two logical steps: (a) determine filter coefficients, (b) filter the streams.
6. Each microphone picks up both the target, and the noise source(s). The filter combines these multiple signals so as to strip the noise from the target.
7. Coefficients are derived from the data (*bottom-up* data driven analysis) and from user input (*top-down control*). It is useful to postulate two distinct modules: data-driven analysis to produce N clean streams, and top-down *selection* among them.
8. Processing works best if the array includes at least one microphone *close to the target,* and one microphone *close to each major interference source*. This is most likely to be the case for *distributed arrays*.
9. Classic algorithms usually assume a relatively compact "antenna" array of microphones for which performance is more severely constrained. It is useful to look beyond these classic algorithms.
10. The multichannel filter that suppresses the noise also affects the target which is *spectrally distorted* (spectral coloring, phase distortion, temporal smearing). However spectral distortion can be milder for distributed than for compact arrays. We should thus concentrate on interference rejection.
11. Perfect interference rejection is attained in principle in two cases: microphone close to target (in which case its signal is clean) and microphones close to all interfering sources (in which case their filtered signals can be subtracted from other microphones). Performance will be less ideal in practice, but it is expected to be better than for a compact array.

12. Major issues of wireless networks are: *latency* of transmission, *bandwidth*, *synchronization* of clocks between nodes, *flexibility,* and *resilience* to failure.
13. Wireless transmission must be *faster than sound*. Protocols such as Bluetooth have too high latency (at least ~40ms), but alternatives with lower latency are available.
14. Bandwidth requirements are massive, but distributed processing helps reduce them. Once the filter coefficients are determined, only one stream need be transmitted from each node. More streams may be required to estimate the filter, but their latency requirements are less stringent.
15. Synchronization of clocks between nodes is an issue, but solutions exist.
16. The system must be resilient to failure of nodes or links. Ideally it should be "opportunistic", capable of reaping benefits of available resources (e.g. additional nodes, hand-held devices, infrastructure). In the extreme it should also be able to fall back on the standard on/in ear hearing aid configuration.

The need

Hearing impairment is a major health issue that is becoming more severe as the proportion of elderly in the population increases. The prevalence of debilitating hearing impairment is in the range 5-8% for developed countries, but it increases very sharply with age, implying that the overall prevalence will grow as populations age. Together with objectively measured hearing impairment (loss of sensitivity), reflected in these statistics, "hidden hearing loss" results in an increased difficulty in understanding speech in noisy environments, despite a "normal" audiogram. The only way to address the problem is to enhance the *Signal-to-Noise Ratio (SNR)* at the subject's ears.

Single-vs-multichannel enhancement

Two approaches are available: *single channel signal enhancement* (the signal from a single microphone is processed based on signal characteristics specific to signal or noise), and *multichannel acoustic scene analysis* (signals from multiple microphones are combined to improve SNR based on differences across microphones).

Many single-channel signal enhancement schemes have been proposed, but the result is often disappointing. The problem is that target and noise signals are intricately mixed within the signal analysis domain (e.g. time-frequency representation), so that one cannot be removed without severely degrading the other. Typically, "ease of listening" may be somewhat improved, but intelligibility is usually not. New approaches based on Machine Learning might yield better results, but the degree of improvement (as measured in dB SNR improvement) is expected to be limited, especially at low SNR.

In contrast, multichannel acoustic scene analysis can provide much greater SNR improvement, at least in principle. For example an algorithm that cancels an unwanted source (for example Generalized Sidelobe Cancellation) can, provide infinite SNR improvement of a target relative to that source. Cancellation works by subtracting signals from one another such that the contribution of the unwanted source is set to zero.

Early devices such as the "hearing horns" of our grandparents perform a basic form of acoustic scene analysis, as do the directional microphones available in some hearing aids. However modern hardware (micromachined microphone arrays, wireless transmission, computing power) and new algorithms greatly expand the range of what can be achieved.

The appeal of acoustic scene analysis lies in (a) the magnitude of the potential SNR benefit, and (b) the degrees of freedom among which to search for a solution. The latter implies increased likelihood to find a good solution, however the diversity of solutions can also lead to confusion and a fragmented solution space.

A good review of microphone array signal processing is Bertrand (2011).

Constraints and obstacles

A technological solution faces multiple constraints related to *acoustics* (noise sources, reverberation, limits on microphone positions), *hardware* (microphone noise, processing power, bandwidth, latency, physical size), *cost* and *marketability* (the solution must be desirable for the user and profitable for the provider). An additional concern is *privacy* (acoustic scene analysis could enable evesdropping). The variety of potential solutions can be an obstacle, in that it complexifies the solution landscape, and exposes the promoter of a given solution to the risk of being superseded by some other solution.

