The sensation of sound is a thing sui generis, not comparable with any of our other sensations. No one can express the relation between a sound and a colour or a smell. Directly or indirectly, all questions connected with this subject must come for decision to the ear, as the organ of hearing; and from it there can be no appeal. But we are not therefore to infer that all acoustical investigations are conducted with the unassisted ear. When once we have discovered the physical phenomena which constitute the foundation of sound, our explorations are in great measure transferred to another field lying within the dominion of the principles of Mechanics. Important laws are in this way arrived at, to which the sensations of the ear cannot but conform.

Lord Rayleigh, The Theory of Sound 1877

These words were written almost a hundred years ago, just before the dawn of what has been euphemistically called The New Physics, and at a time when nothing was believed completely understood until it had been reduced to the level of a mechanism. This was the first paragraph in what was to become the greatest single work in acoustics. It is still a basic reference.

The message of Rayleigh*s first paragraph is very clear, yet it was lost on many who were so impressed by his mastery of mathematics that they relegated the sensation of listening to a station below that of the mechanics of sound. Today much of audio continues to be mechanism oriented in its thinking. Our blueprints to understanding this mechanism are composed of meters, oscilloscopes, graphs, mathematics, and inferences based on the way physical objects react. The mechanistic view is so deeply engrained that there are many who never question the accuracy of the pieces of the blueprint for evaluating how well the complete audio system does its job.

As a result of this, audio components, and occasionally whole systems are *measured* thoroughly and exhaustively as though they are

simple mechanisms. Ranking such systems for quality then consists of analyzing the measurement data and ordering the results on a numeric scale. One percent distortion is thus considered less desirable than a third percent distortion. A more uniform frequency response is considered to be better than an irregular response, and so on.

The problem comes when you get down to listening to these systems, for the numbers may have very little correlation with how you might rank them based on what you hear.

A great deal of bile and acrimony has been generated over the years by this fact, with the result that much of audio has been split into two views: the objective view and the subjective view. The subjective view may best be summarized *don*t bother me with measurements, how does it sound?* ; while the objective view is, *if you golden ears hear something, it doesn*t conform to the laws of physics*.

A great many people in the audio profession choose to ignore the controversies which such differences of opinion can create. It is almost as if they hope that by ignoring the problem it will go away. But the problem will not go away. It persists.

I would like to share with you some of the results of my own personal research into this mystery of why our measurements do not always agree with what we hear. The technical basis for this material has appeared in several papers in the Journal of the Audio Engineering Society. There is no need to discuss technicalities here. Instead I will summarize the salient points which are of interest to those who share the experience of listening to sound rather than analyzing it.

WHILE A GOOD MYSTERY NOVEL CAN BE COMPROMISED BY PREMATURELY REVEALING THE PLOT, I THINK IT MIGHT BE HELPFUL IN THIS DISCUSSION TO TELL YOU SOME OF THE RESULTS BEFORE DESCRIBING THE ANALYSIS INTO OUR OWN SPECIAL MYSTERY OF SOUND.

PERHAPS THE MOST IMPORTANT RESULT IS THAT BOTH THE SUBJECTIVE VIEW AND OBJECTIVE VIEW ARE CORRECT. WE DO HEAR THINGS WHICH ARE NOT NOW UNIQUELY MEASURED, ALTHOUGH OUR MEASUREMENTS ARE ACCURATE AND WHAT WE HEAR IS CONTAINED TO SOME EXTENT IN THEM. THIS MEANS THAT WHAT WE HEAR CAN ULTIMATELY BE RELATED TO SOME MEASUREMENT. BY TAKING AN OBJECTIVE VIEW WHERE FORM HAS MORE SIGNIFICANCE THAN HAS EVER BEEN CONSIDERED BEFORE, SEVERAL APPARENT PARADOXES OF SOUND MAY BE RECONCILED. FOR EXAMPLE, DISTORTION IN AN AMPLIFIER IS NOT THE SAME THING AS DISTORTION IN A LOUDSPEAKER AS WE NOW MEASURE IT, AND A GIVEN AMOUNT OF DISTORTION IN AN AMPLIFIER WILL BE MORE OBJECTIONABLE THAN THAT SAME AMOUNT IN A LOUDSPEAKER. AMONG OTHER THINGS THIS MEANS THAT IF SOMEONE TELLS YOU HE CAN HEAR THE DIFFERENCE BETWEEN TWO AMPLIFIERS OR PHONOGRAPH CARTRIDGES WHILE LISTENING THROUGH A LOUDSPEAKER SYSTEM WITH 5 PERCENT DISTORTION, YOU WOULD DO WELL TO BELIEVE HIM BECAUSE IT IS ANALYTICALLY POSSIBLE. OTHER EQUALLY INTERESTING RESULTS DROP OUT OF THE ANALYSIS, BUT THESE EXAMPLES SHOULD GIVE THE FLAVOR AND PERHAPS WHEY YOUR APPETITE FOR WHAT IS TO COME.

NO ONE CAN HOPE TO UNDERTAKE ANALYSIS OF SUBJECTIVE SOUND UNLESS HE CAN EXPLAIN WHAT IT IS HE IS ANALYZING. YOU*VE GOT TO KNOW THE TERRITORY. SO LET ME START BY EXPLAINING WHAT SUBJECTIVE AUDIO AND IN PARTICULAR THE LISTENING EXPERIENCE MEANS TO ME.

The subjective sound experience is an illusion - a sound image. It is a personal mental experience of real or apparent acoustic presence. The principal way we can communicate this experience to others is through verbal description. The words we use in this description are generally chosen to evoke comparable images of experience and describe the acoustic illusion in terms relating to known sources of sound.

We fully recognize, for example, that the sources of sound in stereo reproduction come from two discrete places. Yet we can experience the illusion of phantom sound images spread out between the speakers. And in a really fine reproducing system the illusion can be startlingly good.

Subjective sound is a total experience. In addition to the physical acoustic stimulus we receive from the reproducing system, our subjective judgements draw upon other factors such as emotions, experience, prior training, and a fusion of impressions due to other sensory stimuli like vision and touch. Cohesion of representation for the acoustic image may utilize periods of time rather than instantaneous moments and may involve the inter relationship among parameters. In this way it may be similar to the enjoyment of melody as either an appreciation of articulated relations among a set of components or an overall melodic contour.

It is clear that we are not going to be able to hook meters onto people and measure subjective perception. Rather we must utilize descriptive terminology which relates to shared common experience.

By now it might appear that I have put myself out of business for objective measurement of audio systems. Not so, as we*ll see in our next discussion.

The acoustic image of subjective sound is a very real thing to each of us. As a matter of fact we cannot distinguish the difference between an apparent sound and some actual sound capable of giving rise to the same subjective impressions. The Holy Grail of audio is reproduction which causes an identical subjective affect as that of a previously established actual sound.

Suppose than that I postulate an equivalent actual sound which produces the same subjective image as that which we experience. This actual sound does not exist, but if it did, and you heard it, the subjective image would be identical to that you experience for the subjective sound we want to measure.

