

Spaces Speak, Are You Listening?

Experiencing Aural Architecture
Barry Blesser and Linda-Ruth Salter

Auditory Spatial Awareness Space as an Acoustic Arena

An acoustic arena is an intermediary between acoustic science and social science. Architects, acousticians, and engineers make decisions about spatial geometry, construction materials, and building technique, all of which influence the size, shape, and aural attributes of various acoustic arenas. Sociologists, anthropologists, and psychologists analyze how those who occupy these arenas react to them in terms of mood and behavior as well as a sense of privacy or social cohesion. Aural architecture bridges these two disciplines. An acoustic arena has both social and physical properties, serving as a shared concept for both disciplines.

Most descriptions of spatial boundaries arise from visual appearances and social markers, cultural signals that delineate a transition not just in social function, but in political rights. Acoustic arenas do not respect those transitions. When the windows of a private house are open during a summer afternoon, the acoustic arena of activities in the public street extends well into the private spaces of the house, and to a lesser extent vice versa. Yet ownership and social rights associated with both the house and street remain independent of the state of the windows. If you are the owner of a private space, you control who can enter and what they can do, but when you open the windows, you relinquish your control over the access of sonic events. The sounds of public life freely enter a private space, and an animated family discussion becomes part of the public arena, heard by any passerby.

An open window fuses visually and socially distinct spaces into a single arena. The social consequence of an acoustic arena is an acoustic community, a group of individuals who are able to hear the same sonic events. Within such a community, an individual who broadcasts some signal or information makes a sonic connection to everyone within the arena. The broadcaster can change membership in the acoustic community only by changing the size of the arena. We whisper to make an acoustic arena small and private, and we shout to make it large and public, thereby determining who is inside and who is outside. Using an inverted definition of a private acoustic arena, Leo Beranek (1960) describes it as a space where excluded conversations are inaudible.

In the strongest manifestation of a private acoustic arena, acoustic privacy is bilateral: outsiders cannot hear broadcasts emanating from within, and insiders cannot hear broadcasts emanating from outside. Given the importance of acoustic arenas, the following discussion explores the social consequence of public and private arenas. The concept of an acoustic arena applies equally to environments of all sizes and types: small private rooms, concert halls, large townships, and natural soundscapes. We expand our understanding of aural architecture by considering not only buildings and auditoriums that were designed according to a specific criterion, but also natural and accidental environments occupied by people and other mammals.

Human beings are only one of many species that evolved a sense of territory based on the size of their acoustic arena. Marc D. Hauser (1997a), in his analysis of animal communications among numerous species, described the complexity and importance of vocal signaling in a shared acoustic environment. Broadcasting vocal signals in a complex environment, such as a forest, is one of the most effective means of communicating because the acoustic horizon can be far larger than the visual or olfactory

horizon. Many species therefore evolved specialized auditory biology and social systems, adapting to their specific acoustic environment, to their acoustic geography—nature's aural architecture.

Early humans first adapted to nature's acoustic geography: open savannas and mountain ranges. Modern humans adapt, in a weaker way, to the acoustic architecture of urban centers and of enclosed dwellings and gathering places. Both natural and fabricated environments are relatively constant and difficult to change, but by changing their vocalization behavior, those who occupy them adapt, whether as individuals, groups, or species. Every acoustic arena is an application of the principle that social groups create or select an environment, which in turn, determines the resources of their acoustic arena. The vocal behavior of a social group creates an acoustic arena as a geographic region that supports an acoustic community. Large arenas allow for larger acoustic groups spread over a larger area.

No single acoustic arena illustrates, or even manifests, all possible uses of a space. Use depends on the prevailing cultural values. At a basic level, acoustic arenas can be sorted into three categories—natural, private, and public. Natural acoustic spaces, at least historically, were shared by competing species. Use of private acoustic spaces, because of controlled design and limited access, is often the prerogative of those with resources and power, both financial and political. Public acoustic spaces, with sonically porous boundaries that connect several physical spaces into a single acoustic arena, are influenced by a multiplicity of occupants, designers, and owners. Whether in natural, or private, public acoustic arenas, occupants adapt their behavior to the properties of the arenas available to them.

