from Companion to Contemporary Musical Thought
Vol. 1, J. Paynter, T. Howell, R. Orton &
P. Seymour, eds. London: Routledge, 1992.

20

Electroacoustic Music and the Soundscape: The Inner and Outer World

BARRY TRUAX

The relation of music and technology is becoming increasingly paradoxical. Within musical thought, computer technology allows the composer to conceive musical ideas from strikingly new perspectives. At the same time technology allows music to be increasingly embedded within the 'media environment' as a commodity with more exchange value than use value. Music frequently functions as a surrogate environment that interposes itself between us and the external world, imposing a relationship instead of expressing it. Music as an expression of the human spirit is locked in a struggle with music as the product of corporate control.

This chapter will examine the contemporary role of music first from the outside, in terms of music within the soundscape, and then from the inside, in terms of the soundscape within the music. In the first section I will trace how the traditional and complex relationship of the listener to the environment has been transformed by the social production of the electroacoustic listener as a consumer within a surrogate media environment. What originated as a complex relationship in the acoustic world reflecting a coherent unity has evolved into a merely complicated relationship based on inherent contradiction in the electroacoustic world. This evolution poses a dilemma for both the contemporary listener and composer, namely how to live with contradiction.

In the second section, I will explore the new soundscape that is evolving within the practice of music and technology. Most of this presentation will be based on my personal experience as a composer whose creative work is

intimately entwined with the development of software systems for composition and sound synthesis (Truax 1985). This experience has led me to conclude that what is significant about the composer's use of technology is its ability to change the process of compositional thinking, not merely its output. One example of this will be illustrated by my experience both with interactive computer music composition systems and with the domain of granular synthesis. This unusual method of sound construction links the micro-level world of sound 'grains' at the millisecond level with the macro level of perception and results in new compositional concepts relating sound and structure (Truax 1988).

The practice of electroacoustic music is producing, at least in some of its manifestations, a new kind of listener and a new soundscape, while simultaneously being equally concerned with producing consumer-orientated music and products. However, I see the more positive role of electroacoustic music as a search for unity and coherence within an electroacoustic environment based on inherent contradiction. Whether music can resolve such contradictions and lead us into a more balanced relation to our environment remains to be seen; it is this quest that will concern us here.

MUSIC WITHIN THE SOUNDSCAPE

Listening and Acoustic Communication: Complexity within the Acoustic Environment

From the contemporary perspective of ecology we can understand the traditional soundscape (i.e. the acoustic environment prior to industrialization and electrification) as a balanced ecosystem. Unfortunately it seems that we are only able to conceptualize that balanced relationship as an ecological one because we have since lost it. We are now confronted with such obvious incongruities in the soundscape that we are able to understand what it was we once took for granted. The necessity of the ecological concept springs from the context of loss, or at least from the present threat to survival. The question for us now is whether a new balance can be regained. Can we – with consciousness – be part of a new ecosystem?

The conscious awareness we seek can be encapsulated by the term 'design'. It comprises the models, the knowledge and the techniques that we understand and can put into practice. Murray Schafer characterizes soundscape design as a compositional challenge that involves everyone, not just composers (Schafer 1977). In my own view, design must also include contemporary technology, the same technology that is largely responsible for the imbalanced situation we now find ourselves living in. Technology is the external embodiment of our knowledge, and we are only hiding from our responsibility for its effects if we blame it for our problems. At this point in the helix, we must choose between trying to reverse direction – an impossibility in my view, given that consciousness cannot be reversed – and using our awareness and our technology to bring about that return but on a different, more conscious level.

Our concepts of design must be informed by an understanding of the past. That past cannot be regained, only re-created by a conscious restructuring of relationships. But the process of restructuring can benefit from a knowledge of how the soundscape functioned prior to industrialization and electrification, even if we cannot restore those conditions. Based on the foundation work of the World Soundscape Project (1973, 1977a, 1977b, 1978a, 1978b) at Simon Fraser University, Burnaby, Canada, in the early 1970s to document and study the acoustic environment, I have evolved a 'communicational model' to describe how soundscapes function (Truax 1984). Unlike the models of the traditional disciplines that study sound, which are based on concepts of energy transfer and signal processing, the communicational model is based on the concept of information exchange.

At the centre of the acoustic communicational model is the listener, because listening is the primary interface where information is exchanged between the individual and the environment. The auditory system may process incoming acoustic energy and create neural signals, but listening involves the higher cognitive levels that extract usable information and interpret its significance. Listening is also multi-levelled because it can involve various degrees of conscious attention. We normally think of listening as always attentive, and in fact this is the most analytical form that listening takes. However, most of the time, we process acoustic information more at a background level without attention being focused on it. This information provides the environmental context of our awareness, the ongoing and usually highly redundant 'ground' to our consciousness.

However, background listening is still a sophisticated cognitive process, involving feature detection, the recognition of patterns and their comparison to known patterns and environmental 'signatures'. Moreover, background listening can trigger conscious attention to be focused on an incoming sound when there is sufficient need or motivation from the listener. We have all

experienced the recognition of a particular sound, such as a voice, footsteps, or a car door closing, even when we are not listening for it. The sound attracts our attention either because it is unfamiliar ('who could that be?') or instead, precisely because it is familiar and because it has significance ('oh, I need to talk to that person'). I call this situation 'listening-in-readiness' because it involves both background and foreground listening strategies. It requires a favourable acoustic environment for information to be available (a good signal-to-noise ratio in technical terms), and an active cognitive processing of patterns and their comparison to known ones.

These different levels of listening involve analytical attention being paid to short-term details in the foreground case, and holistic or Gestalt pattern recognition in the background case. These two complementary strategies are often described as the respective provinces of the two hemispheres of the brain, and music is known to involve either or both such strategies, depending on the listening context and the listener's training or competence (Bever and Chiarello 1974). But I suspect that our traditional experience with environmental sounds predates the development of these cognitive abilities with music, both for the individual and the species as a whole, and therefore soundscape experience is fundamental to all forms of listening. However it is understood, listening is at the centre of the complex relationship between the individual and the environment.

I have suggested elsewhere that instead of listening being the end stage of a series of linear energy transfers from source to listener, we can understand it within a system of information exchange where sound mediates the relation of the listener to the environment (Truax 1984). In the traditional soundscape, that mediation entwines the listener in a unity with the environment. Hildegard Westerkamp has characterized the relationship as a balance between input and output, impression and expression, listening and sound-making (Westerkamp 1988). The information we take in as listeners is balanced by our own sound-making activities which, in turn, shape the environment. Acoustic ecology mirrors and complements the social and biological ecology in traditional societies.

