

The Importance of Being Digital

Paul Lansky

3/19/04

tART-lecture 2004

commissioned by the tART Foundation in Enschede and the Gaudeamus
Foundation in Amsterdam

Over the past twenty years or so the representation, storage and processing of sound has moved from analog to digital formats. About ten years audio and computing technologies converged. It is my contention that this represents a significant event in the history of music.

I begin with some anecdotes.

My relation to the digital processing and music begins in the early 1970's when I started working intensively on computer synthesis using Princeton University's IBM 360/91 mainframe. I had never been interested in working with analog synthesizers or creating tape music in a studio. The attraction of the computer lay in the apparently unlimited control it offered. No longer were we subject to the dangers of razor blades when splicing tape or to drifting pots on analog synthesizers. It seemed simple: any sound we could describe, we could create. Those of us at Princeton who were caught up in serialism, moreover, were particularly interested in precise control over pitch and rhythm. We could now realize configurations that were easy to hear but difficult and perhaps even impossible to perform by normal human means.

But, the process of bringing our digital sound into the real world was laborious and sloppy. First, we communicated with the Big Blue Machine, as we sometimes called it, via punch cards: certainly not a very effective interface. Next, we had a small lab in the Engineering Quadrangle that we used for digital-

analog and analog-digital conversion. (A few years earlier it had been necessary to drive forty miles through central New Jersey—not a pleasant ride—to Bell Labs in order to use their converters.) Our lab included two large digital tape drives. One was a vacuum column monster that read 800 BPI (bits per inch) tapes, holding about 20 megabytes. It sounded like a jet engine when the vacuum columns were engaged. The other was an unreliable 1600 BPI mechanical drive. We wrote the digital tapes on the mainframe and hand-carried them to the lab. The first computer made by Hewlett Packard, a refrigerator-sized 2116 minicomputer with 64k memory, ran the drives and the converters. The noise in the room was deafening when it was all powered on. Our 16-bit D-A converters were made by Hewlett Packard. Our A-D converters were 12-bit and custom-made. At the end of the chain were Ampex and Scully 15-ips tape recorders,

I was in my 20's, had good high-frequency hearing and distinctly remember marveling at the purity of the signal coming directly off the converters, as best I could hear it through the noise in the room. I noticed and reluctantly accepted the gentle high-frequency hiss that the tape machine added. In the mid-1970's my first computer piece, *Mild und Leise*, was issued commercially on LP thanks to a competition held by the International Society of Contemporary Music/ League of Composers[1]. It was a 19-minute work made entirely synthetically (there were no sampled sounds), and it fit neatly as track 2 on one side of the LP, following a short two-minute piece. It was not surprising to me that in addition to the tape hiss there was the added noise of the needle dragging through the grooves, and neither was it surprising that the quality of sound was noticeably better the first time I played the disc than the second or subsequent times. I even learned to accept the mild pre and post echoes of the sound during quieter moments when print-through from adjacent grooves bled through. What was surprising was that the quality of the reproduction got somewhat worse toward the end of the piece. I mentioned this to my father, a recording engineer for Capitol Records at the time, who told me that I was hearing 'inner diameter distortion'. The angle of the needle to the grooves grew more acute as the inner grooves were reached, creating distortion. This was very distressing.

So the act of transferring digital sound into the analog domain was fraught with traps and pitfalls at each stage. But the story doesn't end there. During those days we were working with very low sampling rates, typically 14khz, in order to

accommodate the slow speed of the tape drives (and also save computer time since it could take hours to create a few moments of sound). Then, as now, at the end of every digital-analog conversion is an analog low-pass smoothing filter that eliminates the mirrored signal between the Nyquist frequency (half the sampling rate) and the sampling rate. Our filters were hand-designed in our engineering shops and were a delicate and sensitive conglomeration of capacitors and transistors. But, in order to completely eliminate any frequencies above 7khz the filters had to start to slope down at about 5khz. Thus the birthing pains of the digital signal were significant even before it hit the tape recorder. We subsequently worked out a version of over-sampling commonly used in CD players these days. The signal was up-sampled to twice the sampling rate, allowing an analog smoothing filter to be applied at a higher frequency, well above the signal we cared about, and the digital signal was first low-passed by a digital filter, while still in the digital domain, which we were able to design in software with a much steeper cutoff, thus giving us audible frequencies well over 6khz! (A lot of work for 1khz worth of frequency range.)