In our technological society, mechanical and electronic interventions have largely obviated the need for social cooperation in regulating the public arena. The earlier social rules for creating and controlling sonic events become less relevant when everyone exists within his or her own isolation chamber. Simply put, there is less need to regulate sounds outside of encapsulated spaces. Technology has produced high-quality private acoustic arenas, making public acoustic arenas less relevant, whether these are quiet or noisy. Simultaneously, the function of the public acoustic arena has been replaced by other means for achieving social cohesion, mostly in the form of electronic communications. But, even though we have far greater control over our electronic than our acoustic connections with others, these do not have the intimacy and immediacy of an acoustic community in a public arena.

Acousticians such as William J. Cavanaugh and Joseph A. Wilkes (1999) have sufficient knowledge to create nearly complete sound barriers, from recording studios and home theaters to mansions for the rich and famous, and partial sound barriers even for public environments, such as highways and airports. Over the years, new materials, installation techniques, and manufacturing methods have resulted in incremental advances in the technology of sound isolation. Any competent acoustician can design a sonic isolation barrier that approaches the theoretical limits of physics. The decision to produce an arena of one type or another is only a matter of economic and cultural choice. In some cultures, physical boundaries are sonically porous, and acoustic arenas depend on social agreement. For example, in Japan, paper screens serve as walls; in tropical islands, windows and doors are always open to allow the air to circulate. In the United States, doors are often hollow and have intentional gaps at the bottom. This contrasts with some Germanic countries where habitable spaces have rubber gaskets on solid doors, tight seals on windows, and thick concrete walls.

Although, to a lesser extent than physical boundaries, sound absorption can also subdivide a space into multiple acoustic arenas by creating virtual partitions. A concert hall, using a minimal amount of sound absorption, remains a single acoustic arena. In such a space, a noisy individual disturbs everyone. In contrast, a large living room with deep-pile rugs and well-upholstered furniture supports many independent conversations in separate arenas. They are private because sound absorption suppresses reflections that would fuse with the direct sound to make it louder and propagate farther. In such spaces, conversation is possible only when the speaker is facing the listener. Sound-absorbing surfaces allow the occupants to dynamically partition a space into separate arenas, whereas sound-isolation barriers prepartition a space without the active involvement of the occupants.

Somewhat paradoxically, a high level of background noise also partitions a space into many small acoustic arenas, creating a matrix of tiny virtual cubicles. For example, in an industrial factory, workers communicate with each other in private acoustic arenas that may have an acoustic horizon of a few inches from mouth of speaker to ear of listener. At all greater distances, factory noise masks the conversation. There can be hundreds of small acoustic arenas in such a space, and each of them is as private as an arena created with acoustic isolation barriers. There is, however, a major difference in comfort between a small acoustic arena created with high background noise and one created with sound isolation or absorption. Given the choice, few would choose a high noise level as the preferred means of creating a small acoustic arena. Using sound isolation or absorption is more expensive, but the arenas are more pleasant.

The history of human societies can be viewed through the prism of their acoustic arenas and acoustic communities. Like air, water, and land, acoustic arenas are the factory whistles to signal the start or finish of a work shift. Towns were organized around these soundmarks, and no one outside its arena enjoyed social cohesion with the community. Public acoustic arenas were valued for their ability to integrate individual into the social fabric of their community.

In his extensive study of bells in the nineteenth-century French countryside, Alain Corbin (1998) showed that self-esteem, emotional well-being, civic pride, and territorial identity all depended on hearing the town bells. When citizens heard the chiming of the bells, they felt rooted within a cultural geography that could easily be walked. Soundmarks provided local cohesion, a contrast to the modern concept of citizenship in a sovereign nation composed of millions of individuals spread over millions of square miles. Competition among towns and communes occasionally resulted in stealing one another's bells, and legal confrontation over the right to ring the bells resulted in riots. Corbin (1998) summarizes their attitudes with the well-known platitude "A town without bells is like a blind man without a stick."

Because the arena for a soundmark determined the scope of the town, those geological formations that would support sound propagation determined which regions could be absorbed into the township. Sound propagates farthest in valleys, which act like sonic conduits, and least over mountains, which cast acoustic shadows. As aural architecture on a grand scale, sonic geography controlled the social fabric of early rural communities. In the early twentieth century, when urban growth polluted the natural soundscape with noise, trolley lines rather than nature's sonic conduits defined social cohesion and its community boundaries. Transportation arenas replaced acoustic arenas. The public acoustic arena survived, but on a reduced and less personal scale.