New developments that facilitate acoustic scene analysis

MEMS technology allows large numbers of miniature microphones to be assembled in dense arrays at low cost. Silicon-based technology can potentially facilitate designs associating microphones and signal processing integrated on the same device, paving the way to ad-hoc networks including very large numbers of microphones assembled into nodes with local processing.

Steady increases in processing power (Moore's law) expand feasability limits on processing tasks, at lower power, smaller size, and with decreasing costs. Floating point support simplifies algorithm design and improves accuracy. Progress in wireless technology makes large ad-hoc microphone arrays more feasable, although latency and susceptibility to interference are issues that need addressing. The recent widespread availability of portable and hand-held devices, each equiped with microphone(s), high-power processor, wireless communication abilities, and high-level operating system functionalities, offers a potential platform for developing acoustic scene analysis ideas.

Progress is also being made on the algorithmic front, with a shift of emphasis from compact arrays of well-defined geometry, to widespread ad-hoc arrays, and from the task of localization, or beamforming based on location, to that of data-driven signal enhancement.

Major hurdles

There are several major hurdles on the path towards a useful acoustic scene analysis solution. They include *hearing-aid specific* constraints including size and power consumption, *acoustic* constraints related to sound field complexity, reverberation and noise, *implementation* constraints including wireless latency, synchronization between nodes, power supply, *product* constraints including cost and marketability.

Hearing aid-specific constraints

In a traditional hearing aid, all elements are included within a compact device behind or within the ear: microphone(s), loudspeaker, processor, and power supply. This implies stringent constraints on size, robustness with respect to the biological environment, aesthetics, power consumption, microphone geometry, acoustic attenuation between speaker and microphone, etc. If an external device is associated with the hearing aid, these constraints are relaxed, while additional constraints are added: the on/in ear device(s) must be capable of wireless transmission, and an external device is required together with all of its own constraints.

This is a crucial design choice: a *self contained* on- or in-ear device, or else an on- or in-ear hearing aid associated with an *external device* (or several). The self-contained option has major advantages, but its constraints severely limit the benefit it can provide. If one or more external devices are added the solution space becomes much wider, and a much greater benefit is attainable. A crucial question is how to evaluate the cost/benefit tradeoff between these two options. Whatever the outcome, here we assume the latter option, involving one or more external devices. This does not mean that the hearing aid can *only* work in this configuration: ideally the device should gracefully switch from one mode to the other, so that the availability of an external device is always perceived as an advantage.

Acoustic constraints

SNR can be improved acoustically in three ways: (a) by placing or selecting a microphone in a favorable position, close to the target and/or far from interference, (b) by the use of directional microphones, and (c) by combining the multiple signals from microphones within an array. Approach (c) subsumes (a) if the array includes appropriately spaced microphones. Here we consider only this approach.

The effectiveness of multimicrophone processing depends on the nature of the *acoustic field* (number of sources, reverberation, noise, movement, temporal and/or spectral sparsity, etc.), *sensor noise* and linearity (for example of AD conversion), and *sensor number and geometry*.

A diffuse field, such as results from reverberation or diffuse noise, is difficult to remove by array processing. Levin et al (2015) show that the directivity factor in diffuse noise, averaged over all look directions, is equal to the number of sensors (minus the number of nulls), which typically translates to a modest benefit in dB. Spatially localized sources can be handled more effectively if the number of microphones M exceeds the number of sources N. The presence of *reflections* increases the effective number of sources, whereas *sparsity* of their activation (in time and/or frequency) decreases it. The limit of large N corresponds to a diffuse field. Moving sources add to the complexity.

Microphone noise may be an issue, particularly for MEMS microphones. Wind noise is an issue outdoors, rubbing against clothes an issue for body-worn devices, solid-transmitted noise an issue for e.g. table-top devices. Quantization or coding noise may be an problem, particularly if cost or bandwidth constraints force to use low resolution ADCs or low bit-rate compression. Such channel-specific noise cannot benefit from multivariate processing.

The more microphones, the better. More microphones allow more spatially coherent sources to be isolated, with shorter impulse responses (Benesty et al 2007), and they allow better rejection of diffuse noise (Levin et al 2015). In a distributed scenario, increasing the number of sensors makes it more likely that at least one sensor will be close to the target, or to a major noise source that needs to be factored out.

Implementation constraints

Implementation constraints include power consumption, processing power, wireless transmission bandwidth, reliability and latency, device clock synchronization, as well packaging, size, aesthetics and reliability constraints. As pointed out earlier, technical constraints are more severe for the on/in ear device than for external devices, although multiplying the devices multiplies the problems to be solved.

Power consumption is a well-known issue for an on/in ear device. It is made more severe if wireless transmission is required, although downloading tasks to an external device can also potentially save power. The external devices too need power, and the scenario of an ad-hoc array of multiple devices is made less attractive by the prospect of having to change batteries or recharge all these devices. Power harvesting may become an option if power efficiency improves.

