How to Build an LSD

By Roy Leamon

Table of Contents

0. Introduction....3
1. Concepts....7
2. The Analog Audio Playback LSD....15
3. The Analog Video Playback LSD....29
4. The Analog Audio-Video Playback LSD....41
5. Analog Audio and Video Capture and Recording LSDs....44
6. Hybrid Digital-Analog Recording and Playback LSDs....50
7. The Fully Digital LSD....62
8. Programming a Digital LSD....89
9. LSD Layers (still to come)
10. LSD Shapes and Structures (still to come)
11. The Business of the LSD (still to come)
12. The Future of the LSD (still to come)

0. Introduction

Welcome to my book about the LSD. After many years of working on the concept, and almost as long trying to figure out what to do with the idea, I finally hit on simply writing a book about what I have discovered.

Who am I, and what is the LSD? Let's try to answer the second question first.

The LSD is a strategy for using technology to create a device that reproduces sounds and images at resolutions never before experienced. "LSD" stands for *Light Sound Dimension,* and is so named because it has the potential to convey the impression that the audience has entered a new dimension of reality based on an unprecedented wealth of visual and aural information.

It can be like a movie theater on acid, appropriately enough. But the LSD is not a single device, in a single size and shape. It's an approach, a protocol, for building a limitless variety of devices, each of which can have its own properties. For instance, one LSD may reproduce sound, another may record video, while others may do both and more.

The words "record" and "reproduce" have specific meanings depending on the state of the art in reproduction and recording technologies. Thomas Edison's first audio recordings lead to what we have today, but sounded quite different! The LSD is *another* redefinition of the state of the art in both recording and reproduction, in both audio and video. It leaps way past the way things are done now. As you read these pages please bear in mind that what you are used to thinking of as "recording" and "reproducing" reality is rendered moot by the LSD.

But what *is* the LSD? Why is it so different, and special?

The LSD uses *arrays* of components to generate a much, much higher resolution of depiction, in both audio and video, than humans have ever seen before. Indeed, the LSD is theoretically capable of *exceeding the ability of humans to perceive audio and video information*. When the world you are experiencing is as real, even more real, than the one you are used to, who is to say reality itself has not been breached, even conquered?

A grandiose claim, to be sure. But let's not get ahead of ourselves. The following pages will only do small justice to the concept, but by the end of them you should have a clear understanding not just of what the LSD is and how the forms of it I have discovered so far can be built, but of its potential to sweep away everything that's come before. Suffice it to say that in my mental travels through this new world, some of the devices I have envisioned are almost *too* powerful, *too* intense; the aural and visual worlds that could be created in them too *beautiful* to grasp for more than a moment at a time.

As you proceed, I hope you will take a moment to try to imagine the audio-visual world an LSD can create. An entire chapter is devoted to programming the LSD, but it is only the merest start. If for no other reason, I hope the LSD will be built so I can experience the amazing things artists can do with one.

Finally, who am I? Probably not the first person who would spring to mind as a candidate for coming up with all of this. I majored in English, not Engineering; I love my Mac, but I don't know how to program it; I appreciate good music but have never recorded it in a studio; I don't know how to work a soldering iron or design a building. I'm also not an entrepreneur in the usual sense, and don't want to be megawealthy. I have always worked in rather menial jobs to put food on the table and keep a shelter over my head.

In other words, I am not ‘qualified,’ in the usual sense of the word, to suggest that the LSD is as revolutionary as it in fact is -- I don’t have expert credentials. I hope and encourage the people out there who *are* experts to pick the thing apart and let the idea’s flaws — and virtues — be known.

Most importantly, though, I didn’t ask to be inspired to focus obsessively on the LSD, which is one of the reasons I have held it close for so long. You don’t receive visitations in the night of ideas this powerful without feeling humility, fear, and some reluctance to go stand on the streetcorner and shout at strangers about it. But the vision has been relentless, and the weight of it too great not to let it finally go. I will not live to see my furthest designs built, but hope to experience at least *something* close to a real LSD before I die.

So if you feel compelled, after reading this book, to go out and start building an LSD, or to contribute to my efforts to develop the digital version, please go right ahead. Let’s get started! I am happy to answer any questions I can if you’ll write me at roy@infinitebaffle.com.

This book is the distillation of 10 years’ work. I hope it will help explain to my friends what I’ve really been up to all this time! My only regret is that my father, who loved and supported me every inch of every step of the way, regardless of how ‘successful’ (or mostly *not* successful) I was, did not live to see this published, but most especially, to see an LSD actually built. He would have loved it.

To him I dedicate this work.

Roy Leamon III

Austin, Texas

1. Concepts

What is the LSD? Perhaps the easiest way to get a feel for the concept is to understand some basic principles behind it.

When we look at a television, or watch a movie flicker on a screen, we are unconsciously accepting a basic premise – that the depiction of reality on that rectangular screen, with those few speakers, is acceptable, despite the fact we don't actually *see* or *hear* in such a limited way. Like a horse wearing blinders, we see a patch of image, always in front of us, always bounded by a frame. We might as well be wearing nearly-soundproof earmuffs at the same time.

Consider stereo, which has dominated audio reproduction for decades. Stereo recordings consist of two channels, or *streams*, of sound, designated Left and Right. Stereo has served our purposes well; the rich variety of music that has been created in stereo is astounding when you think about it. But it's also astounding to imagine what artists could do with a much greater number of channels.

The screen's frame, and stereo's limitation, impose an order that is a source of both discipline and imprisonment. The aesthetics of motion pictures, for instance, are rooted in the nature of film's presentation. Would movies be edited as they are if the screen completely surrounded you? Having to confine one's gaze to a single forward screen, or a few speakers, forces choices that would no longer have to be made if the artist was creating for an LSD. We cut movies, and mix audio, to fit through narrow portholes.

But our eyes are at the front of our heads, you say, and we only have two ears, so isn't any extra information wasted? To answer this, I want to take you on a walk through a beautiful forest. But before we go, you must choose between two ways of experiencing the forest. The first is to put your head inside a hollow cardboard box with a rectangle cut out of one end, such that appears you are looking through a hole 27" in size diagonally. A piece of glass in that rectangle reduces the resolution of what you're seeing well below what your eyes are capable of seeing. The box also has 3-foot-long tubes sticking out from the left and right sides, which diminish all the sounds of the forest to two streams of audio but take away your ability to locate the sounds' sources with any more precision than their being to the left or to the right.

The box is lightweight and ventilated, so you'd feel very comfortable wearing it on your walk; aside from your senses of sight and sound being significantly diminished, you'd hardly know you were wearing it.

Your second choice is to walk through the forest with no box on your head, your ears and eyes free to hear and see as you were born to do.

Which option would you choose? What if I told you you couldn't turn your head while wearing the box either, unless I chose to do it for you?

That's what I thought. But we *do* choose to see and hear the world in this way—wearing the box—every time we turn on a television or sit in a movie theater. Again, within those limitations, amazing sounds and images have been created that have fed our culture for almost a hundred years now. But now that we have the means to create something much richer and freer, is there any reason we should not break free from these bonds that confine us so?

An LSD is a device that can liberate us in this way. Although it could never replace an actual walk in the forest – LSDs are not intended to reproduce scent, taste or touch information – it can record and play back the image and sound information so much more vividly than any other technology you've experienced that your feeling of *being* in the forest will be much greater than you can imagine.

How is this done? As I write this in 2002, the technology used in the device is based on using the speakers and screens (monitors, televisions, etc.) available today. As audio and video technologies progress, it's likely new components will achieve the same effect but with much-diminished infrastructural needs and even greater quality.

Today's technology is amazing—the power available in commonly-available electronic components that you can purchase at any consumer electronics store is far beyond what anyone in the early part (or even middle part) of the 20^{th} Century could probably have imagined. Computing now allows an entirely digital signal path for audio and video, and large numbers of analog and digital components can be bought for relatively little money. The awesome, cheap, and programmable power of today's technology is what makes the LSD conceivable *and buildable*.

We're used to thinking of using this power for personal use only. We might hook computers together in a network, but individuals still use each computer. One of the most exciting moments I had while discovering what an LSD is came when I realized there's no reason we can't *combine* these 'personal' components to create a much larger, unified device. The combination of components is called an *array*.

The LSD uses arrays of components to produce a much larger and more intense effect than any single component ever could by itself. These arrays are built out of the best components available (or affordable). Each component in the array can handle its own *discrete* channel or stream of audio or visual information.

The fact that each information stream is discrete is vitally important to understand. When I write of combining screens, say, in an array, you might imagine I'm talking about something like a video wall, which is a stack of four or more televisions set up to show a single image. The difference between a video wall and a real LSD is that the video wall displays a single stream of information that has been electronically divided into substreams. A four-screen video wall shows a *quarter* of the original signal on each screen. The LSD doesn't do this – the goal is to have each screen display an *undivided* stream that is displayed at the *highest resolution possible* for that device. The same holds true for audio – each speaker should receive a pure full signal, rather than just a portion of one.

The arrays in an LSD can take many shapes, as determined by the overall structure of a particular LSD. You might build your LSD to resemble in size and shape a motion picture screen, but you could also build one as a box, a circle, a tube, anything really. You can combine these shape primitives to create very sophisticated audio-visual architectures.

And you don't have to start out knowing exactly what you want to do (though it probably helps!). You can begin with just a few components and add more as time goes on—your LSD can *grow* from a small seed. You could start out with a few dozen speakers, for instance, and eventually add more and more, with no predefined upper limit to how many your LSD contains.

LSDs can also grow in *quality*. As better speakers are produced or become affordable, they can replace older, lower-quality speakers. Even better, though, the LSD does not require obsolescence. The older speakers can simply be displaced to other parts of the LSD to grow it even more.

The arrays in an LSD might be built of tens, hundreds, even thousands of components. These components could be speakers, televisions, computer monitors, microphones, video cameras, digital still cameras, amplifiers, computers and so on. Depending on which components are used, the arrays can generate and/or record sound and/or image in an inward and/or outward direction.

The simplest LSDs, and the most limited in function, are analog, meaning the information traveling to the output devices is generated or delivered by varying pulses of electrical power over wire, as opposed to purely digital information. I'll cover analog LSDs first.

Next comes analog-digital *hybrid* LSDs, in which the signal may be generated or routed by computers but travels over analog signal paths for output on analog components.

The last category is the fully-digital LSD. This is the holy grail, the apogee, of LSDs, at least as I've imagined them. In a digital LSD, the sense information is as purely digital as possible, and the sound and image recorded and displayed can be manipulated or even created from scratch by computation.

So what does an array of components in an LSD allow an artist to do? For one thing, it allows him or her to manipulate the *location* at which an image or sound will appear. The LSD's arrays form a canvas that mimics the locational possibilities found in nature much more closely than any previous medium. With only one screen in a system, an image doesn't have many options for where it's going to appear; a stereo system only gives a sound two options for where it can play. If I design an LSD that outputs 100 channels of sound over 100 speakers, though, I can decide which sound goes to which speaker, and then move the sound to another speaker or set of speakers, and decide how that move would happen. The same holds true for screens.

The ultimate expression of LSD technology is a completely immersive experience that enables an audience to hear sound from any location at any time and to see image in any location at any time—and allows the creator to *manipulate* those locations *over* time.

As if that weren't enough, I want to close this chapter with probably the most important, and, for me at least, the most mind-bending quality the LSD possesses: *the perceived resolution of an LSD is incredibly high.* Much higher than any of us has ever experienced with a device built by human hands.

To understand how this works, you must accept the fact that when we talk about 'resolution' with any device we are really talking about *perceived* resolution. Hold a baseball at arm's length, then throw it 20 feet away. It looks smaller now, right? The same holds true for pixels and sound sources. The actual resolution of a component – say, 1600 pixels horizontally across a 17-inch monitor – appears differently when your eyes are one foot away from the monitor versus ten feet away.

An LSD takes advantage of the relationship between distance and actual device resolution to wring the highest possible *perceived* resolution out of its arrays. The bigger the array, the farther you would generally stand from it, and the 'finer' the perceived resolution is.

The resolving power of the human eye has been estimated at 17,000 pixels per inch. Built correctly, an LSD is capable of achieving—and surpassing—this resolution *as perceived*, without requiring you to build a monitor that has a *native* resolution of 17,000 pixels per inch. What effect a display device that operates at the very limit of human visual perception would have on an audience is difficult to say, but I think it'd be interesting to find out.

We'll get into visual resolution in much greater depth in later chapters. Audio resolution is also 'finer' with larger devices; a sound moving across an array of 1,000 speakers will seem to move much more fluidly than if it covered the same distance, in the same amount of time, over only 10 speakers.

The idea of a thousand speakers arrayed around you might provoke some alarm; wouldn't it be unbearably loud? It would be if you cranked every speaker to its limit (rock on!). Our current schemes of audio reproduction involve pumping the entire world of sound through just a few speakers, so we're used to turning up the volume so we can hear each element in the mix as distinctly as possible. In our thousand-speaker array, each speaker handles much less of the work of producing the overall sound, so that rather than producing a *louder* sound, a thousand speakers would produce a *bigger* sound, one that is spread across a larger area—even *soft* sounds can be *big* on an LSD. And since each speaker is handling a much more manageable portion of the overall audio 'image' than in traditional scenarios, the sound is likely to be cleaner as well.

The richness, subtlety and beauty possible when a single unified sound—consisting of hundreds of subtlely different channels—is played over a thousand-speaker array is almost too frightening to consider. Combine it with a supremely high-resolution visual image, and things start getting fun.

2. The Analog Audio Playback LSD

My first work on the LSD resulted from an attempt to bring back the dead – the Grateful Dead, that is. I had seen them in Las Vegas in 1993, and returned home to find myself wondering what would happen when they stopped playing: how could such a grand experience be recreated?

If you aren't a Deadhead, the grandness of the experience will be hard to understand; most attempts to describe it fall pathetically short. But it inspired me to start thinking about sound, and the way we perceive sonic reality, and it soon became clear simple stereo, or even multichannel sound as Dolby Laboratories has popularized, falls way short of what we not only can hear, but *should* be hearing.

One of my early efforts involved building a cube-shaped frame about three feet tall out of 2 x 4s, stringing the sides with a few strands of rough twine, and attaching to the twine three Walkmans driving two small battery-powered speakers each, for a total of six channels of sound. Using probably the cheapest MIDI setup ever made for a personal computer, I created a 30-second long six-track composition that would've gotten me kicked out of any music school. I recorded two channels at a time to cassette tapes, popped the cassettes into the Walkmen, and carefully positioned myself inside the cube.

While this takes just a minute to write, we're talking probably six hours work just to get to this point. The moment of truth was at hand, but unfortunately the truth I discovered was not the one I was hoping for – without an extra set of hands, I would not be able to start all three Walkmen at the same time, and any hope of even a roughly synchronized rendition of my composition was lost.

Not having had high hopes of musical glory in the first place, I started two of the Walkmen and quickly started the third. And with that, the LSD was born.

To build an analog audio LSD, probably the simplest LSD to build, you must first lay out a design. How big will it be to start? What components will you use, and how many will you begin with? What will hold the components? What audio program will be played over the array, and how will you create it? Will the components be synchronized, and if so, how? What method will you use to start, pause, stop, and rewind (if necessary)?