Bringing sounds from the analog world into the digital world was an equally painful process and it was actually more perilous since the hardware was poorer and noise entered the picture on the way in, and never disappeared.

My point in describing this process is to develop some perspective on the two sides of the fence that we lived with in those days. The power of digital signal processing was at our fingertips, but our ears could only hear a poor approximation of it.

It was really not until the early 1990's that things changed much. In the mid 1980's our work moved to personal computers, still pitifully slow compared to today's machines. In the early 90's computing and audio technologies converged: recording studios switched to digital production, and eventually the compact disc became a standard storage medium for data as well as audio. People began to talk about the demise of 'vinyl' and the debate on the relative virtues of analog versus digital sound that had been brewing for ten years came to a boil. Today there still remains an ultimate moment when every sound must be converted to an analog signal, even if it is just at the level of a speaker cone, but now this only happens once in the life of any audio reproduction.

Despite my early trials with digital and analog sound it is not my intention here to weigh in on one side or another in the argument about the virtues of analog versus digital sound. The heat generated by this discussion has far exceeded the light. Audio technologists rage on about the differences and sound theorists grasp at flimsy straws. In a 1997 *Musical Quarterly* article, Rothenbuhler and Durham, for example, go to the absurd extreme of stating

The crucial difference between phonographic [analog] techniques and digital recording and playback is that the digital storage medium holds no analog of either the original recorded signal or the resulting playback. The digital storage medium holds numbers—data. These numbers are related to waveforms by a convention arrived at in intercorporate negotiations and established as an industry standard; but they could be anything.[2]

I suppose the ‘intercorporate negotiations’ they’re referring to were the ones held by shepherds 50,000 years ago while they were sampling the motions of heavenly bodies and deciding how to plot them: perhaps the first digital representation of a continuous signal[3]. They go on to assert that analog recording has a ‘physical relation’ to the sound and thus has a continuous link to the originators of the signal, while digital representation is merely a ‘measurement’ of the sound. What reaches our ears from digital signals, therefore, are not the voices of the past, but reports on the voices of the past. The supposed virtue of analog representation lies therefore in its physical connection with the dead. Their point is charming even though it is based on a misunderstanding.

There are certainly things to be said on both sides of the argument concerning the accuracy of either technology, although I would assert that digital technology, if not already, will soon be capable of being indistinguishable from the best analog reproduction. The ultimate question in the digital/analog debate is therefore, to my mind, not one about sound quality. What is most interesting is what effect it has on what we are able to do with sound.

There are two important areas I’d now like to consider, both a function of digital signal processing (DSP). The first is the new status of the idea of an original and a copy. The second is the distinction between hardware and software. Each of

these areas carries implications for the meaning of music.

It seems obvious that having a digital representation of a signal means that there is no longer any distinction between an original and a copy. With care, a digital signal is capable of being duplicated exactly an infinite, or at least a very large number of times. Every 'copy' of a compact disc contains exactly the same set of numbers as the original digital master, and is capable of producing the same signal an infinite, or at least a very large number of times. There is no artifact introduced by copying or processing at any stage, nor any degeneration of the storage medium caused by the act of listening. There is no trace of a history. My dismay at the successive stages of degeneration of my precious mainframe-created music in the early 70's is gone. Today, everyone who hears my music on CD, for example, is capable of hearing a perfect performance, identical (to within the quality of their audio equipment) to what I hear when the signal is first created. I no longer have the privilege of best possible point of audition, which is now shared by everyone, and there is only a one-stage transition from the digital to the analog domain, and that of course is necessary because we don't yet have digital ears.