Transmitting multiple audio streams between multiple nodes requires a considerable *bandwidth*, although inter-stream correlation allows for efficient compression (c.f. work by Alexander Bertrand). Furthermore, once a filter solution has been derived, each node needs to transmit *only one* stream to others or to the hearing aid (c.f. work by Alexander Bertrand). Multiple streams may however need to be transmitted to calculate the inter-stream correlation parameters from which this solution is derived. Large bandwidth requirements increase the risk of *interference* from other devices that share the wireless spectrum, which in turn requires that the algorithms be robust to such failure.

A major issue in wireless audio is *latency*, due to the concatenation of stages that include digital signal processing, packet assembly, transfer and decoding. Latency is a potential problem for the user, and for signal processing. For the *user*, at the largest latencies the sound may appear out of sync with visual cues, and for shorter latencies the user might in some situations hear both direct and wireless-transmitted sound with a temporal lag. *Signal processing* may be compromised if node-to-node latency exceeds speed of sound. Specifically: the sound stream picked up by a microphone may be freed of interference from a distant source by subtracting the signal of a microphone near that source, but this works only if that signal arrives before the acoustic wave. Latency of standard Bluetooth is rather large and variable (~150 ms). Specialized versions of Bluetooth (aptX) boast a shorter latency of ~40 ms (equal to latency of acoustic propagation over ~12m). This is reported to offer tolerable audio-visual synchronization in video and games, but it is not accepted by musicians for example for wireless microphone or guitar links, and it would severely restrict the possibility of real-time processing of streams picked up over the network of microphones.

Specialized lower-latency digital protocols are available with latencies as low as 1ms (see NHK paper in Web Resources section), or 0.5 ms (see Comfort Audio doc paper in Web Resources section) and analogue transmission, FM, Near Field Magnetic Induction (NFMI), or infrared also allow low latency. A critical design choice is between a standard such as Bluetooth that is widespread and available on many devices, but severely limits processing options, and the more dedicated solutions. One option may be to use a dedicated link such as NFMI over short links (e.g. between HA and external device) together with Bluetooth for links over longer distances and communication with other devices. The issue is entwined with that of *robustness to jamming*, and *power consumption*.

Device clock *synchronization* is an issue for processing audio streams from nodes with different clocks, as the time alignment between streams may be unknown, and it may even drift with time. The issue is serious for source localization and segregation based on sensor array geometry, but less serious for data-driven segregation methods that can to some extent realign the time axes automatically. Considerable efforts have been devoted to solve the synchronization problem, either by designing algorithms insenstive to it, or by realigning the clocks based on acoustic cues, wireless synchronization, GPS, etc. It is probably safe to say that synchronization is a nuisance, but not an unsolvable problem.

With a network of nodes with wireless communication, *reliability* is a major issue. The system needs to be able to adjust to failure of a node (e.g. battery), or a link (e.g. wireless jamming), or changes in quality of acoustic signals (e.g. wind or loud sound near a microphone). This might require maintaining a list of alternative topologies, and switching smoothly between their outputs in an opportunistic fashion.

Product constraints

The technical solution must fit additional constraints such as cost, marketability, etc. A complicating issue is the variety of technical choices, parameters, and configurations that can be envisaged. Scenarii range from a single hearing aid, or pair of communicating hearing aids, to a network including also one or many remote devices, some of which may belong to the user, others belong to other users, a wider shared service such as a “smart” building, and so on. This diversity makes it hard to find a unique “sweet spot” that optimizes performance and tradeoffs, and that is immune to the emergence of other solutions that may confuse the technological landscape and fragment the market. On the other hand, creatively addressing other uses (such as "hearable" devices for the non-impaired) may widen the market and overcome some of these issues.

Control

The acoustic processor can, in principle, enhance any among a number of acoustic sources. The user needs to choose which one. COCOHA's remit is making this choice based on brain signals, but this is one among many possible control mechanisms, others being for example a graphic and/or tactile user interface, or a gesture-based controller, etc. A question of interest is whether the acoustic processing algorithm (e.g. beamformer) requires or can benefit from top-down control, or whether it is better to assume a simple one-among-N selection process at the output of a purely bottom-up process. Arguments of modularity and simplicity favor the latter modular design, but computational requirements might lead to choosing a top-down control solution. In this discussion the modular design is assumed: the acoustic processor uses only bottom-up acoustic information.

Privacy

An effective solution of the SNR enhancement problem may raise privacy issues. These may be exacerbated for solutions targeted at a wider public of "normal hearing". This is a societal / ethical question, but it might lead to certain technical choices to mitigate the problems.

A perspective on distributed microphone processing

Microphone array processing is a well-developed field, but there has been a shift of emphasis from relatively compact arrays with well-defined geometry, towards *distributed ad-hoc arrays* with unknown geometry. The concept of "beamforming" is blurred if there is no well-defined location from which to beam.