That is a pretty wild concept but that equivalent actual sound is something which CAN be measured and described. I have no plans at all to measure this *thing*. All I want to do is establish that there is some mathematical model which can stand for the subjective image, at least in part.

In the Audio Engineering Society Journal I took a more formal approach to get to this point. I presumed that if two or more people can describe the same subjective experience in words which both can recognize, then what they describe exists, no matter whether it shows in our measurements or not. That takes more than a little nerve because it apparently says our mathematics is wrong. I then asked what words we use to describe the subjective experience. After sifting through a fair number of terms it became evident that there was a commonality of a special sort. First of all we describe things in terms of a number of dimensions. There are a number of coordinates

OF REPRESENTATION WHICH WE USE. FOR EXAMPLE, THE DIRECTION FROM WHICH A SOUND APPEARS TO COME MUST BE INDEPENDENTLY SPECIFIED AS WELL AS ITS INTENSITY AND TONAL PROPERTIES. THE FACT THAT AN OBOE IS PLAYING A IN NO WAY IDENTIFIES ITS APPARENT POSITION IN SPACE OR HOW LOUD IT IS PLAYING.

THAT*S PRETTY OBVIOUS. BUT THE THING THAT HITS YOU RIGHT BETWEEN THE EYES IS THAT MOST OF OUR AUDIO MEASUREMENTS ARE ONE DIMENSIONAL. I WILL GET BACK TO THIS LATER, BUT THIS IS WHY WE SEEM TO BE ABLE TO HEAR THINGS WHICH ARE NOT MEASURED.

HAVING ESTABLISHED THAT THE SOUND EXPERIENCE IS DESCRIBABLE AND MULTI-DIMENSIONAL, THE QUESTION REMAINS, IS THERE SOME WAY OF MATHEMATICALLY HANDLING THIS EXPERIENCE? THE ANSWER IS, YES, BUT IT IS NOT THE MATHEMATICS WHICH WE HAVE USED ON AUDIO PROBLEMS FOR LO THESE MANY YEARS. THE TYPE OF ANALYSIS WE MUST NOW CONSIDER IS A BRANCH OF TOPOLOGY. WE MUST THINK OF A GEOMETRY OF SOUND WHERE FORM HAS A MORE VITAL MEANING.

NOW WHEN I SAY GEOMETRY I DON*T MEAN SIMPLY SPATIAL CONFIGURATION. I MEAN A GENERAL GEOMETRY INVOLVING EVERYTHING WE NEED TO DESCRIBE THE ACOUSTIC IMAGE. THIS INCLUDES THE TONAL, SPATIAL, TEMPORAL, AND INTENSITY PROPERTIES. THE WHERE, WHEN, AND HOW OF THE SOUND ILLUSION.

THE REASON FOR WANTING A MATHEMATICAL MODEL FOR THE SOUND IMAGE IS THAT WE CAN OBJECTIVELY HANDLE SUCH A MODEL. AND QUITE POSSIBLY WE CAN HOPE TO MEASURE SOME OF ITS PROPERTIES. SO LET*S TAKE A BIG JUMP AND REPLACE THE MENTAL ILLUSION OF SUBJECTIVE SOUND WITH THIS HYPOTHETICAL GEOMETRICAL MODEL.

IN RAYLEIGH*S TIME THE ONLY HOPE OF UNDERSTANDING SOMETHING WAS TO REDUCE IT TO A MECHANICAL MODEL. THAT IS WHAY RAYLEIGH MEANT IN

HIS OPENING PARAGRAPH. THIS WAS THE GOLDEN ROAD TO UNDERSTANDING. AROUND THE TURN OF THE CENTURY A SMOLDERING MOUNTAIN OF THINGS-THAT-DID-NOT-QUITE-FIT EXPLODED INTO A VOLCANO THAT SHOWERED DOWN ON THE DOMAIN OF MECHANICS. THE DETAILS ARE WELL KNOWN, BUT MUCH OF THE GOLDEN ROAD LAY BURIED, NEVER TO BE SEEN AGAIN.

THE REASON I MENTION THIS IS NOT THAT ACOUSTICS WAS AFFECTED, BUT THAT A MORE LIBERAL VIEW OF ANALYSIS EMERGED THAN COULD HAVE BEEN ANTICIPATED BY RAYLEIGH. NO LONGER COULD AN ANALYSIS BE CONSIDERED UNACCEPTABLE IF IT SEEMED WEIRD AND YOU COULDN*T TIE IT TO A MECHANICAL MODEL. THE ONLY TEST AN ANALYSIS HAS TO PASS IS WHETHER OR NOT IT IS FORMALLY ACCEPTABLE MATHEMATICALLY AND IF WHAT IT PREDICTS ACTUALLY CAN BE OBSERVED TO HAPPEN.

WITHOUT CONSIDERING THE ONTOLOGICAL SIGNIFICANCE, ONCE WE HAVE IDENTIFIED THE MATHEMATICAL MODEL WITH THE EXPERIENCE, WE CAN LIKEWISE TAKE ANY MATHEMATICAL TRANSFORMATION OF THE MODEL AND TIE IT BACK TO THE EXPERIENCE.

THIS MEANS THAT THE ELECTRICAL VOLTAGES GOING TO THE SPEAKER TERMINALS MAY MATHEMATICALLY BE CONSIDERED TO ACTUALLY <u>BE</u> THE ACOUSTIC IMAGE. IT JUST HAPPENS TO BE TRANSFORMED TO A DIFFERENT SET OF COORDINATES. THE SPIRAL GROOVE ON THE RECORD IS ALSO THE ACOUSTIC IMAGE, AS IS THE MOTION OF THE STYLUS ON THE PLAYBACK CARTRIDGE.

ALL OF THIS MIGHT SEEM SO DARN ABSTRACT THAT IT INTRUDES ON OUR MEANING OF THE SOUND IMAGE. ACTUALLY THIS IS THE UNSPOKEN PHILOSOPHY BEHIND THE MEASUREMENTS WE NOW MAKE ON AUDIO SYSTEMS. WHEN WE VIEW THE OUTPUT OF AN AMPLIFIER ON AN OSCILLOSCOPE WE REALIZE THAT THE COMPLICATED BEHAVIOR ON THE SCREEN SOMEHOW REPRESENTS THE SOUND WE HEAR.

In that case the only thing we can see is an indication of a voltage changing with time. We are like an observer by the side of a railroad track watching a swiftly moving train passing us on the way to the loudspeaker terminal. Our perspective is one dimensional because that is the only view allowed to us of the sound image at that point and with that piece of equipment.

Once we voice our philosophy and identify what we are doing, it is possible to bring to bear fantastically more powerful tools of analysis than we could have hoped to achieve by just thinking about the relationship among things. The subjective view stops when it gets to the word descriptions. Theirs is not the province of mathematics. On the other hand, the traditional objective view is considered complete when the mathematically proper description of a given part of the audio system has been achieved. Words are not considered part of the mathematics. Since we are trying to establish a link between the two views, we must go beyond the borders of each and into a no man*s land. Next up is a look at this no man*s land.

