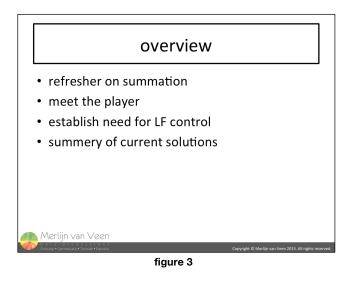figure 1

This is a transcription of my 40-minute presentation on low frequency control during the 59th AES conference on sound reinforcement in Montreal (figure 1). Needless to say that 40 minutes is barely enough time to scratch the surface on such an involved topic. I therefore decided to focus on establishing the need for low frequency control, something I felt I could do within 40 minutes, and not attempt to cover each intricacy of the various solutions at our disposal. I would very much like to thank Arthur Skudra for reviewing and editing this transcription.
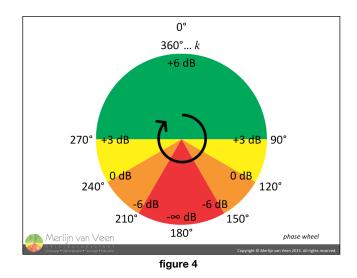


figure 2

I would like to thank Nathan Lively for talking me into this, the people of the AES, specifically Ben Kok and

Peter Mapp for liaising and finally my self-appointed mentors: Bob McCarthy and Mauricio Ramírez.



figure 3

During this primer we'll start (figure 3) by briefly refreshing on summation to make sure we're all on the same page. Next we'll meet the player commonly referred to as a subwoofer. We'll continue with establishing the need for low frequency control, the essence of this introduction and finally have a look at current typical solutions.



figure 4

This is the phase wheel (figure 4). It shows us the amount of summation or cancellation when adding two correlated signals of equal magnitude generating the same frequency with a phase offset.
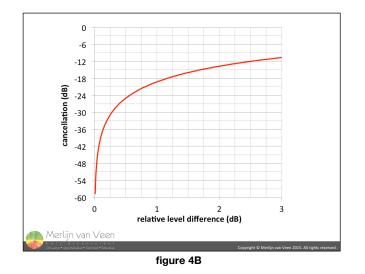
Correlated signals are characterized by a causal relationship such as the one between a parent and a child. Evidently if we play Beethoven's "5th" and AD/DC's "Back in black" side by side at equal loudness, the increase in level won't be 6 dB. However if we play Beethoven's "5th" twice side by side, it won't sound twice as long, but 6 dB louder. And if we play the same piece twice in series, it will sound twice as long but not louder. But what about every iteration between 100% and 0% phase overlap?

In the case of elementary sinusoids, the basic building blocks of (audible) audio, the phase wheel shows the progression for two signals of equal magnitude over phase offset as long as both signals are present. Most notable values are 0° of phase offset which results in a 6 dB increase, 120° and 240° of phase offset mark the break even points where nothing is gained or lost and finally 180° of phase offset which causes "perfect" cancelation.

This trend repeats itself over $k$ cycles as long as both signals are present and maintain their magnitude. Two thirds of all possible outcomes result in summation of as much as 6 dB whereas one third of all outcomes result in cancellation of as much as minus "infinity".



**figure 4B**

I would like to point out though that in the real acoustic world, cancelation in excess of 30 dB at a 180° of phase offset is considered a good result and requires 0,25 dB accuracy or less (figure 4B).

A bigger difference in relative level offset at 180° of phase offset will significantly reduce cancellation.
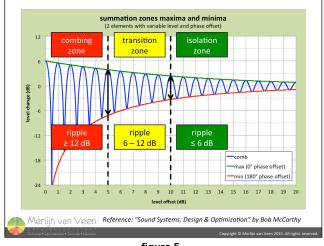


**figure 5**

When level offset is added on top of phase offset (figure 5), we're in desperate need of simplification. Bob McCarthy's summation zones [1] distill all iterations between relative level and phase offset down to 3 possible outcomes.

The green line in the chart shows all instances of 0° of phase offset over different level offsets. The red line shows all instances of 180° of phase offset over level offset. These lines represent best and worst case scenarios respectively. The values in between are the result of all remaining amounts of phase offset between those two extremes over level offset. Like the phase wheel, this trend repeats over $k$ cycles as long as both signals are present.