Historically, for the average person without servants to act as messengers, living in a private acoustic arena meant social isolation. In contrast, a large public acoustic arena provided social inclusion. Schafer (1978) quotes a resident of a small town who remembers from the early twentieth century the importance of a large acoustic horizon, and the value of identifying horses by the sound of their steps: “The iceman had a couple of very heavy cobs . . . the coalman had a pair of substantial Percherons that always walked . . . the dry-goods store had a lightweight horses . . . and the Chinese vegetable men had very lazy horses.” In a town with acoustically porous living spaces, you could hear the fishing boats returning to harbor, the children walking home from school, the rattling of leaves in the wind, and the dog fighting with the cat. You would know that it was time to visit your neighbors when you heard their wagon returning from shopping. Sitting at home, and without moving from your chair, you were intimately connected to the activities on your street.

As part of acoustic ecology, this is but one example of how a sonic environment creates a connection and cohesion among people. In her review of Steven Feld’s documentary soundscape series, *The Time of the Bells*, Rachel Lears (2005) broadens the concept of soundmarks by mentioning the role of bells on bicycles, in carnivals, ceremonies, churches, government, and sheep farming. But she ignores the role of acoustic geography, even though it determines the scope of the arenas for those sounds. Thus the size of the acoustic arena for sheep bells, which effectively determines the protected grazing area for the herd, varies with the terrain, being larger in valleys and smaller on hillsides.

Modern society has a mixed attitude toward the size of public and private acoustic arenas. Radio, television, newspapers, e-mail, and telephone have replaced the public acoustic arena as ways to maintain social connections on a large scale. Cities are so noisy that residents treasure private acoustic arenas, often at the cost of feeling isolated, lonely, and anonymous. In contrast, within a modern household, a family arena exists when all family members sit together or keep their doors open. Our household, like many others, has no doors for any of the common rooms. When sitting in my office with the doors and windows closed, I am fully isolated from the public acoustic arena, whereas when I move to the backyard, I am fully immersed in the activities of my neighbors, the local squirrels, and the neighbor’s cat. Some companies place workers in a single large acoustic arena, with only managers having private acoustic arenas. Similarly, later discussions on musical spaces illustrate a cultural progression from the shared acoustic arenas of public performances in churches and concert halls to the private arenas of sound reproduction in homes and automobiles. Headphones produce the most private of all acoustic arenas.

To summarize: the principles of acoustic arenas apply directly to the aural architecture of all spaces. To create an attractive space, be it a courthouse, school, civic center, family residence, or house of worship, an aural architect must also incorporate contemporary attitudes toward acoustic arenas and acoustic communities.

Social Spheres and Acoustic Arenas

A scarcity of acoustic arenas, as with all limited resources, provokes competition among the groups of people using those arenas. The social dynamics of human groups determine the outcome of that competition. Although stronger groups capture a larger percentage of available acoustic arenas than weaker groups, laws and social conventions provide cooperative mechanisms for regulating particular kinds of arenas. For example, concert halls have strict rules that give musicians the exclusive right to create sounds, whereas

taverns have weaker rules that give any patron enjoying food and drink the right to sing. Airport agencies specify where and when airplanes can fly. Households have rules about the volume of television sound. In contrast, a self-indulgent motorcyclist riding through a neighborhood usurps the right to make noise, interrupting hundreds of conversations in less than an hour. Injecting noise of whatever kind into an acoustic arena is nothing more than the exercise of sonic power: social or political, autocratic or democratic, supportive or destructive.

When they design a space, traditional architects exercise power ultimately as potent as that of social or political agents in determining, however unwittingly, the size, use, and attributes of the acoustic arena of that space. For their part, those who occupy or live within a space have a dynamic, bilateral, and continuing relationship to space within their own aural architecture of created acoustic arenas. In contrast, the relationship of traditional architects to their spatial creations is severed at the completion of the project.

Whereas a traditional architect creates an acoustic arena in a space by erecting boundaries that are sonically impermeable, the occupants of that space create equivalent arenas by asserting their social or political right to generate sonic events. From the perspective of the acoustic community, voluntary silence and physical barriers produce equivalent arenas. Especially in the previous century, many societies passed laws to control nuisance noise, such as that made by vendors, barking dogs, radios, carpet beaters, and street musicians. In some cities, they attempted to enforce quiet on Sunday to emphasize the solemnity of a day devoted to religion. We have all been surprised at the large size of public acoustic arenas on a quiet Sunday morning. Like the airways, public acoustic arenas are common resources owned and thus to be regulated by the people.