So far in this analysis I appear to have left a big loose end flying. I started from the subjective view and moved over to the objective view by introducing topology. What about all that swell mathematics of Fourier transforms and such that the objective person considers de rigueur. Does it go down the drain? The answer is no. When you look very carefully you find that the time domain and frequency domain methods are absolutely correct but are actually part of the topology we had to introduce. Rather than being the only possible ways of describing a system, the time and frequency domain expressions are simply two out of an infinite number of ways. If an Englishman and a German were the only persons who ever carried on a conversation they could easily be convinced that no other language were possible. Imagine their surprise when they are ushered into the United Nations.

I had stated earlier that the fact the sound image is one dimensional in the way most people measure it, and multi-dimensional when we perceive it, is pivotal to explaining why our measurements seldom correlate with what we hear. In order to illustrate what I am saying, let me use an analogy in the form of a *thought* experiment. Mathematically it is possible to take a two dimensional object, like a painting, and cleverly transform it into a one dimensional form. Suppose I had found a way to code the Mona Lisa and have it presented on a string. It*s all there, and you can transform this string back into the original painting with absolutely no loss of detail.

I hand you this string which I state is the Mona Lisa. Before you can react, I hand you another string which I say is Gainsborough*s Blue

BOY AND ASK YOUR OPINION ON THE GENIUS OF DA VINCI VERSUS GAINSBOROUGH AND WHICH WAS A BETTER MASTER OF COLOR AND DETAIL.

AS FAR AS YOU ARE CONCERNED THESE ARE NOTHING MORE THAN TWO WEIRD LOOKING PIECES OF STRING. EVEN IF YOU ACCEPT ON FACE VALUE THAT I WASN*T A CANDIDATE FOR THE FUNNY FARM, YOU COULDN*T POSSIBLY EVEN IDENTIFY WHAT THESE WERE, LET ALONE PASS JUDGEMENT ON NUANCES OF DETAIL. ALL THE INFORMATION IS THERE, BUT IT IS NOT IN A FORM YOU CAN RECOGNIZE. IT ISN*T UNTIL I TRANSFORM THESE ONE DIMENSIONAL STRINGS INTO TWO DIMENSIONAL PAINTINGS THAT YOU SUDDENLY IDENTIFY WHAT THEY ARE.

IF I DIDN*T STOP BUT CONTINUED TO TRANSFORM THIS TWO DIMENSIONAL FORM WHICH YOU RECOGNIZE INTO A THREE DIMENSIONAL BLOB, YOU WOULD AGAIN FIND IT MEANINGLESS. THE ONLY THING THAT MAKES SENSE AS A PAINTING IS AN OBJECT WITH THE PROPER NUMBER OF DIMENSIONS.

OK, NOW SUPPOSE I HAND YOU A PLOT OF THE ACOUSTIC IMPULSE RESPONSE OF CARNEGIE HALL FROM STAGE CENTER TO SOME SELECTED SEAT AND ASK YOU TO LOOK AT IT. I THEN HAND YOU A PLOT OF THE IMPULSE RESPONSE TAKEN IN LINCOLN CENTER AND ASK YOU WHICH HALL HAS BETTER ACOUSTICS FOR A VOCALIST STANDING AT THAT SPOT ON EACH STAGE. I DON*T THINK YOU COULD TELL ME. I COULD HAVE USED THE STEADY STATE FREQUENCY RESPONSE RATHER THAN THE IMPULSE RESPONSE, BUT THE RESULTS WOULD BE THE SAME BECAUSE THIS IS ONE DIMENSIONAL DATA. IT IS SOUND PRESSURE AS A FUNCTION OF POSITION ON THE TIME AXIS (OR ON THE FREQUENCY AXIS) OF THE PIECE OF PAPER. IT IS CARNEGIE HALL PRESENTED TO YOU ON A STRING AND, JUST AS IN THE CASE OF THE MONA LISA, DOESN*T HAVE ANY MEANING TO YOU EVEN THOUGH ALL THE DATA IS THERE.

I use the analogy of the two dimensional painting of the Mona Lisa transformed to a one dimensional string because it is so absurd you*ll remember it. However, the concept of topological transformation which I AM trying to get across is very real. A hologram is a special three dimensional to two dimensional map which is gibberish to you until an inverse mapping is performed by coherent illumination. Only then will you be able to see a three dimensional image of a scene. In just the same way a microphone is a map from a multi-dimensional sound field to a one dimensional signal. And just like a holograph, the microphone signal doesn*t mean a thing to your senses until you transform it back into a multi-dimensional sound experience. In order to illustrate some of the points brought out by the type of analysis I am describing I will continue with the painting on a string analogy. This is a pretty useful model and besides, I can*t think of any analogous five dimensional thing which I would need in order to be completely accurate about sound.

If I cut out small pieces of the string and splice it together, it is tempting to say that not much was done to the Mona Lisa. Actually, you might be horrified with the results. Both the enigmatic smile and the background scene which appear in distinctly different places on the painting now share these same pieces of string which I have snipped out. The resultant distorted painting may have a disembodied smile hanging in air like Alice*s Cheshire Cat. Then again all that may happen is a blurring of detail. Yet all I did was change only small details on the string.

SHIFT THE SCENE NOW TO A *ONE DIMENSIONAL* AMPLIFIER. IF I HAVE A SLEW RATE PROBLEM SO THAT FOR BRIEF MOMENTS IN TIME THE SIGNAL IS *SNIPPED OUT*, IT MAY NOT LOOK LIKE MUCH OF A CHANGE IN WAVEFORM. ACTUALLY, THE VOCALIST AND CLARINET IN OUR SOUND IMAGE SHARE THESE SAME PIECES OF TIME IN THE AMPLIFIER. THE RESULTANT DISTORTED SOUND MAY HAVE THE VOCALIST SMEARED TOWARD THE POSITION OF THE CLARINET BECAUSE OF THIS PROBLEM. THEN AGAIN THE EFFECT MAY BE SOMETHING SO SUBTLE YOU DON*T NOTICE IT, DEPENDENT UPON WHAT KIND AND HOW MUCH DISTORTION THERE IS.

OR FOR ANOTHER ANALOGY SUPPOSE I TAKE EVERY PLACE ON THE STRING WHERE I SEE THE COLOR RED AND I CHANGE THE SATURATION. RED ON THE STRING DOESN*T NECESSARILY MEAN RED IN THE PAINTING. THE STRING NOW LOOKS GROSSLY DIFFERENT, YET WE HAVE TO LOOK HARD AT THE PAINTING TO NOTICE THAT ALL WE DID WAS ALTER THE OVERALL BALANCE IN A WAY INDISTINGUISHABLE FROM A SMALL CHANGE IN THE AMOUNT OF LIGHT FALLING ON THE PAINTING. SHIFTING NOW TO OUR AMPLIFIER, A SUBSTANTIAL AMOUNT OF PURE THIRD HARMONIC DISTORTION DUE TO A SAG IN GAIN WITH LEVEL CAN LOOK PRETTY BAD IN THE SPECIFICATION SHEETS, YET IT MAY HAPPEN THAT THE SOUND IMAGE IS DISTORTED IN A WAY SIMILAR TO A SMALL EXAGGERATION OF THE NONLINEAR DISTORTION ALREADY IN OUR HEARING PROCESS. THE NET RESULT IS THAT IT SOUNDS CLEAN.

