20

ACOUSTIC SPACE, COMMUNITY, AND VIRTUAL SOUNDSCAPES

Barry Truax

Introduction: A Dual Processing Model

The human auditory system has a remarkable ability to perform what I call "dual processing"—that is, a process of simultaneously extracting two distinct types of information from a sound event. I am not referring to the common distinction of separating sound from its structural pattern of organization since the latter occurs over a much longer timeframe. Instead I am referring to the perception of a single sound event, or short gestural pattern, that yields complementary types of information. One of the most basic examples of such processing occurs when we identify both the nature of a sound source and the energy input that set it into vibration. The latter is called the excitation function in acoustics, and although the resonant properties of most objects are relatively fixed, different kinds of excitation can bring out differences in the perceived timbre of a sounding object. Early psychoacoustic experiments on the perception of everyday objects found that listeners could easily identify both some quality of the sound source, for instance, the hardness of a material, while at the same time obtaining information about the nature of the mallet that struck it or the process of excitation, such as bouncing or scraping (Freed 1990; Gaver 1993; Warren and Verbrugge 1984). Trevor Wishart (1996) has further identified situations where there is a single energy input, or repeated sequence of inputs, which results in the inherent qualities of the sound source being emphasized, compared with continuous energy input that results in a perceived gesture (or "imposed morphology" as Wishart puts it) whether that of, for example, the wind or some form of human or mechanical energy. This latter situation can be compared with the process of modulation, where information is encoded (and later decoded) into the pattern of change of a given parameter of a carrier, whether it is an audio signal, an electromagnetic wave, or the air stream passing through the vocal tract. Listeners clearly have a detailed knowledge of what kinds of gestures are associated with natural, human, and mechanical energy sources.

A second example of dual processing occurs with speech where we simultaneously extract semantic meaning from sequences of phonemes, while also being aware of the paralanguage with which the speech is delivered. We often describe the latter in terms of pitch inflections, timbre, dynamic changes in loudness, tempo and meter, patterns of stress, and most importantly, the use of silence—exactly those attributes which are used to describe a musical melody. Paralanguage is essentially how something is said, not what. The phonemes themselves may be regarded as digital units that combine to form words (but with smooth transitions between them) whereas the paralinguistic aspects are analog in nature, subject to continuous variation. It is interesting to note that paralanguage tends to remain identifiable even when the speech itself is difficult to understand, or distorted in some fashion. We readily interpret the paralinguistic aspects of speech as reflecting the mood, intent, and sincerity of the speaker, for instance, as well as the relationship between speaker and listener. However, such interpretations are highly culturally specific, and it may be dangerous to take them over into a cross-cultural situation where their meaning may be quite different, even if certain basic elements are similar.

A third example brings into focus the impact of audio technology on the listener over the past century as it deals with the simultaneous perception of the content of a reproduced sound along with the quality of that reproduction. The famous Edison "tone tests" during the early mechanical reproduction of music (Thompson 2002) demonstrated the audience's inability to "tell the difference" between the live voice of an opera singer (one contracted with Edison, of course) and its reproduction on disc, despite the obvious limitation of frequency bandwidth at that time. I argue that, since listeners had never before heard what we regard as good or bad reproduction, it was sufficient for them to merely identify the singer and the music. In other words, they had no auditory competence based on analytic listening to judge reproduction quality separately. Not surprisingly, the nascent audio industry seized on this concept of "fidelity" as a selling point to educate these new consumers and convince them that every technical advance was worth paying for, a strategy that has extended up to, but not including, today where compressed audio in the MP3 format has promoted sound reproduction of a lesser quality than what could theoretically be made available (Sterne 2012).