For both architects and occupants, silence reveals more about the social and cultural aspects of acoustic arenas than sounds. Silence is far more than the absence of sound, a definition that considers only the physical properties of sonic vibrations. Rather, silence may be understood as an active choice by the creators of acoustic arenas: the occupants and the architects. The absence of sonic events—silence—is important because it leaves the acoustic arena available for low-level sonic events that add nuances to communications. Silence creates large acoustic arenas as a common resource, whereas loud sound consumes that resource. Only the highest-quality acoustic arenas, with very low background noise, communicate silence.

A few examples illustrate the social and psychological complexity of silence. It can signal: a cessation of both social and natural activity, a state of psychological tranquility, a powerful emotion that transcends speech, a cooperative agreement to respect the public soundscape, a silent prayer communicating with a deity, a preoccupation with inner thought, a punitive response to social or political transgressions, or an acceptance of the right to be left in peace. Such nuances of communication are severely degraded when the aural environment falls victim to intruding noise.

The level of background noise determines the quality of an acoustic arena and the reliability of its auditory channels. A silent environment creates the best auditory channel; a noisy environment the worst. The sonic properties of the channel determine what messages can be transmitted. Communicating with ringing bells from a bell tower is more reliable than communicating a public announcement by voice, which is more reliable than communicating emotional intimacy by subtle tonal inflections. A noisy acoustic arena only allows for basic communications, such as a bell sound, because noise degrades the

subtler aspects of vocal communications and social cohesion. For example, even at the most fundamental level, oral communication becomes more stressful when noise masks the short silent interval that distinguishes a voiceless consonant from its voiced counterpart, such as a p from a b. Similarly, noise prevents signaling with a hesitant pause, which may signal the speaker's lack of confidence, or with a sudden cessation of speech, which may be intended to coerce an unwilling response from the listener. Such signaling requires a silent background.

Unlike practitioners of vocal religions, Quakers value silent prayer as a way to distinguish that activity from the profanity of ordinary speech. They regulate silence using strong rules that forbid transgressing on the religious commons (Bauman, 1983). Group silence is the ultimate manifestation of social cohesiveness because silence can exist only if all members cease from speaking—total deference to the group's values. When silence dominates, vocalized prayer takes on special meaning: voices framed by the boundaries of silence rather than lost in an ocean of sound. Silence is the central component in many religions and rituals (Szuchewycz, 1997).

Teachers, judges, priests, and tyrants all have the power to silence others. To be silent in the face of authority can show either deference or defiance. The asymmetric relationship between those who give orders and those who must obey is always demonstrated by who controls access to the soundscape. The common command "Silence!" demonstrates political power because it defines who is allowed to express a point of view. Adam Jaworski (1993) called these interactions "the politics of silence and the silence of politics," and Wreford Miller (1993) stated that silence, or the lack of it, has been politicized in modern society to the point where the sounds themselves matter little.

Acoustic arenas are commercial as well as political. In exploring the history of background music over the last half century, Hildegard Westerkamp (1988) observes the unchallenged right of commercial organizations to exercise control over individuals in their acoustic community. For them, an acoustic arena is private property to be leased by the highest bidder. Marketing literature from companies that sell music services to commercial enterprises is explicitly blunt. With training in behavioral psychology and human engineering, the founder of Muzak claims that you will "see the difference in customers," and injected music will "teach your cash register to sing with the foreground music from AEI" (Westerkamp, 1988). Airport lounges, even as semipublic spaces, saturate their occupants with television advertising. Waiting passengers may avoid attending to the visual component of that space, but they cannot block its aural counterpart. Television sound creates a sufficiently large acoustic arena for its message that other acoustic arenas are reduced in size. Just as sponsors or owners may commission traditional architects to design acoustic arenas by manipulating acoustic parameters, they may also design these arenas themselves by injecting background and foreground sounds or by enforcing rules about who else can also inject sound.

More commonly, ownership rules of an arena are created informally when two or more individuals congregate for a social interchange. Territorial bubbles appear as if by magic around a group of individuals if they begin to interact, and the group quickly acquires rights to the arena. When encountering such a social bubble with its implied acoustic arena, outsiders are reluctant to intervene or to create sonic events (Lindskold et al., 1976). The strength of ownership rights to an acoustic arena depends on the distance between individuals, their perceived status, and the nature of their interactions. Cultures assign implicit rights to acoustic arenas, and there are complex unwritten rules governing the size of an arena being claimed.