The range of answers to these questions point to the tremendous diversity of LSD designs and experiences possible. An LSD in a small room will have a very different effect than an LSD erected in a large open space. An LSD with ten speakers generates a lot less information than one with a hundred speakers, but depending on the space the LSD is built in, ten might be a better number. The type of audio program you're intending to play over the LSD makes a big difference too, and whether or not the components in the array are synced will have a lot to do with what kind of program you choose.

Let's start with the absolutely simplest LSD to make – a non-synced array with, say, 20 speakers. The device will be driven by VCRs, which are cheap, generally reliable, and can play for a couple of hours without pause. Also, if you get the same brand and model of VCRs, it's possible you can start, pause, stop and rewind them all at the same time using a just one or two remote controls. They are not likely, however, to stay in sync over the course of playback. If you have the money and expertise, you could use time-code-synchronized video decks like the ones used in video editing. They're expensive, but you could probably get precisely synchronized sound out of your array if you use them.

For now, though, let's focus on simple off-the-shelf solutions. Each VCR will be connected to a standard amplifier, which will drive two speakers each. (If you buy different brands of amplifiers, you can control each one individually, which would help you 'tune' the LSD using volume and equalization settings.) We are driving 20 speakers, so we'll need 10 VCRs and 10 amplifiers, the cable to connect them, a good power source, and 10 videotapes containing our 20 channels of sound.

How do we build those 20 channels of sound? A variety of ways come to mind; I'm sure there are many more. Probably the simplest way is to use one of the many products for personal computers that let you compose, record, and play back audio tracks of various kinds. You could also record ten sets of live stereo sounds. Create a 20-track composition, then lay off two tracks at a time to each successive videotape.

Don't forget while you're composing, or while you're laying the tracks off, that you can move a sound from one track to another, or more literally, you can rig what's on the videotapes such that a channel moves from one speaker to another in your array. And don't forget the power of silence, either. Assaulting the audience with 20 channels of sound all at once is certainly one strategy, but starting off subtlely, with, say, a single sound that bounces around a little from speaker to speaker, will help ease your audience into what your LSD is capable of doing, and builds in a little suspense as the number of channels playing at once keeps rising. Don't forget that as rough as it is, your simple 20-

channel analog LSD is capable of doing a lot more than anyone in your audience has probably ever heard before from a sound system.

What structure have you decided on for your LSD? The simplest is probably a circle; find enough space where you can set out 20 speakers in a ring around you, with enough distance between them to make their discrete outputs truly distinct from one another. Is your LSD going to be inside or outside? Of course there are trade-offs either way. Inside is protected from the elements, probably less bothersome to neighbors once you start cranking your device, and power is easier to access. The downside of an indoors LSD is that the space has to be big enough to hold all that equipment and a room may not permit a good clean strong sound without your altering the room acoustically to absorb sound reflections.

Outdoors isn’t protected from the elements, and it’d have to be fairly remote from people who’d mind all the noise you’re going to be making. You also need a safe source of power. But so long as there aren’t a lot of hard reflective surfaces nearby to bounce sound back at you, each channel is likely to come across more discretely outdoors.

You could build a wooden shelf-like structure to hold the components, or just lay them on a piece of plywood on the ground. If you plan well, it won’t take long to set your LSD up and take it down when you’re done.

You could also lay out your array as a wall, or make four towers of five speakers each, with space between each cluster. You could suspend some or all of the speakers overhead. You could place them in different rooms in a house and turn the entire house into a large speaker box.

Is each speaker the same as the others? Possibly, but not necessarily. Your audio program is likely to contain a range of frequencies from high-pitched sounds to low-pitched ones: a flute, French horn and tuba, for instance. Or a soprano vocalist, an electric guitar and a bass guitar. Or weird computer noodling that goes all over the place.

Considering these differences in pitch, it's very likely you'll want to choose your components accordingly, to vary the types of speakers you use according to what kind of sound will be coming from them. If the output is going to range widely, a good general speaker with adequate frequency response will work fine. If a speaker is going to be used exclusively for high- or low-pitched sounds, then choosing a tweeter or subwoofer (respectively) that's very efficient at reproducing those sounds makes sense. I'll discuss speaker selection and differing frequencies in a bit, when I get to LSD stereo spatialization.

So you've set up your components in the array of your design. You're powered up, you have the videotapes cued correctly in the VCRs, and you've verified you have good connections all around. The moment of truth has come. Lift your remote control (if you have a team armed with remotes, do a countdown), say a little prayer, and hit Play.

While it's perfectly fine to keep your eyes open while you listen, you might want to get into the habit of closing them so you can listen without the mental distraction of visual information.

Your first attempts to program your LSD are not likely to be your strongest, and the work you are having to do to create 20 channels of sound is not insignificant. A degree of frustration is inevitable. Remember that you are exploring a new medium, and discovering what works and what doesn't. Considering how awful most movies still are, after all these decades of 'professionals' making them, don't put a lot of pressure on yourself to master the LSD in your first few tries.

So it didn't sound as good as you wanted. Feel free to change the configuration of the array as necessary – move speakers, adjust volumes, even add to or reduce the number of speakers you're using. Moving the entire system to another space might help it sound better too.

Try recording different kinds of audio programs as well. There's a certain randomness that's inherent with a non-sync system; play with that a bit (and certainly don't try to attain sync unless your LSD is built for it)—and remember that at least *two* channels in your system are in sync with each other. Play with location as much as you can without going crazy trying to figure out where a sound should appear, move to, and end, and certainly don't *over*do the location effect.

When you finally get your LSD rocking, invite friends over to give it a listen, and try to incorporate their feedback into your design and program. While it's perfectly fine to create an LSD solely for personal satisfaction, there's a lot to be said for creating LSDs that are enjoyed by groups of people, especially as the devices get bigger and more sophisticated than the analog one you've built.

In fact, you may have enlisted friends to help build your LSD, or it's a communal effort that belongs to all of you. Many people have an old stereo system or VCR they either don't use or will replace some day; why throw out or give away such perfectly good components, or sell them at a steep loss, when they are perfectly adequate for starting, or adding to, an LSD?

The blueprint I have described for creating a 20-channel analog audio LSD is one of many, many routes you can take. You could bypass separate amplifier components and use self-powered speakers. You could use any number of analog sources, and they don't even have to be consistently the same within a single LSD: tape decks, CD players, DVD players, boom boxes, radio receivers, live musical instruments and vocals. Heck, you could even use a Walkman.

All of this is fine and good, and a lot of fun to play with, but at some point you will want to play regular prerecorded music over your LSD, probably from a CD. How in the world are you going to do that?

As I mentioned earlier, analog LSDs are the least flexible. You pretty much have to design your LSD to do whatever specific task you want, and if that task changes, your design may have to change radically. If we were to start from scratch we could build an LSD intended mainly for prerecorded music, and it would look somewhat different from what we have so far.

But let's assume we're using our 20-channel rig. Since we are playing from a single CD, we don't have to worry suddenly about synchronizing our VCRs; they're bypassed when playing prerecorded music (in this scheme at least). Instead, we need to get the signal from the single CD to all those speakers.

We'll start out easy – we just want to split the signal enough times to get output from every speaker in the array. I realize this is sort of the same thing as the video wall issue I pooh-poohed in the last chapter, but remember, this is just a start, and we're dealing with a non-native LSD audio source.

We could simply take the two outputs from the CD – left and right – and feed them into an amplifier, which could then drive ten more amplifiers using Y-connectors. You'd have to watch out for distortion as the signal is amplified more than once. You could also use the Y connectors right off the bat, and drive the 10 amplifiers directly from the CD, but your signal strength might not be strong enough to give each amplifier enough to amplify cleanly.

Another method – I've actually done this -- is to drive multiple speakers from a single amplifier speaker output by *splitting Monster speaker cable* into as many threads as you need. I was surprised that it worked for me as well as it did.

Either of these methods – Y connectors splitting signal to all your amps, or splitting actual speaker cable to drive more speakers from a single amp – gets the stereo signal to your speakers. This is great in its way, but it sort of sucks considering you have an *LSD* here, not just a simple stereo.

With a little clever design we can make those 20 speakers work for us to make even two channels sound much better than they would on a normal, traditional stereo. Remember that fundamental to the LSD world is the notion that you are *dimensionalizing* audio and video information. What we want to do is dimensionalize the stereo sound. I'm suggesting that we use a simple technique to make the each channel seem bigger, cleaner, and more interesting to listen to than it normally is. That way we can feel like our 20-speaker LSD is *improving* stereo rather than just trying to recreate it.

We're used to hearing two channels, left and right, but we aren't used to hearing them spread over space so that differences *within* each channel are perceptible as the differences *between* channels are. So rather than create a box that simply splits two channels into 20, how about one that does that and something more. You may have guessed by now what it is – it's a hack of a box that you can already walk into any consumer electronics store and buy. (I was an idiot and did it myself, but could easily have been electrocuted or badly shocked. Do *not* do this unless you know what you are doing!)

What you want is a 20-band graphic equalizer. (If your LSD had more speakers, you could use graphic equalizers that have more bands.) A graphic equalizer takes the incoming stereo signal and splits each channel according to frequency from lowest to highest, so that within the equalizer you actually have 20 separate channels (10 for left, 10 for right) that can be boosted or lowered in volume. After passing through the volume adjustment area, the signals are mixed back down to left-right channels and sent to the amplifier.

My hack involves intercepting the signal after it's been split by frequency. It can be intercepted either before the volume adjustment component or after. I used my hacked box with self-powered speakers that had volume controls on them, so I didn't necessarily care to have the ability to adjust volume at the equalizer box. (Okay, so I didn't know what I was doing and intercepted the signal at the wrong point, and have to rationalize it somehow.) After cutting each of the 20 wires and adding wires for the ground, I drilled holes in the front of the box and attached 1/8" stereo plugs, and attached each wire (with a ground wire) to the plugs. When I was done, I had a box with 20 separate outputs divided by left and right, with each channel divided into 10 frequency bands.

Why do we want to output 20 channels of sound that are differentiated by frequency? Because it dimensionalizes the stereo sound. If we didn't do this, every speaker would play the same thing as every other speaker carrying that signal. With our EQ hack, we actually get ten different signals per left-right channel, or 20 discrete channels in all. And it just so happens we have a 20-speaker LSD. Serendipity!

To make this work, you'll take your CD player and run its signal through the hacked EQ. From the EQ you now have 20 separate outputs. Since each amplifier handles two streams, pair up wires in a way that makes sense (or randomly even) until you have 10 pairs. These are fed into the amplifiers, which drive the speakers. If you're using self-powered speakers, you can run each of the 20 channels to its own speaker.

I'd probably arrange the speakers in a sort of left-right tower format, with the bass frequency speakers at the bottom of the tower and the treble at the top, at least for now. Depending on your LSD's setting and the components used, your strategy may differ.

This leads to the promised mention of matching speakers to frequencies. If your LSD is geared solely towards reproducing dimensionalized stereo sound such as what I've just described, you'll want each speaker to be as closely matched to the frequency of the output it will be handling as possible. That means bigger subwoofer-type speakers for the lower couple of rungs, smaller tweeters for the highest rungs, and good woofers in the middle.

What we're in fact doing is traditional speaker design on, well, acid. A normal multi-cone speaker cabinet has a *crossover network* inside it that splits the incoming signal according to frequency and sends each split to the cone best suited to handle it. LSDs playing frequency-dimensionalized stereo (or 5.1 format for that matter) are doing the same thing, but it sounds a whole lot better when you give the music a chance to open up and spread out even more.

This concept is the most advanced one I have actually built, in 2001 when I lived in a small house in California. (Aside from picking up a Spanish-language AM radio station, the hacked equalizer box worked great.) My final count on speakers was 27 pairs. They were all self-powered computer speakers using 1/8” connectors.

None of the speakers cost more than $20 a pair. While they weren’t the best speakers in the world, the fact I wasn’t requiring a single pair to carry *all* of the sound actually helped the overall sound seem clearer, which validated my notion of multi-speaker arrays providing *bigger* sound that isn’t necessarily *louder*.

I nailed shelf brackets into the walls on either side of the room. I strung twine between the brackets, and tied the speakers to the twine, so I had two suspended speaker arrays of 27 speakers each on opposite ends of the room. One array got the left channel, the other the right channel. Starting at the lower-frequency speakers, I connected the signals from the EQ box to the speakers. Since I had more speakers than signals, I used Y connectors to send some of the signals to more than one set of speakers.

This led to some interesting fun with determining which frequencies sounded better when spread over more speakers, and which frequencies don’t need as many channels. As a rule, I found that the high-middle and low-high frequencies – but not the very highest – benefited the most from having more speakers playing their signals. This makes sense from a psychoacoustical point of view; sound grows less directional as the frequency is lower (which is why your subwoofer can be placed ‘anywhere’ in a room), and having a more vivid sense of the physical location in space of the higher frequency sounds gave the overall sound a greater sense of fidelity.

In other words, the overall sound was more 'there,' more alive and spacious, than simply using two speakers (which I did try for comparison). The dimensionalization using frequency splitting worked, I am happy to report, though as soon as I had fired the system up I was ready to tear it down and rebuild it bigger (and preferably outside).

But I haven't done that yet; I've subsequently moved, and all of those speakers sit now in boxes. The equalizer met a sad fate when I accidentally touched the power supply with the wrong wire (it's a good thing I didn't meet the same sad fate; don't go opening electronics and hacking their insides unless you know what you're doing, and ALWAYS UNPLUG ANY COMPONENT YOU ARE WORKING ON). Another attempt to hack a different brand of EQ resulted in complete failure. The best approach is to pay someone to do it right – to hack an EQ box to produce good, clean dimensionalized sound without picking up an AM radio station!

The hacked EQ isn't the only analog signal processing you can introduce to your audio streams. Although much audio production work is done on computers now, "outboard equipment" that fits in a rack takes audio signals and does things to them using chips and filters. Since the audio postproduction world is generally moving to a purely computer-based workflow, good used rackmount equipment should be easy to find and relatively cheap.

What can you do with a piece of sophisticated signal processing hardware like this? Well, in addition to (or instead of) simply splitting audio according to frequency, you could add reverb or echo or both to sets of speakers in your LSD.

You could build speaker 'stalks' or towers and arrange them along the sides of a long tall room or open area. The middle stalks could play unfiltered sound, or frequency-dimensionalized signals. The stalks behind you could be filtered using reverb, while the stalks in front could be filtered using echo and delay. With enough outboard equipment you could change the filters used on these signals in real time as part of a performance.

Other types of effects are possible in a purely analog LSD. A box could be built that moves signals from speaker to speaker in response to a simple controller. (The analog box that runs the Audium in San Francisco is apparently something like this.) The type of equipment modern DJs use in dance clubs could fit in very nicely with a modestly functional analog audio LSD.

The audio source doesn't have to be a CD. It could be live music, but it could also be generated by a computer, with all its signal processing capabilities invoked, and even with some intention towards manipulating location based on where a sound occurs in the frequency spectrum. To illustrate this last point, suppose we have a sound that's in rung 3 of the left channel. If we kick up the sound's frequency to rung 5 and route it to the *right* channel, we've 'moved' the sound by manipulating the frequency band where it occurs. You wouldn't necessarily need a computer to do this, but it might be a lot easier with one.

Remember that even with 20 speakers (or 54, as I was using), we're at the very lowest range of the number of speakers needed to make an LSD sound like, well, an LSD. The speaker count could be 200 or 2,000 – just repeat each unit of components as necessary, and add more amplification where necessary to keep the signal good and strong.