Classic algorithms assume a relatively compact array and their goal is usually to optimize the "directivity pattern" around this array so as to emphasize the target source and attenuate off-target directions (e.g. "sidelobes") corresponding to interferers and/or diffuse noise and reverberation. The notion of directivity makes less sense for a distributed array. With an ad-hoc distributed array, the geometry of the array is often not known in advance. There exist algorithms that allow the array to be "calibrated" (positions, gain, etc.) automatically, however knowledge of spatial positions may actually not be necessary if the acoustic analysis algorithms are signal-driven. Classic algorithms put equal emphasis on maximizing interferer suppression and minimizing target distortion. However, as argued below, target distortion is less of a problem for distributed arrays, for which the emphasis can be put more squarely on interferer suppression.

Whatever the method, acoustic scene analysis involves applying a *multichannel filter* (typically FIR) to the signals from the microphones. The analysis filter can be applied in the time domain by adding microphone signals with appropriate delays and coefficients, or in the short-term Fourier domain by applying the equivalent transfer functions.

In the time domain:

$$y(t) = \sum_{k,\tau} g_{k,\tau} x_k(t-\tau)$$

where y(t) is the filtered signal, $x_k(t)$ are the microphone signals, and τ is delay. The $g_{k,\tau}$ are the filter coefficients. Each microphone signal itself is related to acoustic sources $s_j(t)$ by:

$$x_k(t) = \sum_{j,\tau} h_{j,k,\tau} s_j(t).$$

The coefficients of the source-to-microphone impulse response $h_{j,k,\tau}$ reflect the room acoustics and source and microphone positions, and the coefficients of the microphone-to-output impulse response $g_{k,\tau}$ are determined by the acoustic scene analysis algorithm. The aim of the algorithm is to approximate:

$$f_j = \sum_k g_k \circ h_{j,k} = 0 \text{ for } j \neq j_0 \text{ (perfect rejection)}$$

$$f_{j_0} = \sum_k g \circ h_{j_0 k} = \delta \text{ (pure delay)}$$

where j_0 is the index of the target source.

In general these objectives cannot all be attained together, in particular there is a tradeoff between interference rejection and *spectral distortion*. For example if the g_k are chosen to minimize interference, the target-to-output impulse response f_{j0} is usually not a pure delay, implying spectral distortion. With compact arrays, the distance (and thus the gain) between source and microphone is roughly equal across microphones, and this tends to produce severe spectral distortion. Distortion can be milder for distributed arrays.

As a simple example (2 sources, 2 microphones, anechoic propagation), if each microphone has equal gain from both sources, the interferer s_1 is cancelled if:

$$y(t) = x_1(t) - x_2(t-\tau_a)$$

where τ_a is the delay that compensates for the difference in propagation delay from s_1. We then have:

$$y(t) = s_2(t) - s_2(t - \tau_b)$$

where τ_b equals τ_a augmented by the difference in propagation delay from the target s_2. The recovered signal is thus related to the target by a filter with deep zeros at frequencies n/τ_b, n=1, 2,... (comb filter). In contrast if gains are unequal the output is instead:

$$y(t) = s_2(t) - \alpha s_2(t - \tau_b)$$

where α is the ratio of gain ratios, $\alpha \neq 1$. For ratios different from 1 the notches of this filter are shallower and its effect mild. For distributed arrays the source-to-microphone gain ratios are likely to be different from 1.

Assuming similar trends for more complex scenarii (>2 sources, >2 microphones, non-anechoic propagation), spectral distortion should be milder for distributed than compact arrays. In any case spectral distortion has relatively mild perceptual effects. Thus for distributed arrays we can *ignore spectral distortion of the target and focus purely on interference rejection*.

To understand the potential benefit of a distributed array, it is worth considering two limit cases. In the first, a microphone is colocated with the target: SNR is infinite and processing consists merely in choosing the right microphone. In the second the target microphone is less favorably located and picks up sound from one interfering source, but a second microphone is colocated with that source. In this case, processing involves (a) estimating the transfer function between these two microphones and (b) applying this transfer function to the signal of the second microphone and subtracting it from the signal of the first. In both cases, unlimited SNR is obtained with negligible distortion. This ideal situation will be approximated more or less faithfully depending on the actual layout of microphones.

If a microphone is proximal to an interferer, the impulse response between it and the target microphone(s) can be estimated, and the interference removed perfectly. This holds (in principle) whatever the length and complexity of that impulse response. If the microphone is distant from the source, the interferer-to-microphone impulse response must be estimated and inverted, which typically requires an analysis filter with long impulse responses, although the presence of multiple microphones eases this requirement (MINT theorem, see Benesty et al 2007).

Interference that comes from spatially localized sources can be removed in this way, at least in princple. Interference that cannot be localized (e.g. diffuse noise) cannot be removed in this way. SNR improvement for diffuse noise hinges on the proximity of at least one microphone to the target, possibly augmented by delay-and-sum effects allowed by having *several* microphones close to the target.