As interesting as these examples may be, this chapter chooses to focus on one of the most subtle but pervasive examples of dual processing by the auditory cognition system, that is, the simultaneous perception of a sound event and the acoustical space in which it is produced. The two are so intertwined that we often ignore the influence of the physical space on the sounds we pay attention to, unless it has a unique or pervasive acoustic character because of symmetrical reflections (e.g., whispering gallery, parabolic reflectors, or canyon effect), exaggerated amounts of reverberation, or the unique situation of the anechoic chamber where the very absence of acoustic reflections creates a disorienting sensation for most listeners.

The modern science of acoustics over the last century has broadly treated the spatial aspect of sound in two contexts: propagation in a free field, and the behavior of sound fields in enclosed spaces, the latter being the basis of architectural acoustics. This work has resulted in a significant body of theoretical and applied literature, including many approaches to the complex problem of modeling the acoustical properties of actual and proposed spaces (Blesser and Salter 2007: ch. 6 and 7; Vörlander 2008). Although the acoustic complexity of real spaces may exhibit subtleties that require further research, the general principles involved seem well established. However, the perception of acoustic space—how we interpret sound as creating a sense of space—is not as well understood. Perhaps the greatest impediment is our reliance on visual models of physical space that are relatively stable and detailed, giving us the impression that space is a fixed entity through which we can move. The practice of architectural design is similarly characterized by an emphasis on the visual aspects of space, with few schools until recently giving any thought to the acoustic aspects of design.

How does the auditory perception of space differ from its visual counterpart? And how are the two related? The most fundamental difference is that the auditory perception of space depends entirely on time, meaning that it is in a constant state of flux. I will argue here that the time domain is central to two related aspects of auditory space—the space or "volume" within a sound (i.e., its perceived magnitude), and the sense of space created by all of the sounds within a soundscape (Truax 1998), what Blesser and Salter (2007) call "aural architecture." Clearly I am putting the emphasis on the human perception of auditory space as to how we interpret acoustic cues, which therefore is the domain of psychoacoustics. However, my goal is broader than that, because I will argue that the perception of acoustic space, and our perceived orientation within it, is a central concern of acoustic ecology, an emerging field of study whose main concern is the relation of the individual to an environment as created by sound, and by extension, the relationship between a community and its soundscapes.

Time, Volume, and Space

Modern architectural theory often suggests that what we build is not simply placed "in" Cartesian space (which is assumed to be uniform in all directions), but rather that what we design and build "creates" space. Similarly, I am suggesting that sound *creates* auditory space and that the sounds we hear create perceived volumes within that space, which even in some cases of total immersion become the space itself. How is this concept related to traditional acoustics? The answer can be found in the most basic acoustical concepts related to motion and time, namely frequency and the speed of sound (Truax 1999). Laypersons often confuse these two concepts, frequency and speed, because they both refer to the temporal behavior of vibratory motion. In classical acoustics, they are related by the concept of wavelength, at least for simple harmonic motion, as illustrated by the equation:

$$f = c/\lambda$$

where f is the frequency of vibration in cycles per second or Hertz, c is the speed of sound in feet or meters per second, and λ is the wavelength of the vibration in feet or meters. Given that the speed of sound is constant for a given medium at a certain temperature, frequency is inversely related to wavelength, with high frequencies having short wavelengths and low frequencies having long ones. Frequency can be thought of at the micro level as the rate of change of phase of a vibration, whereas the speed of sound is its rate of propagation through the medium, which for air is relatively slow, at least compared to light, being around 1,100 ft/sec or 330 m/sec. A useful rule of thumb is that sound travels about a foot in a millisecond, keeping in mind that all frequencies travel at the same speed, meaning that complex vibrational patterns travel coherently from source to destination and create an analogous vibration at the eardrum.

What does this basic equation from acoustics have to do with acoustic space? Quite simply, it is the complex pattern of simultaneous vibrations or frequencies within a sound that creates its sense of perceived volume or internal acoustic space. We hear the complex resonances of a steel ship's hull being struck as coming from a much larger object than a small wooden boat, for instance; its perceived magnitude is clearly greater, even if its overall intensity level is the same. On the other hand, it is the brain's ability to detect the small time differences caused by a sound reaching our ears by different paths, such as those caused by reflections, that creates a sense of an external acoustic space through reverberation (Truax 1998). In contrast, the speed of light being extremely fast compared to sound means that the light we perceive coming from all objects in our world (but not from the stars!) seems to arrive instantaneously at our eyes.