At the conclusion of its study of five European villages, the World Soundscape Project postulated that the soundscape of each, to the extent that it reflected its traditional function, was characterized by variety, complexity, and balance (World Soundscape Project 1977b). Variety characterizes the amount of potential information in the environment, complexity the processing of that information by the listener which determines the listener's relationship to the environment, and balance the set of constraints that regulates and stabilizes the system. The result is the acoustic community, a communicational system in which sound plays a formative role. Music is an intense and specialized form of sound-making, both individual and social, through which that coherent relationship to the environment can be expressed.

Attali (1985) suggests that music reflects the need for order in the face of the violence of noise and affirms that 'society is possible'. He also points out that the control of music – and hence of all sound-making – is the concern of those in power, those responsible for social order. Music therefore reflects the prevailing economic order. Even more dramatically, Attali hypothesizes that changes in music-making may predict and precede changes in the economic order! Could the mediating role of sound mean that new forms of acoustic communication presage a new social order?

In addition to Attali's dialectical relationship of music and noise, there are other implications of noise on the acoustic ecosystem. Traditional acoustical engineering approaches have concentrated on the negative effects of noise, ranging from the physiological damage it can cause, to psychological effects, disruption of communication, and social trauma. Within the acoustic communicational model, we can understand noise as a destabilizing force. Instead of the traditional objective definitions of noise as an aperiodic waveform or a sound with high intensity level, or the subjective catch-all of noise as 'unwanted sound', we can focus on its role in disrupting or obscuring acoustic information. The simplest way in which it does this is by masking, that is, by making other sounds more difficult to hear.

Within our model of the acoustic community, noise is anything that upsets the balance of the system or its constraining forces. It is anything that reduces variety and increases redundancy, anything that obscures the clarity of a sound or reduces the definition of the acoustic environment. It is anything that simplifies or weakens the relationship of the listener to the environment, anything that reduces the desire to listen or the opportunity to make sounds. In practice what noise creates is an information-poor environment, one in which there is little desire to listen because there is so little that is meaningful to listen to. In the worst case, noise impairs the physical ability to hear and stresses the physiological functioning of the entire body. More commonly it simply reduces the meaningfulness of the aural experience and the sense of self and place.

Perhaps most significant of all is the way in which noise as the cause of degeneration of an acoustic community paves the way for profound changes, at both the individual and social level. Faced with the destruction of meaning from one perceptual source, the individual seeks a replacement in another, usually the visual. But the visual sense has developed to complement the

auditory, not to replace it. It objectifies and detaches, contrary to the auditory sense which involves and connects. By the end of the nineteenth century, the forces of industrialization and urbanization had profoundly changed the social and economic order of Western Europe and North America. They were also threatening to change the acoustic environment of most communities irrevocably, not only with the sounds of industry but also with those of motorized vehicles, both on the ground and in the air. Into this acoustically precarious situation a radically different type of sound phenomenon was about to be introduced – electroacoustic sound. Music and the soundscape would never again be the same.

Electroacoustic Communication: The Listener as Consumer within the Surrogate Environment

It has been claimed that the introduction of the phonetic alphabet was a major turning point in the history of communication. By allowing speech sound to be objectified and stored, it eroded the oral tradition and paved the way for literacy (Ong 1982). Although literacy has had a profound effect on human consciousness in that it allowed and encouraged inner speech, together with the development of abstract thinking and the concept of the self as an individual separate from the environment, it has not had a comparable effect on the soundscape. On the other hand, the electroacoustic revolution has had a profound effect on both. By allowing any sound to be objectified, stored, and reproduced in any other context, electroacoustic communication has not only eroded literacy but has also paved the way for a complete redefinition of what constitutes the acoustic environment and the listening experience.

Whereas industrialization changed both social conditions and the sounds that populated the soundscape – both of which threatened the balance of the acoustic community – electroacoustic communication provided a completely new element: contradiction. In the natural environment, contradictions cannot occur; a species in contradiction with its environment simply dies out, unless of course it can adapt. Within electroacoustic communication, contradiction is an inherent and inescapable fact.

The essence of the electroacoustic process is that it transforms acoustic sound, with all its physical characteristics and limitations, into a different energy form, namely an electrical signal which has very different characteristics (such as travelling at the speed of light) and none of the same limitations. In other words, the electroacoustic process breaks all the constraints of the physical acoustic world, but in so doing it creates all its contradictions.

In the acoustic world, sound is constrained by being tied to its context, in relation to which it derives at least part of its meaning. In the electroacoustic world, sound can be taken out of its original context and put via a loudspeaker into any other, where its meaning may be contradictory to that environment. The contradiction in any electroacoustically reproduced sound is that it exists in this time and space but it is from another; it exists in this context but is not necessarily appropriate to it. The sound comes from an inanimate source, the loudspeaker, but can we treat it as an inanimate environmental sound if we recognize it as speech or music? It is both, and the contradiction cannot be resolved. Schafer (1969) termed this phenomenon 'schizophonia' to indicate its 'aberrant' quality.

In the acoustic world, sound can only last a certain time before its energy dies away; it can only travel a certain distance, or have a certain loudness and, in general, certain timbral qualities. In the electroacoustic world, none of these constraints apply. With amplification and electronic sound synthesis, the behaviour and characteristics of sound are almost completely arbitrary. It can travel nearly any distance, be many places simultaneously, sound arbitrarily loud and be transformed in timbral character until it is unrecognizable.

One of the most significant contradictions in electroacoustic sound, in terms of its implications for the listening process, is the ability to repeat a sound exactly. Attali (1985) sums up the most essential characteristic of both monopoly capitalism's stockpiling of commodities and music in the twentieth century with the same word: 'repeating'. In the natural acoustic world, a repetition of a sound is never exactly the same, nor can a sound overlap itself; to do so would be a contradiction. Variety and difference are inherent in the acoustic process, and since information is based on the perception of difference (Bateson 1972), the acoustic world is full of information. Exact electroacoustic repetition, which is not only possible but commonplace, eliminates difference and hence reduces information. In the mixing studio it is a simple matter to repeat a sound exactly or combine it with itself up to an arbitrary density of layers. What kind of sense is the brain supposed to make of what previously had been a contradiction in acoustic terms?