The difference between the maxima and minima is known as "ripple". It's a metric to indicate the scope of change in level by comb filtering that will occur whenever equal timing between multiple signals or equidistance to multiple sources can't be kept. It applies to both the electrical and acoustical domain.

Relative phase offset will determine which frequencies will be affected and relative level offset the scope of change in level.

---

[1] Sound Systems: Design and Optimization 2nd ed. by Bob McCarthy, Focal Press

Small time offsets only affect high frequencies. It takes proportionally bigger time offsets to affect lower frequencies. Relative level offset comes second or first depending on how you look at it and determines the scope of change in level caused by the temporal misalignment.

What the chart shows us is that regardless of phase offset, ripple in the order of 12 dB or more is to be expected for relative level offsets of 5 dB or less. This range is called the combing zone. On the opposite end of the chart for relative level offsets of 10 dB or more, ripple in the order of 6 dB or less is to be expected and the audibility of comb filtering will be negligible. This area is called the isolation zone. The area in between is called the transition zone and predicts ripple in the order of 6 dB to 12 dB.

In other words, level offset is the steering mechanism in reducing the audibility of comb filtering that's going to occur, whether we like it or not, when more than one signal or source is used.
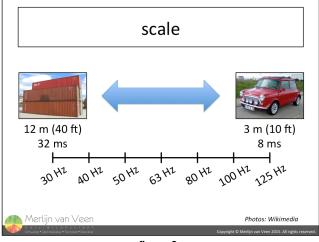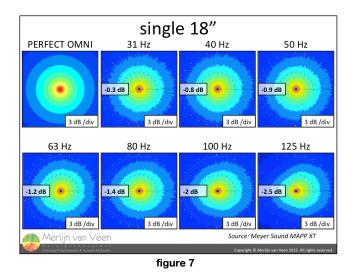


**figure 6**

A typical subwoofer has an operational frequency range of 2 octaves spanning form 30 Hz to 125 Hz (figure 6). To put into perspective the scale we're dealing with, imagine wavelengths ranging from a 40-foot long intermodal shipping container to a 10-foot long Mini Cooper respectively. This order of magnitude makes these frequencies hard to absorb and control. They abide to the same rules of physics but the solutions often become proportionally big.

According to popular belief, subwoofers are thought to be omnidirectional which they technically are, because the -6 dB criterion used for determining coverage angles when it comes to omnidirectivity is arguably the loosest specification in pro audio. This criterion allows for a 2:1 ratio of omnidirectivity, ranging from -6 dB to 0 dB in the back and if you want to split hairs, even +6 dB in the back.



**figure 7**

Close inspection however of i.e. a single 18" driver (figure 7) shows front to back level differences of as much as 3 dB due to self rejection and pistonic behavior. This might sound trivial, but these differences matter when multiple subwoofers are arrayed together. Especially in inverted stack gradient configurations where the rear facing subwoofer(s) must be matched in level to the front facing subwoofer(s) to achieve optimal cancellation.

**figure 8**

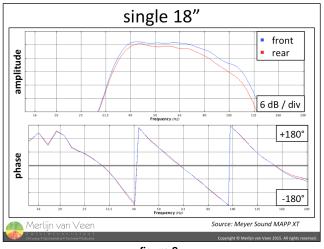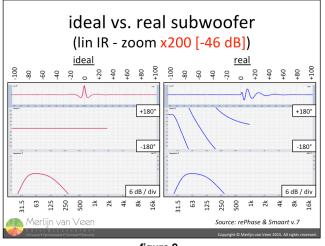Transfer functions (figure 8) in the frequency domain tell a similar story.



**figure 9**

In the time domain I would expect the linear impulse response of an ideal subwoofer on my analyzer to look like the top plot in the left hand screen capture of figure 9, a perfectly symmetrical response. This is underscored by the flat phase response in the plot below. I imagine this behavior would deliver the most impact. Like dropping a bag of sand all at once on top of somebody instead of cutting a hole in the bag and slowly pouring the contents over him. The latter metaphor however, has closer resemblance to the behavior of a real subwoofer, illustrated in the right hand screen capture of figure 9.