I COULD GO ON AND ON WITH ALL SORTS OF EXAMPLES, BUT I THINK THE POINT SHOULD NOW BE CLEAR. YOU CANNOT ALWAYS SAY THAT A HIGH DISTORTION IN A ONE DIMENSIONAL MEASUREMENT AUTOMATICALLY MEANS A HIGH DISTORTION IN A MULTI-DIMENSIONAL SOUND IMAGE. THIS DOESN*T MEAN THAT OUR PRESENT HARMONIC AND INTERMODULATION TESTS ARE WRONG. BUT IT DOES MEAN WE MUST BE A GREAT DEAL MORE CAREFUL IN HOW WE INTERPRET THEM THAN IS NOW OUR PRACTICE. JUST HOW CAREFUL WILL BE COVERED NEXT.

Suppose we now take a closer look at the meaning of distortion in our unfolding mystery of why measurements may not correlate with what we seem to hear. What do we mean when we say something is distorted? Does it mean deviation from fidelity to the original? If the form of a signal is changed by a device so that what comes out doesn't look like what went in, isn't that distortion? Sometimes it might be, but what about a microphone. It takes sound pressure changes and makes voltage changes out of them. That certainly changed the form but we don't like to call a microphone a distorter because it does what we design it to do.

That may be a little far out so let's concentrate on considering devices such as amplifiers where the output and input are both the same thing which can be measured by the same type of instrument. Let's put in a reference signal, like a sine wave,and look at the output. Suppose the tops of the sine waves are rounded off - isn't that distortion? It certainly is, because instead of a single sine wave output we now have a harmonic set of sine waves. But what if the output appears to be a perfect sine wave at the frequency of the input signal, is such an amplifier distortionless? Are you really sure? What if we had used one of the newer breed of compressor amplifiers which reduces gain with average increase in level but does not inject harmonic distortion in the audible range. Program dynamics are now affected so that a given percentage increase in the input is not proportionately reflected in the output. In professional application they are used with care and protect recordings from excess program dynamics. But an over adjustment can be highly distorting, yet there is no measurable steady state distortion. This demonstrates that it

isn't enough to measure for harmonic distortion at one signal level, we must measure as a function of signal level as well.

All right, let's take compressors out of the picture and see if a perfect sine wave output represents a distortionless reproduction. Now the output amplitude tracks the input amplitude. Is it finally distortionless - especially if the gain is constant for every frequency under test? Are you really sure? Don't place any money on it. We haven't even begun talking about phase and time delay.

If the phase out of the amplifier is shifted by a lag of ninety degrees for every possible frequency, it represents what is known as a Hilbert transform of the input. Feed a square wave into the amplifier and the output will look simply terrible with sharp peaks where the square wave level changes occur. If lack of fidelity between output and input is a measure of distortion then this must be distortion.

Or, suppose the phase is well behaved for all frequencies up through middle C in the musical scale, but has a characteristic above that pitch value which indicates a time delay of 50 seconds. This "distortionless" amplifier could now treat you to a musical spectacular where the facile right hand of a pianist continues to play long after the phonograph cartridge is lifted from the record.

All right, let's knock off all the amplitude and phase variations with level and frequency and postulate a good old fashioned wire with gain. Now at last we are distortionless. Or are we? We still haven't really considered the final sound image. The amplifier is now distortionless - as an amplifier - but the sound image can still be imprecise. If the

gain to the left channel is greater than it should be, the sound image will still be distorted by being laterally shifted and warped. Or, if both channels are adjusted for identical gain but one of them is out of phase with respect to the other then we know the sound image will be warped. Finally, if the amplifiers are perfect, have the same gain, and are both in phase - what is left? Two things remain: proper level so the image is neither too loud nor soft, and absolute phase, so that explosive sounds are represented by overpressure rather than underpressure.

My purpose in pursuing this line of reasoning is to demonstrate that when it comes to subjective sound the only distortion that has any meaning is deformation of the final sound image. When we force ourselves to consider a geometry of sound where form has significance, we must chase the signal through all of its various manifestations until it emerges as something related to the illusion of sonic presence.

I'll admit this forces us to tear ourselves away from a well engrained habit of detailing the performance of <u>pieces</u> of the audio system. But stop and think for a moment. Isn't the final sound illusion what we are <u>really</u> trying to improve in audio. When we start to think about the subjective affect of distortion we must always relate our measurements to the final sound image. In order to do this we must break some very old habits of analysis - we must look at the whole system rather than pieces of the system.

Having introduces the concept of a geometry of sound let us now make some use of the analytical tools which this gives us. The sound image can have a number of dimensions, dependent upon where in the audio

system we choose to analyze that image. If the image is one dimensional as we may look at it on an oscilloscope coming out of an amplifier, and five, or more, dimensional as we perceive it as an illusion, then what happens when distortion is introduced? The answer is simple. The affect of distortion in one level of dimensionality is to redistribute the representation among coordinates in another level of dimensionality.

Remember the string analogy? A simple change in a one dimensional form can cause a cross coupling of representation among coordinates in a higher dimensionality. In order to give this a name let's call it representation distortion.

As an example of this, the conceptual position in space occupied by an oboe during the reproduction of a symphony doesn't depend upon what notes are being played, when they occur, or how loud the instrument is played. We also come to expect that the space and tonal properties of that oboe should not change simply because other instruments are playing. Each instrument has its own special representation in our sound image.

What happened when that sound image existed as two separate voltages in our stereo amplifier? If either channel of the amplifier can warp the final sound image, then we have distortion. If in particular the instantaneous values of those voltages is dependent upon the complexity of what is happening at that moment (and for times around that moment) then the coordinates of representation of the final sound image will become tangled up with each other to some extent.

The oboe with independent space, time, tonal, and loudness representation will share the same amplifier channel with the entire string section.

As the strings play, the oboe will be influenced because of amplifier non linearities. The result is representation distortion where the oboe in the subjective sound image is smeared.

It isn't necessary for two or more instruments to be playing in order for crossmodulation of coordinate energy to occur. The oboe will, all by itself, have representation distortion if a non linearity exists in the audio chain. The very first multidimensional representation distortion which was identified as caused by one dimensional non uniformity is time-delay distortion. Next up on the agenda is a discussion of time-delay distortion.