Closely linked to our sense of perceived volume of a sound, is our ability to trade off size with distance. Our visual ability in this respect is well known when we have an experiential reference. Knowing the usual size of a human, we assume that a person who appears smaller must be more distant. A photograph or painting that includes such familiar forms works similarly as long as accurate perspective is maintained, undistorted by a lens or by a painter who, for instance, uses foreshortening of distance. On the other hand, a more abstract texture in a photograph, such as in an extreme close-up, might be mistaken for an aerial photograph of a desert landscape. Similarly with sound, we perceive the source of a sound to have a certain size or volume, and if its loudness decreases, we assume it is more distant, not giving off less energy unless its spectrum is weaker.

In discussing the speed of sound propagation, I noted that all frequencies travel at the same speed, hence vibrations are transmitted coherently with no disruption in phase (although if the source is a loudspeaker tweeter and woofer, this coherence can no longer be taken for granted). However, every interaction of the sound wave with the medium of transfer, the distance of transfer, and most importantly, with its encounters with solid obstacles, or within an enclosed space, changes the relative strength of the various frequencies involved because of absorption and resonance. A sound outdoors in a relatively open space will not have any low frequency boost that the same sound, such as a voice, would have in an enclosed room. Therefore, when we hear that voice over a telephone, which doesn't transmit frequencies lower than about 300 Hz, the voice sounds more distant. Therefore, what we arrive at is an intertwining of sound with physical space. Every sound we hear carries information about the vibrational pattern of the source *and* the physical space through which it has traveled. It is one of the most amazing abilities of the brain that it can decipher both kinds of information simultaneously.

A single reflection of a sound wave can produce an echo if the returning sound doesn't fuse with the original and occurs around 100 ms after it (50 ms is the theoretical limit of the auditory detection of echoes, but architectural acoustics accepts any early reflection arriving within 80 ms as reinforcing the original sound and not contributing to reverberation). Multiple reflections, such as in an enclosed or semi-enclosed space, create reverberation, the accurate simulation of which continues to challenge those designing digital signal processors (Blesser and Salter 2007). Not only are the reflections numerous temporally, but they are also complex in the frequency domain, the totality of which might be called the acoustic "signature" of the space. If we record a sudden, short broadband sound in a space including the resulting reverberation (what is called the "impulse response"), we can simulate any other sound being perceived to be in that space by the process called convolution (Roads 1996). When we convolve the given sound (preferably recorded in a dry space with little environmental coloration) with the impulse response of a space, the result is that our given sound appears to be located in that space, because it is colored with the space's frequency response and reverberant decay. What convolution does is to multiply the two spectra (or frequency content) together such that any frequency that is strong in both is very strong in the output, and conversely, any frequencies that are weak are strongly attenuated. Again, sound and space are linked.

An interesting extension of impulse convolution is called auto-convolution where a sound is convolved with itself, thereby emphasizing its strongest frequencies, removing weak frequencies and doubling the length of the sound. The auto-convolved sound can then be convolved with the impulse response of a space, as in my work *Temple* (2002), in which case the sound hovers somewhere between the original sound and its reverberated character, creating a fusion of the internal space within the sound with its external acoustic space (Truax 2012).

The Perception of Acoustic Space

What this close connection of sound and space means is that sonic events in time are required for us to hear acoustic space, whereas we imagine an architectural space to be independent of who or what is present within it, though its perception clearly depends on the lighting conditions. For a blind person, then, the cessation of movement or activity means that that aspect of the world "disappears." The auditory world is entirely dynamic and can never be static. Sound requires motion (within a certain range of audible frequencies), audible sound is the result of that motion, and that motion creates space when perceived. Tim Ingold (2007: 11) perceptively suggests that sound "is not the object but the medium of our perception. It is what we hear *in*."