A typical bass reflex (ported) design will exhibit 4[th] order high pass filter behavior at the tuning frequency (between port and driver[s]) by mechanical-acoustical design alone. Combined with electronic (IIR, minimum phase or analog) high and low pass filters to protect the subwoofer from physical damage and band limit it to it's intended operational range, a typical subwoofer can easily exhibit 720° and 270° of phase shift within a 2 octave interval at the lower and higher frequency limits respectively. This is why within the pass band of most subwoofers different frequencies arrive at different times.
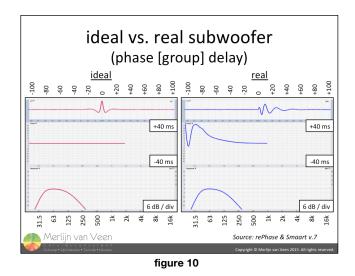


**figure 10**

Figure 10 shows the same comparison but instead of phase, phase (group) delay is shown. The ideal, time invariant, subwoofer on the left shows no phase (group) delay, whereas the real subwoofer on the right is stretched out over more than 35 ms throughout its operating range. You could also say that the virtual depth of the driver(s) or cabinet is stretched out over a distance of 12 meters (40 feet).
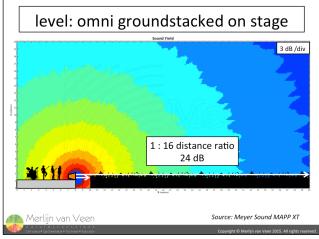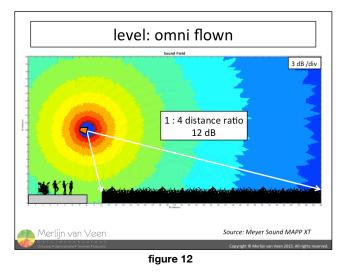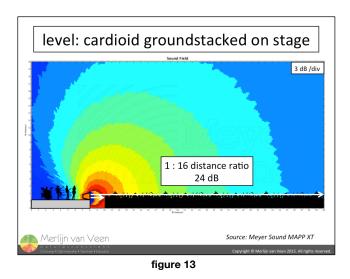
figure 11



figure 12

Now that we've observed the basic properties of a typical subwoofer let's establish the need for low frequency control. We'll start by looking at level first.

When we regard a subwoofer as "omnidirectional" it will be pretty much immune to rotation. We can orientate it any way we like, but it will make little to no difference. We're solely at the mercy of the inverse square law. So when we stack the omnidirectional subwoofers on the floor or stage (figure 11), depending on the depth of the venue, the audience will experience far less than acceptable front to back level differences due to the distance ratio. This practice will also introduce the maximum amount of low frequency stage wash.

Moving the subwoofer up in height (figure 12), away from the front of the audience, improves the front to back distance ratio considerably and reduces the level variance. Simultaneously the amount of stage wash due to the increased distance is also reduced.



figure 13

A cardioid subwoofer or configuration on the floor or stage (figure 13), with at least a 60 cm (2 feet) air gap in front of or besides the stage, will significantly improve the situation on stage, but unfortunately not in the audience that resides within the propagation plane on axis to the subwoofer.
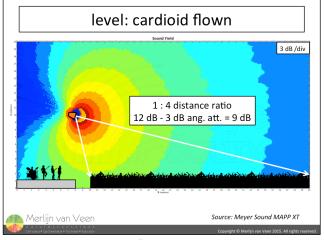
**figure 14**



**figure 15**

Moving the cardioid subwoofer or configuration up (figure 14) will provide the same advantages as the flown omnidirectional subwoofer with the added bonus of angular attenuation because the former is no longer immune to rotation due to its inherent directivity pattern.

In terms of preventing environmental noise pollution omnidirectional subwoofers are evidently a poor choice over cardioid subwoofers or configurations.