What is pitch? Boy, now there is quite a question. One gets instantly tangled up in psychoacoustics trying to answer such a query. But suppose I ask how one may begin to measure some physical property of a signal which would be agreed upon as representing pitch. What then would we measure? The answer seems to pop out that we should measure the frequency of a signal which represents a tone. But is that mathematically correct?

Let's get technical for a moment. The thing that is mathematically identified as frequency is a coordinate in one of the two one-dimensional representations we employ for signals. The other possible representation uses time as its coordinate. The representation of a signal as a function of frequency is obtained from the time representation by a Fourier transformation. This is a mapping from the time domain to the frequency domain, and conversely. If you follow the rule in going from a time representation to a frequency representation, one of the things you must do is integrate from minus infinity to plus infinity in time.

This never bothers a mathematician, but if you stop and think about it, if by "time" you mean that thing which is measured by the wall clock, then all eternity has to transpire before you mathematically can talk about frequency. You are forbidden to talk about "time" and "frequency" in the same description.

This means that a glissando can never be described as a tone which glides through frequency with time. You cannot, strictly speaking, state that a tone has any frequency value at any particular time. This is a very strange thing - especially to someone who loves music. Try to write a musical score under these groundrules and see how far you get. A great

deal of nimble toe dancing has been done over the years in trying to reconcile what we hear with the concept of frequency. And these usually boil down to mysterious ex cathedra pronouncements involving uncertainty principles and non commuting operators.

Let's knock off all that stuff and look at things from the standpoint of a geometry of representation. The frequency representation, from the standpoint of the dimensions of the coordinates used, is one-dimensional. You can sketch out representations on a frequency line, or a complex frequency plane, or a frequency cube, if you wish, but you still have coordinates dimensioned as Hertz.

Once you identify the coordinate called frequency, you can choose to use another equally accurate one-dimensional representation. You can, in the parlance of topology, map a frequency description into a time description. It is no accident that the proper dimensionality of Hertz is reciprocal seconds (which is what we used to call it), because the other possible description is seconds.

Pretty abstract. But where in all of this is pitch? Suppose I now say that you can map either of these one-dimensional representations into a two-dimensional form. Remember our Mona Lisa? Now what do we have? If you play the mathematical game properly you have one coordinate of this two-dimensional representation which has the dimensionality of reciprocal seconds. The other coordinate has the dimensionality of seconds. Ah Ha! Now, we have something which involves both seconds and reciprocal seconds in the same description. We have a two-dimensional Mona Lisa which we got from a one-dimensional string.

Now things are getting interesting, because if we take a signal which we can agree fits all our subjective ideas of a pure pitch tonal, and ask where it falls in the two-dimensional representation, it comes as the coordinate which is dimensioned in reciprocal seconds. What, then, is the subjective equivalent of that other coordinate which is dimensioned in seconds? It turns out to be time delay - or as some of us conceive it, relative time.

Let's pursue this a little further. What do we <u>really</u> mean subjectively when we think of "time"? Don't we really mean relative time. If not, then where is zero time on an absolute scale? Is it the birth of Christ, ..or last Tuesday, .. or tomorrow morning?

I must unfortunately inject some of my own philosophy here, so please bear with me. Object if you will, but hear me out. We choose a mathematical model for describing things which best fits our observations. Time, or really relative time, has a strong human identification. We have a mathematical model which seems to fit, so we use it and say it is the time functional. Once we do that we inherit the alternate model which is a frequency functional. We get so hung up on the mathematical model that we sometimes refuse to let go of it even when common sense says it doesn't fit everything. One of the most difficult lessons to learn is when to grab for the brass ring - and when to let go. One time to let go is when you realize frequency is not what we subjective perceive as pitch when at the same time we have such a strong association with "time's arrow".

Enough philosophy. Let's look at what we now have in this mystery of subjective and objective audio. A network transfer function is a

cause and effect relationship between what we identify as an input and what we identify as an output. The transfer function of any device, whether it is an amplifier, a phonograph cartridge, or a loudspeaker, can now be expressed in three possible ways - based on what we have just found. We can express it as a frequency response, or as an impulse response, or as a joint pitch-delay response.

When we put a signal into a device (it could be a Brahms Symphony into a loudspeaker) what emerges can be analyzed in any of these three ways. It can be analyzed as a steady state frequency spectrum which includes everything from the opening passage through the finale. We assume nothing happened prior to the symphony and eternity stops with the last note when we do that. Or we can analyze the complex (yes, complex, not just real) instantaneous values perceived. Or we can analyze the signal in terms of the relative arrival times as a function of tonal pitch.

The measure of how perfect the device does its job is made by comparing the output against the input. The frequency response and impulse response of a device are the old standards we have come to know and love. But what is the nature of the pitch-delay response? It turns out that a non-perfect network does something strange to a signal when you analyze its pitch-delay behavior. A perfect signal going in will emerge as a multiplicity of signals, each one of which has its own delay as a function of pitch and with its own intensity. Many instead of one. A perfect signal is dispersed in time delay as a function of its pitch. Instead of a signal emerging at the proper time, many equivalent signals emerge with varying time delays. <u>That is why this is called time-delay distortion</u>.

Time-delay distortion was first predicted mathematically and given a name in 1968. Once one knew what to look for, it was immediately perceived and measured. It is now officially recognized as one of the perceivable veiling distortions in reproduction, although the true topological nature of this distortion was not identified in technical literature until 1973.

Time-delay distortion is one type of representational distortion, which we discussed earlier. What is happening is that a distortion in frequency response causes the emergent signal to have a cross coupling of signal energy in the dimensions of pitch and relative time. The signal is smeared in the time of arrival of pitch components.

CLICK! Another piece of the mystery is falling into place. The sound image we perceive has pitch and relative time as two of its independent coordinates. When a device distorts the sound image, we now see that one effect is a time smear of pitch components. Time-delay distortion. If you imagine that the sound image is blurred for certain tones, you are correct. Do we measure that when we test the audio components? The answer is yes, but we may not be able to recognize what it means when we see it as a frequency response or an impulse response distortion. The information is there, but we have to recast it into a form recognizable as an ultimate time-delay distortion. Just as in the case of the Mona Lisa it didn't mean much on a string. But express it as a two dimensional object and it suddenly makes subjective sense.

The simple step of mapping from a one-dimensional to a two-dimensional representation thus begins clearing up some aspects of the mystery of audio. The analytical details of this map and its representation on what is called

a Delay Plane is a bit too technical to go into here. But some of the predicted results are significant enough to be worth considering.

For example, a particular type of network known as an all-pass can make mince meat out of a square wave signal. If you are a firm believer that a good square wave response is a sine qua non of fidelity, then such a device would be instantly cast away as unlistenable. The poor square wave starts out with a great rise time, but goes all over the place in a most unmannered way thereafter. The frequency response looks great - in amplitude- but the phase is anything but linear. The problem is that if you place such a network in the audio chain, its affect may be virtually inaudible. How come? One answer is afforded by the pitch-delay representation. What is happening is that there is a single value of time delay for each pitch component. Rather than being smeared into a multiplicity of images, the final sound image is distended in the delay of pitch components analogous to the view seen through a fun house mirror. Articulation and clarity are not significantly modified, but the time delay relationships among pitch components is altered in a smooth manner.