What psychoacousticians are still investigating is our auditory ability to sort out the complex vibration that arrives at each ear, which is the result of several sources of vibration being added together. Somehow, based on that analysis, we perceive a coherent "auditory scene" populated with identifiable entities at various distances (Bregman 1990). One strategy involved in this process is binaural localization, which refers to our ability to detect the direction of a source. Differences in time of arrival at the two ears for low frequencies, and intensity differences for high frequencies, are the primary cues. However, even subtler cues are present that distinguish when sounds are in front, as opposed to coming from behind, and when they are higher or lower than ear level. These cues result in a subtle coloring of the sound in the upper frequencies by the external ear flaps, or pinnae. Ridges on the pinna create small reflections of an incoming sound which, when combined with the direct version, results in an attenuation of certain high frequencies and a slight boost to others.

A major component of auditory scene analysis—sorting out complex vibrations into separate sources—is the detection of coherent patterns of those sources. A voice coming from a certain direction will have a pattern that is different from one coming from a different direction. The ability to follow one or the other voice at will, or to switch attention rapidly between them, is called "cocktail party effect" (Truax 1999). It involves the brain's ability to suppress one pattern while enhancing another, and assumes that the coherent vibrations that emanate from the same direction are probably from the same source. On the other hand, later arriving vibrations, such as the reverberant tail of a sound, are random, uncorrelated vibrations that are interpreted as indicative of the space where the voices are located. Or to use the language of this chapter, the uncorrelated sound creates the sense of acoustic space, within which correlated patterns with their own sense of volume and distance, are interpreted as sources (Truax 1998).

The auditory sense of sources and space is usually confirmed by visual cues, but not necessarily. The World Soundscape Project (WSP) was founded by R. Murray Schafer at Simon Fraser University in the early 1970s to study acoustic environments, and in one of its recordings from a small town in Italy in 1975, there are three distinct sources that can be easily discerned, each creating their own sense of acoustic space. In the foreground of the piazza are men talking, their voices brightened by the surrounding, reflective surfaces which also adds a degree of reverberation suggesting the size and shape of the physical space. Simultaneously, an unseen choir is heard coming inside a church facing the square, its muffled sound indicating both its distance and the sense of being heard through the walls. Also simultaneously occurring are the bells of another church on the other side of the village and not visible in the square, providing a sense of a distant horizon to the complex acoustic space of the recording. Both the physical space and the social space of the community are revealed in the acoustic space of that soundscape.

Listening to any recording, of course, requires us to try to identify the soundscape without a visual reference. However, even with a good quality stereo microphone (or multiple mikes), the soundscape has been subtly distorted, similar to how a camera lens re-presents a scene. In most cases, the auditory space is flattened out, just as photographs usually flatten out distant objects in a scene. Sounds coming from behind the recordist may be repositioned in front by the listener, since the mikes do not have the binaural colorations expected by the ear. Binaural or "kunstkopf" recording provide a more vivid sense of an acoustic space, but have to be listened to on headphones for the space to be externally localized. But, even if subtle or less subtle distortions of the actual soundscape are present in the recorded versions, the auditory system can usually produce a reasonable mental representation of the original space.