When such 'contradictory' phenomena as the radio, telephone or gramophone record were first experienced, the wonderment at the 'magic' of technology in breaking normal constraints seems to have overpowered any uneasiness at the underlying contradiction. When we consider how poor

the quality of early forms of sound reproduction was, we may wonder why listeners were so convinced they were hearing the original. Was it simply because there was sufficient resemblance or could they not analyse the difference? Since then both scientific and aural analysis have been based on the electroacoustic reproduction and repetition of sound. As a result, our sense of 'sound quality' is conditioned by the fact that, thanks to electroacoustic technology, we have heard many poor as well as good examples of reproduced sound.

Today, frequent exposure to electroacoustic sound – and a survey of my students indicates that they hear in excess of fifty hours per week of reproduced sound (Truax 1984: 154) – has lessened the sense of magic and made the contradictions at least conventional if not actually banal. The same seems not to have occurred on such a wide scale with new forms of electroacoustic music unless the music imitates conventional instruments and musical styles. Known sounds that are reproduced become conventional fairly quickly; new electronic sounds seem strange when first introduced (e.g. tones replacing bells in telephones) but again, they too become quickly accepted, presumably because their meaning is already known. However, non-derivative electroacoustic music finds much slower acceptance because it requires the development of a new frame of reference, a new language, in short, a new system of communication.

What has been accepted with remarkable ease is electroacoustic sound as an environment in daily life. Whether it is the ubiquitous background music (rapidly becoming foreground music) in public places, the radio or television as accompaniment to daily activities in the home or car, or the Walkman and portable radios as an accompaniment attached to your person, electroacoustic sound, including a lot of music, is seldom far away. Already in the early 1960s, Mendelsohn (1964) found that among New York City radio listeners, the majority used radio as an accompaniment to daily activities, as a way to bracket the day and 'lubricate' social relations by having common experiences to talk about, as well as to influence mood and behaviour, a trait that the Muzak Corporation has been exploiting since the 1930s (Cardinell 1948).

I suggest that this use of electroacoustic sound as an everyday accompaniment creates a 'media environment', and I suspect that it functions as a surrogate environment for most people, one that seems richer and more meaningful than the alternative 'natural' one which may be noisy, lonely, boring, uninformative or alienating. Into these 'gaps' in life comes sound that will at least seem to fill the deficiency, the radio that will get you through the rush hour while you are trapped in traffic, the record player that masks the sounds of neighbours, the television that provides distraction for the tired and bored and 'company' for the lonely.

The problem with such surrogates is, first, that because they don't alleviate the real problem, they create a psychological dependency on its antidote, and secondly and more seriously, that most of these media environments are intimately connected to the commercial world. As a result, the listener becomes the recipient of countless advertising messages as well as the consumer of audio products and services. As part of the same process, music becomes a commodity like any other, with more exchange value than use value. On the broader political-economic level, American popular music also becomes the forerunner and most significant symbol of American products and values in a world increasingly dominated by corporate interests.

Interestingly enough, even if few composers seem to have learned anything from environmental sound experience, it appears that advertising designers have. Most of their audio advertising techniques, even if not consciously designed to work like environmental sound, certainly seem to act the same way. We have already noted the importance of background listening in the traditional soundscape where sounds common to an environment can be processed without conscious attention being focused on them. The frequent repetition of advertising works in much the same way, even when the listener is not actively attuned to it or the sound itself is low level. The frequent presence of background music in a particular environment builds up much the same pattern of association.

I call this type of listening where attention is focused elsewhere 'distracted listening' (Truax 1984). The problem for the advertiser is how to communicate to a person in this kind of situation. Factual information cannot be conveyed with any reliability of recall, hence advertisements seldom include such specifics. Instead, advertisers surround their product name with the desired image, supported by stereotypical music and voices, with attention 'hooked' by easily recognized tunes, motifs, slogans, rhymes and other 'catchy' items. These sounds and images strike what Schwartz (1973) calls the 'responsive chord', in that the listener supplies the appropriate meaning or association. After many repetitions, the product will be embedded in the listener's memory with a variety of images and emotions attached to it. When a need arises or when a decision to buy is being made, the sight or name of the product is sufficient to evoke the appropriate response. All of this occurs, indeed best occurs, without analytical attention or logical thought being applied (in fact, paid attention to, an ad usually seems absurd and the listener may become defensive about the message).

With the rise in environmental noise and the loss of community as a source of meaning, electroacoustic sound has fertile ground in which to operate as a surrogate. Once commoditized, music has become inextricably linked with this process and the multi-billion dollar industry that goes along with it. The individual's psychological dependence on the media environment and industry's continual need for new products and services guarantees that the electroacoustic relationship will be long-lived, even if imbalanced. Yet Attali and others claim that a crisis is developing, both economically and socially, and change is inevitable. He even suggests that certain developments in contemporary music foretell this change: 'Music, the ultimate form of production, gives voice to this new emergence, suggesting that we call it composition' (Attali 1985: 134).

From the Complex to the Complicated: A Crisis for the Composer and Society

The electroacoustic revolution has completely transformed the traditional soundscape and the role of the listener within it. Impressed by the scope of technological development, we are often tempted to say that we live in a 'complex' age. We assume that technological sophistication leads to progress and then are surprised when that technology creates more problems than it solves. The nuclear arms buildup and the problems of environmental pollution have convinced many that this 'progress' has brought us to the edge of destruction of the planet.

Instead of regarding our technological age as being complex, I think it is more appropriate to understand it as being complicated. The essential difference between the two descriptions is that, to me at least, the term complicated implies contradictions, a lack of coherence that complex systems exhibit. To be complicated, a system poses problems, perhaps unresolvable ones, and if so it is unstable. Such instability may lead to adaptive behaviour if the system can reorganize itself, but it can also lead to breakdown if it doesn't.

Technology poses problems, and electroacoustic technology, as I have shown, by being based on inherent contradiction, poses some quite difficult ones. One of the most difficult stems from the fact that we do not wish to give up what we perceive as the positive effects of technology, which we have come to depend on, even if they are accompanied by negative implications. In fact, the technological 'habit' of most people is such an integral part of their 'life-style' that any criticism of technology is often taken as a personal attack. It is difficult to stand back from such personal involvement to gain the larger perspective and perceive alternatives.