The current tumult in the recording industry is entirely due to this fact. Even though mp3 files, a compressed form of digital signals, are inferior in audio quality to the original digital version, the fact that they too can be copied indefinitely with perfect accuracy, further intensifies the debate. Everyone owns the original, and the propriety of ownership is no longer bounded by physical laws, only by legal statutes. The pathetic efforts of the recording and media industries to curtail this by introducing degeneration, either by physical means or by software will ultimately fail. They are attempting to reshape the digital world in the image of the previous analog generation and fail to see the shifting ground beneath their feet. This leads to concerns over intellectual property and copyright, which I will evade here by choice, except to say that there will be an ongoing battle between opposing forces and this will merely slow, but not stop the realization of the consequences of digital formats. A digital format is not a result of 'intercorporate negotiations', but rather a mathematical object whose manipulation is the business of computers. Decoding the numbers on a DVD or compact disc, while requiring some advanced knowledge and programming skill, is not rocket science. Anything in a digital format can ultimately be read as a string of bits, and reproduced by anyone[4].

The profound difference between reproduction in the digital and analog domains obviously lies, then, in the extent to which analog representations trace a history of the actual copying process. When we view or hear something that is an analog copy we are simultaneously experiencing an intervention between the original and our moment with the object. Someone at some point had access to an earlier generation and copied it onto some medium, adding artifacts in the process. When you see a Xerox copy of a newspaper clipping it is clear, for example, that what you are seeing is the result of the specific actions of an individual. The original peers through the artifacts introduced in the process of copying. When experiencing a digital representation there is not usually any intervention between us and the original: what we are experiencing *is* the original.[5]

The elimination of degeneration caused by copying leads us to a world in which signal processing carries few penalties in terms of the creation of noise and artifacts. This means that the space between imagination and realization has vast new potential. Recorded sound is no longer a fragile object, subject to rot, print-through, copy noise, distortion and inaccurate reproduction. The consequence of this is that editing, mixing, transforming and otherwise modifying sound has an entirely new dimension and creates new realms of potential. The quality of any processing will basically be subject to degeneration caused only by the attributes of the original. Nothing is added as a result of external factors. (There are, of course, limitations introduced by bit-depth and sampling rates, but most processing is now done in floating point format, and higher sampling rates move frequency domain distortions well out of the range of human hearing.) The tradition of *Musique Concrète*, for instance, rested at first on a very flimsy foundation. At every stage potential for degeneration entered the picture and exerted a strong constraining influence on what was possible. The finished product, moreover, was necessarily of another generation and decidedly inferior[6]. The genius of early electronic works by Stockhausen, Varèse, Berio, Schaeffer and others lies partly in the ways they circumvented the limitations imposed by analog processing. And although constraints are always part of the challenge to the ingenuity of composers (whoever would have thought one can make great music on a harpsichord), the limitations of potential degeneration shrank the conceptual domain in which they were working. Nevertheless, a gradual transformation of our perspective on the potential meanings of recorded

sound begins at about the same time.