The analog audio LSD is a great introduction to the concepts that drive all LSDs, and the effects possible even with relatively simple, inexpensive, off-the-shelf components point to how powerful LSDs can get when more sophisticated components are used to drive them. But if you want to get into the LSD realm, with enough care and planning – and a healthy dollop of imagination – you can create an analog audio LSD that will stay amazing for a long time.

3. The Analog Video Playback LSD

Building an analog video LSD – or any other type of LSD, for that matter – follows the same basic principles used to build an analog audio LSD. You build an array of many components as far away as necessary to make the perceived resolution of the components as high as possible.

The *signal source component* for an analog video playback LSD can be a VCR, VCD or DVD player, a video camera, or any other means of getting video signal to the output component. Even a computer with the right video card could be used.

The *output component* can be any sort of television or digital projector, or a computer monitor if full-screen video can be routed to it.

As with the audio LSD in the previous chapter, if you use VCRs or DVD players that will respond to the same remote control, it may be possible to start and stop all of these components at once, for a very rough sync. (How closely they'd stay in sync through the rest of the program is a big variable.) Or at least you might be able to start all of them at once without having to go from machine to machine.

Synchronization between analog signal source components might be feasible if you invest in professional-quality videotape decks, but perfectly interesting video LSDs can be built with off-the-shelf components that play either in rough sync or with no sync at all.

The design and layout of your video playback LSD can follow any whim, depending on the effect you want. Among other designs, I've imagined the following types of video LSDs: simple rows of televisions; multiple overarching (or partially overarching) curved screens; a single large flat screen you'd watch like a movie screen; a sphere; a dome; a ring or series of stacked rings; and on and on.

The information you are presenting will help you decide on your design: from a single color filling each screen to a single large image that is either still or moves. With an analog system, however, you have to record each and every one of those videotapes or DVDs if you choose to display a prerecorded program. If you are using live video feeds your upfront work is reduced significantly. (We'll get into recording LSDs in Chapter Five.)

Your cheapest and easiest route to take with prerecorded analog video is probably an LSD unit – i.e., the basic building block of your array – containing a television and a VCR. Since you're only working with video and not audio at this point, you don't need to worry about adding speakers or amplifying sound. It's probably easier to go with a standard size of televisions across the array, unless you like the effect of your program playing back on a variety of sizes. Don't forget, however, that smaller televisions display the same number of pixels as larger ones *in a smaller space*, so it's possible their perceived resolution might be higher. When using a mix of sizes, you may even want to cluster the smaller televisions together where, for instance, the resolution of the displayed image should be highest.

The as-yet unresolved problem with video LSDs of any kind are the screens' bezels, meaning the plastic frame that runs around the screen itself. Some televisions and computer monitors, especially flat-panel displays, have minimal bezels; some projection televisions, such as those used in certain models of video walls, have no bezels at all. But if you're building a video LSD, you're very likely going to have to deal with the unhappy fact of bezels separating the displays more than you'd like for them to.

A few ways of dealing with this unhappy fact suggest themselves. First, choose components with the least amount of bezel around the edges as you can. Second, overlap them slightly so one component's bezel obscures the adjacent ones' as much as possible. Third, *if and only if you know what you're doing*, remove the component's case and either mount the screen naked or put it in another cabinet that has less bezel. Fourth, mount a piece of glass with slight magnification (i.e., a lens) over the front of each screen so the image goes to the edge of the glass.

The most important tool for minimizing the bezel effect, though, is to build as large an array as you can, and place it far enough away from the audience that the bezel interruption *appears* minimal. After all, if you look closely at most tube-based television screens you'll see a black screen or mask separating each pixel (or more accurately, each red, green or blue pixel); it disappears as you stand farther away. If your video array isn't big enough to make the bezels appear somewhat minimal, it's probably not big enough in the first place.

Actually, the most important tool for minimizing the bezel effect is to make the programming so compelling that, like those Japanese puppet plays in which the puppeteers are onstage but wear black outfits, the audience soon ceases to notice it!

How big should your video LSD be? While using a handful of televisions will definitely work, you don't start getting into interesting places until you have dozens or even hundreds of display components in your array. Assuming you're using standard NTSC video with standard televisions, your resolution for each screen is only 320 x 240 pixels. To achieve the same 1,600-pixel horizontal resolution as on a good computer monitor, you'd need five televisions in a row; to achieve anything approaching LSD-quality resolution, you'd want at least dozens of televisions in each row.

Remember that LSDs are additive; if you get a row of five televisions and VCRs fired up, you can keep adding more – stacking rows on top and adding columns left and right or wherever you wish.

I would like to see a 30 x 30 television LSD. But I'd *love* to see a 300 x 300 television LSD. And to see a 3,000 x 3,000 television LSD would be more than any of us could probably handle!

Where are you going to get your televisions and VCRs? Search the want ads in your local paper or on the Internet. Go to thrift shops, garage sales and used electronics stores. Buy them off your friends when they upgrade to newer ones. New or refurbished VCRs can be had for relatively little cost, and as people buy DVD players it's entirely possible the number of used VCRs for sale will swell and the prices will come down even more.

Where are you going to build your analog video LSD? Unlike audio LSDs, you don't have to worry about sound waves bouncing around a space, though you do need to make sure the array won't be polluted by ambient or direct outside light. An old movie theater might be a good place for an LSD video wall; a space the size of a basketball court would work for a four-sided video LSD. If you can seal or otherwise protect the components from the elements, you can also set something up outside. If your LSD will only be showing live video, be sure your cameras are going to be recording something interesting and, I would suggest, constantly changing.

Though not as powerful as a fully digital video LSD, an analog video LSD can produce amazing images, even with little or no sync between components (assuming prerecorded content is used). Even if you have a relatively small analog video LSD, you can still have a lot of fun with it.

One of the best ways to get into visual LSD concepts — and to get a video LSD up and running as quickly and easily as possible — is to use live video. Unlike a video wall, which splits a single video source among several monitors (and degrades the image quality of the original signal), to be a true video LSD, each monitor needs to display video at its highest native resolution (320 x 240 in the case of NTSC video).

For live video, this means a square array of 10 x 10 televisions will be fed by a square array of 10 x 10 video cameras. We'll have an NTSC resolution across the array of 3,200 x 2,400 pixels (or 7.7 megapixels). This is quite low, relative to larger arrays and to digital LSDs, but it's sufficient for our purposes for now.

You *could* use HDTV sets and cameras to bump the resolution up significantly, but you'd probably be spending enough money to cover the funding of the *digital* LSD's development!

Suppose then you have a square 10 x 10 array of monitors -- to feed live video to each one, mount the cameras *in* the television array, or create a second frame or structure, even a mobile one, to point the cameras at whatever subject you wish.

Mounting the cameras in the television array itself gets us into some funky territory where the LSD's mindbending qualities start to become apparent. First, realize you have not one but *two* ways of mounting them – facing the same direction as the monitors, or facing the opposite direction.

What would these two camera orientations display on the television array? This might seem easy for you to figure out, but since I'm sort of slow, I had to rig some televisions and video cameras to make sure of the answer.

Remember that when a video camera is pointed away from you, it "sees" as your own eyes see – objects moving left to right across its field of vision are displayed as moving left to right on a monitor. But when the same camera is turned *towards* you, the opposite happens – objects moving left to right are displayed as moving *right to left*.

With cameras pointed in the same direction as the screens (i.e., pointed at you), and each camera mounted at the television closest to it, as you walk left to right in front of the array you will observe yourself entering from the *right* side of each television — your image will enter from the far side of *each* column of screens in succession, meet you halfway, and exit behind you. Put another way, the image doesn't move *with* you across the array, but approaches and recedes in the *opposite* direction of your movement.

This is only mildly interesting; after all, you can get the same effect at the video camera counter at an electronics store. What I really want to create is an array that displays a single unified image. To solve this problem when the cameras are pointed at you, you must route the cameras' signals differently than you might initially think.

To get a single 'logical' image with a 10 x 10 array containing cameras pointing at you, the signal from the Column 1 cameras must be routed to the corresponding televisions *in Column 10.* Column 2's signals are output at Column 9, and so on. If you had an odd number of columns in your array, the center column's camera signals would be routed to the center column's monitors.

What kind of image do you get with a live analog television array that has its signals routed this way?

Why, you get *The LSD Mirror.* No, it's not the same as an optical mirror. In fact, it has a distinct advantage: optical mirrors flip reality, and therefore don't show you what you really look like, whereas the LSD Mirror *reflects you as you actually appear to other people*. For instance, I part my hair on the left side. When I look in the mirror, the image reflected back to me shows my hair is parted *on my visual left*. If a woman is admiring my coiffure, though, she'll perceive that my hair is parted *on her visual right*. Optical mirrors lie; the LSD Mirror shows the truth. (To be honest, the same can be said for a single video camera and monitor as well, though it's probably a much smaller 'reflection.')

So what would it be like to walk in front of a 10 x 10 LSD Mirror?

If you start at the left side and walk to the right across the LSD's field of vision, parallel to the array, your image will enter the array on the *right* side and head towards you. You will meet it at the middle of the array as you keep walking; you will exit at the right, while it will exit at the left. If you stand in the middle facing the array and lift your left arm, the image of yourself will lift the arm that is on *your visual right*, as though you were looking at another you.

Could the LSD Mirror have any practical application? Well, I'd sure like to know what I actually look like before I go out into the world in the morning! I'd also like to know how clothes actually look on me at the clothing store. When I'm getting my hair cut, I'd rather see my appearance as it will look like to other people than the way it'll look only when I see myself in a mirror. And, less practically but perhaps more important, it's bound to affect the mental image I carry around of myself to see my entire body *truthfully* -- as I appear to others.

Let's move on to the other configuration, in which the cameras are mounted on the *opposite* side of the wall from the direction the televisions point. In other words, the television array is facing you, but the camera array is behind them facing away.

What would you have now? Why, you'd have *The LSD Invisible Wall.* You're seeing everything on the opposite side of the wall, *as if neither array were there*. Well, okay, you can see the televisions, but hopefully your perception of them will fade somewhat as you realize in effect you're *seeing through them* to the other side, that the array is essentially transparent. Since there are no cameras on *your* side, you're completely invisible to anyone on the other side. And remember, you're seeing the other side at full NTSC resolution, meaning each television is displaying at its highest level of detail.

When the cameras are mounted this way, their signals are sent to the television *closest* to them, *not* to the television on the opposite end. As my friend walks left-to-right (relative to me) on the other side of the array, her image is displayed on the Column 1 (left) monitors and jumps to each column as she moves across, exiting right. You don't have the same cross-array wiring problem as you had when you mounted the cameras on your side of the array. That's because the cameras are now oriented 180 degrees opposite from before, pointing *away* from you, which flips their signal — relative to you — and makes things work in a more 'natural' way.

You can build a single-sided wall with cameras on this or that side quite easily. But suppose you built a *double-sided* wall with cameras and televisions on *both* sides.

Depending on how you route the signals, this could result either in two LSD Mirrors, or a *completely* invisible LSD wall, which would allow you to see the other side and for the other side to see you, but without physical contact (considering there's a bunch of equipment separating you).

Cameras don't even have to be mounted *in* the television array. You could set them up in their own array structure in an entirely different location. And you don't have to have just one wall. You could always *start* with one wall, then add left, right and back (even ceiling and floor) arrays later. If your array were in the form of a cube and had live video feeding in from all four (or even six) sides, you could be completely protected from your surroundings physically but completely immersed in it visually.

Setting up the video cameras will take some finesse no matter how you mount them (more to come on LSD Cameras in Chapter Five). Once you've got your camera array set up, you'll need to turn and zoom each camera to the appropriate point for it to capture only so much as it should, with little or no overlap between it and the camera next to it. Check the output on the screens as you go along, and don't forget to take the bezels into effect. Zooming your camera array in and capturing a very *small* area to display on your 10 x 10 array of televisions should be just as possible as capturing an image that's as large as, or even larger than, the array itself. Not having built and 'focused' a camera array yet, I can only venture to say it's likely a delicate and time-consuming task that will require some patience.

Playing back *prerecorded* images with our analog visual LSD follows a process similar to the techniques for playing back prerecorded sound on our analog audio LSD in the last chapter — take a 10 x 10 camera array and record its input onto VHS or digital video tape (or whatever you want), then, if not playing the video back directly from the cameras themselves, burn 100 videotapes (or DVDs) and put them into the corresponding VCRs (or DVD players) in your television array, take a deep breath and hit Play.

The cameras used to record signal for your television array don't even have to be video cameras. You could capture a single scene using an array of digital still cameras, or even analog film cameras, then create 100 videotapes that contain the still image parceled appropriately among the televisions. Converting an array of still cameras' output to work on an array of televisions takes a little work. If using film, you'd have to scan the processed film (slides or prints) into digital form or have a PhotoCD made from it. Then (regardless of its origin) use a computer or video editing equipment to divide the overall image into 320 x 240 (i.e., NTSC resolution) chunks, and record each chunk to videotape or whatever medium you're using. Depending on the size of your array, you might even have to downsample the scan to fit. You'd also have to take into account how many pixels the bezels rob you of.

With the higher resolution images possible with still cameras, you'd likely need a smaller array of them — or no array at all, but only a single camera — to capture enough pixels to display at full resolution on a 10 x 10 array of televisions or larger.

What you *don't* want to do is record too *little* information for output on your array. That gets us back into the crappy video wall situation, in which good solid signals are subdivided and the output resolution accordingly decreased.

If your video LSD is only handling still images, the sync between components isn't necessarily crucial. You could record chunks of a single still image onto all those videotapes and give each of them a duration of 30 seconds, or two minutes, or whatever, then switch to another image; the array will change like flashcards at a football game, with some screens turning a bit early, most around the same time, and some a bit late, all depending on how well the sync has been maintained by the individual components (and assuming you started them all at the same time!).

The effect of showing a single still image on a 10 x 10 array of televisions might be pretty cool. But increase the array's size by adding more televisions, and the perceived resolution can grow much higher, depending on the number of televisions you use. How much higher can it go? Remember, we're using televisions. (Computer monitors have much higher native resolutions, but we'll save that particular intensity for later.) A single NTSC television displays 320 pixels wide by 240 pixels tall. We multiply these to get the total number of pixels per screen. We'll use 'megapixels,' which are millions of pixels, to keep things simple.

At 320 x 240 pixels per screen we have a total of 76,800 pixels. Divide that by 1,000,000 to get 0.077 megapixels (rounded up) per television. (For the sake of comparison, a 1,600 x 1,200 computer monitor displays total 1,920,000 pixels, or about 2 megapixels.)

A 10 x 10 array of televisions contains 100 televisions. To find out the size of the array in pixels, take our 0.077 megapixels per screen and multiply that by 100. The result? 7.7 megapixels. That seems pretty big by computer monitor or digital camera standards, but it's still small by LSD standards.

I'd prefer a really *big* LSD, so I'm going to build my array to contain 100 x 100 televisions, or 10,000 total screens. This might be as tall as, say, a 10-story building, but I dream big, and a 10-story-square screen would be *awesome*. To calculate the size of the display, multiply 0.077 times 10,000, and we get 770 megapixels, or 0.77 *giga*pixels.