I have argued elsewhere (Truax 1984: ch. 13) that two key factors in any solution are the need to regain control and the need to experience alternatives. The crisis is similar for both the public and for the professional composer – the lack of personal control and the lack of alternatives. The two are closely related. The audio industry and the mass media, by gearing their products to the widest possible audience (or 'market' in their terms), ignore the needs of the individual. Two familiar problems for most electroacoustic composers are, first, the lack of access to the media and to other means of distribution, and second, the dependence on the companies that manufacture computer music instruments for the means of production. Despite the inherent flexibility of software and programmed control, the currently available computer music products exhibit a depressing degree of similarity and closed or 'black box' design.

The personal computer has been developed and marketed with a strongly democratic appeal; likewise, personal computer music systems have promised to bring the power of large-scale digital music systems into the realm of personal ownership. To a certain extent this has happened, but the new products (and one is always waiting for the 'next' one) invariably come with musical limitations and a 'closed system' mentality. The reasons seem to be two-fold, the first being the necessity of the mass market-place to base its products on the lowest common denominator of market appeal, and the second being the proprietary secrecy that stems from the competitive nature of the industry. The result seems to be that, with the worldwide spread of the same 'democratic' technology, everyone has the 'freedom' to sound the same!

The situation is particularly ironic with digitally based technology because of its inherent flexibility. Hardware costs and limitations are quickly ceasing to be the most constraining factors for users; ironically, it is the software that poses the greatest limitations because of the conventional musical models it implements. Software by its very nature is open-ended and flexible. It is notoriously difficult to standardize, to transport to other hardware configurations, to copyright or to profit from, but the industry has managed to do all of those things. The root of the problem is that if companies provided truly open-ended software, they would go out of business – they would no longer be needed. Individuals could become self-sufficient because they would have the means to evolve and develop their own ideas, instead of being the user of others'. For industry to prosper, it must reduce differences and perpetuate dissatisfaction.

On the other hand, I'm not sure that most individuals actually want that kind of open-endedness. It would mean taking a great deal of responsibility for one's own development. It would mean active involvement and participation, taking an initiative, constantly being in a state of learning, and worst of all, risking failure. We have not been conditioned by the market-place to accept this scenario readily. However, if we do not embrace this path, we will continue to live, with a greater or lesser degree of discomfort, with contradiction in a society that is increasingly vulnerable to breakdown. At this critical juncture does music itself provide a glimpse of any answers?

THE SOUNDSCAPE WITHIN THE MUSIC

We have seen in the first part of this chapter that music is increasingly a part of the everyday environment. In many cases it becomes the key feature of the soundscape, even if it is not meant to be listened to. It works its effects on us wherever we go and becomes associated in our memories with environments and products. In essence, it becomes the 'ambience' of the media environment. Once habituated to its omnipresence in public, we purchase it for private use in order to duplicate its effects or to signify our social role. The more serious stockpile the commodity in the hope that time will be found to listen to it in future. In the microcosm of our personal lives we replicate the macrocosm of the social and economic world.

It is interesting to consider the music that resists this incorporation into daily life, music that refuses to become ambient, music that demands involvement, music that we cannot passively consume but which requires us, even as listeners, to construct its meaning. The fact that such a music still exists gives us hope that music as a human phenomenon remains alive, that music has not become merely environment. For each person, the music that serves this function will be different. It is a personal discovery and may be found anywhere, in the conventional or in the experimental, in the wilderness or in the city, outside our culture or within it, with technology or without, through sound-making or listening in depth. Wherever it is found, it will involve thinking through sound.

For me the most vital form which that music takes is electroacoustic music, because it is in that medium that I am simultaneously composer, performer and listener. The mingling of what were traditionally these three separate functions gives this type of musical experience an intensity which is less present for me in other forms. I propose to examine the soundscape within this music in order to discover in the microcosm the characteristics of the new macrocosm that we seek.

Technology and the Compositional Process: Changing the Instruments of Thought

The practice of electroacoustic music leads us back to first principles. It is constructed from the ground up, as it were. Nothing is given, everything must be reconstructed and reconceptualized. As a result, everything must be done with knowledge and conscious awareness. The world this music creates is artificial but it is still a world, and to function properly that world must be coherent. Contradiction arises most often through the interface of the electroacoustic world to the real world.

One of the most fundamental characteristics of electroacoustic music is that the composer designs the sounds themselves, their organization into larger forms, and ultimately the space or environment which the composition creates (Emmerson 1986). That space is inherent at all levels of the composition, in the sounds and their organization, as well as in the performance or 'diffusion' of the sounds in an acoustic space. This latter aspect is often a source of contradiction – the space implied by the sound structure is inconsistent with the quality of reproduction and the acoustic characteristics of the space in which it is performed. The art of diffusion is to make the inner and outer spaces coherent and complementary.

Likewise, if electroacoustic sounds create space, they also create time. By comparison to the traditional view that sounds exist in time, contemporary ideas about cosmology suggest that the movement of objects creates our sense of time. Sound, possibly more so than visual phenomena, has always had the dynamic role of creating time, and music can be thought of as a particular means of organizing and shaping subjective time. The electroacoustic composer has the unprecedented ability to control time at the micro as well as the macro level.

Constructing the electroacoustic world requires knowledge, a great deal of knowledge. I have already described technology as embodying human knowledge, and the most powerful embodiment today is in the computer. It is one of the paradoxes of electroacoustic music that it would not exist without such external embodiment, hence its profound appeal to the inner world of imagination. I suspect that our age-old fascination with technology has a lot to do with that interplay of the inner and outer world that it represents. Technology as the embodiment of knowledge (whether in the

'soft' form of the ancient I Ching or the 'hard' form of the modern computer) is a mirror of human consciousness. We are simultaneously drawn to it and fearful of what we might find, and so it is not surprising that the computer today provokes such strongly ambivalent reactions.

The computer is a powerful tool for controlling complexity. It does so on the basis of a specific model of a given complex task and this model is implicit in the software used to address that task. Therein lie two critical questions: what is the source of the model, and what is the locus of control? Both questions challenge us in fundamental ways. The model reveals how we conceive of the task at hand. If we rely on traditional models of how music is constructed – for instance, according to instrumental music practice – then we shouldn't be surprised if the music that is produced is an extension and reflection of that practice. There may be some pragmatic value in using a machine to produce that kind of music, but all that it reflects back to us is what we already know. In many cases even that may not truly belong to us because it has only been inherited, not revealed as the result of our own search.