The emergence of digital signal processing helps complete an evolution that began many years ago concerning the meaning of recorded sound. The visionary originators of *Musique Concrète* saw that recording was potentially much more than merely a notation of past events. Coming about fifty years after the first recordings they began to imagine musical ways of cashing in on the act of hearing a recording as a primary experience. We see a gradual acceptance of the loudspeaker as a kind of musical instrument beginning at this point. Interestingly, film sound theory provides a useful perspective on some of these same questions. The role of sound in film helps develop our perception of the meaning of the loudspeaker. James Lastra, in *Sound Technology and the American Cinema*^[7] distinguishes between ‘identity’ and ‘non-identity’ theorists. Briefly, the former believe that recorded sound has the potential to be indistinguishable from the live sound, while the latter feel that recording destroys a large part of a sound’s essence and presents a different order of experience than the original. The debate has its origin early in the history of film when sound technicians were divided about whether or not the aural perspective of the audience should be as close as possible to that of a potential auditor actually in the scene. Practice in film evolved quickly to generally abandon any attempt at total realism and as a culture of film viewers we have become accustomed to disparity between the physical structure of a scene and its audio representation. We’ve long since learned to associate the sounds of struck coconut shells with horse’s hooves, for example, as well as accept the disparity between the apparent physical characteristics of some architecture and the acoustic signature of sound presumably occurring in that space. While filmmakers nod toward realism in this respect it is generally the case that their efforts are half-hearted. Lastra bridges the gap between the identity and non-identity theorists by focusing on our perception of the inner narrative of a sound text. Our attention oscillates between meaning, timbre, texture, rhythm, syntax, pitch, creating a complex weave in which the total package matters less than the aggregation of the individual characteristics of perceived sounds. Quoting Lastra:

There is never a fullness to perception that is somehow ‘lost’ by focusing on a portion of the event, by using the event for certain purposes, or simply by perceiving with some particular goal, say understanding, in mind.^[8]

Film theory helps intensify our perspective on the nature of recorded sound by illuminating questions about our individual interactions with, and understanding of it. In other words the space created by recording is malleable: though fixed in spectral terms we react to it individually and idiosyncratically. It goes far beyond merely being a notation of an event.

In essence then, a recording can create what could reasonably (although unfortunately) be called a virtual world and we as listeners have become acculturated to peering into that world, accepting and disregarding its limitations and its contradictions of reality. It is instructive to remember the astonishment with which each generation, beginning with Edison, greeted the new technologies: viewers fled the image of an oncoming express train; Edison's 'tone tests' challenged listeners to distinguish between a real singer and what we would now consider a terrible recording of one; the wonders of stereo reproduction and now the marvel of multi-channel digital sound poke at the potential for recorded sound to ultimately be indistinguishable from the real thing. (Whether or not this is ever possible is beside the point. We certainly are approaching that goal). In each case we formerly peered into a world that had some sort of curtain around it, and was only a weak approximation of reality as we know it. As the technology improves the curtain becomes more transparent. But rather than try to cast recording as an incomplete representation of reality it is more useful today to imagine that there are two realities, the experience of recorded sound and the experience of live sound.

Just as in film, unrealistic juxtapositions, overlaps and manipulations of time often contradict our real-world experiences but we accept them as collage-like interpretations of reality, so too in recording we accept things as musical fact that would be impossible to experience in the real life. Most professional recording today is designed to create an illusion of reality that differs substantially from an original experience. Not only are errors removed so that the performance is at least letter perfect, but in most cases the point of audition is entirely artificial. An orchestral recording is going to entail dozens of microphones and many channels, enabling the editors to tweak the balance among many different aural perspectives at any moment. Not even the best seat in the house has this vantage point. The difference between hearing a good orchestra recording and hearing a live orchestral performance is vast, and probably not reasonably comparable. It

is impossible to assertively state that either is superior, they are simply different. In popular music it is probably more common than not that a given recording consists of innumerable overdubs.[9] The difference between a recording and a live rock concert is even vaster than in the orchestral instance. There is no way to even approximate the body-shaking experience of a stadium-sized amplification system on a home CD. The rock concert model is, in many ways, an emulation of the CD, which often precedes the concert tour. The orchestra recording, on the other hand, goes to some lengths to create a virtual image of a concert experience, unrealistic though it may be. Most other recordings fall somewhere along this spectrum.

Recorded sound has thus emerged as a complex and rich arena of musical potential. It is no longer merely archival. In the past fifty years many composers have been attracted to it as a primary medium, and the emergence of digital signal processing has inaugurated a new stage in this evolution. I will now walk through a series of examples to demonstrate the power of DSP for musical purposes.