That's a lot of information to be throwing around, but building a 10,000-television analog video LSD is not much different from building a 100-television one, or even a 10-television one. How hard would building such a device be? Not *that* hard, I suspect, with a little intelligence in design and construction. And you can start out, as I've said, small.

What kind of image would you show on such a screen? I must leave that to you to imagine. Please do remember that a 10 x 10 television array is just about the lowest-resolution array you'd even want to consider building. While you might feel somewhat immersed inside a six-sided television array, using higher-resolution output components such as computer monitors takes things to much higher level. When we get to the digital LSD, things will grow even weirder and more wonderful, but ultimately the incredible nature of LSD visual reality is due to its extremely high perceived resolution, and televisions are only the beginning.

As with the analog sound LSD described in the last chapter, though, analog video LSDs are simple and relatively inexpensive, and you can achieve cool things like the LSD Mirror and Invisible Wall relatively easily.

4. The Analog Audio-Video Playback LSD

Now that you're up to speed on how to build analog LSDs, it should be fairly clear how you'd combine audio and video components to generate both kinds of sensory information on the same device simultaneously.

The LSD unit, or building block, used to create an analog audio-video playback array could consist of a television, two or more speakers, an amplifer, and a VCR or DVD player. (You could also use a computer to provide audio and video output.) The VCR's video output can go directly to the TV, and its audio output can be fed to the speakers through the amplifier. If you're using an A/V amplifier, your video can go through it as well.

If you're using televisions, you could use their speakers for audio output instead of adding more. And video cameras can record sound as well as image. Personally, though, I like the idea of having more control over the quality and type of audio input and output, so I'd probably use better speakers, especially for low frequencies. It's even possible I would remove all the extraneous material from my televisions – cabinets, speakers, cable tuner, etc. – and just use the video display tube and circuitry mounted in a thin-bezeled frame.

When adding speakers to a television array, though, the cursed bezels can end up being our friend; you can place the additional speakers in the bezel area so the television screen isn't obscured. (You can leave small computer-style speakers in their cases, but you'd want to remove larger cones from their cabinets.) You don't necessarily need to put speakers at every juncture in the screen array, but do make sure you have sufficient speaker coverage to match dimensions of the audio that's been captured and give a good LSD location/direction effect.

You can also separate your speaker and television arrays so that hiding the cones among the televisions isn't even an issue. One example of this would be the Home Theater on Acid concept, which would basically replicate the layout of a 5.1-channel system (ideally adding dimensionalization using the LSD graphic equalizer hack described in Chapter Two). With a large television array up front, and speaker arrays at each of the 5.1 locations (the subwoofer can be anywhere, or you can have multiple subwoofers at multiple places), nesting speakers in the television array wouldn't be necessary.

But that's just one option; your analog LSD can take any shape or size you can dream up. To my mind at least, the most interesting marriage of sound and video in an LSD allows sounds to originate from specific visual locations, and interspersing speakers among the televisions is the way to do this.

Using *digital rear projection* is one intriguing way to hide the speakers. It's expensive, but it allows you to hide the speakers *behind* the screen, next to the video projector itself. This comes up again with the digital LSD; I'll leave off giving a more detailed description until then.

As with the other analog LSDs, syncronization in an analog audio-video playback LSD is a factor to consider. With high-end videotape decks, sync is possible, but with low-end decks, and, one assumes, DVD players, the sync is only as good as your ability to start everything at once and the components' ability to play back at a constant rate.

Interesting effects are possible despite the drawbacks to using analog devices. The effects LSDs that are solely audio or visual achieve individually become even more powerful when you combine image and sound on a system. Since development of the purely digital LSD would take a while, working in analog is probably not a bad idea – it will allow you to get the basic concepts down, to explore programming strategies, and to start to turn people on to what the LSD actually is.

5. Analog Audio and Video Capture and Recording LSDs

By now you should have a fairly clear understanding of what an analog playback LSD is, but how to get signals to all of those monitors and speakers may be less clear. (More important is figuring out how to get *compelling* signals to them, but I'll cover that in detail in Chapter 8 on Digital LSD Programming.)

The device used to capture and possibly record audio or video (or both) for LSD playback *is also an LSD* — it uses the same principles of arrays, location, and perceived resolution as a playback LSD. Although we're covering analog capture at the moment, the techniques are essentially the same for digital capture LSDs.

Analog audio on an LSD can be live or prerecorded. Any type of sound that's currently recorded on traditional media can be played back on an LSD – voice, music, natural sound, electronica, etc. Adding LSD dimensionalization and other kinds of analog massaging can make the sound more interesting.

Your analog audio LSD may consist merely of microphones, amps and speakers, with live throughput from input to output. Depending on your design, this configuration may resemble a PA as you'd find in a club or concert venue. If I were designing this kind of LSD, though, I'd steer away from the PA design (especially with its feedback issues) and get into more interesting uses that take better advantage of what the LSD is about.

One example would be to have two soundproof rooms, one lined with microphones and one lined with speakers. (By "lined" I mean either all four sides or all six surfaces.) You'd need enough of each component to ensure each sound, and the location of each sound, is captured and played back adequately.

What sort of microphones should you use? I haven't actually built one of these devices, but I would guess that monophonic directional microphones would work best, with each mono stream fed to a single speaker in the other room. It's possible a stereo mike with a cardiod pickup pattern would work just as well, but you'd need to make sure the orientation of the speakers matches the orientation of the microphone's pickup pattern.

Such a setup would yield interesting sounds if you 'play' the microphone room intelligently. You could ride a bike with a bell on it in circles in the microphone room, and your audience would hear it circle them in the speaker room. You could have actors speak lines, or musicians play different instruments, as they walk around, moving laterally as well as nearer and farther from the arrays. You could also set up a small speaker array at the center of the room pointing *outward at* the microphones, and play any sort of sound you want, with the effect of potentially dimensionalizing the sound even more, since the microphone array could break up the sound it's recording into more location streams than the original possesses. *Moving* the speakers in the center array *during playback* would result in a live manipulation of a given sound's location for the other room's audience.

Another design strategy for live sound is to keep the speaker room as it is, but turn the microphones *outward*, probably in an outdoor setting. I call this type of recording device an *LSD cluster*, in this case a microphone cluster. Each microphone would point outward from a common center. Ideally the shape of the cluster would match the shape of the playback array, but it doesn't necessarily have to.

The microphone cluster could be placed anywhere, so long as the live signal can reach the output array without generating feedback. You could hang such a cluster in a place that's not safe or is otherwise impractical for humans to be, but where the sounds are compelling – or in a place where the presence of a human would create a disturbance, especially in nature.

The point of view of the listener when hearing sound generated by a microphone array depends on the orientation of both the recording array and the playback array. In the original example, the microphones in the second room were all pointing inward at the sound sources. A microphone cluster points outward. If you stand in the center of the room with the inward-pointing speaker array and switch back and forth between sounds recorded by inward and outward microphone orientations, you would experience each orientation's sound somewhat differently. (Take a moment to try to imagine what the differences might be.)

The two-room live sound LSD could also contain microphone and speaker arrays *in each room*. This is somewhat analogous to the two-sided LSD Invisible Wall. Feedback issues would no doubt prevent you from live recording and playing back sound in the same room, so you'd want to play back recorded signals either later in the same room, or live in the other room's array.

You could also design your *playback* array as an outward-pointing cluster as well, with a somewhat different playback experience than the inward-pointing speaker strategy. (Again, take a moment to try to imagine how this would sound, and how it would sound different from an inward-pointing speaker array.)

Recording the sounds you obtain with an LSD microphone array is simply a matter of adding a recording-capable component to the chain – a video camera, cassette recorder, Nagra, VCR, or even computers. You could record each microphone's input even as it's being played live on a speaker array, or the recording device could be a discrete LSD that records what it captures for later playback on another system elsewhere.

Each microphone's level should be balanced with the others in the recording array – the loudness of a sound in playback will play a large role in determining how close or distant the listener perceives it to be. To accurately monitor the audio input of a recording LSD, you'd need a playback LSD with the same number of channels and spatial orientation. You could always mix the input to stereo so you can use headphones, but it won't give you an accurate *locational* portrayal of all the signals you're actually getting. Another strategy is to be able to select any pair of microphones (or a single stereo microphone) at a time to check what it's recording and adjust line levels as necessary.

One solution to the monitoring problem is to use a soundboard that will accept as many inputs as you have and allow you to listen to a stereo mixdown, and/or monitor inputs a pair (or more) at a time. Just be careful not to mix down what you're actually recording!

Without good synchronization, analog recording of actual live sound is likely to suffer, although the effect of playing back unsynced (or gradually de-synchronizing) sound might be sort of cool in some circumstances. Prerecorded programs that don't depend on synchronization are much more likely to prove compelling.

Analog *video* capture and recording works much the same way audio does. If you take our two-room scenario and replace the microphones with video cameras and the speakers with televisions, you'll have a device as compelling visually as the original was aurally. As with the LSD Mirror described in Chapter 3, since your cameras are pointing inward you'll need to route the camera signals appropriately so the television array is displaying a unified image rather than a series of discrete columnar images.

The orientation of the cameras and televisions will affect the way the audience perceives the visuals to a much greater degree than with audio. If all four walls of the capture room contain cameras pointing inward, and the signals are routed to an array of televisions also pointing inward, then the image being captured will be repeated on all four walls (and the ceiling and floor, if set up that way). The cameras themselves will also be captured.

This is contingent on the cameras being zoomed and focused to capture the room at a 1:1 size ratio. If they are zoomed and focused to concentrate on a much smaller point in the center of the room, say, it's possible the object being captured will play back *around* the audience in a manner that can perhaps best be described as 'being inside the outside' of the object. Let's imagine a room is completely coated with an LSD camera array, and each camera is zoomed and focused to capture a small portion of a basketball at the room's center. The live signal played back in the room lined with televisions will be an *inversion* of the basketball, as though the basketball's skin has been stretched around the outside of the room but faces *inward*.

If, however, the cameras are pointed *outward* — arranged, in other words, as a camera cluster — and their signals are displayed on an *inward* television array, the viewer will perceive himself to be observing the world around him in a much more natural way. This could take place either inside the room, with the camera cluster residing at the room's center, or outside the room.

Just as with the audio rooms, the video rooms can be set up with enough components to allow each room to capture its contents and play them back in the other room simultaneously.

You can record the video you are capturing for later playback, and just as with audio there is a disadvantage to playback if the video is not synchronized. On the other hand, a still image can be captured and recorded on videotape such that synchronization doesn't especially matter; as described in Chapter 3, an array of televisions changing from one still image to another with less-than-perfect sync would resemble a flashcard section at a football game, and might actually be sort of cool.

Finally, there's no reason a bank of audio or video recording components has to be in a strict array. A "recording LSD" could capture patches of image or sound at different locations, with the captured signals played back by a strict output array or even one that has its own sets of subarrays. A recording LSD might consist of dozens or hundreds of microphones at dozens or hundreds of locations. It might take some doing to unify the signals they capture for playback, but who's to say it wouldn't be loads of fun?

6. Hybrid Digital-Analog Recording and Playback LSDs

Much of the content in the preceding chapters may have frustrated you, as the limitations of analog LSDs are quite severe relative to more sophisticated digital incarnations. But your struggle to this point has lead, I hope, to a fairly good understanding of what an LSD is. The means of arriving at "a device that is an LSD" are varied – you can use all sorts of components to create high (or at least higher) resolution sensory information in a multi-channel array.

If we still lived in the 1970s and wanted to build an LSD, analog versions would be our only route, and they be a lot more expensive than they are now. The notion of using a computer to augment or even drive a 1970s LSD would be laughable. But because we live in the aftermath of the personal computer's ascendance, it's hard *not* to imagine building an LSD without a computer being involved in some way or other.

The simplest way to bring in a computer is to have it produce or record signal, as I've alluded to (parenthetically) up to this point. Computers can play DVDs, VCDs, even video streams directly off their hard drives. They can also, of course, play stereo or even 5.1-channel sound. While they don't generally include amplifiers (some current Macs do), speakers used with computers are usually self-amplified. Computers can also drive traditional analog amplifiers.

One immediate advantage of using computers to drive sound and image is that it is all digital. A videotape has to be rewound, and will degrade much sooner than a hard drive will. It is much quicker and cheaper to copy computer files from one machine to another than to duplicate videos from a master. And the quality of native digital audio and video is better, cleaner, than analog sound, though your LSD might include analog elements to bring a certain warmth that digital, it is generally agreed, lacks.

More importantly, though, computers are programmable. They can be networked to communicate with each other. We will see only bits of the importance of this in hybrid systems, but these two concepts are worth bearing in mind as you read this chapter, because the following chapter – on the Digital LSD – would not be possible without them.

I might use computers to drive a small analog LSD because it's easier to store lots of sound and image data without the hassle of a big stack of videotapes. An Internet connection will let me download (or upload) even more stuff. I can connect lots of components to a computer, even at the same time: a digital (or analog) video camera, a digital still camera, a microphone, two or more speakers, an analog amplifier, a synchronized video tape deck, etc.

This is all fine and dandy, but if I can only get 5.1 channels of sound out of a computer, and I want 100 or 1,000 channels of sound on my LSD, it's going to be awfully expensive to use computers simply as a sound source. Instead, I should buy a bunch of relatively cheap amplifiers.

But then I still have to get my sound *to* those amplifiers. Do I really need to have 50 CD players for a 100-channel LSD? Or 17 5.1 DVD players for a 101-channel LSD? Wouldn't it be much better if I could drive all my amps with just one or two computers?

Why, sure it would. It might end up costing just a bit more than the 17 DVD players or 50 CD players, but guess what you gain from it?

SYNC.

I tried my best to convince you in the analog chapters that you could do cool things without sync. That's entirely true. But *with* sync, you can do more. Much more. In fact, the "LSDness" of an LSD only starts to come out, I think, when things are absolutely synced. Voices that match moving lips. Downbeats that happen at the same time across the array. Sounds that move like large aural structures; images that grow, shrink, move, swirl.

With sync, the LSD world can turn.

One decidedly cool way of programming a synchronized audio LSD is to take the original *unmixed* multi-track recordings used to make a song or album. Often this is at least 24 distinct tracks, each containing a different part of the music – vocals, percussion, lead guitar, etc. These original tracks are mixed to stereo for normal recordings, but record companies – ever on the watch for new sources of income – could sell the unmixed individual tracks for use on LSDs, enabling users to make what we're used to hearing in stereo come alive again in a significantly more dimensional way.

How would we build a hybrid LSD? Let's start with audio, and assume we're using two computers as sound sources for a 120-channel analog audio playback LSD. Each computer would drive 60 channels, which are amplified by a 20 6-channel amplifiers (a 5.1-channel amplifier is really 6 channels total). Each of these amplifiers is connected to one of the computers by way of an *LSD digital-analog converter*.

Yes, another LSD box. (Are you listening, entrepreneurs?) This time it's not about spreading out the *sound* of the audio, as the dimensionalization box was, but of distributing digital audio signals to multiple analog targets – the amplifiers.

This works much the same way a TCP/IP network (like the Internet) works. Each packet of information is headed to a specific address on the network. In our case, the packet of information is sound information. (We'll get to video in a bit.)

What information do we need to get to the amplifier to make it generate the sounds we want?

Wave amplitudes. Sounds are waves, or disturbances, in the air. When air moves suddenly, a sound is made. A speaker cone moves many times a second to create sound, *but only along a single axis.* So the sound information we need to send is simply a number representing how far the speaker cone has moved from a stationary zero point.