I think it's safe to assume that, even in 2015, the stereo subwoofer setup is still the de facto standard. So let's look at the relative level offset between two subwoofer positions using omnidirectional subwoofers and its effect on the spatial distribution of tonal and ripple variance. After all the most prominent reflection is the other subwoofer.
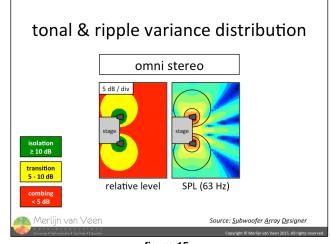
The left plot in figure 15 shows a plan view of the spatial distribution of the summation zones discussed earlier. Because of the omnidirectional subwoofer's inherent immunity to rotation, areas of isolation can only be created by brute force. The majority of the audience falls within the red combing zone with ripple in excess of 12 dB, except for the audience members along the geometrical centerline of the venue, equidistant to both subwoofer positions. The minority in the yellow transition zones will experience ripple in order of 6 dB to 12 dB and only a few are "saved" by brute force in the green isolation zones, extremely close to either subwoofer position. Over there, the ripple will be less than 6 dB.

The right plot in figure 15 reflects this subdivision of the audience. The contrast in sound pressure level between the power alleys and valleys is apparent and substantially larger within the red combing zone in comparison to the transition and isolation zones where relative level offset comes to the rescue. The first power valleys off-center will manifest themselves in the frequency response as 1-octave wide (50% of the operational bandwidth) gaps. The second power valleys as 1/3-octave wide (17% of the operational bandwidth) gaps.
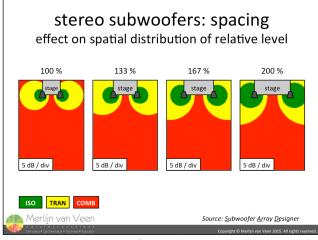
**figure 16**



**figure 17**

Increasing the physical spacing between the subwoofer positions (figure 16), even by as much as 200%, does very little. The majority of the audience will remain within the red combing zone.

The final arguments in favor of low frequency control will focus on indoor reverberation. Typical medium format speakers are quite capable of maintaining their nominal coverage pattern from 1 kHz and up. Larger format speakers are capable of extending this control by as much as 2 octaves towards the lower end of the audible spectrum. Line arrays are even better at maintaining low frequency control in the vertical plane depending on the length of the array. Even so, all these systems exhibit one governing trend; they lose directivity when the reproduced wavelengths are relatively large compared to the dimensions of the speaker or array.
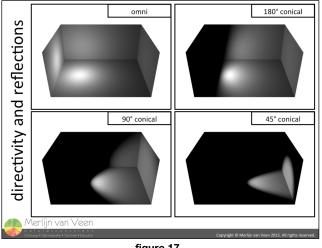
This mid and high frequency control allows us to avoid exciting the room. By not illuminating the walls and ceiling with sound (figure 17), the amount of reflections and therefore the perceivable intensity, not the decay time, of the reverberation is reduced in comparison to the direct sound. On top of that, absorption by air, a distance related phenomena, attenuates the high frequency reflections even further as they propagate through the venue.
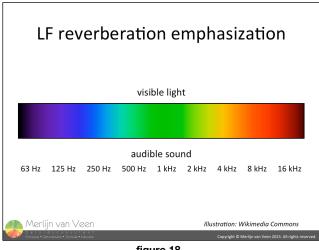


**figure 18**

Because of the relatively large wavelengths of low frequencies, which make them hard to control and absorb, we're virtually incapable of avoiding excitation of the venue. If we compare audible sound to visual light (reversed order and it's 1-octave wide range spread out over 9 octaves) in figure 18,

the reduction of mid and high frequency reverberant intensity by pattern control is like attenuating the green to red part of the visual spectrum, leaving us with nothing but ultra violet.



**figure 19**

A black light, low frequency reverberation, party (figure 19). The mid and high frequency pattern control effectively emphasizes low frequency reverberation. This reverberant energy, smeared out over time, tramples over the rest of the spectrum, reducing intelligibility and clarity.
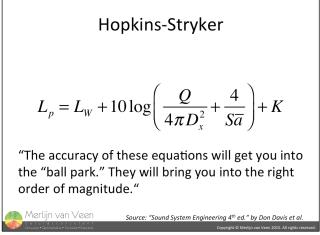
### Hopkins-Stryker

$$L_p = L_W + 10\log\left(\frac{Q}{4\pi D_x^2} + \frac{4}{S\overline{a}}\right) + K$$

"The accuracy of these equations will get you into the "ball park." They will bring you into the right order of magnitude."