Much of the work that many of us have done in electronic music invests heavily in the idea of hearing recording as a primary experience, and in a world in which we peer at a kind of virtual reality through the windows, or lenses, of loudspeakers. We cash in on the suspension of disbelief that recording and film has taught us to employ: that even though what we perceive contradicts our experience of reality, we still accept it as a newly formed version of reality. It is into this virtual world, influenced by our acculturation of recorded listening that the most interesting meaning of digital signal processing arises. It is here that the worlds enabled by DSP manifest their real potential.

Many, but not all of the things we do in the digital domain are at least theoretically possible in the analog domain. Analog filters can do the work of digital filters but at considerably more cost and with less efficiency and flexibility. My primitive example above of using oversampling to move a signal out of reach of sloppy analog anti-aliasing filters, and using finely tuned digital filters to do the work instead, is a simple case. The power of software, on the other hand, enables us to work with sound in ways that are impossible to imagine using analog tools. Thus the space behind the speaker now becomes one that is vastly more pliable than before, and as observers and creators of virtual worlds

we have significant new abilities.

One of the most interesting aspects of digital signal processing research is in the area of separation of the various components of a sound: pitch, duration, timbre. A substantial difference between working in the analog and digital domains is the ease with which, in the latter, we can create and store analytical data about various aspects of a signal. Modeling thus becomes a vastly more flexible and powerful enterprise. No longer wedded to a direct link between these parameters composers are free to imagine idiosyncratic uses for each, in combination and in indirect ways.

Some of my earliest work in computer synthesis involved a modeling technique known as Linear Predictive Coding (LPC). This technique is specifically modeled on the mechanisms of speech and has its origins in the Channel Vocoder, developed by Homer Dudley at Bell Labs in the 1930's [10] This was an analysis/synthesis machine that separated out the components of speech into excitation function (glottal folds) and filters (head, chest, nasal passages). During the 1960's and 70's researchers at Bell Labs developed LPC as a digital technique along similar lines. Briefly, in the analysis process multi-pole filters are constructed for roughly each 1/100th second of speech to approximate the vowel formants at that moment. In the synthesis process an artificial excitation function, easy to synthesize, is fed through these filters. Various means are used to detect voiced/unvoiced speech in the analysis and in resynthesis white noise is used in place of a pulse to create plosives, fricatives, etc. By independently varying the rate at which the filters change, and the frequencies of the pulse excitation-function we have effectively separated pitch, timbre and tempo. In other words the crippling stranglehold of the tape recorder, in which changes in tape speed uniformly affected pitch, timbre and speed no longer holds. In addition to this, Kenneth Steiglitz devised a method of altering the formant characteristics of the filters allowing us to change the apparent dimensions of the original source: changing a violin into a viola, a man into a woman, etc. [11] Example 1a demonstrates different LPC resyntheses of a single spoken phrase. Examples 1b and 1c are from my early work *Six Fantasies on a Poem by Thomas Campion*.

Audio Examples 1[12]

<http://www.paullansky.org/beingdigital/example1a.mp3>

<http://www.paullansky.org/beingdigital/example1b.mp3>

<http://www.paullansky.org/beingdigital/example1c.mp3>

The Fast Fourier Transform, (FFT) a quick digital shortcut for the Discrete Fourier Transform is another way to separate out various components of an audio signal. Working digitally in the frequency domain with windowed frames enables us isolate and arbitrarily modify frequencies, timbres and speeds. A rather brilliant use of this is demonstrated by the *Shapee* program of Christopher Penrose.[13] In this example the frequencies of an arrangement of the Brahms Lullaby are mapped onto the timbres of a vocal work by Perotin.