The issue of control similarly reflects our preconceptions, tinged in this case with psychological implications. We need the computer to control complexity, but we are afraid of losing control to it. However, the computer brings our need for control to a crisis point by showing that it leads to absurdity. Even if we decide to specify all events in a composition, the computer demands that we specify all parameters of those events. And then if we specify all parameters, the complete digital specification of the audio signal requires many thousands of sound samples per second – clearly an absurdity to control directly. The computer ultimately reduces to absurdity the traditional linear model that music originates with an all-powerful composer, passes through an obedient performer, and ends its journey as a communicated message to the passive listener. To escape the inherent contradiction of this model, the composer must rethink what the process of composition means with a computer.

What the computer does is to allow us to change the process of thinking with sound (Truax 1986). It becomes a new instrument for musical thought, not merely for new sounds. It provides a framework, an extendible one, for thinking about music. By virtue of its open-endedness, it allows the experience gained from its use to be channelled back into further software development. Whether the strategies used are deterministic or stochastic, whether the process is top-down or bottom-up, whether the details are determined by algorithms or by direct specification, the computer must be guided by the composer as to whether the results are musical. The concept of control by the composer gives way to the role of musical guidance, the composer as a musical strategist with the tacit knowledge to evaluate the musicality of the results of the process.

The result of this different approach to musical process is inevitably a new musical language and a new electroacoustic soundscape. Its characteristics implicitly reflect the new system of elements which have resulted in its birth. It has been evolved and negotiated by the composer, not imposed. It creates its own sense of space and time. Its imagery may have some basis in the real world but in general it seems as if a curtain has been drawn aside to reveal a different and completely coherent universe, to which the composer is but a conduit. The music is not a static entity, but within technical constraints may have a dynamic existence such that it changes from performance to performance and may even be adapted to a different medium. It creates an environment that requires the listener's involvement to establish its meaning, a meaning that is at once collective and personal. By contrast, the masterpiece tradition of the nineteenth century constructed an immutable musical edifice, solidly grounded on the foundation of tonality and to be visited in reverence in order to receive a reflection of the composer's inspiration. Today's successful electroacoustic work resembles more a participatory environment in which information is exchanged between listeners and the soundscape.

From Micro to Macro: Complexity and the World of Granular Synthesis

The digital technique that, in my experience, goes back the furthest to first principles is granular synthesis. It has been suggested as a computer music technique by Xenakis (1971) and Roads (1978) and is based on the production of a high density of small acoustic events called 'grains' that are less than 50 milliseconds in duration and typically in the range of 10–30 milliseconds. Typical grain densities range from several hundred to several thousand grains per second. Such high densities of events have previously made it difficult to work with the technique because of the large amount of calculation required, and therefore until recently few composers have experimented with it. Using a digital signal processor controlled by a microcomputer, I have been able to implement the technique with real-time synthesis and incorporate it within an interactive compositional environment (Truax 1988).

In taking us back to first principles, granular synthesis questions what

those principles really are. It challenges the traditional Fourier model of sound, which states that any periodic waveform may be reduced to the sum of a set of harmonically related sine waves each with a different amplitude. The first problem with the basic Fourier theorem is that it ignores time; it implies that the constituent sine waves have infinite duration. In order to deal with actual musical sounds, each cycle of the sound may be analysed to establish the time variance (or envelope) of each component harmonic (Mathews and Risset 1969). However, such sounds are not periodic as they begin, having instead an 'attack transient' or burst of noise as the instrument is set in motion. It is interesting to note that this short period of time as the sound begins, when the amount of acoustic variation is the greatest, is the most critical for the identification of the sound. It is also the part that contains the greatest amount of information. If that part of the sound is cut off, the rest is barely recognized as coming from the instrument in question.

The second problem with the Fourier model is that it corresponds poorly with the psychoacoustic reality of how our auditory system actually processes musical sound (Roederer 1975). Although we can perform Fourier analysis, i.e. distinguish separate harmonic components, we can do so only for the lower harmonics whose frequencies are sufficiently far apart, and only under conditions where the sound is prolonged such that we have time to focus on those harmonics. In other words, we have to simulate the unchanging conditions implied by the Fourier theorem for it to be applicable. In more realistic situations, with sound densities in music involving several notes per second in a melodic line, possibly combined with chords and various accompanying instruments, the auditory system uses much more efficient means to identify and track particular instruments.

In 1947 the British physicist Dennis Gabor proposed an 'acoustical quantum' as the fundamental unit of sound that incorporates both frequency and time because 'it is our most elementary experience that sound has a time pattern as well as a frequency pattern' (Gabor 1947: 591). In other words, the quantum is the shortest duration of sound that will activate the auditory system. It is an 'event', not merely a fixed stimulus. Although the techniques of granular synthesis depart from the ideal Gabor grains, they still involve the principle of building complex events from seemingly trivial micro-level, enveloped quanta called grains.

In fact, what is most remarkable about the technique is the relation between the triviality of the grain (heard alone it is the merest click or 'point' of sound) and the richness of the layered granular texture that results from their superimposition. The basic characteristics of the grain are preserved in the macro-level texture, and changes to the basic shape or content of the grain will alter the resulting texture. However, the macro-level is also more than the sum of its parts because of the way in which the grains are organized with respect to each other. For instance, if the component grains all have a similar frequency, the result is a musical pitch, but as the frequencies of the grains are spread over a larger range, the result becomes a broad-band noise, much like the natural textures of wind and water. These sounds create a remarkable sense of space and volume. Similarly, if the durations of the grains are all similar, the frequency spectrum and timbre of the resulting sound is much richer because of the phenomenon of amplitude modulation. In addition, when a short delay is placed between grains and then increased, the fused granular texture begins to 'pull apart' as isolated events emerge perceptually from the texture and eventually establish rhythmic relations. Once global amplitude contours are added, we have replicated all of the basic acoustic properties of sound.

One of the primary characteristics of granular synthesized sound is its dynamic quality. A static, unvarying sound is virtually impossible to achieve. It is the exception, not the rule. This quality is reminiscent of the natural soundscape in which everything is in a state of flux. It was not until the electroacoustic revolution that an entirely static sound was produced by the electronic oscillator. Predictably, such sounds fell on the ear only to produce boredom. The auditory system is oriented towards change and quickly habituates to redundant stimuli. The result of this habituation, a process called adaptation, is that the loudness of a steady tone decreases until theoretically, with infinite exposure, the sound disappears! We also react to steady, high-frequency tones with irritation. They seem piercing and unnatural. Nature has fortunately provided for transience in such sounds (such as birdsong) and a pattern of decreasing energy in the upper frequency range to make them palatable.