**figure 20**

Another important metric indoors to be mindful of is critical distance, which is the distance to the source where direct and reverberant sound power levels are equally as loud.

The Hopkins-Stryker equation (figure 20) allows us to estimate the sound pressure level ($L_p$), for a given sound power level ($L_W$), directivity factor ($Q$), surface area ($S$) and average absorption coefficient ($\overline{a}$) over distance.

The two fractions between parentheses determine the sound pressure levels of the direct and reverberant sound respectively. Direct sound is dependent on directivity factor and distance (inverse square law) and the reverberant sound on surface area and absorption coefficient. Notice how the latter isn't distance related.
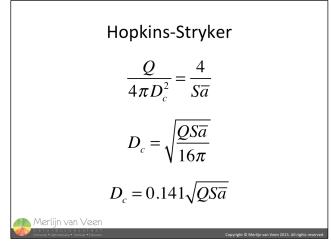
### Hopkins-Stryker

$$\frac{Q}{4\pi D_c^2} = \frac{4}{S\overline{a}}$$

$$D_c = \sqrt{\frac{QS\overline{a}}{16\pi}}$$

$$D_c = 0.141\sqrt{QS\overline{a}}$$

**figure 21**

If direct and reverberant sound are equally loud at critical distance we can estimate its range by rearranging the terms between parentheses (figure 21). The outcome is highly interesting because it shows the numerical weight of both directivity factor and absorption coefficient. Doubling the directivity factor is equally efficient as doubling the absorption coefficient!

Unfortunately the absorption of low frequency power is virtually none. It takes an absorption coefficient of 0,5 to absorb 3 dB of sound power, very unlikely with materials like concrete and metal holding the venue up. The rigidity and stiffness of these materials pretty much rules out help in the form of diaphragmatic action. The walls won't flex and subtract sound power. This assures us the building won't fall down on us.
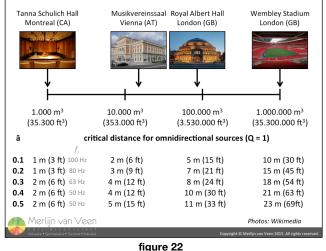
But if the absorption coefficient is expected to be close to zero and the directivity factor for omnidirectional subwoofers is 1, then the only parameter left is surface area. We need lots of low absorbent square meters or feet to compensate.

This gives larger venues under similar circumstances an advantage but in smaller venues trying to increase the directivity factor of the subwoofer(s) is pretty much the only viable option left.

Normally I'm not a big fan of critical distance because as the directivity pattern of loudspeakers changes over frequency so does the critical distance. Anyone with access to a dual-channel analyzer (i.e. Smaart), can locate these positions for a given frequency quite easily. Because one part signal and one part noise in the form of reverberation, the condition at critical distance, will have a coherence value of about 50%. But because the subwoofer's directivity factor stays virtually constant throughout its operating range, the estimated critical distance can be quite insightful.
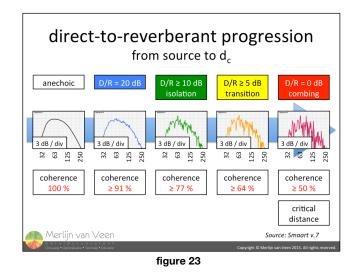


**figure 22**

If we look at some critical distance estimates for various volumes[2] (figure 22) using a directivity factor of 1 and modest values for absorption coefficient, the results are depressing. It's readily apparent that

---

[2] The venues in this example are chosen to give a sense of volume. Their acoustical properties are of no concern in the interest of this exercise.
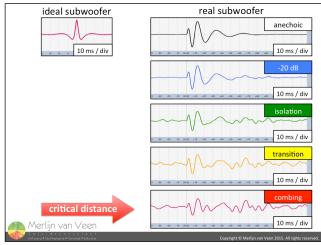
your typical omnidirectional subwoofer is pretty much dead on arrival. And as volume increases it doesn't get much better given the distances that need to be crossed to reach the back of the audience.