Understanding the social rules for acoustic arenas requires the concept of social distance, as embodied in the term social sphere, which then becomes the means for evaluating arenas. The sounding of a foghorn is a public broadcast intended for everyone with the expectation that its acoustic arena will be large, whereas a whispered comment is a private communication intended only for an intimate companion with the expectation that its acoustic arena will be small. Social expectations determine the properties, especially size, of an acoustic arena, and social behavior then adapts to available arenas. For example, if an acoustic arena were large enough to signal emotional nuances over a great distance, its large size would conflict with expectations of privacy. Similarly, a small public acoustic arena conflicts with the need to broadcast public information to a large population. To be socially useful, acoustic arenas and their properties must match the cultural norms governing social spheres.

Whereas physical distance is measured in meters or feet, social distance depends on the social context. The social anthropologist Edward T. Hall (1966) divided social distance into four spheres: (1) the intimate sphere, which ends at about half a meter (1–2 feet) and is reserved for intimate friends and relatives; (2) the personal sphere, which ends at about 1 meter (3 feet) and is reserved for acquaintances; the conversational sphere, which ends at about 4 meters (12 feet) and is reserved for oral interchanges with strangers; and the public sphere, which is determined by the acoustic horizon and is impersonal and anonymous. How we experience a person, object, or sound depends on these distances, which Hall called “proxemics,” the experiential manifestation of anthropological distance, which varies from culture to culture.

For each of these four spheres, a culture provides implicit ownership rules for the corresponding acoustic arena. Rules for the intimate sphere are rigid—lovers do not permit outsiders to enter. Strangers encountering an intimate sphere are likely to fall silent or speak softly. Rules for the public sphere are malleable—the social consequences of transcending sonic norms are minimal. Other spheres are intermediate cases between intimate and public spheres. Aural architecture fails when there are conflicts between social spheres and acoustic arenas. For example, individuals in a conversational sphere spanning a distance of 4 meters (12 feet) cannot coalesce if the arena diameter is only 3 meters (9 feet). Similarly, with an acoustic horizon of only 4 meters, sonic events in the public sphere are inaccessible.

To illustrate an application of social spheres, let us use proxemics to evaluate a chamber music concert. Musicians are located on stage in their conversational sphere, whereas listeners are located in their audience seats in the public sphere. Even if the management provides audience seating on the stage, some listeners are uncomfortable in a socially inappropriate sphere. The performers own their conversational sphere. But in the nineteenth century, performers and listeners often sat together in a small chamber, comfortably sharing a common conversational sphere. Today, if you put on binaural headphones by using spatial synthesizers, an audio engineer can place a virtual musician two inches from your left or right ear, well within your intimacy sphere. When such technology creates additional freedoms to move the location of a sound source, conflicts between the social and artistic expectations of the appropriate sphere may suddenly appear.

Proxemic distances are useful for evaluating the relationship between social spheres and acoustic arenas. If society does not provide the appropriate acoustic arena, then the corresponding social sphere is unavailable, and the corresponding social activities

are not possible. Availability of an appropriate acoustic arena, in turn, depends on the aural architecture, which itself is a combination of acoustic design and the social rules for regulating sonic events. Aural architecture is not only the physical design of a space, but also part of a complete social system. We can only appreciate the importance of aural architecture when we recognize the interwoven relationship between spatial awareness, social behavior, and the design or selection of a physical space.

Only listeners with motivation, dedication, and aptitude become expert at transforming the acoustic attributes of objects and geometries into a useful three-dimensional internal image of an external space. As with training sonar operators to identify underwater objects by how they modify incident sound, acquiring expertise of any form of auditory spatial awareness requires hundreds of hours of practice. Why would someone invest so much effort to acquire this proficiency?

Some listeners obviously benefit by having this ability. Musicians and composers include spatial attributes as a component of their art; acousticians depend on spatial awareness for designing concert halls; and audio engineers create spatial illusions with synthesizers. Listeners who must move around in places without light are likely to acquire the some basic abilities to recognize open doors, nearby walls, and local obstacles. But of the many groups of listeners who use auditory spatial awareness in their personal and professional lives, those with a visual deficit have the strongest motivation: hearing is a way to orient and navigate space, and their reward in developing their spatial awareness is the possibility of leading a normal and fulfilling life.

By itself, blindness never improves hearing. The auditory acuity of blind people as a population is average, spanning the same range of abilities to be found in the general population. On the other hand, some blind individuals are indeed motivated to enhance their spatial abilities far beyond the average. Practice is the most important predictor for achieving a high level of proficiency. With sufficient practice, some become expert, often displaying skills that are so extraordinary as to border on the magical. Such individuals illustrate what our species, in the limit, is capable of achieving. We are how we live—there is no generic human being.