In terms of the issues of compositional control raised earlier, it is clear that granular synthesis requires a shift in thinking. It is obviously impossible for the composer to specify each individual grain, given that there may be thousands of them per second. Hierarchic levels of control are absolutely necessary, and at each level the composer simply specifies key control variables on the basis of which all of the specific data is calculated by the program. Interestingly enough, the structure of the software that realizes the technique is different from conventional programming approaches that tend to be linear and deterministic. The granular software, in contrast, more closely resembles an organic system in which parallel processes are working independently, linked by access to common variables representing 'the state of the world' and by messages passed between processes. The standard

debugging technique of stopping the program and stepping through its instructions one at a time is meaningless in this case; once stopped, the program loses its most essential characteristics, its time behaviour and the interaction between levels. In this kind of system, the user functions as the source of control messages that guide the overall process without directly determining it.

There are also many fascinating parallels between the world of granular synthesis and the subatomic level of physical matter. First, of course, is Gabor's idea of the quantum of sound, an indivisible unit of information from the psychoacoustic point of view, on the basis of which all macro-level phenomena are based. However, more specifically, Gabor postulated that the parameters of this quantum were frequency and time. The student of acoustics is often taught that these are two distinct and unrelated parameters, but of course, on a deeper level of understanding they prove to be reciprocal to each other, both physically and psychoacoustically. There is also an 'uncertainty principle' relating them, analogous to Heisenberg's uncertainty principle about particles at the subatomic level. He demonstrated that the more precisely you determine the position of an electron the less you know about its velocity, and vice versa. This is because velocity is the rate of change of position. Similarly, the smaller the time window in acoustics, the larger the uncertainty (i.e. bandwidth) in frequency. Frequency, as the rate of change of phase, is the reciprocal of time.

In another analogy to quantum physics, time is reversible at the quantum level. The quantum equations are symmetrical with respect to the direction of time, and, similarly, the quantum grain of sound is reversible with no change in perceptual quality. That is, if a granular synthesis texture is played backwards it will sound the same, just as if the direction of the individual grain is reversed (even if it is derived from natural sound) it sounds the same. This time invariance also permits a time shifting of sound derived from the real world, allowing it to be slowed down with no change in pitch. Grains are taken from the sound sequence at specific points but the rate at which these points advance through the sequence need not be at the original clock rate. When 'clock time' is stopped, the grains can move backwards through the sound with no difference in sound quality. All of the characteristics that establish the directional experience of time occur at the macro level. Granular synthesis is in a unique position to mediate between the micro and macro time worlds.

The basis of granular synthesis in the seemingly trivial grain has had a powerful effect on my own way of thinking about sound and music. It clearly juxtaposes the micro and macro levels, as the richness of the latter lies in stark contrast to the insignificance of the former. Moreover, the range of densities obtainable, from the low levels associated with human gestures through those perceived as rapid and virtuosic, culminating with entirely fused textures, suggests a compositional continuum ranging from human scale to abstract. Finally, in terms of sound and structure, it is clear that the two are inseparable with this technique. The macro-level structure can only be described in terms of its component sounds, and the resulting sound complex is definable only in terms of the structural levels that characterize its organization.

To date, I have composed three works using the granular synthesis technique. The first, *Riverrun* (1986), was based on synthesized grains (Truax 1987b). The fundamental paradox of granular synthesis – that the enormously rich and powerful textures it produces result from its being based on the most 'trivial' grains of sound – suggested to me a metaphoric relation to the river whose power is based on the accumulation of countless 'powerless' droplets of water. The opening section of the work portrays that accumulation, as individual 'droplets' of sound multiply gradually into a powerful broad-band texture. The dynamic variation implemented by changing variables continuously allows the piece to create a sound environment in which stasis and flux, solidity and movement coexist in a dynamic balance similar to a river, which is always moving yet seemingly permanent. The piece, I find, also captures some of the awe one feels in the presence of the overpowering force of such a body of water, whether in a perturbed or calm state, and as such it seems to create a different mode of listening than does conventional instrumental or electroacoustic music.

Two more recent works, *The Wings of Nike* (1987), for computer images by Theo Goldberg and two soundtracks, and *Tongues of Angels* (1988), for oboe d'amore, cor anglais and four soundtracks, are based on the granulation of sampled sound using very short fixed samples. In the case of the first work, these are male and female phonemes, and in the second piece the samples are derived from the live instruments involved in the performance. Despite the brevity of the source material, very rich textures and complex rhythmic patterns can be obtained from it. The pitch and timbre of the resulting sound are determined by the source material unless the grain duration is too short, in which case a broad-band spectrum results. However, the overlay of up to twenty simultaneous versions of such sound, each with its own variations, produces a 'magnification' of the original sound, as well as the possibility of gradual or rapid movement through its micro-level characteristics.

The degree of magnification involved can be appreciated when it is

realized that three of the four movements of *The Wings of Nike*, lasting approximately 12 minutes, were derived from only two phonemes, each about 150 milliseconds long! The stereo tape is a mixdown from an eight-track original which includes four stereo pairs of the granular material, and therefore the vertical densities of sound are around eighty at any moment, and the horizontal densities range from quite sparse through to 8,000 events per second at the very end. The sounds are heard at approximately their original pitch combined with versions an octave up or down, except towards the end of the first movement where a slow downward glissando is heard.

In the case of the male phoneme ('tuh'), timbral changes are possible as the grains move from the aspiration at the beginning of the sample, through the consonant to the pitch of the voiced vowel and finally past it. Microscopic timbral changes that normally go unnoticed become evident with the repetition of the overlapping grains. The analogy in the work is between the sampled visual image of the statue (the Winged Victory, or Nike) that forms the basis of the computer-generated images and the sampled vocal sound as the source of the tape. Each is based within the human dimension but the transformation techniques extend them towards the supra-human, as at the end of the first movement (Album).