The grey values in the first column for the 1.000 m$^3$ (35.300 ft$^3$) example represent the lower frequency limits. In the frequency range below these limits, the venue is governed by room modes that inherently behave non-statistical and render the Hopkins-Stryker equation unsuitable. I would like to point out that these values represent gradual and not absolute razor sharp transitions.



**figure 23**

Now that we have a way of estimating critical distance let's look at the implications along the path from subwoofer to critical distance.

Close to the subwoofer the direct sound, affected by distance (inverse square law), dominates over reverberant sound which is independent of distance. There's good signal to noise ratio between the both which shows up as high coherence on the analyzer. However as we move away from the subwoofer, reverberant sound increasingly gains market share over the direct sound which decays at a -6 dB per doubling of distance rate. The signal to noise (reverberation) ratio decreases and coherence drops. In the frequency domain (figure 23) this manifests itself as a gradual steady increase in ripple.

**figure 24**



**figure 26**

In the time domain (figure 24) the decrease in direct to reverberant ratio appears as time smearing in the linear impulse response plot. By the time we've arrived at critical distance the initial impulse response has grown a 100 ms tail of equal magnitude! And remember, this already occurs after a fraction of distance (figure 22).

Compare this response to the response of an ideal subwoofer at the top left of figure 24 and think again about impact and bags of sand! The majority of the audience is well beyond critical distance and is listening predominantly to indirect reverberant sound.
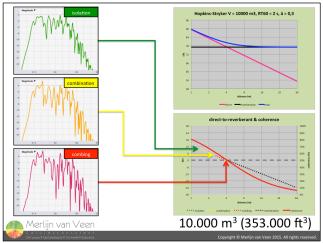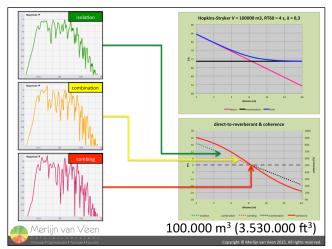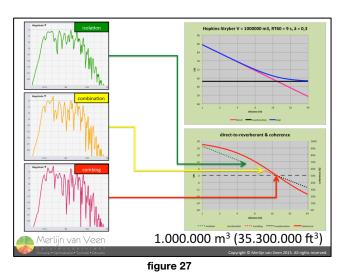


**figure 27**

Figures 25 through 27 show the typical progression, without further modifiers ($M_e$, $M_a$ and $N$)[3], for an arbitrary sound power level, an optimistic constant absorption coefficient of 0,3 of various volumes over logarithmic distance.

In the top right plots, the pink line represents direct sound pressure level and drops with -6 dB (one division) per doubling distance (one division). The black line represents the distance-independent, reverberant sound pressure level. Critical distance is where both lines intersect (equal sound power level).



**figure 25**

[3] Sound System Engineering 4th ed. by Don Davis et al., Focal Press page 222

The blue line represents the combined sound pressure level.

The bottom right plots show the direct to reverberant ratio over logarithmic distance. The line has been shaded in accordance with the summation zones. The red line represents coherence, based on signal to noise (reverberation) ratio.

Beyond critical distance, the reverberant sound increasingly gains more market share over the direct sound, and coherence (SNR) drops below 50%.

The direct to reverberant trend over distance is identical for all volumes in these examples. The difference is in the values that grow over volume.
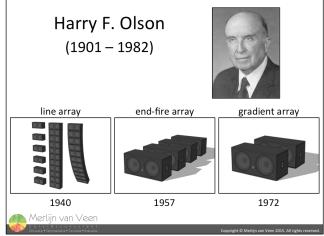


**figure 28**

Now that we have established the need for low frequency control, especially indoors where room acoustics become the primary obstacle, let's look at an overview of the most typical current solutions (figure 28).

We can make a distinction between line arrays and forward arrays, and as far as I can tell they are all described by Harry F. Olson[4] by as much as 75 years ago.