Acoustic Ecology and the Design of Acoustic Space

If acoustic ecology is concerned with the relationship of the individual listener and communities of listeners to their environment as mediated by sound, then the individual and collective perception of acoustic space must play a fundamental role. Perhaps the most basic role is that of orientation. The habitual sounds we experience daily both reflect and confirm our sense of physical space, as well as our place within it. Individuals and communities have a definite sense of "what belongs" in their acoustic space, and what kinds of noise are "invasions" of that space. The World Soundscape Project has referred to such intrusions as "sound pollution" as distinct from noise pollution which is generally defined by degrees of harmfulness and risk. In other words, familiar sounds and their temporal patterns define and characterize our sense of place. Even subtle changes to the habitual pattern (which we usually take for granted) may be noted; for example, "it seems too quiet here today" when we sense that something is missing, or the opposite, "something special must be going on." The characteristic ambience of a given space adds to the "feel" or "atmosphere" of it, even if we would be hard pressed to define what contributes to that character. Often it is what the WSP calls the "keynote" sounds—those that are in the background of our perception but typify a space the most. Foreground sonic events, or "signals," may provide specific information that we know how to interpret, even if fleetingly, and culturally important sounds recognizable to all in a community can be termed "soundmarks" (Truax 1999).

I have suggested that an information-rich, balanced soundscape contributes to the sense of an acoustic community, one where sound plays a formative role in the definition and life of a group of people, no matter how their commonality is defined (Truax 2001). Such an acoustically defined community will likely exhibit a large variety of sounds, many of which are interpreted by locals with a complexity of contextual information, and the resulting layering of sounds is balanced by a variety of spatial, temporal, and social forces that make it functional. Sound will also define what is the boundary of the community, whether the scale is small or large, by distinguishing between what is "local" from what comes from the "outside." In one study of an acoustically defined Vancouver neighborhood that is bisected by a busy thoroughfare, some locals referred to those passing through as "the others."

In other words, traffic moving in one set of directions was regarded as "other," and that moving in a different direction was "local"—and in fact, the latter was characterized by a greater pedestrian component along with slower moving cars. The two sonic components of the soundscape thus created a mental map to the locals that reflected these intersecting spaces, and in fact most of the people interviewed could draw a version of such a map (Paquette 2004).

The "enemy" of the acoustic community is not so much noise per se, but rather any element that lessens the clarity and definition of an acoustic space, or dulls people's inclination to listen. In other words, the acoustic community depends on information exchange, and anything or any habit that detracts from or inhibits that exchange weakens the sense of community. Bland, uniform sounds that lack character or are not perceived to be on a human scale might be the most obvious culprits, such as broadband noise from ventilators or machinery or excess amounts of traffic. Although such sounds may be acoustically complex in some sense, they are usually perceived as lacking in information or character, although as already noted they may come to be recognized as "keynotes" in the community. Even worse, they frequently mask other, more individualistic sounds, thereby reducing variety. In the language of ecology, a few dominant species with little diversity crowd out numerous smaller species that are able to co-exist. Just as the loss of genetic diversity is a problem, so is the loss of aural complexity.

Besides the effects of orientation and the communication of information, an acoustic space can also encourage various types of interaction. An early study of the soundscape of Boston termed the positive character of such interactions as "responsive spaces" (Southworth 1969). In other words, the fundamental acoustic principles of reflection, resonance, and absorption, all of which contribute to the sense of acoustic space, are the main variables which can be designed to promote (or deter) human interaction. The details and variety of approaches to this topic are beyond the scope of this chapter, but perhaps a brief look at the extremes of the continuum will clarify its nature. At one end we have the "free field" where there is little or no reflection because of the lack of any barriers to reflect the sound (though in real situations there is always the ground). That extreme end is the anechoic chamber where absorption is maximized and reflection minimized, and usually this type of acoustic space is disorienting to the individual because there is no interaction, no feedback, and essentially no acoustic space. The other end of the continuum is the "diffuse sound field" which maximizes reflection (or resonance if the space is smaller) and minimizes absorption. A marble-lined space, an indoor swimming pool with highly reflective glass and water, or a gymnasium with polished floors and high ceiling are common examples. Sound comes from everywhere and nowhere; the acoustic space is omnidirectional and often as equally disorienting as the anechoic room, except for the opposite reason. If one did not have to act or communicate in such a space, one could enjoy the womb-like envelopment, but otherwise the eyes have to be alert for orientation, verbal communication is almost impossible, and noise levels tend to become exaggerated. As noted earlier, the sound *is* the space, and vice versa. In between these two extremes lie the truly interactive acoustic spaces where reverberance and envelopment are balanced with the needed sense of clarity and definition.