In the second movement, the Scherzo, the statue is given a rather androgynous head, and the phonemes in the tape part are detached and realistic. However, each image accumulates again into more elaborate patterns that mask the source. In the third movement (The Illuminated Nike), frequency modulation grains are used. Just as the image of the statue undergoes colour variations in illumination, different types of modulation 'colour' the carrier frequencies of the tape sounds producing timbral variations. In the last movement, the Coda, both the visual and aural elements are progressively multiplied (the wing feathers and drapery in the visual case, and the syncopated phonemes in the musical accompaniment) until they fuse. A new kind of machine emerges that releases the earth-bound images into upward flight.

Speed, rhythmic patterns and density changes such as gradual accumulation or evaporation of the sound are other compositional strategies that granulation makes available. Speed and density, for instance, can change from isolated events, through to fast repetitions (each with slightly different timbral characteristics), and finally result in fused textures. Stochastic patterning of the grains can occur over the same range of densities as well, giving the sense that the sound is coalescing or evaporating.

In the interaction of the tape with the live performer in *Tongues of Angels*, this progression from speeds and patterns within the performance range of the instrumentalist to those achievable only by the computer creates much of the underlying dramatic tension of the work. We sense the struggle of the performer to match the virtuosity of the tape part and to transcend the limitations of human performance. The transcendence finally materializes at the end of the work when, after a gradual accelerando in both parts from syncopated rhythms through to fused textures and continuous gestures, there is a dramatic downward glissando on the tape. Despite the fact that all material was derived from the instruments heard live in the performance, the effect is suddenly that of a choir of voices, a transcendent image that inspired the biblically derived title of the work. Perhaps it could also be a metaphor for the technique of granular synthesis that seems to break into a new sonic domain.

I have described these compositions in some detail to give a sense of how dramatically this synthesis technique has changed both the musical thinking process and its results. The computer is an absolutely essential element in the system that has allowed this to happen. It is an interesting footnote to the granular synthesis story to consider why I should have been the first to implement it in real-time synthesis and hence to explore its potential with ease. It was technically feasible to do so about ten years earlier than my 1986 implementation, even though it is still technically impossible with most commercially produced MIDI-controlled synthesizers. I suspect that the main reason is the shift in basic concept which the technique demands. Technology always seems to be limited most by our own preconceptions and in turn it limits what is 'thinkable' with that technology. There is no precedent for the granular concept in instrumental music, hence a mind-block to its implementation. One wonders how many other techniques we have similarly excluded.

Resolving Contradiction: Towards a Re-integration of the Inner and Outer Soundscape

In the first part of this essay I argued that music in the outer soundscape has become devalued as a result of the electroacoustic revolution. Functioning primarily as a commodity and a surrogate environment, it is situated within a system based on complication and inherent contradiction. In the second part I argued that the inner soundscape of non-derivative electroacoustic music has resulted from the same revolution. That soundscape, as a new system of acoustic communication, is based on complexity and inner coherence. The task facing us now is to re-integrate the inner and outer soundscapes.

The search for complexity is an attempt to resolve contradictions and to seek unity and coherence, whether in the musical or social worlds. The technique to be used is that of design. I would like to think that composers could take a guiding role in that search because of their technical knowledge and because of their deep awareness of the aural experience. Perhaps the simplest place to begin is to see the act of composing as an analogy to and expression of that search, with the composition as its dramatization. The experience of such works could resonate within other individuals for whom the search is slightly below the surface of their awareness. That experience within the controlled environment of music could lead outwards to become applicable to, by analogy, the problems of the external world.

However, I do not see many composers being concerned with such matters. No matter how paradoxical the situation becomes, they seem to have little concern for the acoustic environment or the effects of media, preferring instead to complain about their victimization by these forces or about the lack of interest by the public in their esoteric concerns. From the ghetto of the new music community we hear a squabbling over crumbs that alternates with a mood of despair. Most indicative of all, we see an attempt to bring about a return to 'old values', to tonality and harmony, the principles that once worked and that presumably are the only ones that can still communicate to an audience that, in the meantime, has become accustomed to tradition as another commodity.

For composition to have any relevance to the problems I have described here, it must be based on a different world-view, a different educational process, and a different relation to technology. The purpose of all of these is integrative not exclusive, and it requires a change in our system of values. For the composer it may mean forgetting about the primacy of 'concert music', the most difficult shift of all in values. Our need for more concert music is surely less than our need for trained people to work in education, in media and technology, or in social and environmental areas involving sound. And yet we resist such departures from the field of 'pure' music because we have been taught to think of these other applications as second-class. Who has taught us to think that way?

In other words, I think we know what the answers are but we resist them psychologically. We need equality but we cling to the power and privilege to be found in hierarchy and inequality. According to the old ways of thinking, concert music is better than popular, instrumental better than electroacoustic, composing better than teaching, art more important than environment. Even if we resist those values, they are still embedded in our psyche as much as any other form of ideology. Unfortunately they link up with the much more untenable values of racism, sexism, patriarchy and imperialism that are similarly based on power and domination. These are the values that have brought society to the brink of collapse and they must be eradicated.

The alternative is subtle. By equality we do not mean the reduction of complexity found in liberalism's equation that everything is the same, that all ideas and cultures have equal value, and that music is the 'universal language'. This philosophy gives the appearance of a lack of prejudice while maintaining the status quo of power. It reduces difference instead of acknowledging it. The differences between cultures, ideas, and men and women are real and can never be eliminated. What must be eliminated is the notion that one is superior to the other. We can start by acknowledging that fact in music. Power and dominance was expressed in nineteenth-century music through the hierarchy of the tonal system, in which conflict had to be resolved by a reassertion of the tonic. Our music can reflect other models, such as the yin–yang concept of the complementarity of opposites, in which musical energy can derive from the interplay of difference, not its annihilation (Truax 1982, 1987a).

Attali characterizes his vision of the new social and economic order by the term 'composing'. He sees a decentralized economic system in which a network of individual and small group efforts replaces the monolithic, centralized forms of modern capitalism, the musical precursors of this change being the musics found on the fringes of the commercial mainstream. As a composer with a typical lack of understanding of abstract economic matters, I found his formulation surprisingly true to what I observe happening among musicians today, particularly the more technologically orientated. There is little reliance on government, institutions or business because of their inflexibility in dealing with individual needs. Instead, self-publication, self-promotion, home studios, and even software development in some cases, are becoming commonplace. Communication links are informal, though often global in scope, and sensibilities are highly regional, though often sophisticated in outlook. Large institutions are suspect and organizations tolerable only to the extent that they serve the network of individuals, not the accumulation of power. There is no centre; music has moved elsewhere.