Put another way, this number is a digital representation of an analog event -- the pulse of energy that shoots down a standard copper speaker wire to drive a speaker cone. (The pulse of energy is translated into the cone moving a commensurate amount, and thereby generates a sound wave of the desired amplitude.) But because it's digital, our signal can contain *more* than just the wave's amplitude – it can also contain a *location value* and an encoded *sync tick*. (Although the computer may handle sync exclusively, some day speakers might have embedded processors that will let them read the sync tick and stay even more time-aligned.)

Combine the amplitude we want the speaker cone to jump with location and time information, and we have LSD audio data encapsulated in TCP/IP packets moving over 100mbs Ethernet cable.

How many signals can we generate down an Ethernet connection? At 100mbs, the standard speed at the moment, we can generate lots and lots. And that number will only grow larger as we get into Gigabit Ethernet and beyond.

The LSD digital-analog converter has an Ethernet plug in one side and whole bunch of analog (perhaps even optical) connections on the other. Inside the converter a simple thing happens. The digital signal is split out to as many streams as it contains, according to the location assignment information embedded in each packet, and each stream is converted or decoded, using a digital-analog chip, to a simple analog signal traveling down a 1/8" RCA or optical cable. How many speakers this drives may depend entirely on the amplifier and type of signal being sent to the converter.

How hard would it be to create a converter such as this? Could't be *that* hard; similar types of converters are found in digital speakers, CD players, and other components that translate digital signals to analog signals. Would you be willing to fork over a little dough to buy as many of these devices as you'd need to drive dozens of discrete-channeled speakers?

Software is needed too. After all, the world is generally rooted in stereo, though some set-ups will allow you to hook up 5.1 systems to the computer directly. The simplest way of looking at is to imagine a mixing board on your computer screen, with each channel or stream of sound occupying one of the control strips – you can affect the volume, echo, reverb, etc., of each channel independently, over time even, with the computer recording each move you make.

Each sound stream can be assigned, by you, to a specific speaker (or speakers) in your array. A given stream might be routed to Speaker A, but over a set amount of time be moved to Speaker Z. The 'set amount of time' could range from instantaneously (i.e., a jump) to a much longer period. The signal might move in a straight line across intermediate speakers, or travel in circles around the device, with each orbit moving closer to the final destination row. Or it could move randomly, or in any other fashion the program would let you assign.

So the software we're speaking of will route multiple signals to multiple outputs, and allow changes to outputs over time. The software should also be able to prepare sound streams for output – loading streams, syncing them up, previewing them, etc.

Developing such software might sound like a lot of work, but there are already any number of audio applications that do almost as much already – DJ mixing software, for instance, that happens now to be oriented toward stereo output, but with only modest modification could allow this type of multi-channel location manipulation.

Or think of it this way – an MP3 server, or a version of Quicktime Streaming Server (which is open source, by the way), that sends tens, maybe hundreds of channels of sound out over Ethernet, maintaining sync even while allowing channel location assignments to change willy-nilly, on the fly, in real time. I am not aware of this capability being part of any type of sound-serving solution, but would it really be that difficult to add?

Take a moment to step back and absorb what we're talking about. Rather than having a VCR or DVD player for each and every speaker or television in your array, a single computer could drive dozens of output components *in sync*, without stacks of media that have to be loaded, without rewinding, without remote controls you hope will start each component at the same time. With a computer-based LSD you can stop and start at will, move sounds around as you like, adjust volume for one or more tracks at a time without scampering around the array, replace one track with another, and so on, all the while maintaining sync.

What could you do with such a device? Besides the notion that whatever you recorded with your analog recording LSD could be rendered in sync after being assembled on a computer, you have the potential of introducing to the mix (so to speak) anything a computer can do to sound, which is a *lot*. Even stereo dimensionalization algorithms (relicating the LSD hacked EQ described in Chapter 2) could be run *in real time*, with each subchannel sent to the appropriate speakers in the array.

At a certain level of sophistication, it is important for the LSD software to be *topology-aware* – in other words, to 'know' how the particular LSD it is 'printing signal' to is constructed. I envision an initial configuration routine that establishes a 'logical LSD' in the software's data space, represented onscreen in such a way that you can 'paint' sound, location changes, and effects onto your arrays.

What a powerful device this would enable *for audio alone.* When we get to video analog-digital hybrid LSDs, the barn doors start to strain against their hinges… because with the introduction of digital control into an analog video playback LSD, things get, as with audio, much more interesting.

We focused on televisions as the primary output component in our analog visual playback LSDs. An analog-digital hybrid device could use televisions as well, at least initially; to keep things clear, I'll use televisions in the following example. Read carefully, for I believe this will demonstrate the kind of power we are talking about when we introduce digital control to an analog LSD.

An NTSC video stream has a resolution of 320 pixels wide by 240 pixels tall. I'd like to build an analog LSD that consists of 320 *televisions* wide by 240 *televisions* tall.

Yes, this is 76,800 televisions. Maybe they cost $100 each, so I've spent $768,000 on televisions alone. But soon you'll understand how *cheap* this is for what we'll be getting.

The first thing I want to do is recreate a 320 x 240 pixel NTSC television signal such that *each pixel* is represented by an *entire television* in my 320 x 240 television array. In other words, *a single television screen built out of televisions*.

To do this, I shouldn't need much more than a single computer to throw out NTSC signal at that resolution. After all, a computer can display NTSC video quite easily, and at the moment that's all we're asking of it.

The LSD video signal, like the LSD audio signal, contains a bit more information than NTSC. Each signal consists of pixel color (instead of audio wave amplitude), location and time. The color may change 30 times a second, but sending these three numbers as packets of information over Ethernet to that many televisions that quickly shouldn't be too tricky.

But wait a minute -- what converts the signal from the digital signal generated by the computer to an analog signal that drives the television to generate the right color?

Why, an LSD digital-analog *video* converter, that's what. Like the audio version, such a device would take the single NTSC color that each television is supposed to receive at a given moment from the computer and translate it into the analog signal that the television understands. Again, this doesn't seem like a particularly complicated thing to do. After all, doesn't a common digital cable box do exactly the same thing, though only for one television at a time?

Using our LSD digital-analog video converter box, let's display a unified, synchronized image across the 320 x 240 television array. We'd have to stand fairly far back from this array in order to see the unified picture, but still, 76,800 televisions all firing at once in such a coordinated manner is sort of cool – and standing that far back would minimize the visual disruption caused by the bezels around the screens.

But wait, you say. One color per screen at a time? Or more accurately, one *pixel* per screen at a time? Isn't that kind of a waste, since *each television* can show a resolution of 320 x 240?

You bet it's a waste! Our simple analog video LSD, which had a VCR for each television, could show full-resolution NTSC video on each component in the array. The point of adding computers to the mix, beyond at *least* attaining the same resolution we were getting with the purely analog device, is to introduce capabilities only a computer can supply.

The converter, then, should be able to handle a full NTSC stream of information for each television it outputs to. I don't know the math, but surely 100mbs Ethernet is capable of handling multiple streams of NTSC video. Gigabit Ethernet will handle even more. You'll likely need more than one computer for our 76,800-television array, but you certainly won't need 76,800 computers.

To introduce our audience to what our array can do, let's start with each television displaying a single pixel so we show them what is basically a very large television. But then divide each television screen into quarters so it shows *four* pixels. Divide that yet again into a 4 x 4 grid (16 pixels), and on and on, until each *television* is showing all 76,800 pixels it can in the 76,800 *television* array.

Take a moment to shut your eyes and imagine what this looks like. Start from one pixel per screen to 4, 16, 256, then in modest steps up to 76,800, *repeated 76,799 times*.

Though it's still at a much lower resolution than we will eventually get to, you are to be congratulated for reaching a milestone -- *you're now seeing LSD video*.

How many pixels are we slinging around? 320 pixels per television, with 320 televisions, is 102,400 pixels across. 240 pixels per television, with 240 televisions, is 57,600 pixels deep.

102,400 x 57,600 pixels is a total of 5,898,240,000 pixels. Let's just round it off to *six gigapixels.*

That's a pretty big chunk of visual information. The perceived resolution of such a device, depending on how far you're standing from it, would be pretty high relative to other media we're used to looking at. In fact, using televisions in this way we're slightly *above* the pixels-per-inch resolution of an IMAX movie screen—and we are in fact showing up to 57 times *more pixels* than an IMAX displays! (See Appendix X for a comparative resolution chart.)

How do the computers at the heart of such a beast store the information needed to drive it? If the data format is efficient enough, you can actually hold a lot of NTSC data on a DVD. Even a software DVD decoder like a Mac's does fine with full-screen video; suppose it was decoding and sending out three or four (or 10 or 20) streams of NTSC data in real time?

Or suppose it was reading the information off its own hard drive, or out of its own RAM, instead of off of a DVD? Indeed, if we are simply talking MPEG-2 video files that the computer is reading and decoding, the files could be stored anywhere, including being streamed across the network. The computers in a hybrid LSD could be augmented by another layer of computers that *computes* what each screen will display – as opposed to the base layer, which simply gets signal out onto the screens.

But failing that, with enough computers we can at least get 6 gigapixels of prerecorded signal out to our array. How hard would it be to capture 6 gigapixels of image with an array of 3-megapixel digital *still* cameras? (You'd need a couple thousand cameras, but prices are falling, right?) You could change the images on your television array as quickly as you can download the new image.

If each of our televisions has at least one channel of *sound* it can play as well, we'd be able to play 76,800 channels. It might take more computers to generate it, but adding that much sound to this vast visual array would be quite grand.

This has brought us to near the end of the chapter. Let's take stock of where we have gotten to. We are now speaking of 6 gigapixels of screen resolution – potentially *live* screen resolution – and tens of thousands of sound channels. While bearing in mind that we are still at the *low* end of the LSD experience, we are already, at least as I see it, in pretty interesting territory. But if *you* still aren't convinced there's something to this LSD business, I very much urge you to read the next chapter.

In our hybrid device we are seeing the beginnings of something quite powerful in terms of the capability of the system to model reality, in a very believable audio-visual sense, in a synchronized presentation. The computed audio-visual LSD signal resembles that of an extremely powerful video game console (out of arrays of which, I should mention, it could be built).

Build your television array in a circle, rising up to a center point that faces down at you—a dome, in other words, like an indoor sports arena—and your LSD will surround you with an immense audio-visual experience displaying much more information, even with relatively cheap televisions, than you'll get in an IMAX theater.

Can you now appreciate the grandeur of the dimension you are entering? Perhaps not, but if you can't, don't blame the device; until it is built, I only have words.

7. The Fully Digital LSD

As fun as it has been to describe analog and hybrid LSDs to you, the fully digital LSD is so powerful and revolutionary that anything less advanced is bound, in hindsight, to seem disappointing. I wish the following statement were an exaggeration, but after many years, it continues to withstand my skepticism: *A digital LSD is so powerful that it redefines the assumptions we make about reality, time, and consciousness itself.*

How could it be so powerful? A sufficiently large and 'resolved' digital LSD is a new dimension of experience that we simply cannot imagine. Our eyes have not seen, and our ears have not heard, anything like it. It allows humans to create experiences that *transcend reality itself*, insofar as it surpasses the ability of our senses to perceive.

That's a pretty big claim to make. Let's see if I can back it up.

In the previous chapter we created a 320 x 240 array of televisions resolved to the fullest native capacity. This gave us a 6-gigapixel display.

Computer monitors, however, are capable of displaying a much higher resolution than 320 x 240 – 1,600 x 1,200 pixels in most cases.

To compare monitors to televisions, let's not change the *number* of screens in the array, let's just change the *type*. In other words, an array that is still 320 x 240 cathode ray tubes, but made out of computer monitors instead of televisions.

We can easily show a 320 x 240 pixel NTSC image across our entire computer monitor array, at one pixel per monitor, just as we did initially with our televisions. We can also show a different 320 x 240 image on *each* monitor, meaning we can show as much as we were showing on the television array at its highest resolution. That gets us our 6-gigapixel screen.

The television array described in the last chapter could use 27" (diagonal) televisions, and the computer monitor array could use, oh, 19" monitors. There is already an increase in the resolution of the output if I show one NTSC stream per monitor, because the dots are appearing on a *more compact* screen. In other words, I'm showing a more resolved image because I'm displaying the same number of dots on a smaller physical space. In other words, it's 'tighter.'

But the real magic happens when we up our resolution to the full 1,600 x 1,200 pixels for each monitor.

With 320 monitors across, and 1,600 horizontal pixels per monitor, my array is 512,000 pixels wide.

With 240 monitors down, and 1,200 vertical pixels per monitor, it's 288,800 pixels tall.

Multiply 512,000 pixels by 288,800 pixels, and you get 147,456,000,000 pixels. Divide this number by a billion to convert the number to gigapixels, and you get 147. Yes, *147 gigapixels* in a *smaller* space than the 6-gigapixel television array.

There's something important to be said at this point about the way resolution works. When you have a 100 x 100 pixel grid, and in the same space you suddenly put a 300 x 300 grid, you are not making number of dots in the grid three times higher, you are making it *nine* times higher. Do the math: 100 x 100 = 10,000, while 300 x 300 = 90,000.

Following this line of thought, if the horizontal resolution of a 27" television is about 15 pixels per inch, and the horizontal resolution of a 19" monitor at 1,600 x 1,200 is about 114 pixels per inch, then the number of pixels per *square* inch is *not* 7.6 times higher, it's *57* times higher. (See Appendix X for the math.) You have 57 times more dots in a square inch of monitor (13,061) than you do in a square inch of television (229 dots).

In addition, a computer monitor displays visual content using a *progressive* scan – each line of the image is drawn top to bottom in one pass. Televisions show *interlaced* images, meaning every *other* line of the image is displayed top to bottom in one pass. It takes *two* passes to create a completely television screen full of image, whereas a computer monitor does it in *one* pass.

What does this mean? Native computer graphics appear much steadier, with much less flicker. In addition to the higher resolution, that's why you can stare at a computer all day from arm's length or even closer. To put it in film terms (and somewhat simply), the *frame rate* of a computer monitor showing native computer video signal is *twice* that of a television, since it shows a full frame per pass instead of a half-frame per pass.

So with a fully digital LSD, we not only can display a vastly more detailed image, but a much steadier one too.

Have I convinced you what you're seeing on a purely computer-driven array is enormously more detailed, more rich, than anything you've seen before? No? What shall I compare it to, then?

How about the best commercial solution (so far) the motion picture industry has offered us: an IMAX movie theater. The television array came to within 60% of the resolution per inch of an IMAX theater screen (even while showing a great deal more information). But the digital LSD using computer monitors blows the IMAX away. It is capable of displaying 65 times *greater* resolution than an IMAX screen.

Can you imagine what such a display looks like? Of course not! It's unfathomable, astonishing, unprecedented! To see a large digital LSD array cranked up to full resolution would quite simply blow your mind.

Let's take a look at how a digital LSD is built. We've seen how analog LSDs use analog components and analog transmission of signal to record or play back audio and/or visual information. A hybrid LSD uses digital storage and a hybrid of digital-analog transmission. A digital LSD uses digital storage and digital transmission.