There is evidence that those who practice a sensory or motor skill for thousands of hours change their brain wiring. Neurological studies, discussed in chapter 8, show that the cortical regions that process specific auditory cues are larger in conductors, musicians, and those with visual handicaps than in other people. Enhanced auditory spatial acuity is entirely a property of specialized sections of the brain that have been trained to interpret relevant audible cues. Listeners strengthen their neurological structure by repeated auditory exercise, just as athletes strengthen their muscles by physical exercise. Although the superb physiques of Olympic swimmers are plain to see, we cannot see the correspondingly superb “physiques” of “auditory athletes,” except by observing their behavior while engaging in life’s activities.

Cognitive strategies for decoding spatial attributes use such cues as the difference in time, amplitude, and spectrum between the sounds arriving at the two ears, as well as detection of changes in the expected spectral and temporal attributes of familiar sounds. Although some acoustic cues are specific to interpreting spatial attributes, most cues are unrelated to spatial acoustics. The cues that distinguish a p from a b, or a violin from an organ, are unrelated to space. Learning to hear space is mostly a matter of inventing a cognitive strategy that can decode the specific cues that arise from the acoustical behavior of objects and geometries in the world. From a physiological

perspective, we all hear the sonic attributes of objects, but, absent training, we neither attend to their aural cues nor invent cognitive strategies for interpreting them. Although placing your hand a few inches from one ear illustrates that the hand's presence is audible, to translate audibility into a conscious sense of a hand, with its corresponding size, location, and skin surface, you must adopt a unique cognitive strategy. Far more difficult to detect, a traffic sign at a distance of a few meters also produces a set of cues that allows a skilled listener to detect the sign's existence and shape. At best, even when highly developed, auditory "seeing" of space (echolocation) is comparable to extreme visual nearsightedness, identifying physical objects that are relatively nearby or comparatively large. Small or remote objects simply do not produce aural cues that can be interpreted by any human being.

Echolocation is directly relevant to aural architecture because it conclusively demonstrates that our species has the neurological endowment to make judgments about objects and spatial geometries just by listening. Yet most aural architects, both amateurs and professionals, are unfamiliar with the native ability of human beings to hear space. A listener using a cognitive strategy to transform auditory cues into an image of a space, by sensing the doorway to the bathroom late at night, for example, is experiencing the navigational spatiality of aural architecture.

Experts at Hearing Objects and Geometries

Although history is replete with anecdotal accounts of blind persons "seeing" space, it was only in the mid-twentieth century that this ability came to be understood as an auditory skill. Curiously, the auditory ability of bats and dolphins to navigate without vision or smell was also discovered at about the same time, and it is now known that other species have a residual ability to sense their environment in the dark. Rather than thinking of this ability as a curiosity, such as sensing magnetic fields or infrared light, scientists now recognize that hearing space is more common than first imagined. Even though most animals and people with adequate vision and available light have little need to enhance their residual ability to hear space, it remains a viable alternative for supplementing vision.

The scholarly language to describe orienting and navigating in a space through hearing is ambiguous and confused. For example, the literature incorrectly uses the term echolocation (locating by means of self-generated echoes) for all forms of spatial awareness.

This name originated from studies of bats and dolphins, which have a synchronized means for vocalizing and then decoding the responding echoes. Currently, the term echolocation applies to sensing spatial attributes with any kind of sounds, not just with self-made ones, whether by vocalizing, clicking fingers, or tapping canes. Background noise, for example, may provide sufficient sonic illumination to "see" aspects of a space. Moreover, the concept of echolocation, as now understood, also applies to acoustic cues other than echoes. Terminological confusion arose because the phenomenon of spatial awareness was recognized long before its physical and perceptual basis were understood.

One of the earliest written records of face vision, the early name for echolocation, was recorded by Denis Diderot (1749), who described the amazing ability of some blind individuals to perceive objects and their distances. Two centuries later, as part of his work at the Perkins School for the Blind, Samuel P. Hayes (1935) reviewed and cataloged the evidence for echolocation from a scientific rather than philosophical perspective.