Attali, however, underestimates the role of technology in this process, seeing it only as the expression of power. I have tried to show that music technology, and the computer in particular, while being the present servant

of power, has another, more subtle side to it, one leading to freedom. Using technology is by no means the only way to proceed and I am not sure that we really understand all of its implications yet. But I am convinced that it does show a path forward.

The inherent contradictions of electroacoustic technology cannot be resolved, but they can be transcended. The initial effects of technology are disruptive of traditional systems of acoustic communication, and the power that technology wields has been appropriated by dominant institutions. Our task now is to regain control of technology at the individual level and through conscious design to reintegrate it into our lives and our environment. Only then will it become a liberating force and an extension of our consciousness.

ACKNOWLEDGEMENT

This essay is based on a series of talks given at the conference, Musica/Complessità (Music/Complexity) held in August 1988 in the convent of SS. Annunziata near Amelia, Italy. I would like to thank the organizers of this conference, Michela Mollia and Walter Branchi, for the opportunity to synthesize the diverse aspects of my work in music, technology and the soundscape, as well as all of the conference participants for the interaction and enthusiasm that gave it life. My attendance at this conference was supported by the Department of External Affairs, Ottawa, and the equivalent agency in Italy to whom thanks are also due.

REFERENCES

ATTALI, J. (1985) *Noise: The Political Economy of Music*, Minneapolis: University of Minnesota Press.

BATESON, G. (1972) *Steps to an Ecology of Mind*, New York: Ballantine.

BEVER, T.G., and CHIARELLO, R.J. (1974) 'Cerebral dominance in musicians and nonmusicians', *Science* 185: 537–9.

CARDINELL, R.L. (1948) 'Music in industry', in D. Schullian and M. Schoen (eds), *Music and Medicine*, New York: Schuman.

EMMERSON, S. (1986) 'The relation of language to materials', in S. Emmerson (ed.) *The Language of Electroacoustic Music*, London: Macmillan.

GABOR, D. (1947) 'Acoustical quanta and the theory of hearing', *Nature* 159 (4044): 591–4.

MATHEWS, M. and RISSET, J.C. (1969) 'Analysis of musical-instrument tones', *Physics Today* 22 (2): 23–30.
MENDELSOHN, H. (1964) 'Listening to radio', in L.A. Dexter and D.M. White (eds) *People, Society and Mass Communication*, London: Macmillan.
ONG, W.J. (1982) *Orality and Literacy: The Technologizing of the Word*, London and New York: Methuen.
ROADS, C. (1978) 'Automated granular synthesis of sound', *Computer Music Journal* 2 (2): 61–2; reprinted in C. Roads and J. Strawn (eds) *Foundations of Computer Music* Cambridge, Mass.: MIT Press, 1985.
ROEDERER, J.G. (1975) *Introduction to the Physics and Psychophysics of Music*, 2nd edn, New York: Springer.
SCHAFER, R.M. (1969) *The New Soundscape*, Vienna: Universal Edition.
——(1977) *The Tuning of the World*, New York: Knopf.
SCHWARTZ, T. (1973) *The Responsive Chord*, Garden City, NY: Anchor.
TRUAX, B. (1982) 'Timbral construction in *Arras* as a stochastic process', *Computer Music Journal* 6 (3): 72–7.
——(1984) *Acoustic Communication*, Norwood, NJ: Ablex.
——(1985) 'The PODX system: Interactive compositional software for the DMX–1000', *Computer Music Journal* 9 (1): 29–38.
——(1986) 'Computer music language design and the composing process', in S. Emmerson (ed.) *The Language of Electroacoustic Music*, London: Macmillan.
——(1987a) '*Sequence of Earlier Heaven*: The record as a medium for the electroacoustic composer', *Leonardo* 20 (1): 25–8.
——(1987b) *Digital Soundscapes*, Cambridge Street Records, CSR-CD 8701, 4346 Cambridge Street, Burnaby, BC, Canada V5C 1H4, and Wergo Records WER 2017–50. Note: this Compact Disc includes *Riverrun*.
——(1988) 'Real-time granular synthesis with a digital signal processor', *Computer Music Journal* 12 (2): 14–26.
——(1990) 'Composing with real-time granular sound', *Perspectives of New Music* 28 (2): 120–34.
——(1991) 'Capturing musical knowledge in software systems', *Interface* 20 (3–4): 217–33.
——(1992) 'Musical creativity and complexity at the threshold of the 21st century', *Interface* 21 (1): 29–42.
WESTERKAMP, H. (1988) 'Listening and soundmaking: A study of music-as-environment', unpublished MA thesis, Simon Fraser University, Burnaby, BC, Canada.
WORLD SOUNDSCAPE PROJECT. The Music of the Environment Series, R.M. Schafer (ed.), Vancouver: A.R.C. Publications.
(1973) No. 1 *The Music of the Environment*.
(1977a) No. 3 *European Sound Diary*.
(1977b) No. 4 *Five Village Soundscapes*.
(1978a) No. 2 *The Vancouver Soundscape*.
(1978b) No. 5 *Handbook for Acoustic Ecology*.
XENAKIS, I. (1971) *Formalized Music*, Bloomington: Indiana University Press.

21

Composition with Machines

CURTIS ROADS

The earliest computer music activity – *computer-assisted composition* (see Hiller and Isaacson 1959) – is still misunderstood by the public. This essay traces some experiences with machines programmed to represent musical structure and execute compositional processes.

Computer-assisted composition implies the existence of software tools that aid the composer in creating and manipulating representations of music. In the past, computer-assisted composition referred to large composition programs written for batch processing computers. The composer would specify a set of initial data, and the program would generate an entire score according to the initial data and encoded musical rules.

Today, with interactive machines, the concept of computer-assisted composition has broadened. A composer can choose from a catalogue of programs that lend assistance with any number of musical tasks: from transcription systems that convert a keyboard performance into music notation, to sequence recorders that recall and manipulate a stored phrase, to interactive patch editors used to tune a digital synthesizer to digital sound mixing systems that are conducted via graphical scores. Furthermore, one can program the machine to assist with a specific compositional task. This labour can be made more efficient by a programming language that contains the appropriate musical constructs and interaction tools.

Underneath the languages and the graphics of the machine is a thick layer of intercommunicating automated procedures that manipulate rep-