Audio Examples 2

<http://www.paullansky.org/beingdigital/example2a.mp3>

<http://www.paullansky.org/beingdigital/example2b.mp3>

<http://www.paullansky.org/beingdigital/example2c.mp3>

This technique is similar to convolution, in which two signals are multiplied by each other in the frequency domain, resulting in a merging and exchange of characteristics. There are an endless number of examples of convolution, and of transformations using frequency domain rather than time domain characteristics. The Phase Vocoder, developed at Bell Labs in the 1960's and subsequently by Mark Dolson at UCSD in the 1970's, is a further development of Homer Dudley's Channel Vocoder. Much of the work by composers such as Chris Penrose, Paul Koonce and others has been based on this technique. Unlike LPC, the Phase Vocoder does not attempt to understand the particular physical characteristics of a signal to begin with; it separates out frequency and amplitude characteristics according to FFT derived data rather than formant regions. Wavelet Transform analysis[14] is an approach that is more sensitive than a simple FFT in that it is attuned to the difference between short-term high frequency changes, and slower moving low frequency events while an FFT will view both in the same temporal domains.

Working in the frequency domain provides extraordinary power, but some time domain approaches are equally suggestive. An approach that I have found interesting is to map envelope and frequency characteristics from a source sound to another, artificially created sound. In my 1988 piece, *Smalltalk*[15], I mapped the amplitudes, rhythms and frequencies of casual conversation onto plucked

string filters. All that was necessary to do this was a frequency and envelope analysis of the original signal.

Audio Example 3

<http://www.paullansky.org/beingdigital/example3.mp3>

Some years later I made a few pieces that extended this technique. In *Now That You Mention It*[16] the same parameters are mapped onto piano sounds

Audio Example 4

<http://www.paullansky.org/beingdigital/example4.mp3>

and in *For the Moment*[17]

Audio Example 5

<http://www.paullansky.org/beingdigital/example5.mp3>

they are mapped in a different sense.

The power of software lies not only in the extent to which aspects and qualities of sounds can be teased apart, but also in the freedom it provides to arbitrarily patch amongst and between sound's dimensions. These mappings are useful concepts in fleshing out a virtual world in sound. Just as objects in virtual reality can have aspects that are familiar to us but behave in unexpected ways, so too in the audio domain the apparent products of physical actions, speaking, striking something, pulling a bow, can appear logical but act irrationally.

Every analog synthesizer had a random number (white noise) generator. The works of Morton Subotnick and others are rich testimony to the value and use of random control voltages. In the digital domain, however, random numbers are generally created by formulas that approximate white noise. There is generally not going to be much difference between the spectra of random signals in the analog or digital domains except in one respect: they are precisely repeatable in the digital domain. While this may at first seem contradictory it can be quite useful. When used in non-audio time, i.e. to control choices of notes, timbres, envelopes, etc. being able to reproduce results is extremely valuable and casts the use of random numbers in an entirely different light. A number of my pieces,

beginning with *Idle Chatter*[18] in 1984, use randomness as a generator of foreground detail.

Audio Example 6

<http://www.paullansky.org/beingdigital/example6.mp3>

I typically use ‘random selection without replacement’ to control the distribution and qualities of textures. In 1998 I composed a piece for sampled piano called *Heavy-set*[19] that used a model of the right hand of a pianist. The model worked by randomly selecting intervals to play, deciding when to play two notes instead of one, choosing directional changes, periodically, deciding to play some notes louder than others, and making various other decisions that were modeled on the kinds of thoughts an improvising pianist might have while playing. The system was thus capable of producing an infinite, or at least a very large number of pieces that were similar in harmonic terms (since this was not chosen randomly) but different in detail. The detail is the fine grain but nevertheless has a consequential effect, and thus a significant part of the compositional process involves choosing among randomly seeded synthesis runs. Here is an example of a passage that repeats a phrase that seemed especially effective.