A rather polite all-pass network will in fact have a delay alteration much less than that due to many other members of the audio chain - particularly the loudspeaker. Yet a square wave response can look terrible. This doesn't mean that such a distortion is inaudible. Not by a long shot. But perhaps it can best be described as benign, since the magnitude of the energy stays constant in the dimension of pitch.

Now, let's take a little magnetic tape print-through. You've heard the

effect where a loud chord is heralded in a soft passage by increasingly louder ghost chords which occur once for every revolution of the tape reel. It can destroy the appreciation of a good performance, yet the signal level may be so low as to be in the noise of most measuring instruments. Not so in the delay plane where the network representation is now a series of delayed outputs. It looks about as bad in this representation as it sounds.

These are by no means the only manifestations of time-delay distortion, as we will discover next.

We have seen from previous discussions that the way we measure audio components may be perfectly correct but not have much identification with human subjective concepts. This arises because the parameters of measurement are not those of perception. We hear a multi-dimensional experience. But we measure at the places where this multiple dimensioned sound image has been transformed into a one-dimensional signal. We measure the string, but enjoy the painting.

The first step in bringing measurement and listening together is to learn how to understand some of our one-dimensional measurements in terms of a two-dimensional representation. From this we can begin to understand how to expand into something like the multiple dimensions of listening.

Let's take a measured component in the audio chain and see how to approach it as a two-dimensional modification of the sound image. A conventional one-dimensional measurement is steady state frequency response. The other one-dimensional measurement is that of impulse response. It might come as a shock to many circuit analysts, but BOTH frequency response and impulse response should be treated as a complex quantity, either as two conjugate components or as an exponential form having a magnitude and a phase. There appears to be something within us which rebels at a complex time representation, however since frequency has such poor human identification we willingly shove our worries under the rug by accepting a complex frequency response. When you really get down to measuring for the audio image it is necessary to lift the rug. There is nothing magic. The complex representation is nothing more than a recognition of the details of energy density partitioning.

With that out of the way, let's look at what an imperfect frequency response can do to the sound image. We measure frequency response by applying a sine wave to a network and comparing what comes out against what we put in. Strictly speaking, the sine wave must last for all eternity, but few of us can afford to wait that long. So we measure for a period of time long enough for "steady state" conditions. Which means all the nasty turn-on transients which resulted from applying the test signal have dissipated themselves to a level below our measuring capability. Right there is one clue in our audio mystery, because much of our music could also be considered transients. But we measure for steady state response because, being one-dimensional, we cannot consider time in the measurement.

Have we thrown away any hope of reconstructing the behavior for music because all we can describe is response versus frequency? Not at all. But you must record both the amplitude and the phase of the output sine wave in order to state what happens uniquely.

Now let's look at the same network in terms of its two-dimensional representation of pitch and delay. Suppose the network we are testing is sealed in a black box so neither of us can see what is inside. All we know is that the box has an input and an output. We measure it and determine its frequency response. Suddenly, in a puff of smoke, a genii appears and hands us another black box with input and output terminals, and asks us to measure it. We do so and find that the genii's box has precisely the same response as our own. Suspecting something funny, we measure our own network with an impulse so as to determine its other one-dimensional response. Lo and behold the genii's box also has the same impulse response.

With a smirk, the genii now hands us a can opener and we open both boxes. They are not at all alike inside. Our own box had things we could recognize like resistors and capacitors. But the genii's box contained a number of funny looking delay lines and networks, all of them hooked in parallel.

If we snip out one of these parallel sub-networks with side cutters and measure the frequency response, a surprising thing is observed. The network has no amplitude change with frequency but does have a peculiar phase angle change with frequency. The genii's box contained parallel all-pass networks and a recombining network so the output was on one terminal.

What is the transfer function of each all-pass? It is a special pitch dependent delay line. Every signal of fixed pitch fed into the box (and a sine wave is such a signal) is split into a number of parallel delayed paths and emerges as a superimposed multiplicity of delayed pitch components. This is our two-dimensional representation we were looking for. This splitting and summing is of course what we call time-delay distortion.

Sound pretty weird? Not really. Let's take a very simple example to show how a well known one-dimensional network representation can be dealt with as a two-dimensional thing. Suppose our first network is a simple resistance-capacitance low-pass filter. It's frequency response is constant in amplitude out to near what we call cutoff, where it is down 3 dB, then it continues to drop, approaching a 6 dB per octave slope with increasing frequency. The phase shift is zero at DC but begins lagging with increasing frequency and finally approaches a 90 degree lag at very high frequencies.

Now let's look at the genii's box. It contains two parallel networks. One of them is a piece of wire, and the other is what is called a first order all-pass lattice with a frequency of maximum delay at DC. The wire is dropped through a divider to have a gain of 1/2. The all-pass delay line also has a gain of 1/2 and is phased so that input and output are in phase at DC.

Now let's see what happens to the genii's box when we feed sine waves in. At very low frequencies the no-delay piece of wire and the delay line add together to pass the signal very well. In fact unity gain at DC. As the frequency increases, the delay line begins to have enough affect on the phase shift of the sine wave that its output now begins adding out of phase with the wire path. At what we call cutoff frequency, they are ninety degrees out of phase, but because each has a gain of 1/2 the resultant signal is down 3 dB and only lags 45 degrees. Finally as we increase the frequency, the delay line component increasingly approaches an out-of-phase condition with the wire path and the result approaches nullification. The frequency response is identical with out RC network.

That's no big deal. But, whereas in the case of the frequency response we were forbidden to consider dynamic musical signals, now with this pitch-delay expansion we can meaningfully describe such things.

Those who are experienced in the technicalities of signal processing are probably highly agitated by now and might tend to look upon this discussion as a variant of the old shell game. Didn't I just say that we can't talk about frequency and time in the same description? And yet I boldly postulate a special type of network having a dispersive delay and describe its frequency response. Have faith.

In the previous discussion we introduced the concept of time-delay distortion as a consequence of the topology of a two-dimensional representation of the acoustic image. In this discussion we are independently introducing time-delay distortion as a consequence of network theory. This duplication is worthwhile. Why? Because the first step in breaking with tradition is always the most difficult and needs all the assurance one can give it.

A very bad mistake was made during the time of Lord Rayleigh in analyzing the time delay of propagating systems which have a non-uniform amplitude response for frequency. It arose out of pre-relativity measurements of the speed of light. I won't go into a discussion of it here because it is a very long and involved story and includes some of the most famous scientists of all time - Rayleigh, Gibbs, Einstein, Voight, Sommerfeld, Brillouin, Michelson, etc. But the important thing from the standpoint of trying to come to grips with the mystery of audio is that group delay NEVER represents the true time delay in a minimum phase system. There are, however, special networks for which we can truthfully describe time delay in terms of how long it takes after we put a signal in before something comes out. These are all-pass networks. The network concept that provided the key to time-delay distortion was one which showed that most of the networks of the kind found in audio can be duplicated by special parallel combinations of all-pass networks. This is the genii's box. Thus even the lowly RC low-pass network can be considered to split the input into two pieces which recombine after suffering their respective pitch dependent time delays.