That's not to say you couldn't use analog components in part of your array. LSDs are *additive* in size and quality, as we've discussed. If you build an analog LSD, add computers to create a hybrid, then decide to go digital, there's no need to toss your analog components unless you really want to; just put them farther away in the case of visual components (so the perceived resolution stays in line with the closer, *finer* components), use them for less-critical parts of your array, or create another *layer* in your LSD structure that plays to the strengths of the analog equipment. (More on layers in Chapter X.)

But for the sake of illustration, let's assume fully digital components: computers, monitors, digital video or still cameras, digital sound capture, and digital speakers.

The digital LSD uses computers to record and 'print' sensory information. Unlike a hybrid, the signal path all the way to the end point should be as digital as possible; all signals are recorded or originate as files on a computer-based medium. For reading data, this could mean a CD or DVD, a pre-downloaded file that is read off the hard drive, a file streamed over a network, data held in RAM, etc. For writing data, this means all information is written to a computer-readable audio or video format.

A single computer will drive as many input or output components as possible in the array. As mentioned briefly in the previous chapter, one layer of computers may be used to handle ‘printing’ signals to components, while another may handle other tasks -- generating those signals, decoding compressed files, relaying captured signals, maintaining sync, etc.

The digital LSD is, in fact, a parallel computing cluster intended to collect or generate a unified sensory experience; the computers act as a single entity in an array just like any other component does. Each computer handles as much ‘resolution’ – its data processing capacity – as it can; each computer records or generates signals at its highest level of throughput.

A typical cheap ($300-500) Windows-based computer can, these days, provide literally hours of high-resolution audio and video by using its hard drive, add-in card, RAM, optical drive and network capacity to its fullest. You can add several USB, Firewire, audio and/or video cards so the computer can drive and/or record from multiple components. Dedicated audio devices like Mark of the Unicorn’s Firewire 828 box might work as well -- it allows you to input and output eight discrete signals, and more than one 828 can be driven by a single computer. Although I’m a Mac guy, the commodity-level pricing of Windows machines makes them especially attractive for use in a digital LSD, though a good case might be made for running Linux or Unix on them rather than Windows. You could still use a Mac to control everything, or use older, cheaper PowerMacs to drive input or output components.

To build your digital LSD simply buy or build from scratch a typical small-office network. You’d end up with a server, 10-20 computers, Ethernet cabling, and a router/switch. Add to each computer whatever cards or add-ons you need to drive components, hook up the components themselves (microphones, USB or Firewire speakers, video sources, monitors), download the appropriate files to each computer's hard drive, and hit Play in your LSD control application.

Yes, we'll need someone to write some software for us. (At least writing software for a hybrid LSD gets us a long way towards the fully digital LSD.) It's easiest probably to drive all your computers from a central server or controller. The software is an extension of the program described in the last chapter: it would handle assembly of the sound and video streams, signal processing and assignments, synchronization, and system integrity checks. It should also be able to generate the signals live via computation rather than rely solely on prerecorded streams. The computers used to output the signals could be called upon to do some of the signal processing themselves, or an intermediate layer of computers might perform the processing and distribute the signals to the player computer slaves.

Indeed, the software used to drive a sophisticated digital LSD is really more of a cluster operating system or media server than a personal-computer application. But it starts out quite simply – each computer that drives output components needs to be able to route multiple signals simultaneously to multiple outputs in sync. This shouldn't be very hard to accomplish; we're really only talking about a multi-channel streaming media server that includes synchronization algorithms.

The USB and Firewire busses can handle dozens of components, and assuming the bandwidth is there, all of these components could be digital speakers, each receiving a discrete, high-fidelity audio signal. The same holds true for video – many modern video cards can write signal to *two* monitors, so with three cards you could conceivably drive six monitors at full resolution. If we had LSD analog-digital converter boxes going as well, we could shoot signals out over Ethernet too.

To illustrate a digital LSD application, let's return to our original analog audio LSD, which was a bunch of speakers fed by a graphic equalizer that had been hacked to create 10 frequency-divided audio streams per stereo channel. There's no reason why you couldn't do the very same thing using a computer and digital speakers – creating, in fact, a digital audio LSD that accomplishes the same thing the analog LSD did. The application would take a stereo sound source and divide it by frequency (just as equalizer programs already do), and rather than re-mix them back to stereo, would route the resulting signals to individual speakers on the USB or Firewire bus. With a computer, though, the number of frequency divisions, and the points at which the frequencies are divided, could be user-adjustable, and could change over time.

The computer-processed frequency hack for existing stereo or multi-channel recordings could be extended much farther by applying sophisticated algorithms to the sound source. The first level of improvement would involve analyzing left and right signals (using stereo as an example) and dimensionalizing *between* the two sources, rather than 'up and down' them by frequency, by determining the degree to which a sound is *shared* between channels. If Sound A is 75% of its total strength on the right, and 25% on the left, a computer processing the signal might be able to determine this and locate that sound at a corresponding horizontal point in the LSD array.

The second level of improvement would be 'intelligent extraction' of each of the distinctive sounds in a recording. (Humans might need to point out to the computer which sounds belong to a given instrument.) In a sense, this is *unmixing* the analog sound, so that rather than having all sounds playing together out of one set of speakers, each individual instrument (or even farther into it, each individual *note*) could be extracted from the mix and assigned a location.

Even more provocative, the speaker assignment for each channel could *change* over time as well, so the audio could move from one speaker to another in a way you could never accomplish with a hardwired analog system.

Once an individual computer is routing multiple sound streams out of a stereo source, routing multiple sound streams out of *multiple* sources shouldn't be that hard to accomplish. These sources could be a digital file or stream originating from a CD, hard drive, RAM, network server, or the Internet, or a combination thereof; they could change on the fly.

The source could also be *computed* sounds, whether they are processed locally by the output computer at the behest of a central controller, or by another computational element.

Combine a bunch of computers capable of reading streams (be they audio or video) from any or all of these sources, and outputting them to multiple digital components *in sync*, and you have a digital LSD capable of quite powerful effects. That doesn't seem *too* hard now, does it?

The digital LSD application is ideally object-oriented, in that each component – audio or video playback or recording, synchronization, location assignment and signal routing, signal generation and processing, etc. – can be upgraded without the entire system having to be reinstalled. It should also be distributable, to allow each task to be processed by any computer on the network with available CPU cycles.

The application should be extensible too – you should be able to load different kinds of interfaces to allow you to pull back to abstract renderings of LSD data, or drill down directly into the actual streaming data as much as you like (much as MIDI sequencers let you do everything from paint notes on a time grid to type in numbers for MIDI values in a spreadsheet). Plug-ins – in fact, industry-standard VST audio plug-ins – should be easily insertable. Applications such as Adobe's Photoshop and After Effects could feed the LSD directly; my idea of a cool interface for the digital LSD involves a set of Photoshop-like tools for painting various sounds, images and visual effects onto a 3-D representation of the device.

But before we get too used to the notion that we're controlling the LSD from a central server or master computer using a single computer screen, let's take a step back and realize just what we *have* with the LSD. It's not just an IMAX theater on steroids, it is itself a *computing device.* Painting effects on a small one-screen model of a large LSD would be ridiculous, and doesn't make sense in light of the fact that *the digital LSD's visual array is in fact an extremely sophisticated computer monitor*. Your interaction with the digital LSD should take place *on* the digital LSD, not through a tiny little porthole like a single computer monitor would be. That's not to say that you couldn't have a smaller LSD that you stand inside and manipulate to affect a much larger LSD that your audience is in.

Like any LSD, a digital LSD can be audio playback only, video playback only, audio recording only, video recording only, or any combination thereof. An audio or video playback digital LSD will feature tens, hundreds, thousands of speakers or monitors in arrays of your design. An audio-video playback digital LSD will combine the speakers and monitors in some way that minimizes the obtrusiveness of the speaker cones; as I mentioned in a previous chapter, the bezels around the CRTs can be your friend when trying to hide speakers amidst the monitors.

A different solution using current technology exists, however. It's expensive, but it has the advantage of completely hiding your audio output components. Instead of CRTs, use digital projectors to back-project images onto fabric screens. The speaker array would sit behind the screen next to the projector, just as speakers sit behind motion picture screens currently. You'd want to put as many digital projectors into your array as possible to maximize the final resolution; the more you put in, the closer the projector is to the screen. Not only would this hide your speakers, it would also eliminate the bezel problem. But it's pricey… for now. Advances in display technology promise exciting new ways of displaying visual information in the coming decades – including *painting* computer-controlled light-emitting substances onto surfaces -- and the LSD can only benefit from them.

Programming the digital LSD is such a big topic that I'm devoting an entire chapter to it later. But suffice it to say for now that everything we've covered to this point about how you could build LSDs, and how you would program them, applies to a digital LSD, as well as much, much more. We can display still images, stereo and 5.1 audio, and multiple streams of audio and video. Like with the hybrid LSD, we can also generate images and sounds completely by computer. We can manipulate the location of the streams, even down to the pixel-by-pixel (or audio-stream-by-audio-stream) level.

And like previous LSDs, you can build your digital LSD in any shape or size you want. You can start out small and build your digital LSD both larger and finer (rougher too) by adding different kinds of components. You can drive analog equipment using the hybridization techniques described earlier at the same time you're driving purely digital components.

A digital LSD can generate sounds and images using purely computational means. But it's likely you'll want to show *reality* in the hyper-real resolution a digital LSD is capable of. To capture reality in a way that will do it justice in digital LSD playback, you'll want to build a digital recording LSD.

Digital recording LSDs work the same way previous recording LSDs worked, except that rather than having to deal with tapes, everything is recorded onto digital media. The main exception to this rule is the use of digital videotape, but as I write this, digital video cameras that write directly to DVDs or Firewire hard drives are beginning to appear. Digital video cameras with multi-gigabyte optical or hard drives will appear in the coming years (or you'll just hook your iPod into one), meaning the chain from input to output is reduced by several steps and the signal remains digital for as long as possible.

To record reality, you'll need an array of cameras and microphones that generate as much detail as possible. Though it may seem like overkill, recording more detail than it is currently possible to play back is not necessarily wasted effort; the same or a new LSD may eventually be capable of playing back more sound and image data than previous incarnations.

Digital recording LSDs can be static structures or mobile ones, with arrays of cameras and/or microphones pointed outward or inward at the world within or around them. One design might be a recording cluster, shaped like a sphere but studded with recording components pointing outward. You could hang such a device from the bottom of an airplane or hot air balloon. You could mount it in the bed of a pickup truck (or larger, depending on the size of your device). You could set up a permanent LSD recorder in a town square, or park, or music hall, or anywhere else you'd want to capture what's going on.

As we've already seen, you could design a recording LSD to point *inward*. One use of an inward-pointing recording LSD would be to capture a football game. Mount a large array of cameras and microphones on one side of (or on all four sides, as well as above) a football field. Perhaps the most effective LSD to play a football game back on would be an outward-pointing rectangular device that replicates the size and shape of the LSD the game was recorded on. Less costly would be a simple LSD wall that shows the most interesting or relevant viewpoint on one or both sides. This device would be less LSD and more traditional, though you'd still get high-resolution sounds and images.

For close-ups of the players and instant replays, you could use a separate *robotic* camera array and display its signals on a section of the larger output device or on its own dedicated LSD. (More on this type of recording array below.)

We encounter a problem when we want to take the many streams of sense data that have been captured by our recording LSD and play them back on a playback LSD – what if there is a mismatch between the number of channels captured and the number available in the playback device? Actually, the number of *streams* doesn't matter so much as the number of pixels and channels.

The solution? The recording LSD must encode its own size, orientation, and resolution in the digital files it captures (or as an accompanying Configuration file). Then a computer-based LSD can *mix or spread* the audio and/or video to ensure it is played back at the highest resolution the output array is capable of. In other words, the recorded signal contains a model of the device it was recorded on, while the playback LSD *maps* the recorded signal to its own output capabilities by mixing signals or spreading them out.

Mixing is of course similar to mixing down multiple audio tracks into two-channel stereo sound. If our source has 100 channels of sound and our output has only 50 speakers, we'll need our computers to figure out where each of the 100 channels should go, doubling up channels in some cases, or allocating them to more than one output component if the original's location doesn't have a matching location in the array.

To give a simple example, let's suppose we have three microphones in a line, set up to record a car passing by. As the car passes, each microphone in succession 'hears' the sound louder when the car is directly in front of it, then softer after it passes. Plot the volume of the three captured channels on a chart and you have three little humps representing what each microphone has heard.

What if our playback LSD, though, consists of only *two* speakers (not much of an LSD, but it's easier to visualize)? How would the computer handle reproducing the sound accurately?

By using *both* speakers for the middle microphone's signal output. (The stereo effect may very well help simulate a third speaker in the middle.) The first and third microphones' signals are handled by the first and third speakers. The middle microphone's signal would be mixed in with the first microphone's signal on the first speaker, but you'd also hear the same amount of middle microphone signal on the *second* speaker, in addition to what Microphone #3 captured.

Playing 2,000 audio signals on a 1,000-speaker LSD playback array would be handled the same way, though the computation involved would be more significant, to put it mildly!

Mismatches in the *shapes* of the recording and playback devices is trickier when it comes to audio, but the digital LSD, with sophisticated-enough software, can give the user options for how to handle them. Let's record a walk in the forest using a spherical cluster of microphones, then play it back on a flat rectangular speaker array. (Having two speaker arrays facing each other would be a lot easier, but let's keep it complicated for once.) The most obvious solution is to map the spherical sound captured much like a Mercator projection of the Earth's surface. At some arbitrary point, perhaps at the 'back' of the cluster as it moves through the forest, you would draw a line to divide the sound streams into left and right. You'd assign the left side of the split to the left end of the flat playback array, and the right side to the right end. The top and bottoms of the cluster signals would keep their orientation top to bottom on the playback array. Just as with the Mercator projection, in converting a curved 3D surface to a flat one you'd have some sounds that are overemphasized (or stretched) while others are de-emphasized or shrunken in significance. The computer doing the conversion could compensate for these distortions to the greatest degree possible.

Another strategy is to use the rectangular playback array as a sort of porthole onto the entire sphere of sound. The user could 'turn' within the spherical sound so that the 1:1

relationship between source streams and playback streams is maintained, but only a rectangular piece of the sphere is played back at a time, and that could change depending on the user's wishes. The obvious downside of this is that not every sound would be played at a given time (unless the user wanted to mix the 'non-viewed' sound into a single channel played relatively softly).

These solutions may sound like unacceptable compromises from a *locational* standpoint, but don't forget that however the mismatch between source LSD shape and playback LSD shape is handled, the resolution of the input and output remain very high, and we're *still* getting much better sound resolution (and dimension) than we'd ever get with two or six speakers…

Video mixing in its simplest form is similar to this kind of audio mixing. With video, we're talking about the number of pixels captured for a given area – say, 1,000 pixels covering 3 feet of space – and having only 500 pixels of display space to render it on. Although diminishing resolution (and therefore losing detail) is less than desirable, it is bound to happen, and reducing those thousand pixels in half is something computers are generally good at (reducing an image's resolution in Photoshop is an example).

Another issue with rendering video is whether or not there's a monitor in that specific spot to render the analogous pixels on. The playback array, for instance, might contain a bezel break in the middle of a captured stream's image. In this case, the image would need to be split, with each resulting piece allocated to the surrounding monitors in the array (along with pieces of other signals). The computer will virtually *stitch* pieces of various video inputs into a single image for a given monitor, and probably blank out the pixels that would have been shown if there were screen space instead of a bezel in a given area in order to keep playback from being distorted.