Audio Example 7

<http://www.paullansky.org/beingdigital/example7.mp3>

These examples create something of a paradox for me. While I’m happy with the results, I also realize that there are probably an infinite number of possible versions of these pieces that will sound similar on the surface, but differ entirely in detail. With the power of computers today it is certainly possible to generate these pieces in real time, thus creating the potential for this infinitude of pieces, and also the possibility that things I prefer will be fleeting and disappear. I’m also faced with the realization that my ‘frozen’ versions of these pieces represent, in fact, just one among many possible instances of the music. Distribution of a work like this in software rather than audio form is presently feasible though not yet practical. It does appear, however, that the day is not far off when this will be practical. This will mean the emergence of a new form of recorded music, different on every hearing.

As listeners, we have long since ceased to be puzzled by an apparent contradiction in our listening worlds. Every recorded sound carries with it an

architectural acoustic signature. This is simply a series of reflected and diffused images of the signal that gives us clues about the space in which the sound was originally recorded, or in which the composers and audio engineers want us to think it was recorded. Indeed, it is rare that a recording ever is released that sounds as if it were recorded in an anechoic chamber. The function of reverberation is to project the image of a sound into a virtual environment that has palpable acoustic characteristics. But every space we use to listen has its own characteristics as well. Therefore we have to factor two environments in creating the audio image.

Artificial reverberation provides a penumbra that intensifies the virtual image of a signal. It helps us to imagine a location, a place, a physical space for it. While its use is frequently designed to smooth imperfections in a signal this smoothing process is the same one provided by a reverberant space. The components of digital reverberation are simple: delay lines, feedback loops, and filters, generally low-pass. While most software reverberators are designed with parameter controls based on the physics of real spaces, wet/dry, reverberation time, percentage of signal being reverberated, location of the source in a three dimensional space, and so on, the specifics of a reverberant space need not be tied to realistic parameters. In other words we can create a space where the apparent dimensions are constantly shifting, or where the resonant character of the space reinforces a particular harmonic texture, and so on. Imagine a room tuned to C major or an echo that rhythmically reinforced a rhythmic motive. In the following example, *Things She Noticed*, from my larger work, *Things She Carried*[20]. the reverberant characteristics of the speech change harmonically from utterance to utterance.

Audio Example 8

<http://www.paullansky.org/beingdigital/example8.mp3>

The lack of any reverberation at all in a recorded signal provides an interesting contrast, however. In most of the electronic works of Xenakis, for example, as well as much of the music of some current electronica artists such as Matmos and Autechre, there is virtually no reverberation at all. Many of their signals never pass through the air, in fact, and come to the loudspeakers with no acoustic signature at all. (This is partly because their sounds resist interpretation as the results of human physical activity.) In this case the world the speakers envelope is not a virtual world in the same sense as the one I have been describing. The

aura of physical reality lies only in the signals themselves, not in an imaginary space they describe. Loudspeakers have now become actual instruments. They are no longer vehicles of transduction as much as engines of creation.

Physical modeling is a relatively recent development in computer music, led by Perry Cook, Julius Smith and others. Rather than attempt to imitate the sounds of real instruments physical modeling uses a toolkit approach in which individual aspects of the physics of real instruments are hobbled together. To simulate a flute for example, (greatly simplified) a pressure wave is fed into a filter modeled on a cylindrical tube open at both ends, and the output is mixed with a returning wave. The model actually behaves like a real flute in that increasing the pressure results in overblowing and results in harmonics. Here is an excerpt from my work *Still Time*[21], in which flutes of various sizes combine.

Audio Example 9

<http://www.paullansky.org/beingdigital/example9.mp3>

Finally, a whole new area of work has recently emerged as a result of the arrival of processors capable of creating sound in real time with some degree of complexity. At Princeton, for example, Dan Trueman and Perry Cook[22] have been designing controllers that range from traditional forms like violin bows, digereedoods, and percussion instruments, to more unusual objects armed with arrays of sensors. These controllers generate sound by sending performance data to a computer which then interprets it and synthesizes and processes signals in real-time. This is another instance of arbitrary mapping but now it is directly connected with physical activity. To add another level of reality they use spherical speaker arrays so that the sound radiates in 360 degrees, as real sound does. Embedding significant processing power in the space between human motion and sound represents an important stage in realizing the potential of digital signal processing. While analog controllers have been around for years, the reconfigurable power of software in this instance completely changes the story.