We can talk about any signal we want, even sine waves, and describe the effect as a splitting - delaying - then recombining of the signal.

This fits beautifully with the acoustic image. A perfect distortionless reproduction produces one and only one acoustic image in its proper place. Any minimum phase distortion in response in the audio chain splits the acoustic image into many images which spread in space dependent on the instantaneous signal components of pitch. The reason we can say they suffer a space spread is that sound in air travels at a constant speed and a little extra time delay is equivalent to the source having moved back from us a proportionate distance. This merging and blurring in space constitutes a distortion on the sound image.

One of our interests in minimum phaseness in a network transfer function depends upon the fact that most of our equalizers are also minimum phase. If we equalize for amplitude response then all of the multiple images of time-delay distortion can merge into one perfect image. If the audio network is non-minimum phase and we attempt to equalize for "flattest" frequency response, then we are left with a resultant all-pass transmission. The sound image is now distended in space, as we stated earlier, like an image in an amusement park mirror.

Time-delay distortion is really ubiquitous in our audio processing. This even includes the multiple sound delays of reproduction in a room. We will defer talking about rooms until later. There is, however, another aspect of time-delay distortion which answers one mystery of audio involving the improvement in subjective cleanliness of magnetic tape recording when all phase distortions are removed.

A magnetic reproduce head is a case of an audio component in which a uniform amplitude response is not sufficient to guarantee distortionless

REPRODUCTION. THE APERTURE LOSS AT SHORT WAVELENGTHS MAY BE PARTIALLY COMPENSATED BY PASSIVE HIGH PEAKING CIRCUITS, BUT THE PHASE CONDITION IS NOT SIMULTANEOUSLY ACHIEVED. IT THEN BECOMES NECESSARY TO ADD A COMPENSATING ALL-PASS EQUALIZATION FOR PHASE. WHICH IN TERMS OF OUR TWO-DIMENSIONAL REPRESENTATION MEANS THAT THERE WAS AT LEAST ONE ALL-PASS NETWORK LEFT IN THE GENII'S BAG AFTER WE TRIED TO CASCADE AN EQUALIZER FOR AMPLITUDE.

AN OLD TRICK FOR IMPROVING THE SUBJECTIVE QUALITY OF FIRST GENERATION MAGNETIC TAPES IS TO DUB A NEW TAPE FROM A REVERSED VERSION OF THE ORIGINAL. THE EDGE OF SOME INSTRUMENTS SEEMS BRIGHTER AND MORE REALISTIC WHEN THIS IS DONE. THIS LITTLE MYSTERY IS EXPLAINED AS NOTHING MORE THAN CANCELLING THE RESIDUAL ALL- PASS BEHAVIOR OF THE IMPROPERLY COMPENSATED REPRODUCE HEAD. WHAT WE DO WHEN WE PLAY A TAPE BACKWARDS IS TO REVERSE THE DIRECTION OF TIME FOR THE PROGRAM MATERIAL. THE CLOCKS IN THE RECORDING ARE RUNNING BACKWARD. TIME-DELAY DISTORTION IN AN IMPROPERLY EQUALIZED REPRODUCE HEAD WILL CAUSE A DIFFERENTIAL TIME DELAY BETWEEN LOWER AND HIGHER PITCH COMPONENTS. THAT IS, REAL HONEST TO GOODNESS CLOCK DELAY AS MEASURED FROM A CAUSE-EFFECT RELATIONSHIP. BY DUBBING BACKWARDS AT AT ONE-TO-ONE SPEED WE ARE MAKING THE REPRODUCE DELAY A RECORDING PREDICTION, BASED ON THE BACKWARD RUNNING CLOCKS IN THE RECORDING. WHEN WE RUN THE NEW RECORDING IN THE PROPER DIRECTION THE AMOUNT OF PREDICTION IS THUS AUTOMATICALLY WHAT IS NEEDED WITH THE DELAY OF THE REPRODUCE EQUALIZATION SO AS TO MAKE ALL COMPONENTS COME OUT AT THE PROPER TIME. AND, CLICK, ANOTHER PIECE OF THE MYSTERY IN AUDIO FALLS INTO PLACE.

It would be possible to go into many more aspects of time-delay distortion, but we have dallied too long and must move on in these discussions since there is so much more yet to be uncovered. So far all we have considered has been the distortion in the sound image due to a linear processing defect. Next up we will begin looking into non-linear affects, starting with a look at what we normally call harmonic distortion. You may be in for a few surprises.

What would be your opinion of an audio product which had the following advertizing claim: "Harmonic distortion at "average" listening level is only one percent; while harmonic distortion before clipping is held to slightly over ten percent".

Would you buy such a product? Or would you reject it as probably too distorted to satisfy a critical listener such as yourself? Think twice, that's a typical sort of measurement on a recording studio master tape machine. Virtually all of your "clean sounding" records originated on a machine with such specifications.

All right, for another example of distortion, let's set up an A-B test between two power amplifiers. One of them has a tenth percent harmonic distortion at full output while the other has distortion so low we can't seem to measure it until the onset of clipping. We have a two-position toggle switch and set up to listen to the two amplifiers. Through all kinds of program material we toggle back and forth and determine that one amplifier indeed seems more transparent than the other. And there are enough people in our jury who agree with the results that we are convinced a genuine listening difference exists.

But wait a minute. We had to listen through loudspeakers. How much distortion do they have? So we set up a microphone and measure the harmonic distortion of the loudspeakers at the listening levels we just used for our A-B tests. Oh boy! they clock in at a neat three percent harmonic distortion. Thirty times the distortion of the one tenth percent amplifier and a thousand or more times that of the other amplifier. How could we have heard any difference?

And if that isn't bad enough, when we take the masking tape off the toggle switch to see which amplifier sounded better - it was the one with a higher measured distortion!

Does this little scenario seem absurd? Well it isn't. These sort of things happen with enough regularity that it is a genuine "whatsit" in our mystery of audio. Trouble is, very few audio professionals are willing to talk about such things in public. It "spoils the image", don't you know.

Generalities are dangerous, but let's take a look at some of the things we observe more often than not concerning harmonic distortion. A given percentage of distortion in a loudspeaker is a lot less objectionable than in an amplifier. The same percentage of distortion in playback of a magnetic tape and playback of a direct cut disc doesn't sound the same. When comparing amplifiers of different design types - such as vacuum tube and transistor - the same percentage of harmonic distortion produces quite different sound image distortion.

There - that should have stepped on enough toes to get someone's attention. Now let's take a look at distortion from the standpoint of a geometry of sound.