Let's start by looking at the forward arrays.

---

[4] Book: Elements of Acoustical Engineering 1940
Book: Acoustical Engineering 1957
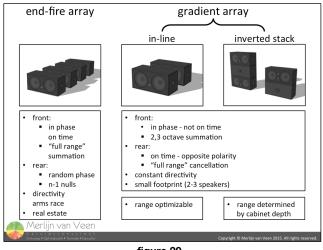AES Journal: Gradient Loudspeakers, 1972



**figure 29**

The forward arrays (figure 29) come in two fashions: end-fire and gradient. The latter can be divided into in-line and inverted (CSA) stack configurations.

The end-fire array is in phase and on time in front of the array. This results in "full range" summation throughout its operational range.

Behind the end-fire configuration, there is randomized phase which results in an amount of nulls equal to the amount of subwoofers minus one.

The amount of subwoofers determines the total bandwidth and amount of cancellation behind the configuration. More subwoofers is better overall, up to a point of diminishing returns.

Regardless of the amount of subwoofers, the end-fire's coverage pattern will narrow over frequency throughout its operational range. The amount of subwoofers determines the global coverage pattern but the lower frequencies will be wider than the higher frequencies. An arm's race.

The end-fire array, unless flown, requires a lot of real estate and should not be put underneath a stage. The stage acts like a boundary and will ruin[5] the coverage pattern. This goes for all cardioid configurations.

---

[5] AES Paper 7971
Subwoofer positioning, orientation and calibration
for large-scale sound reinforcement

The gradient array is one polarity flip away from a reversely oriented end-fire configuration. The amount of subwoofers is limited to 2 for in-line configurations, and 2 or more for inverted stack configurations, depending on the front to back level difference of a single subwoofer.

The gradient array is on time but out of polarity behind the configuration. This results in "full range" cancellation throughout its operational range.

The front of the array is not on time but in phase. This results in a partial summation of a little over 2 octaves.

Notice how end-fire arrays exhibit a perfect solution in front of the configuration and an imperfect solution behind it. The gradient array is the other way around, an imperfect solution in front of the configuration and a perfect solution behind it. That is why the gradient array is also referred to as the mirror image of end-fired arrays.

In contrast to the ever-narrowing end-fire array, the gradient array exhibits near constant directivity throughout its operational range. This turns it into a more suitable building block for more complex designs.

The inverted stack can be designed symmetrically or asymmetrically which relies on the boundary conditions of the floor to keep it symmetrical. The latter configuration therefore can't be flown. The symmetrical inverted stack can be used as a composited module in larger line arrays.

The in-line gradient array requires little real estate, especially the inverted stack configuration.

The in-line configuration allows you to tune the maximum output frequency in front of the configuration by changing the physical spacing whereas the maximum output in inverted stack configurations is predetermined by the physical depth of the cabinet.
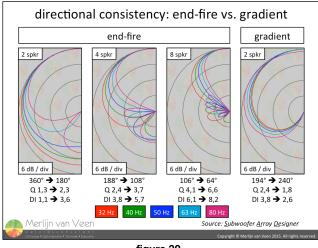


**figure 30**

Figure 30 illustrates the arm's race in terms of directivity for end-fire arrays when the amount of subwoofers is increased. Notice the overall narrowing of the coverage pattern. The lower frequencies however remain wider than the higher frequencies. The gradient array on the other hand exhibits virtually constant directivity throughout its operational range.

| directivity factor ($Q$) | critical distance (%) |
|---|---|
| 1 | 100% |
| 2 | 142% |
| 3 | 173% |
| 4 | 200% |
| 5 | 237% |
| 6 | 245% |
| 7 | 265% |
| 8 | 283% |
| 9 | 300% |
| 10 | 316% |

**table 1**

A look at the directivity factor (figure 30) shows the effort required to achieve practical values (table 1). Especially if we take into consideration that the critical distance formula derived from the Hopkins-Stryker equation has square root in it.
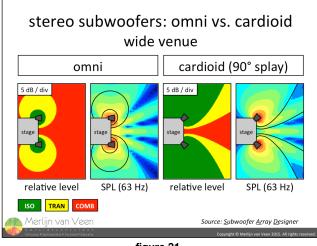
**figure 31**



**figure 32**

Let's look at an example of cardioid solutions in a wide venue (figure 31). On the left hand side we see the omnidirectional approach where the majority of the audience ends up within the red combing zone. But when we change to a cardioid coverage pattern, our subwoofers are suddenly no longer immune to rotation.