Vancouver's spectacular natural setting and its dramatic layout of modern buildings, particularly around the harbor area, provide strong visual imagery for the city, one that is used to attract tourists. As documented on the *Soundscape Vancouver 1996* CD (WSP 1999), many of these visually striking environments are accompanied by bland, technologically derived soundscapes. Sound examples from the CD include the Seabus crossing the harbor, the noisy exhaust fans from the architecturally striking Canada Place, and the bland drones and hums from Arthur Erickson's otherwise dramatic Museum of Anthropology with its marvelous collection of West Coast artifacts. One wonders whether in these environments, the eyes take over and cause the ears to ignore what is accompanying these visual splendors. Fortunately, there are also many examples of planned urban re-development in the city that are designed on a more human scale, not unlike the village model which the WSP encountered in its European study (Truax 2001). Granville Island, Gastown, many parts of the West End, Commercial Drive, and some other neighborhood-based town centers in Vancouver are examples of this approach which have proved popular with the public. In each of these, acoustic space is controlled, at least to some extent, and populated by a wide variety of sounds that function on a human scale, whether made by humans or other sources. Such information rich environments seem to create a positive model of acoustic ecology.

Multi-channel Diffusion and the Soundscape Composition

In an age that seems intoxicated by "virtual reality," we often assume that these artificial visual illusions of space are the only ones, particularly as they acquire increasing degrees of realism. Even if given less public profile, multi-channel and multi-speaker re-creations of acoustic space are just as impressive, and more easily achieve the effect of total immersion, since it is relatively easy to surround an audience with arrays of loudspeakers. Our work at Simon Fraser University over the last decade has shown that this type of aural representation is particularly effective for creating immersive acoustic environments, through what we call soundscape composition (Truax 2002, 2008), including both those which reference actual spaces, and those, such as my *Chalice Well* (2009), that create entirely imaginary ones. Frederico Macdeo (2015) provides a survey of the different roles of acoustic and electroacoustic space in contemporary music and sound art.

The multi-channel approach is an extension of earlier arrays or "orchestras" of loudspeakers where a stereo track was sent to an arbitrary number of speakers with dynamic changes controlled by a composer/performer, usually centrally located. This technique is called "diffusion," a term drawn from acoustics where it refers to the spread of sound in a space. By emphasizing a sound coming from a particular speaker one could create the illusion that the speaker location was the momentary source of the sound, but in general only one, or at most two sounds could be localized at a time.

One of the earliest multi-channel installations occurred at the Brussels World's Fair in 1958, where Edgard Varèse's multi-track work, *Poème Eléctronique*, was projected through 425 loudspeakers attached to the curvilinear walls of the Corbusier designed Philips Pavilion. Early four-channel formats (quadraphonic sound) doubled the number of possible sources, but could only create a coherent sense of space in a relatively small room because the distances between the speakers left gaps in the spatial illusion unless a lot of reverberation was added. Today, the 8-channel configuration works best for medium-sized rooms, as long as the material on each channel is kept uncorrelated, that is, as independent sources such as is the norm in the acoustic world. The spatial layout of these speakers can vary, but the choices are generally circular, equally spaced around the audience, or more clustered in front of the audience, given that our ability to localize is better in front than behind.

Larger, well-equipped halls have extended this principle to even larger numbers of channels and speakers to which independent sources or tracks can be sent. The ZKM at Karlsruhe, for instance, has a rig of 40 speakers in a formation called a *Klangdom*, elevated above and around the audience. The Sonic Arts Research Centre in Belfast has an amazing array of up to 32 channels, two sets of 8 that are suspended at varying heights, another set of 8 around the audience, and another set of 8 beneath the audience but audible through the grid flooring (Figure 20.1). The acoustic panels on the walls are also variable to add or omit reflecting surfaces. These arrays, both at the listener's ear level, and those that incorporate height and depth, are excellent for creating a vivid sense of acoustic space that is totally immersive. With the flexibility and precision of digital control, the composer can literally design a detailed acoustic space, and move the listener through it. It is not an exaggeration to suggest that this approach creates a three-dimensional "aural architecture."