If it's any consolation, having to reduce the number of recorded pixels shown on a playback LSD at least maintains the highest resolution of the output component. *Spreading* occurs when the opposite condition pops up – fewer inputs than outputs. Unfortunately, this means the resolution of the signals that are being output is not as high as the playback LSD is capable of producing, but even still the playback detail should be significantly higher than in other media. Then there's the brute fact that some recordings should simply not be played back on inadequate LSDs, just as low-resolution JPEGs that display fine on computer monitors should not be used for print.

To spread audio, the computer parcels a single channel over multiple speakers, though the algorithms that do this might weight playback to simulate the original's intended location by diminishing the signal's volume somewhat at the outermost speakers.

To spread video, the computer enlarges the original (or decreases the output component's resolution) to match what the output resolution is. This is a bit like taking digital video and playing it back on a television—though the original has a higher resolution than 320 x 240, converting it to NTSC reduces it to NTSC resolution (though often it still looks sharper than video captured only at NTSC resolution).

Writing a program that mixes and spreads audio and video as I've described is not a trivial task, but there's no reason to think it's not possible. Initially it would probably have to be a preprocessing job – download the original recorded data and render it for the specific capabilities of the playback LSD. Eventually, as computers and software become more speedy and sophisticated, it could be done live. Or better yet, throwing *more* computers into an LSD to handle processing this specific task might allow almost simultaneous rendering and display.

The LSD operating system is all about handling the massaging and routing of audio and video streams. Even with a perfect 1:1 match between the input and output arrays, the artist who's producing the final program might *want* to play with the size the overall image or sound is played back at, which would mean enlarging or reducing the original signal anyway. The image of a woman's face, for instance, might initially be displayed at twice its recorded size, then be reduced over time to its original dimensions and finally shrunk to a much smaller point. Tools currently exist that increase the size of an image several hundred percent with minimal degradation of the image's quality.

This discussion of the problems of mapping recorded signals to playback resolution brings us to a rather interesting phenomenon, one that you will need to comprehend before you can truly appreciate the nature of the digital LSD.

In a true digital LSD, the signals, be they audio or video, are fully malleable. They can be manipulated as though they were pieces of stretchy fabric in your hands, made out of dots of image or slivers of sound. To continue the analogy, a motion picture is a single piece of fabric, cut to a specific shape and set in a rigid frame. It is played back the same way each time the film unspools through the projector.

In a digital LSD nothing is sacrosanct about the image's location or nature. Its very *perceptual* qualities can be manipulated in time and space. To illustrate the implications of this, let's suppose we have a video recording LSD array about the size and shape of a motion picture screen – flat, rectangular, and 20 feet tall by 30 feet wide or so. Let's set it up on the Mall in Washington, observing the Washington Monument, the reflecting pool, and various people walking by. We'll record 5 minutes' worth of reality.

The image it captures is "flat," like a movie, and has the same size and shape as the array that captures it. In other words, we have a piece of rectangular sense data that is 5 minutes long. Let's call it "Washington."

The 'depth' of "Washington," in the classical motion picture sense, is a function of the lenses used by the cameras – a 'longer,' or zoomed in, lens captures a shorter depth of field, while a 'shorter,' or zoomed out (or wide angle) lens captures a longer depth of field.

But the 'depth' in another sense is a function of the fact that a frame of motion picture film records a certain image that is projected by a motion picture projector. A movie camera captures light on a frame of film at the focal plane (where the light coming through the lens is focused on), while a movie projector takes the same frame and reconstitutes the image on a screen some distance away by shooting light back through it and out another lens (which also has to be focused, though in this case it's the screen that's the focal plane). Unless the image is manipulated in postproduction, what's captured is what's displayed, in a 1:1 relationship.

Now let's turn to the digital LSD, and try to understand how it goes way beyond movies in this very fundamental sense.

Our LSD version of "Washington" is stored as a bunch of video streams on one or more computers. Each stream is tagged with the location of the recording camera in the overall array, so that the playback LSD will know how to locate its signal correctly.

What LSD shape should we play "Washington" back on? The most obvious answer is that we should use one that is precisely the same size, aspect ratio and orientation as the LSD used to record the piece – i.e., a flat array 20 feet tall by 30 feet wide. But this is assuming, as we do with film, that the qualities of recording and playback need to somehow "match" each other. Remember, "Washington" exists as a bunch of video streams on a computer.

In fact, we could play "Washington" back on virtually any digital LSD, of any shape, that has a sufficiently contiguous array of monitors, *while maintaining the semblance of a flat, rectangular display*. That's because a mature LSD application or operating system *knows* what the playback LSD is capable of, and knows what the original LSD was like, and can translate appropriately.

To show you what I mean, let's play back the flat, rectangular "Washington" on the curved *outside* surface of an LSD sphere.

Your first inclination may be to visualize the "Washington" image stretched around the sphere, as if you were wrapping the sphere with a piece of the stretchy fabric I mentioned earlier. You could do this, of course, but the image would be take on a fish-eye shape, as though it had been shot through a wide-angle lens, and the resolution of the played-back image would vary from lower-resolution at the center to higher-resolution at the edges.

But instead – key point here – on a digital LSD the rectangular "Washington" could appear to *float inside the sphere*, suspended in black space. In other words, unlike a movie projection system, the imaging capability of the digital LSD is *not* dependent on its actual shape! Instead, it is an entirely *virtual* shape, so that just as we can depict a planet floating in space on a flat movie screen, so too can we depict a flat movie screen floating in space on a round object without distortion. The highly-detailed playback array allows us to give the impression of a *deeper reality* than we're used to with motion pictures, and the computer-based nature of a digital LSD means we can manipulate that reality to depict virtually whatever we want.

This is not to say that the audience's orientation to the outside of the sphere makes no difference. There's no doubt a fairly distinct 'sweet spot' where the floating image of "Washington" will look best – meaning perfectly rectangular and distinct in this case. Wander along the outside of the sphere in either direction and the sweet spot will be lost. Nevertheless, the digital nature of a digital LSD means that the shape and size of the visual playback array does *not* require the recorded scene to have been captured by an array of the same shape and size.

The most perfect structure for a digital playback LSD is a dome or sphere that contains the audience within it. This structure creates an immersive experience that can, with a sphere, extend *below* the audience as well. Because both structures are circular and all-inclusive, the arrays can display virtually any LSD program, even ones that are recorded using flat camera arrays.

How would you display, then, a scene captured by a spherical recording LSD in a dome that does not extend the array beneath the audience? One strategy is to have the audience perceive themselves at the *bottom* of the sphere, with the sides curving upward around it to meet at the top. The audience (or someone controlling playback) could 'move' the perceived location within the playback world up or down, even allowing some degree of distortion at the ‘bottom’ so that it can be squished along the bottom edges of the dome and yet remain ‘readable’ for what it is.

The digital LSD recording device, or camera, is an entire industry to itself, with its own line of development that can incorporate advances in the state of the art. An LSD camera might consist of a fixed structure of cameras and/or microphones, arrayed in a fixed grid, and with a fixed orientation to what it’s recording (i.e., all cameras are zoomed and focused on a fixed contiguous area).

Beyond that, however, lies much more interesting territory. First, the array doesn’t have to remain fixed. The structure itself could be mobile, whether attached to a vehicle or self-propelled, and could even be remotely controllable so that someone in an LSD elsewhere could ‘drive’ it around.

Second, the zoom and focus of the components that are recording visual information don’t have to be fixed; they could be manipulated to zoom in and focus on a small object, or zoom out and re-focus to capture a larger scene. To do this, you would build a robotically-controlled array of video cameras. Each camera would be mounted on a robotic swivel, with the zoom and focus of each camera controlled by computer.

Combine the robotically-controlled array with mobility, and you have a powerful recording device that can travel about examining everything in its path as closely as the ‘driver’ wants – with each object rendered in full LSD detail.

Getting into even weirder and potentially more wonderful territory, you could add to your digital LSD recording device something that provides an even finer level of location detail to the mix – radar, sonar, or any other technology that provides distance information. (I’ll use radar for now.) Consider a spherical array of video cameras and microphones pointing outward, capturing the world around it in high visual detail. The fact that a given camera is recording a given section of the entire scene imparts a certain level of location detail to the recorded data. But by adding a component that provides detailed *distance* information for each external object in the scene – via radar – and associating those data with the sounds or images being captured, you might effectively be able to *fully* dimensionalize the recorded information.

If the recording LSD were sitting in the middle of your living room, for instance, the radar would detect the size and distance of each object's surface. Associate that information with the visual and aural information being picked up by the LSD's cameras and microphones, and you have a realistic rendering in three dimensions of the appearance and locations of objects in space.

What would the benefit of this be? Let's suppose move a radar-equipped recording LSD through the living room in a systematic fashion, capturing information about size, shape, color, and sound properties of the room. (Sound properties might include a ticking grandfather clock, for instance.) You would essentially *scan* the room in three dimensions. Later you could manipulate an object's location and appearance discretely, or you could calculate what it would be like to hover at point A or B.

In other words, you would create a virtual reality calculated from an *actual* environment that you could travel through in any way you can think of, much as first-person shooter video games allow you to walk (or run, as the case may be!) through *imagined* environments rendered on the fly by the game device.

To take this one step further, imagine you have captured a room this way, and the room is now displayed on a playback LSD, with you standing in the middle of the LSD and seeing the equivalent of standing in the middle of the room. Hang a location transmitter around your neck, and equip your playback LSD with the means of tracking your movements, and you can effectively walk around the room, and have the LSD reflect your movements, in real time, and at such high resolution you might feel *you are actually there*. The only difference being, of course, that in real life you can't pass through furniture or walk on walls!

I've been writing from the point of view of an *outward* recording device, but an *inward* recording visual LSD captures not only the exterior appearances of each object, but a three-dimensional representation of the object where it is not blocked by other objects. (Though with an inward-pointing dome or sphere recording LSD it's possible at least *one* camera would capture surfaces hidden by other objects.) Indeed, *combining* an inward- and outward-oriented recording LSD *with radar* might allow you to capture a complete representation of an environment with minimal 'scanning' necessary.

Such uses would generate huge amounts of data. How would you store all of it? With analog LSDs, the data are recorded *en masse*, onto an analog medium like tape. If nothing changes from one moment to the next, the image or sound information still takes up "room" on the tape. Also, the entire collection of pixels or sounds is recorded as a series of dots or sound waves no matter what is being recorded. A perfect circle, for instance, will be recorded as a bitmap, with all of the dots used to make up that circle stored for later playback.

In the digital world, however, efficiencies can be realized in the size of the final recorded data by encoding only the data needed to reproduce the original scene. Where a given recording component in the array records nothing, or nothing new, its input can be ‘summarized’ using a simple line of code that says as much – i.e., “no sound waves” or “play same frame for *x* amount of time” – rather than taking up space with frame after frame of bitmap data or second after second of audio data. This type of efficiency encoding is already used in other kinds of digital media file formats.

Beyond this, however, lies more interesting challenges for encoding raw sense data efficiently. I leave the strategies for meeting these challenges to the professionals, but with the huge computing power available in a cluster of today’s computers working as one computing device, and with the advances in number-crunching coming our way in the next decade, it doesn’t seem out of the question to suggest that a ‘vectorization’ of bitmap data would be increasingly possible. (Note that this is similar to the ‘vectorization’ of analog audio data mentioned earlier.)

What does ‘vectorization’ mean? If I take a crisp digital picture of a perfect black circle printed on a sheet of pure white paper, the digital camera stores the entire picture as a series of dots. I could reproduce that image, however, if I wrote a mathematical formula that described a circle and oriented it correctly on a white background. This formula take up 50K of my computer’s storage capacity rather than the 600K a high-resolution JPEG file might take.

Doing the same thing to a complex image is merely a function of the speed and sophistication of the encoding algorithms. An entire scene captured by a digital recording LSD would be processed (whether on the fly or after the fact) so that the final result was a bunch of equations or values. The playback LSD’s computers would then ‘render’ the equations or values – again, either before-hand or on the fly – according to the size and quality of the output array, not to mention the wishes of the person programming the LSD's output.

In fact, by vectorizing the sense data it's probably easier to vary the playback location, and applying filters of various sorts to the playback 'streams' would be much more efficient, especially if the data flowing over the network weren't big chunks of bitmapped data but elegant little formulas instead.

In such a scenario, a 'rendering layer' of computing might be necessary to receive the encoded data, convert it to a bitmap for output, and ship it off to the appropriate component (be it another computer or the final output component) in the array, much the same way a drawing program like Adobe Illustrator converts vector graphics into a screenful or pageful of dots. As computers grow more powerful, it's likely a single machine will be able to handle rendering and final output simultaneously.

One more important aspect of the fully digital LSD is the ability of the computer – be it a single computer or a cluster – to handle its own quality control. Modern operating systems include color calibration technology so that a specific version of the color red on one monitor will match very closely the same red on another monitor. The same quality assurance capability holds true for audio as well; a digital LSD could verify that each speaker in the array is generating sound by playing a tone on each one, and either it (or another computer in the array) picking the tone up with a microphone. By having *all* of the microphones in the array pick up the sound, it could also quality check the microphone array, and compute microphone and speaker characteristics based on the information each microphone receives. A digital LSD could also be configured – perhaps could self-configure – to "know" what each speaker in the array can do, and adjust its output accordingly. It could do this by maintaining a database of speaker types and characteristics and checking each speaker added to the array against it, or by simply 'listening' with the microphone array and measuring the way a given speaker handles a variety of frequencies and tone types.

All of these speculations may sound exceedingly complex to implement, but there's no reason to think at this point that these ideas can't be realized – many of them probably are *already* realized in some fashion or other.

The size of a digital LSD is entirely dependent on the technology used to build it. The general rule I've come up with is that bigger is better, but smaller is just as good if it's *finer*. Let's say we want to build an LSD that operates at the highest visual resolution humans are capable of discriminating – about 17,000 dots per inch, if my source is correct. I don't know how to do the calculation, but it seems clear that to achieve that kind of resolution with today's monitor technology we'd need a lot of monitors at a significant distance from our retinas – a mile away? Ten miles away? (Send me the answer if you know it!)

Future display technology, however, is likely to drive up the dots-per-inch a display device can output. I've seen 200dpi color and 300dpi grayscale flat screen LCD monitors. The higher the native dpi of the display device, *the closer the LSD array can be*, and the fewer components you need in order to build one that maintains our 17,000 pixel per inch resolution.

I predict that some day in the future, probably well after my death, you will be able to walk into a living room-sized LSD and have the same visual experience you'd previously have had to have built a vastly larger device to achieve. And even farther down the road, it will be affordable!

The two downsides I see to a smaller, finer LSD such as this is first, you lose the sound quality in a larger space – bass soundwaves, for instance, need room to propagate. Second, I envision the best LSD experiences to be *group* experiences, and the personal device precludes that to some extent. But who's to say that you wouldn't use your personal LSD *to attend larger LSD events*? After all, if you are seeing the world around you in as detailed a fashion as reality itself, why not, and what, aside from the sound and the physical proximity to others, would the difference appear to be?

To summarize, a digital LSD is a computer-driven, fully digital incarnation of the LSD concept. It offers the maximum resolution of both sound and video possible. Video, for instance, is displayed at much higher resolutions than motion pictures have ever approached. With enough monitors in the array, and a sufficient distance between the array and the viewer, a digital LSD can theoretically display up to, and even surpass, the resolution the human eye is capable of seeing. The implications for how this would affect our fundamental notions of what reality is remain to be seen, but certainly the effects would surpass by far those for any other medium yet devised by humans.