I hope that I've demonstrated that the emergence of digital sound is really a watershed moment in the history of music. It enables us to harness technology in the service of musical adventure in ways that were unimaginable only twenty years ago. The computer is the ultimate instrument of the imagination.

Mild und Leise, whose painful creation I described at the outset is now, ironically, perhaps my most famous piece thanks to the English rock group Radiohead, who sampled part of it and used it as the harmonic backdrop for their song *Idioteque*, on their 2000 album *Kid A*. The irony is increased by the fact that it arose in the digital domain, made its way to LP, was sampled from the LP undoubtedly using digital tools, and included in a composition whose sound world is strongly reminiscent of early analog synthesizers (and in fact uses one for the simulated drum track). Here is the opening of my work

Audio Example 10

<http://www.paullansky.org/beingdigital/example10.mp3>

And the opening of *Idioteque*:[\[23\]](#)

Audio Example 11

(*Kid A*, *Idioteque* track 8)

[1] *Electronic Music Winners*, Columbia-Odyssey Y34149 (1975).

[2] Eric W. Rothenbuhler and John Durham Peters, "Defining Phonography: An Experiment in Theory", *Musical Quarterly*, Vol. 81/2, (1997) 242-264.

[3] Thanks to Kenneth Steiglitz for this analogy.

[4] Encryption may slow this down considerably, but it is a two-edged sword.

[5] In recent years 'plug-ins' for digital-audio editing programs have been created to simulate the sounds and artifacts of LP's and analog devices. Perhaps this is an attempt to reinsert the voice of another individual in the chain from the creator to the listener.

[6] While it is too soon to assess the longevity of digital storage media, the fact that data can be copied with perfect accuracy further separates digital from analog media. Anyone who has tried to play a twenty-five year old tape (there are some that have to be baked in an oven before they can be played only once), or revive an old LP is well aware of the extreme fragility of these storage formats.

[7] James Lastra, *Sound Technology and the American Cinema* (New York, Columbia University Press) 2000.

[8] Lastra, 151.

[9] Jazz, in many of its manifestations, may be the exception. The role of individual virtuosity in jazz performance would clearly be violated by overdubbing and too much editing.

[10] www.cs.queensu.ca/~bradbury/pdf/lpc_paper.pdf

[11] Kenneth Steiglitz and Paul Lansky, "Synthesis of Timbral Families by Warped Linear Prediction", *Computer Music Journal*, vol. 5/3, pp. 45-49, (1981).

- [12] All audio examples can be accessed from a single page at <http://www.paullansky.org/beingdigital.html>
- [13] web.sfc.keio.ac.jp/~penrose/shapee/
- [14] Corey Cheng's thesis at Dartmouth www.eecs.umich.edu/~coreyc/thesis/thesis_html/title.html#titlepage and a paper by Tzanetakis, Essel and Cook at Princeton soundlab.cs.princeton.edu/publications/2001_amta_aadwt.pdf) provide a good introduction to the musical applications of Wavelet Transforms.
- [15] *Smalltalk*, New Albion 030, 1990
- [16] *Conversation Pieces*, Bridge Records 9083, (1998).
- [17] Ibid
- [18] *More Than Idle Chatter*, Bridge Records 9050 (1994)
- [19] *Ride*, Bridge Records 9103 (2001)
- [20] *Things She Carried*, Bridge Records 9076 (1997)
- [21] *Fantasies and Tableaux*, CRI 683 (1994)
- [22] www.music.princeton.edu/~dan and www.cs.princeton.edu/~prc
- [23] *Kid A*, Capitol Records (2002)