In my opinion, the key to designing such a space is to treat the loudspeaker as a point source, and avoid the illusion of what are called "phantom images" that appear between the speakers but collapse when the listener is not placed exactly between them. Just as we can distinguish multiple sound sources in a soundscape (assuming their levels are balanced), so too can we hear the definition of multiple speakers emitting different (i.e., uncorrelated) signals. When some of these channels incorporate a similar sense of reverberation, ambience, or other environmental cues, then those speakers will connect to form an ordered sense of acoustic space. Strategies exist for moving a sound smoothly between speakers (or not), hence adding the possibility of moving sound sources, and/or apparent movement of the listener through different acoustic spaces. Simultaneous "streams" of sound images can

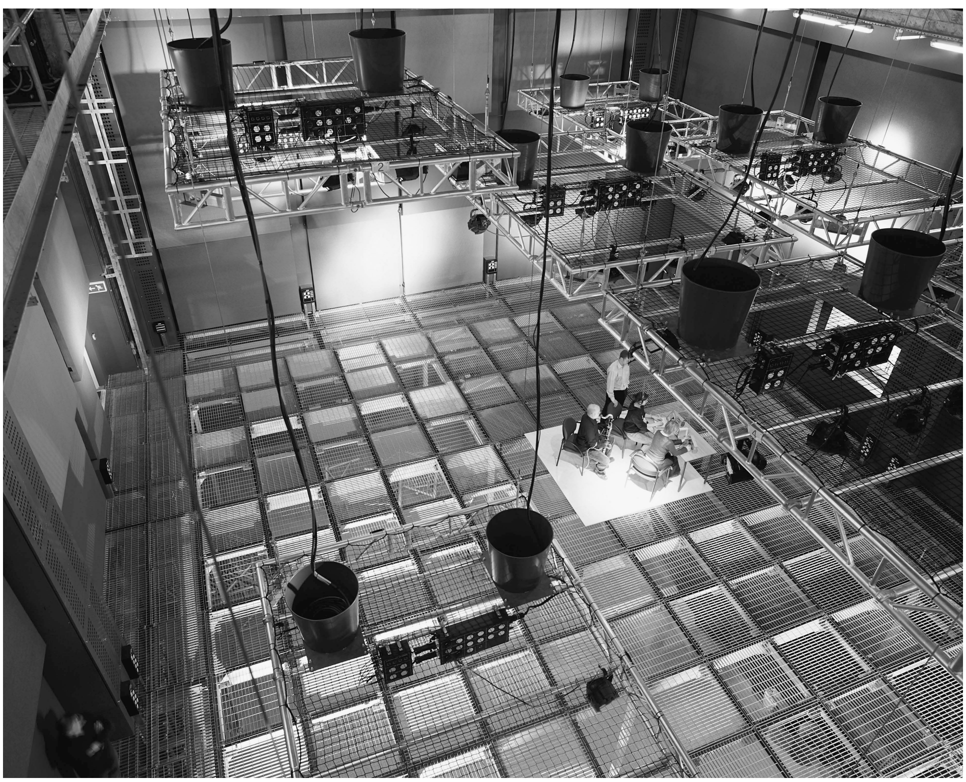

Figure 20.1 The multi-channel theater space at the Sonic Arts Research Centre, Belfast.
Photo: Hall Black Douglas.

also be created, though it is unclear as to how many a listener might optimally follow. The artistic potential of such immersive audio environments is just beginning to be understood and put into practice.