In all cases, the theoretical benefits of array processing are contingent on the ability to *estimate* the optimal coefficients of the analysis filter, which may be difficult due to the complexity of the acoustic scene, and non-stationarities.

A general approach for estimating the analysis filter is to apply a set of M time delays (or else a filterbank with M channels) to each microphone signal, resulting in a multichannel signal with MN channels, where N is the number of microphones. Linear techniques such as PCA or Joint Diagonalization (a.k.a DSS or CSP) are then applied to find linear transforms with specific properties. Linear transforms of time-shifted signals are equivalent to FIR filters, and thus by choosing the appropriate analysis method and criteria we may find the desired analysis filter. For example, the null space of the PCA transform corresponds to the set of filters that cancels *all* sources.

The approach is applicable to non-stationary analysis (for example to take advantage of temporal sparsity of interfering sources) by calculating the covariance matrix (from which the transforms are derived) on restricted temporal intervals. Interval boundaries can be "discovered" by appropriate clustering or segmentation algorithms. Various normalization and whitening schemes can be used to tune this process.

In summary, distributed microphone arrays promise more flexibility and better performance than classic compact arrays. Linear subspace methods are one approach to take advantage of this flexibility.

Source selection and rendering

For simplicity the previous discussion assumed a single output. The same processing can be applied to each of N sources to produce N output streams from which the user can choose. The desired source (or sources) can then be *rendered* at the user's ears, possibly on the basis of spatial information gleaned from the microphone array. Rendering can include the target source, possibly together with some combination of the background sources to provide context or to allow attentional switching.

Rendering is distinct from acoustic scene analysis: the goal is *not* to use the user's (residual) binaural stream segregation abilities to further denoise the acoustic scene. It is also *not* to make acoustic scene analysis "transparent" to spatial cues. Natural spatial cues are likely to be greatly degraded, both by propagation and reverberation, and by acoustic processing. It is not fruitful to try to conserve natural cues if reliable artificial cues can be synthesized.

Distributed processing

A problem with distributed microphone arrays is the *bandwidth* of the many signals that need to be transmitted to the processing node. Transmission costs are exacerbated in the case of a wireless network. This motivates *distributed processing*, by which subsets of microphones are processed within the node to which they are attached, before wireless transmission. Instead of the full set of signals, a smaller number of processed signals is sent to other nodes and/or to the hearing aid. Within-node processing operates on the signals of microphones attached to the node, together with signals from neighboring nodes. A useful concept is that each source is "owned" by the node that "sees" it with best SNR (see work of Sharon Gannot & others), and cleaned by that node and made available for other nodes. Within-node processing can benefit from signals transmitted from those nodes that own the sources that contaminate its own microphones. For this to work well, the transmitted signals must arrive no later than the acoustic signals, and thus *latency* of wireless transmission is a critical factor. Wireless must be *faster than sound*, which is all the harder as the number of streams is large.

A useful observation is that, once the analysis filter has been determined, it is only necessary to transmit *one* stream from each node (see work of Alexander Bertrand & others). This follows from linearity of the filtering process (similar to Kirchoff's law). It may be nonetheless be necessary to exchange other signals between nodes to estimate the between-microphone correlation structure from which the analysis filter is derived. Those additional signals do not share the same latency constraints, and some degree of downsampling or compression might be acceptable.

Flexibility and resilence

The distributed network scenario is diverse. The complexity and topology of the network can vary widely, depending on the design, and on which resources are available at each moment. Processing algorithms should be able to take this into account, drawing in any resource that happens to be available, and they must also be resilient to sudden failure of a node or link (e.g. wireless dropout, or noise affecting a microphone). A general approach is to implement multiple solutions (with & without each node or link) and transition smoothly from one to another when conditions vary. At the extreme, the network can revert to only the hearing aid. This requires reliable and fast failure detection mechanisms.

The ideal building block

To experiment with these ideas, the ideal building block is a device with:

- multiple audio inputs (and/or multiple microphones)
- wireless communication with other nodes to transmit audio and exchange control information
- ability to apply a multichannel filter to streams from local microphones and neighboring nodes
- computational ability to calculate statistics (cross-correlation, etc.) on audio streams

This building block is designed to handle any node of the network, following the concept of “made for all” (MFA). In filtering mode each node needs to transmit a single audio stream to the hearing aid (possibly hopping from node to node) with low latency. In coefficient estimation mode, the nodes need to exchange more streams to allow cross-correlation coefficients to be calculated. Latency requirements are less severe in this case.

Hearing Aids to Hearables

An interesting new development is the emerging concept of “hearable”, an augmentative device targeted at the normal hearing. This can impact hearing aids by (a) creating a new, wider market for technologies essential for hearing aids, (b) overcoming stigmarelated obstacles by marking the device as “cool”. See for example the Bragi-Starkey merger, or the SoundHawk.