In his review, Hayes notes that scientists began to study echolocation only after sufficient anecdotal evidence and personal testimonies demonstrated that it was a real phenomenon. He describes a particularly impressive example of blind navigation he himself witnessed: Martin was a native of New York City and had been blind nine years. He was of a fearless and impetuous disposition, and went about the city without a guide. He passed up, down and across great thoroughfares frequently and only a few times colliding with a bicycle, which vehicle he detested. I was with him on occasions when I marveled at the perfect freedom with which he walked along crowded streets, showing not the slightest timidity, and requiring no aid whatever from me. . . . I was amazed to see him cross Broadway at 14th Street with perfect ease, and imagine my astonishment when he shied around timbers that had been set up across the sidewalk to prop the wall of a building undergoing repairs. He got on and off street cars without a blunder and made his way across narrow streets without betraying his blindness. He used no cane nor did he feel his way with his hands. Had I not known that he was actually blind I would have thought that he was feigning. I asked him how he knew his way and avoided collisions, and he invariably told me that he did not know. He seemed to be guided by what I shall term a miraculous instinct superimposed by a subconscious mental condition. I am inclined to the belief, in the absence of a better theory, that he was directed by what Hudson terms "the subjective mind"! (Hayes, 1935)

The historical literature contains many such testimonies from many periods and cultures. Accepting the introspective comments of those who are adept at echolocation provides the kind of insight that is not yet available from scientific studies, which reveal little about the underlying cognitive strategy for sensing space. These testimonies emphasize several important aspects of echolocation. First, the skill is not conscious, and even those who have a highly developed skill cannot describe how they do what they do. Second, the exclusive use of echolocation for navigation requires great courage. Third, using hearing for navigation, at least at this high level of performance, is unusual; more frequently, blind persons depend on touch with their cane, using echolocation only as a supplement to their tactile sense of space. How blind persons acquire a cognitive strategy for echolocation is still somewhat of a mystery. Ved Mehta (1957), blind from childhood, described his experience of navigational space. Wanting to live a normal life in Calcutta, he learned to jump from banister to banister, from roof to roof, and rode his bicycle through unfamiliar places.

When he later attended the Arkansas School for the Blind, he participated in their echolocation program, which was based on motivating students to avoid the pain of colliding with suspended objects. Teachers simply believed that echolocation could be learned by anyone, and their task was to provide motivation to invest in such learning. Mehta described the environment, not the process of learning the skill.

One day in early spring, all the totally blind students were herded into a gymnasium and asked to run through an obstacle course. Plastic and wooden slabs of all sizes and weights were suspended from the ceiling around the gymnasium. Some of them hung as low as the waist; others barely came down to the forehead. These slabs were rotated at varying speeds, and the blind were asked to walk through the labyrinth at as great a speed as possible without bumping into the obstacles.

The purpose of keeping the slabs moving was to prevent the student from getting accustomed to their position and to force them to strain every perceptual ability to sense the presence of obstacles against the skin—a pressure felt by a myriad of pores above, below, and next to our ears. Some of the slabs were of an even fainter mass than the

slimmest solitary lamppost on a street corner. This obstacle course helped gauge how well an individual could distinguish one shadow-mass from another and, having located the one closest to him, circumvent it without running into yet another. . . . The gymnasium was kept so quiet that the blind people could hear obstacles, although I could not help feeling that I could have run through the labyrinth with a jet buzzing overhead. . . . Forme, going through this obstacle course was child's play. (Mehta, 1957)

Although the details of learning echolocation vary, there is common attitude shared by those who are determined to “see” with their ears. Ved Mehta was not unique. The world-famous jazz musician Ray Charles eloquently describes a similar approach to living as a blind child (Charles and Ritz, 1978): “Being blind wasn’t gonna stop me from enjoying the bike. . . . Somehow in the back of my mind I knew I wasn’t going to hurt myself. Sure, I rode pretty fast, but my hearing was good and my instinct was sharp. . . . On another day Momma asked me to chop wood. . . . I was treated like I was normal. I acted like I was normal. And I wound up doing exactly the same things as normal people do.” A few years later, he went to a special school for the blind, but his attitude toward echolocation was already solidified. “There were three things I never wanted to own when I was a kid: a dog, a cane, and a guitar. In my brain, they each meant blindness and helplessness.” Being sensitive to the nuances of sound in general, he taught himself music and echolocation by listening carefully to the world of sound. Ray Charles never used a cane for navigating a space.