The first, and probably most important,thing to remember is that we should not get hung up on talking about distortion in a component. The only thing meaningful is to chase that distortion through all the steps of audio processing and ask, "what does it do to the final sound image?"

How do we measure harmonic distortion? We start with a pure sine wave signal which we use as a source. Then we intercept the output version of that signal and measure its harmonic content in a variety

of ways. Let's stop here for a moment and reflect on some of the points we raised in an earlier discussion. Most of us seek the easy way out in measuring. It is a great deal easier to say "this device measures four on a scale of ten", than to give a detail specification. By now it should come as no shock that the spectrum of harmonic distortion has an amplitude AND A PHASE of components, and should be measured as a function of drive level as well as source frequency. Sure, it's a lot of work, but in many cases shows differences between components that do not show up on a "one-number" measurement. A better understanding of the mystery of measurement versus listening occurs when we become more precise in our measurement. Let's defer a more complete discussion of that point for another time.

The sine wave is a proxy for our sound image. Harmonic distortion, being of necessity a one-dimensional measurement, is best applied to those components where the sound image is itself one-dimensional. Where could we expect this measurement to be most meaningful? How about amplifiers? Yes. Does it apply to single channels of magnetic recording? Yes. How about phonograph cartridges? Not unless you're very careful. How about loudspeakers? No, not the way we now measure them.

So it seems that harmonic distortion in a loudspeaker as we now measure it, is not really the same thing as harmonic distortion in an amplifier. When we measure an amplifier we are measuring everything that amplifier can do to the portion of the sound image which it processes. So long as we are discussing steady state distortion (we'll come to transient distortion in a later discussion), everything the amplifier

can do to modify the sound image is measured.

When we place a microphone at one position in the sound field of a loudspeaker and measure the harmonic distortion of the microphone signal, we are tacitly assuming that the role of a loudspeaker is one-dimensional. Of course it is not. The loudspeaker interfaces with our normal geometric space dimensions, which an amplifier does not.

This is only a fragment of a clue to the listening difference of distortion because we really must try to understand how the <u>final</u> sound image is affected. In order to begin understanding the geometry, let's consider normal stereo reproduction. That amplifier we measured carried what we call the left channel, let us assume. What do we mean by that? The sound illusion is a whole phantom image spread out between the reproducing speakers, as we generally perceive it. Did somebody take a hot butter knife, slice through stage center of our image and constrain the left half to the left channel? No, only the extreme leftmost part of the image is wholly processed by the left channel amplifier. The rest of the image depends upon the right channel to some extent. All of the space, tonal, and temporal properties of our ultimate sound image is processed as two independent one-dimensional things.

This processing defect we know as harmonic distortion is one manifestation of a nonlinear transfer characteristic. Consider what happens when a musical instrument occupies stage center. Both channels then process the same signal. A sine wave signal will have harmonics, but they will all occupy a stage center position if both amplifiers are identical. An extreme stage left or stage right instrument will also not have its angular position and extent subjectively changed by harmonic distortion.

But a little reflection will show that an instrument anywhere else will have a probable lateral spread of its tonal harmonics. If the signal from the left amplifier is 3 dB stronger than from the right amplifier, the distortion fragments will generally be stronger on the left than the right.

Do you remember the early days of transistor power amplifiers and what was known as "transistor sound", a harsh and brittle sort of unpleasantness in a signal. Much of this was in fact due to low signal level amplifier problems which occured at the crossover transition between power transistors. Consider what might have happened to the geometry of the final sound image.

Most of us seldom think of it, but human "hearing" has quite a bit of nonlinear distortion associated with it. In order to survive in a world with sound intensity differences of greater than a million-million, we have developed somethink like a logarithmic response. Just like the "blind spot" in vision we accomodate to this and are never aware of it. But we know that in most sounds distortion tends to increase with increasing intensity. Louder means more distortion and conversely. Not so in many of the early transistor amplifiers. As you dropped signal level in such amplifiers the distortion rose until the fundamental was nearly "snuffed out" by the crossover notch.

Low level instruments to left of center would have a left channel signal with significantly lower distortion than that of the less intense right channel signal. The sound image was then spread all over the stage with much of the lateral smear tied to distortion fragments.

In addition, the nature of the distortion fragments was such as to constitute very high order harmonic terms, in contrast to what we have to accomodate in normal hearing. The result was listener fatigue, a very brittle sound which exaggerated instrumental attack, and a natural tendency to turn the power amplifier switch to the off position.

There were, and are, other types of distortion. But by considering the geometry of the final sound image as processed under human subjective conditions, harmonic distortion can itself be sufficient to state that such amplifiers would sound bad, even though the conventional method of measuring such distortion gave the amplifiers a clean bill of health.

Getting back to our original train of thought, nonlinear distortion of the sound image when it is processed one-dimensionally will lead to a differential smear of signal energy among the higher dimensions of the final sound image. Depending upon the detail nature of the one-dimensional distortion, the sound image will be smeared in its space, tonal, temporal, and intensity values. This is representation distortion. In some cases, such as with amplifier crossover and slew rate distortion, the image may be laterally smeared with distortion fragments as though you had brushed your hand across a still-wet painting. In other cases, portions of the sound image may distend back from you or accordian together to lose all depth relationships when time and space values are affected. In still other cases, instrumental voices may recess and suppress themselves when temporal and intensity values are affected, etc.

Most of this is available in our distortion measurements when we are careful enough in making them. And particularly when we are careful in the interpretation of these measurements in subjective terminology.

What about our loudspeaker which we so far have left dangling? Harmonic distortion in a one-dimensional amplifier is a total thing. If you've got such distortion it must smear up some of the coordinate representation in the final sound image. No way out. Not necessarily so in a loudspeaker. Harmonic distortion in a loudspeaker just might foul up primarily the tonal properties and leave some of the spatial, temporal, and intensity values relatively unaffected. Harmonic distortion may of course do the same type of thing for a loudspeaker as an amplifier does to the sound image, but that's more the exception than the rule.

When we set up a single microphone and measure sound pressure from a loudspeaker at one point in space, we are not measuring the same thing, as far as the sound image is concerned, as an amplifier measurement. Many of the spatial and temporal properties of the final sound illusion are embodied in where the loudspeaker is placed and in what room. In many cases these tend to overwhelm the modification in representation due to what we measure as acoustic distortion.

This leads to a rather interesting possibility. There may be some acoustic environments in which a certain loudspeaker will sound very accurate. That same loudspeaker in other places may be so-so, or possibly sound noticeably distorted. If you forget about measured component distortion and think in terms of the geometry of the final sound image, another possibility presents itself. It just may be possible that there is a preferred amplifier-loudspeaker combination for a more accurate sound. Such statements made only in the context of impulse response or frequency response measurements would be cause for academic court martial and

automatic banishment from the world of objectivity. But maybe, just maybe, we are beginning to solve some of our mystery by taking a broader view.

Harmonic distortion is only one small piece of the picture. We'll look over some others in our next discussion.