Some existing solutions (including low-tech)

- *Induction loops*/telecoils are highly effective for delivering a high SNR signal, but they depend on appropriate infrastructure and are useful mainly for broadcast signals.
- *Directional microphones* are a common feature in hearing aids. Limitations include relatively modest SNR improvement (~3 dB) and the need to switch between directional and non-directional modes. Other problems include spectral distortion, wind noise sensitivity, etc. Direction control is usually by moving head. An interesting combination is directional on one ear, omnidirectional on the other.
- *Binaural directional hearing aids* (e.g. Siemens Insio) use wireless communication between devices at both ears to perform 2-microphone beamforming.
- A *wireless neck loop* allows a standard hearing aid to pick up audio transmitted by a wireless protocol. Allows connectivity with a wider range of devices (e.g. cellphone), avoids having to include wireless capabilities in the HA.
- Personal *Hearing Amplifiers* pick up sounds from a microphone, send them to HA, earphones or to a neck loop via a wire or wireless. If the microphone is directional the device may be pointed (a simple solution to the control problem!).
- Remote wireless microphones (e.g. Comfort Audio) worn by speakers allow them to be heard by the HA user.
- *Wide Area Assistive Listening Devices*

References

- J Benesty, J Chen, Y Huang, J Dmochowski (2007) On microphone-array beamforming from a MIMO acoustic signal processing perspective, Audio, Speech, and Language Processing, IEEE Transactions on 15 (3), 1053-1065
- Bertrand, A (2011) "Applications and trends in wireless acoustic sensor networks: a signal processing perspective" *Proc. of the IEEE Symposium on Communications and Vehicular Technology (SCVT)*
- D.Y. Levin, **E.A.P. Habets** and S. Gannot (2015) On the average directivity factor attainable with a beamformer incorporating null constraints, IEEE Signal Processing Letters, Vol. 22, No. 11, pp. 2122-2126.

Web-based sources

This list includes a wide and diverse range of resources, both technical and of wider interest (blogs, company sites, etc.). They cover a variety of topics, such as microphone and low-latency wireless technology, multimicrophone algorithms, new concepts and consumer trends, etc. They are not ordered.

- A set of PDFs relevant for acoustic scene analysis: https://www.dropbox.com/sh/mu0ccnfaeutr2dq/AAAJs5SzGsBeOC2InYfGGcx9a?dl=0
- Wireless synchronization over dedicated low-latency ISM channel: http://www.3daudiosense.com/blog/3d-audiosense-beaglebone-black-cape
- Wireless synchronization over ISM: http://wilma.kug.ac.at/index.php?id=15672
- National Instruments paper on synchronization: http://www.ni.com/white-paper/11369/en/. Presentation on synchronization via GPS: https://sem.org/PDF/Veggeburg_Advanced%20Wireless%20Architectures_%20Notes.pdf
- Acoustics-based synchronization: Hon, T.-K., Wang, L., Reiss, J. D., & Cavallaro, A. (2015). Fine landmark-based synchronization of ad-hoc microphone arrays. *Eusipco*, 1331–1335. http://doi.org/10.1109/EUSIPCO.2015.7362600
- Synchronization by sending system time-stamp to ADC input: Lienhart, R., Kozintsev, I., & Wehr, S. (2003). Universal synchronization scheme for distributed audio-video capture on heterogeneous computing platforms. *ACM Multimedia*, 263–266. http://doi.org/10.1145/957013.957067.
- A perspective on MEMS microphones: http://www.memsjournal.com/2015/07/mems-microphones-emerging-technology-and-application-trends.html.
- Discussion on Bluetooth latency: https://www.reddit.com/r/oculus/comments/21lsq5/stay_away_from_bluetooth_headphones/
- Sub-2.3 ms guitar jack: https://www.nordicsemi.com/eng/News/News-releases/Product-Related-News/World-s-lowest-latency-wirelesss-guitar-jack-is-superior-to-a-wired-link
- 5.5 ms wireless audio links: http://www.creative.com/emu/products/product.aspx?pid=18609
- RTX design: ~4 or ~8ms latency wireless audio: http://www.rtx.dk/Wireless_Audio-4069.aspx
- Sony digital wireless microphone, ~3.6 ms: http://www.sony.fr/res/attachment/file/95/1193315636495.pdf
- Near-field magnetic induction: http://www.nxp.com/products/interface-and-connectivity/wireless-connectivity/nfmi-radio-solutions/nfmi-radio-for-wireless-

audio-and-data-streaming:NXH2280UK, http://phys.org/news/2015-02-ultra-low-power-radio-transceiver-enables.html