Conclusion

In this chapter, I have tried to give an overview of a concept of space that is not tied to the visual domain, but rather is created by aural experience. Although the acoustic and psychoacoustic principles on which it is based are mostly well known, our understanding of how humans create their sense of acoustic space based on just two binaural inputs is still fragmentary, and until recently based mainly on speech and music perception, not environmental experience in general where the variables are far more complex. The "architecture" of such spaces can easily be related to the concerns of acoustic ecology and acoustic design. Many involved in the field would say that the design concerns today are increasingly pressing as the forces of technology and urbanization progress. Until recently, one role of music has been to fill existing spaces, designed or otherwise, both to inspire and in the case of "music as environment" (what used to be called "background music") to manipulate those not listening to it. I have suggested here that soundscape composition, as both a musical and communicational form treating "environment as music," can use sound to create acoustic spaces and thus draw attention to our ongoing relationships to the everyday world.

References

Blesser, B. and L. Salter (2007) *Spaces Speak, Are You Listening? Experiencing Aural Architecture*, Cambridge, MA: MIT Press.
Bregman, A. (1990) *Auditory Scene Analysis*, Cambridge, MA: MIT Press.
Freed, D. (1990) "Auditory Correlates or Perceived Mallet Hardness for a Set of Recorded Percussive Sound Events," *Journal of the Acoustical Society of America* 87, pp. 311–322.
Gaver, W. (1993) "How Do We Hear in the World?: Explorations in Ecological Acoustics," *Ecological Psychology* 5(4), pp. 285–313.
Ingold, T. (2007) "Against Soundscape," in A. Carlyle (ed.) *Autumn Leaves*, Paris: Double Entendre, pp. 10–13.
Macdeo, F. (2015) "Investigating Sound in Space: Five Meanings of Space in Music and Sound Art," *Organised Sound* 20(2), pp. 241–248.
Paquette, D. (2004) "Describing the Contemporary Sound Environment" (Master's thesis), Simon Fraser University, www.sfu.ca/sonic-studio
Roads, C. (1996) *The Computer Music Tutorial*, Cambridge, MA: MIT Press.
Southworth, M. (1969) "The Sonic Environment of Cities," *Environment and Behavior* 1(1), pp. 49–70.
Sterne, J. (2012) *MP3: The Meaning of a Format*, Durham, NC: Duke University Press.
Thompson, E. (2002) *The Soundscape of Modernity*, Cambridge, MA: MIT Press.
Truax, B. (1998) "Composition and Diffusion: Space in Sound in Space," *Organised Sound* 3(2), pp. 141–146.
Truax, B. (ed.) (1999) *Handbook for Acoustic Ecology* [CD-ROM], Cambridge Street Publishing, www.sfu.ca/~truax/handbook.html
Truax, B. (2001) *Acoustic Communication*, 2nd ed. Norwood, NJ: Ablex Publishing.
Truax, B. (2002) "Techniques and Genres of Soundscape Composition as Developed at Simon Fraser University," *Organised Sound* 7(1), pp. 5–14.
Truax, B. (2008) "Soundscape Composition as Global Music: Electroacoustic Music as Soundscape," *Organised Sound* 13(2), pp. 103–109.
Truax, B. (2012) "Sound, Listening and Place: The Aesthetic Dilemma," *Organised Sound* 17(3), pp. 1–9.

Vörlander, M. (2008) *Auralization: Fundamentals of Acoustics, Modelling, Simulation, Algorithms and Acoustic Virtual Reality*, Berlin: Springer Verlag.

Warren, W. and R. Verbrugge (1984) "Auditory Perception of Breaking and Bouncing Events: A Case Study in Ecological Acoustics," *Journal of Experimental Psychology: Human Perception and Performance* 10, pp. 704–712.

Wishart, T. (1996) *On Sonic Art*, ed. S. Emmerson, Amsterdam: Harwood Academic Publishers.

World Soundscape Project (1999) *The Vancouver Soundscape 1973/Soundscape Vancouver 1996*, Cambridge Street Publishing, CSR-2CD 9901, www.sfu.ca/~truax/vanscape.html