By introducing a splay angle we make use of angular attenuation and reduce the overlap in the center, while simultaneously creating isolated areas off-center. A look at the right hand relative level offset plot shows that the majority previously in the red combing zone has now become a minority. This is emphasized by the reduced interference in the power valleys between the cardioid solutions. The geometrical and therefore temporal situation has stayed identical, but relative level offset has come to the rescue.
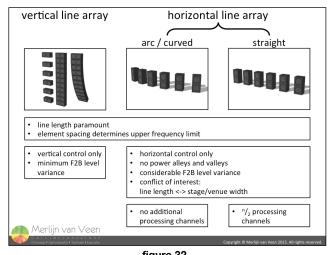
Line arrays (figure 32) are another way of exercising low frequency control. They can be deployed vertically and horizontally. The pattern control occurs exclusively in the same plane.

For successful low frequency control with line arrays, line length is paramount. Long wavelengths combined with narrow coverage angles require long lines. Element spacing determines the upper frequency limit of the operational range before the onset of pattern breakup by under-sampling. Straight arrays however behave dramatically different from curved arrays.

Without further processing, the coverage angle of straight arrays halves with each doubling in frequency throughout their operational range. The arrays aren't capable of maintaining constant directivity.

Curved or arced arrays on the other hand maintain a constant coverage angle throughout their operational range, equal to the angle of the arc section (like a slice of pizza).

Consequentially straight arrays which are intended to exhibit constant directivity, need to be processed as if they were physically placed in an arc.

Vertical arrays provide poor horizontal control but can be very effective in offsetting front to back distance ratios. They'll reduce level variance

throughout the audience and when aimed or steered carefully, they can avoid the ceiling.

Horizontal arrays are often deployed on the floor as a single mono low frequency source that provides horizontal control. When built with omnidirectional subwoofers, arrays in an arc will spread out their energy in front of the configuration and create an undesirable focal point behind the configuration that ends up on stage. Straight delayed arrays with omnidirectional subwoofers on the other hand, are inherently symmetrical. The energy will spread out equally on both sides of the configuration like a Rorschach inkblot test.

In the vertical plane, due to the poor front to back distance ratio from sitting on the floor, the level variance is still at the mercy of inverse square law. But because of the mono approach, the absence of power alleys and valleys is a substantial advantage over the traditional "stereo" subwoofer setup.
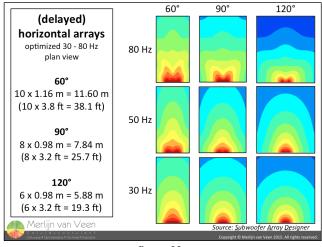


**figure 33**

One particular conflict of interest with respect to horizontal arrays is the desire to extend the line length to the width of the stage or venue. Narrow venues require long lines whereas wide venues require short lines regardless of the width of venue or stage.

Near-field underlap is a local problem, affecting a minority of the audience, and is best dealt with by using local solutions like fill subwoofers.

It makes no sense to sacrifice a macro solution at the expense of the majority of the audience in favor of a minority.

All attempts of maintaining a fixed line length, regardless of coverage angle, like amplitude or frequency shading, effectively reduces the line length over frequency and could be considered a waste of power, something not taken lightly by the sound pressure level preservation society in a part of the spectrum where amplifier power doesn't come cheap.

The rising popularity of "magic bullet" FIR filters offer little help for live sound due to the unacceptable latencies involved at low frequencies.

In conclusion, both vertical and horizontal arrays evidently benefit from cardioid subwoofers or configurations over omnidirectional subwoofers.

Thank you!


For further information on low frequency control consult the Subwoofer Array Designer manual or read the articles "Lack of impact" and "Offset restraint in 2-element cardioid configurations" at my website.