- NHK paper on low-latency (1 ms) digital radio microphones: http://www.nhk.or.jp/strl/publica/bt/en/fe0056-4.pdf
- Comfort Audio low latency (<0.5 ms) wireless: http://www.comfortaudio.com/wp-content/blogs.dir/1/files_mf/ssttechnicalinformationeng11102527.pdf
- Kleer (rival to Bluetooth): http://ww1.microchip.com/downloads/en/DeviceDoc/Kleer_vs_BlueTooth.pdf
- Sceptical comments on wireless audio: http://electronics.stackexchange.com/questions/27280/wifi-limitations
- Impulse ~5ms, used for ear-to-ear: http://www2.imec.be/content/user/File/NEW/Research/Wireless%20Communication/Digital%20baseband/FEB%202014%20IMPULSE%20RADIO%20ULTRA_WIDEBAND.pdf
- Microsemi white paper on ultra low power wireless body networks: http://www.microsemi.com/document-portal/doc_view/127466-white-paper-ultra-low-power-short-range-radio-transceivers
- Phonak FM: http://www.phonak.com/com/b2c/en/products/fm/receivers/microlinkfreedom.html
- Siemens paper on acceptable delay for hearing aids: https://media.sivantos.com/siemens-website/media/2014/07/2009_Hearing-aid-delay.pdf
- Apple Bluetooth earphones? : http://9to5mac.com/2016/01/08/iphone-7-wireless-headphones-beats/, hearing aid? : https://www.hearingtracker.com/blog/apple-hearing-aid/
- Latency on android: http://www.androidpolice.com/2015/11/13/android-audio-latency-in-depth-its-getting-better-especially-with-the-nexus-5x-and-6p/
- Denant (Oticon) invests in wireless: http://hearingmojo.com/william-demant-invests/
- Guest article on wireless at the Hearing Blog: http://thehearingblog.com/archives/5034
- Pros and cons of directional microphones: http://www.audiologyonline.com/articles/directional-microphone-patterns-they-also-1294.
- Review of directional hearing aids: http://www.hearingreview.com/2013/11/the-evolution-of-directionality-have-developments-led-to-greater-benefit-for-hearing-aid-users-2/, Sync or stream: http://www.hearingreview.com/2013/05/it-s-sync-or-stream-the-differences-between-wireless-hearing-aid-features/

- Directional microphones provide objective benefits, but not user satisfaction: http://www.rehab.research.va.gov/jour/09/46/5/Gnewikow.html
- Siemens in-the-ear directional binaural hearing aids: https://www.bestsoundtechnology.com/pro/news/news/the-new-insio-hearing-aid-family/
- Action on Hearing Loss shop: http://www.actiononhearingloss.org.uk/shop.aspx
- Bluetooth neckloop: http://www.actiononhearingloss.org.uk/shop/cm-bt2-bluetooth-neckloop-product-m492.aspx. Other stuff from Action on Hearing loss: http://www.actiononhearingloss.org.uk/supporting-you/products-and-equipment.aspx
- Comfort Audio (wireless microphones, induction loops): http://www.comfortaudio.com/for-users/finding-the-right-hearing-product/product-overview/
- A wok as a directional microphone: https://ronvickery.wordpress.com/2015/07/29/an-ald-that-doubles-as-a-wok/
- IP dispute between GN Resound and William Denant (Oticon) on wireless and other technologies: http://www.hearingreview.com/2016/01/gn-resound-william-demant-settle-patent-disputes/?ref=cl-title, http://www.hearingreview.com/2015/10/danish-drama-gn-wdhs-patent-war/.
- Bragi: http://www.bragi.com, partnership with Starkey: http://www.hearingreview.com/2016/01/starkey-partners-bragi-new-hearable-technology/, https://www.hearingtracker.com/blog/bragi-dash-for-hearing-loss/
- SoundHawk: https://www.hearingtracker.com/blog/soundhawk-review-hearing-loss-expert-analysis/, http://www.soundhawk.com
- Hearables: http://www.nickhunn.com/hearables-the-new-wearables/
- Jacoti: https://www.jacoti.com, https://www.hearingtracker.com/blog/jacoti-and-the-consumerization-of-hearing-technology/
- Cynthia Compton Conley on wide area assistive listening devices: http://www.comptonconleyconsulting.com/blog/how-to-hear-better-in-movie-theaters-lecture-halls-and-in-other-public-areas, and wireless hearing aid systems: http://www.comptonconleyconsulting.com/blog/wireless-hearing-aid-systems-how-to-connect-to-everything-almost
- Hearing loops: http://www.hearingloop.org
- HHTM on the importance of an open wireless connectivity protocol http://hearinghealthmatters.org/hearinprivatepractice/2015/the-importance-of-an-open-wireless-connectivity-protocol-for-hearing-assistance-products/
- ClearOne: http://www.clearone.com/beamformeraudio
- Shifting sands of Hearing Aids and Hearables: http://hearinghealthmatters.org/hearingeconomics/2015/distinguishing-between-hearing-aids-and-hearables/
- Samsung poised to enter the fray: http://hearinghealthmatters.org/hearingnewswatch/2015/samsung-hearing-aid-rumors-heat-up-reports-claim-company-nearing-official-launch-1227/, Earcle: http://hearinghealthmatters.org/hearingnewswatch/2016/samsung-earcle-hearing-aid-hearable-psap-data-leak-0211/,