8. Programming a Digital LSD

If I've done my job correctly, by now you should have some inkling of the power a fully digital LSD offers artists and audiences. It can display images as real as reality itself, and can create a sound environment more rich than you'd ever be likely to find in nature.

What would you do with all this power? What 'program' would you play on it?

The most immediately obvious thing to do is to record reality using LSD techniques and play it back on as detailed a playback LSD as you can build or find. Using the recording and playback techniques described in previous chapters, environments can be captured, rendered and even explored in extremely realistic audio-visual detail. With the introduction of live signal transmission and reception to recording and playback LSDs, and further enhancements involving radar or sonar and mobility in the recording device, the environment that's being captured can also be explored *live*.

One quality of the LSD experience that is important to mention, especially in this context, is that it is safe. You are experiencing an audio-visual world remotely and probably anonymously, so you can go anywhere and do anything the recording cluster will allow. Suppose you built, say, a *submersible* recording LSD, that could be driven around a coral reef, all the while beaming back to you an extremely detailed reproduction of the underwater scene. With cameras pointing all around the sides, bottom and top, you could sink down to the reef and check out the scene (the submersible would need some means of making sure you don't go crashing into things!). You wouldn't need a scuba tank, or even know how to swim – there's no danger of drowning. And you needn't be afraid of sharks or sharp coral either.

A pre-recorded exploration of a coral reef would be compelling as well, but the live transmission of LSD content – *LSD telephony* might be a good term – opens yet another frontier for artistic and commercial exploration of the device's capabilities. Put recording LSDs with live signal transmission in public places around the world, and you can 'visit' there as easily as flipping channels. Doing so takes the notion of webcams to an extremely high level, but the concept is the same. Make these recording LSDs mobile (freely or in a limited way) and you might be able to go exploring the streets of Barcelona or Paris, or the Amazonian rainforest, or anywhere else.

This is just a start, though. To illustrate, let give you as succinct an example of the power of LSD telephony as I have thought of to date. This example also introduces the power of LSD *layers*, or the use of multiple 'coatings' of LSD arrays to achieve sophisticated effects.

Let us build two digital LSD domes, one in your back yard, the other next to a fountain in New York's Central Park. The domes are, say, 15 feet tall.

Your dome contains playback components pointing inward at you – i.e., speakers and monitors. The Central Park dome contains *recording* components pointing *outward* – i.e., cameras and microphones. Because they are domes, every inch contains array components. The recording components in Central Park are sending live signal to the corresponding playback components in your back-yard LSD.

What does this look like to you? Take a moment to imagine it; it's no different from the LSDs already described. The scene displayed around you as you sit or stand in the middle of your dome is an extremely vivid rendering of the fountain, the plaza, the lake, the trees, the people and the sky *at that moment* in Central Park. It's like a live remote feed via television, only the world displayed is much more detailed in image and sound.

If you build these LSDs using only recording components in Central Park, and playback components in your back yard, you cannot interact with the people you see. But if you add capture components to the inside of your LSD, and display components to the outside of the New York one, the people in Central Park *can see and hear you* in as high a resolution as you see and hear them.

The video capture components in your dome will not only pick up your image, they will also pick up the images being displayed on your monitors, so they must be tuned to the frequency of your monitors' refresh rate so you won't get banding. Assuming you have tuned them correctly, the cameras in your LSD will capture and send back what the monitors display at high resolution, meaning that the people in Central Park will see *Central Park* fed back to them. If the fountain is behind the Central Park dome, and displays at the back of your LSD, someone standing in front of the Central Park dome will see a very detailed version of you standing *in front of* the fountain.

In other words, the Central Park dome itself will have in a sense *disappeared*. To an observer in New York, it will seem as though you are standing there, just as it will appear to you as though the observer is standing in front of you.

And remember we have sound as well. The feedback issues are probably pretty thorny; I leave that problem to someone more knowledgeable about them. But certainly there would be a way to be able to interact verbally with the people "around you" at the Central Park LSD.

You are not actually there, but the effect would be so convincing that for your purposes you simply *are*. What would we call a device capable of creating the powerfully vivid impression that you are somewhere else entirely, where you can interact with the people at the other location? Would it be a telephone? A video conference? Those technologies pale in comparison to what's *really* going on – *LSD Teleportation*. While you are not *physically* there, you are *practically* there.

Of course, this teleportation doesn't have to happen between your back yard and Central Park. It could be between your office and another office. Or it could be between your office or home and *multiple* locations – a branch office on one side of your LSD, your child's day care center on another, and a store or sporting event or relative's house or anywhere else on yet another.

Let's not forget that in our example we are only using one side of each dome – the inside of your dome and the outside of the Central Park dome. You could just as easily configure the unused surfaces so that someone *inside* the Central Park dome could be seeing the image and sounds of your back yard as captured by the *outside* of your dome. And with the same mix of components, someone outside your dome could be interacting with someone there too. Or the New Yorker inside the Central Park dome and the person outside your back yard dome could be experiencing entirely different locations independently.

But suppose you don’t *want* to interact with those surly New Yorkers! And don’t want them seeing you or your back yard. Just turn off your recording components and you’ll become an anonymous voyeur.

LSD Teleportation is possible using simple, high-bandwidth audio and video feeds. It doesn’t take into account what computers can do to manipulate recorded reality, and surpass reality altogether. While the ability to record and play back real-world environments at supremely high resolution is a large part of what makes the LSD what it is, the power that computing brings to digital LSDs truly opens up the device to artistic exploration on levels without limit or precedent.

Imagine a canvas that surrounds you with virtually infinite possibilities of painting image and sound. Even a still-motion analog LSD made of televisions presents a field of exploration of color, line, and representation that gives artists a chance to do things never done before. But a fully digital audio-visual LSD goes so much farther that the aesthetics of LSD art will take years, probably decades, to fully explore.

(Is that notion that it takes decades, even centuries, to explore the aesthetics of a medium unprecedented? Of course not – the artistic possibilities of computers, motion pictures, music, television, radio, painting, poetry, novels, etc. are all *still* being developed.)

Let’s leave aside video for now and just talk about sound. Our digital LSD may contain 1,000 speakers, or 10,000, and its shape might be a rectangle, circle, box, dome, sphere, or any other design. We can play stereo sound on it quite easily – orient one channel to one side of the array, and the other channel to the other side.

Even with only two channels, we can do a lot more than we ever could on traditional systems or purely analog LSDs. We can move each channel around: we can fly them across each other laterally, raise and lower them together or independently, widen or shrink the number of speakers each is played on together or independently, jump them from one part of the array to another, clone them at multiple points in the array all at once, clone them in a sprinkle pattern over time or all at once, swirl the cloned channels around a center point, march the clones around the outside edge of the array such that one channel is chasing the other, fade the channels in and out at the same or different points in the array, and on and on, not forgetting that combinations of these movements could be happening simultaneously.

We could have as many channels going at once as our LSD allows, all doing the same sorts of things in much greater variety and complexity. We could process the signals in any number of ways – dimensionalize by frequency and soundfield location (the degree to which a sound is shared across two channels), add echo, reverb, equalization, and any kind of sound filter possible. We could mix a hundred channels down to one, then separate them out again in the space of a second to a hundred speakers. We could have each note of a piece – think “Flight of the Bumblebee” – play on a different speaker around the array. A section of a song might be played by ten different instruments at ten different locations in the array. A song that begins softly could start out at the center of the array, with more of the array used as the song grows louder. The entire LSD could play a single unified sound; imagine a woman’s low whisper captured by 100 microphones and played back – softly, but also very large and vivid – on a 100-

speaker array; now imagine it recorded with 1,000 microphones and played back on a 1,000-speaker array.

In a dome or sphere, sound ‘journeys’ could take the audience on a tour through the real world or something entirely imagined. LSD DJs would have a vast new territory to explore and create complex and compelling mixes in. Acoustical instruments could be recorded and played back with much greater fidelity and subtlety than they are now. The acoustics of the world’s premier concert halls could be recorded and played back either with a performance actually happening there or applied to a performance recorded somewhere else.

Large bands of speakers could be devoted to a single sound, with other sounds sprinkled across other parts of the array. Each instrumentalist could have a section of the array devoted to his or her output, and the sections could be moved and intermixed in any way. A composer might generate music from one area of the array and produce a myriad of other sounds and effects from other parts. Sounds could move closer or farther away, or both simultaneously. The LSD could be programmed to generate a sequence of sounds that unfold around the array in semi-random patterns when just one key is struck.

That’s just a smattering of examples that come to mind, enough, I hope, to convince you that the fields of artistic possibilities in an LSD are much more vast than stereo or 5.1-channel sound could ever be.

Similar effects are possible when location is added to visual playback. Visual storytellers could show a single unified scene over the entire array, or float many scenes simultaneously in different areas (imagine Mike Figgis’ *Time Code* unbounded from a rectangular screen). Not only could these be “flat” in the context of the LSD screen’s shape and size, they could be distant, fuzzy, and still one moment, then close, sharp and moving a minute later. Representational visual content can be life-sized, enlarged, or diminished as a result of either the recorded size or the computed playback size, and of course can be scaled upward or downward at will.

Because we're still stuck in the flat, boring world of motion pictures, computer-generated animation is currently played back on a flat rectangular screen. But there's no reason why a movie like *Toy Story* couldn't be calculated and imaged for a circular, spherical or other playback shape. (Indeed, the *Toy Story* movies were re-rendered for their video release to take into account the difference between film and television aspect ratios.) Along these same lines, a video game powered by a massive LSD computing cluster could render a three-dimensional immersive world in a truly immersive device rather than a flat rectangular one.

Beyond these types of representation is yet another dimension of LSD programming – interactivity with the audience. The way sounds and images are generated and/or displayed could be a function of what the audience itself is doing, whether *en masse* or individually or somewhere in between (teams, for instance). The LSD could track the audience's movements by using location detection. In a large LSD, the program might change to reflect the audience's movements among different areas of the LSD's interior. The audience might be able to control the type of program generated by voting with a simple remote control; they could affect volume, rhythm, location, or any other variable the program would allow to be manipulated.

New kinds of musical instruments that are built around wireless controllers could allow audiences to generate their own programming in the LSD. One audience member, for example, might have a simple wand that produces one sound when it's waved in one direction, another sound when it's waved in another direction, and yet another when it's raised and lowered (with combinations of these sounds when combinations of movements are made). Other audience members could have wands with the same or different programming, and the combination of their effects would create the overall experience.

Audience members could use pointing devices to literally paint colors on the screen array, or paint sounds on the speaker array. Control pads might allow them to alter the colors or images or sounds (or all three) their pointers produce, as well as many other qualities that can be applied to image or sound in an LSD (the amount of time a sound or image event lingers, its location over time, its interaction with other sounds and images, etc.). While we tend to think that sounds would interact with sounds and images with images, an image could affect a sound and vice versa. An abstract image might actually sound the way it looks, with different sounds emanating from different locations in the image.

While groups of people may interact with each other and with the LSD, probably the most intense experience possible in a digital LSD comes you are alone inside one and your location is constantly transmitted to the device. This allows the program to generate an audio-visual world *that moves in relation to your movement*. You could run around the LSD's interior while the computer-generated image of a distant planet's surface moves exactly as it would *if you were actually running around on the planet*. If the LSD were large enough, it could do the same with groups, similar to how a theme park ride like Disney's *Star Tours* works, although without necessarily having a predetermined program.

Just because we are used to the world moving at a dependable rate when we make our way through it doesn't mean an LSD has to do the same thing. I could bicycle around the interior of an LSD and have the screen display a reality passing by that gave the impression I was travelling a hundred or thousand miles per hour rather than just five or ten. I could bicycle my way to the moon, for instance, in ten seconds. Or cross the ocean with a single leap. One LSD I've imagined building for a theme park involves a water ride through a large tunnel that is a tubular, inward-pointing audio-visual LSD. The riders could float down a fairly sedate section of the ride, with a pastoral scene unfolding at a somewhat advanced clip around them, then when they begin to go faster the scene around them accelerates even more. When they come to a modestly precipitous drop down a water chute, the display would show a very sudden fall through space that ends hundreds of feet below where they started from, even though their boat only slid down 20 feet.

The scale of an LSD will affect the types of things you might want to do with it. A personal-sized one that you can't move around in very easily would not work well for the water ride, though you could certainly sit in a recliner (or better yet, a personal motion simulator) and have an exciting whitewater experience regardless – the body will react to physical images of motion by leaning against turns or pulling back at a sudden stop, even if the body itself isn't moving.

A very large LSD could allow not only audience movement, but the construction of a physical environment inside the structure that is itself compelling. Someone could simulate a deserted, half-ruined city on an alien planet that still has working guns fixed to turrets; the array could show a pleasant alien spring day that suddenly grows dark with cloud until a huge flotilla of ships comes out of the sky to attack. The guns might shoot laser-guided artillery images as red streaks that are displayed on the array, with the LSD reacting to direct hits on ships by showing spectacular explosions. If a portion of the *floor* of such an LSD were made of monitor arrays, from the vantage point of the gun turrets the destroyed ships could hit the ground and explode there very realistically.

Lastly, while I have focused on depicting sense information (whether recorded in the natural world or generated by computer), *mind* information can also be displayed on an LSD's arrays. As large computing devices with extremely high-resolution output components, LSDs can be used to display complex calculated information in addition to the types of things we use personal computers for. A huge 3D representation of the minute-by-minute fluctuations of every stock currently traded on the New York Stock Exchange, grouped as desired by industry, market value, increase or decrease in value, earnings history, etc., might be fun to watch for a while (who knows – maybe you'd pick up on patterns that might help you make investment decisions). A large detailed representation of your company's cash flow could help you make management decisions better than the same information (or a reduction of it) in a spreadsheet. In other words, sets of information that are too complex to see at once on a single personal computer screen can be displayed much more comprehensibly on a very large LSD array.

These are all just the barest beginnings of thinking about how to program a digital LSD. The intensely rich visual and aural properties of the digital device yields possibilities for artists that are simply too fantastic to fully imagine; it would be a crime *not* to get a good digital LSD into the hands of those who might create works that go well beyond these simple musings.

Even more important than battles on alien planets and cash flow analyses is, after all, a simple and exceedingly rare aesthetic quality – beauty. We can experience the beauty and ugliness of our own immediate physical world through our senses; we can experience a reduced version of other parts of the world through our media; we can appreciate the inner visions of other human minds through art. Not only can we do all these things with much greater fidelity and power with an LSD, but have many more worlds to experience in an LSD than could ever be made with previous media. A shimmering fantasy of glowing, moving, changing color, accompanied by a richly-detailed but very gentle soundscape. An intricately beautiful world of shape and form constructed in response to the direction we choose to walk in. A magic carpet ride through a garden from a butterfly's point of view. The simple and profound examination of the veins on the back of a leaf.

I see the LSD as a means not only to experience the beauty produced by nature and the human imagination in an immensely heightened way, but to increase our ability to appreciate this beauty. The digital LSD is a powerful means of opening our eyes and ears to more beauty than we can imagine.

9. LSD Layers

10. LSD Shapes and Structures

11. The Business of the LSD

12. The Future of the LSD