http://anewdomain.net/2016/02/10/samsung-earcle-fcc-docs-suggest-samsung-hearable-hearing-aid-plans-exclusive/, Bluetooth HA FCC test results: http://www.scribd.com/doc/298890969/Samsung-Bluetooth-Hearing-Aid-Test-Results-FCC-Application#scribd

- TipTalk (Samsung spinoff): http://www.innomdlelab.com
- Hearing Aid Research Lab: http://www.harlmemphis.org/index.php/about/
- Ear Machine (smartphone as control device): http://hearinghealthmatters.org/hearingeconomics/2013/smartphone-control-apps-continued/
- Williams Pockettalker: http://www.harriscomm.com/williams-sound-pocketalker-ultra-sound-amplifier-heavy-duty-folding-headphone.html. Williams Sound: http://www.williamssound.com
- Soundfocus: http://soundfocus.com
- Kim Cavitt's page has lots of useful pointers: http://hearinghealthmatters.org/hearinprivatepractice/2016/hearing-aid-industry-disruption-2016-gosh-i-hope-so/
- Holly Hosford-Dunn's blog on patents (up to Feb 2015): http://hearinghealthmatters.org/hearingeconomics/2015/us-hearing-aid-device-patent-actvity-for-january-february-2015/ (up to April 2015): http://hearinghealthmatters.org/hearingeconomics/2015/us-hearing-aid-device-patent-activity-for-march-april-2015-2/
- Wired article on Hearables: http://www.wired.com/insights/2014/10/hearables/
- Mutual mobile post on energy harvesting: http://mutualmobile.com/posts/hearables-energy-harvesting
- Earin wireless earbuds review: http://www.theverge.com/2015/11/18/9755726/earin-wireless-earbuds-review
- Doppler Lab's Hear buds: https://www.hereplus.me/
- Wired says "ditch the headset": http://www.wired.com/2009/07/by-headset/
- A Google hearable: http://www.androidheadlines.com/2016/02/google-project-aura-hearable-may-debut-at-google-io.html
- Vapor-hearable: http://arstechnica.com/gadgets/2015/03/indiegogo-hearing-aid-campaign-offers-refunds-after-bait-and-switch/
- Mind-reading wearable: http://www.wareable.com/wearable-tech/looking-inside-the-brain-2219
- Apple AirPods?: http://www.healthyhearing.com/hearing-aid-manufacturers
- Jason Galster on wireless for HAs: http://www.hearingloss.ca/articles/a-new-method-for-wireless-connectivity-in-hearing-aids/
- Jenny Grothe (GN Resound) on wireless: http://www.audiologyonline.com/interviews/interview-with-jenny-groth-m-1341, 5 myths: http://www.hearingreview.com/2010/12/five-myths-about-digital-wireless-hearing-aid-technology/, John Nelson (GN Resound): http://www.csd.jmu.edu/symposium/Nelson%20Ruth%20Symposium%20.pdf

- Donal Schum (Oticon) on wireless: http://www.oticonusa.com/~asset/cache.ashx?id=20464&type=14&format=web
- Onsemi AYRE SA3291 NFMI system for hearing aids: http://www.onsemi.com/pub_link/Collateral/SA3291-D.PDF
- NXP's NFMI chip: http://www.onsemi.com/pub_link/Collateral/SA3291-D.PDF
- Rena Appleby (HIS convention) on the benefits of wireless: http://www.oticonusa.com/~asset/cache.ashx?id=20464&type=14&format=web
- Oticon streamer: http://hearingmojo.com/oticon-made-for-iphone-streamer/
- Wired on the Braggi Dash: http://www.wired.com/2016/01/bragi-dash/
- The Atlantic on wearables: http://www.theatlantic.com/technology/archive/2015/02/what-my-hearing-aid-taught-me-about-the-future-of-wearables/385145/
- Sonos wireless streaming network: https://sonos.custhelp.com/app/answers/detail/a_id/126/~/information-about-sonosnet
- Braggi's not quite up to it: http://documentally.com/2016/02/11/the-dash-from-bragi/
- SWARMs of microphones: http://www.esat.kuleuven.be/sins/.
- The EU Bank gives a 75Meuros loan to GM Store Nord for wireless: http://www.eib.org/infocentre/press/releases/all/2016/2016-013-eib-supports-gn-store-nord-s-hearing-aid-innovation.htm
- Devices based on induction loops: http://www.ovalwindowaudio.com/index.htm