During the ensuing half century, modern methods have evolved for teaching echolocation, but the assumption that it can be taught is still controversial. Many, if not most, schools for the blind have abandoned teaching it. What explains the current lack of interest? In reviewing the literature, I noted that, with the exception of Kish (1995) and a few others, those who teach echolocation are themselves fully sighted, as such, they are very unlikely to develop sophisticated echolocation abilities. In contrast, Kish was blind from childhood, and taught himself echolocation by an intuitive sense of how to acquire that skill. He is now a licensed teacher for orientation and mobility, having created his own teaching methods (Kish and Bleier, 2000). Along with a colleague, Kish founded TeamBat, a program that guides blind teenagers into the mountains on bicycle trips, shown in figure 2.3. The answer to the earlier question is, in part, that echolocation is more a commitment to learn than a teachable skill.

Those blind individuals who use echolocation belong to a unique sensory subculture that has transformed a latent ability to hear navigational space into a high art form. Although there is no question that most listeners possess only the most rudimentary ability to detect spatial objects and geometries by listening, the difference between experts and beginners is only a matter of degree because the underlying cognitive and personal issues are the same.

Like ear training for musicians (Ottman, 1991) and for audio engineers (Moulton, 1993), learning echolocation also involves attending to the subtlest auditory cues. Unlike such training, however, echolocation involves an additional step—using a cognitive strategy to convert binaural cues into spatial images. Those cues originate from a multiplicity of transient sound sources interacting with a range of moving objects and surfaces. Consider the number of sounds and surfaces on an urban street. The cognitive strategy for echolocation must process all of them. Acquiring this ability therefore requires an individual to practice in a real sound field in a real space. For this reason, echolocation is best learned as part of daily life in a real-life environment, unlike other forms of ear training, which can take place in a studio or classroom. It is difficult, if not impossible, to artificially

create or record teaching examples that faithfully replicate realistic sonic environments. The ability to create an internal picture of external objects and geometry is greatly enhanced when strong motivation, greater than average skills, and an extended opportunity to practice are present. For blind individuals, enhanced echolocation ability correlates with several key factors. Engaging in echolocation, if begun in childhood when brain substrates are evolving, can readily adapt neural structures to become optimized for different purposes. A child without any residual vision is simply more likely to discover hearing as an alternative means for navigating a space if permitted to do so. Because practicing echolocation includes the risk of injury, the child needs to be comfortable taking risks, and the child's parents must avoid excessive protectiveness. In fact, participating in activities that normally assume adequate vision is the best predictor of acquiring auditory spatial awareness for navigating, as attested by the personal examples of Daniel Kish (1995, 2001), who categorically rejected the guidance of those who urged him to learn to use the cane, and Ved Mehta (1957), who moved about the streets of Calcutta without supervision. Investing in auditory spatial awareness is always a free choice that any of us can make, but few do.

Even though there are numerous examples of individuals who learned echolocation, the rehabilitation literature is, at best, ambivalent about using hearing rather than the tactile sense for navigating space. When large numbers of soldiers returned from World War II with visual disabilities, formal training programs became a priority, and echolocation was an obvious technique (Bledsoe, 1980). After prolonged controversy and passionate debates, rehabilitation workers involved in helping blind soldiers eventually concluded that tactile navigation—using a cane—was simply easier to teach. Many soldiers could not, or would not, learn to sense subtle auditory cues and invent cognitive strategies. Some schools for the blind explicitly taught auditory spatial awareness, which is fundamentally different from navigational skills (Campbell, 1992), although such teaching proved problematic because most rehabilitation professionals were themselves sighted and could not teach from personal experience. Scientific studies of blind persons using echolocation do not reveal the underlying cognitive processes. As a generalization, cognitive strategies are learnable but not necessarily teachable; for those who cannot echolocate, such strategies have little, if any, practical value in daily life.

The literature on echolocation actually illustrates a larger principle: sensory skills are acquired, rather than innate; they are based on personal utility and lifestyle. Blind persons with the ability to echolocate are an obvious example of a sensory subculture that has the ability to use a specific cognitive strategy to interpret spatial cues arising from one aspect of aural architecture: navigational spatiality. In contrast, professionals who are actively engaged with other aspects of aural architecture, such as designers of concert halls or composers of music, become very adept at other cognitive strategies for interpreting other spatial cues.

Insights into the sensory and cognitive aspects of echolocation contribute to our understanding of aural architecture. And for this reason, it is worth shifting the discussion from anecdotes to research. By the mid-twentieth century, explaining the intractable phenomenon of echolocation became a scientific challenge. As with many perceptual phenomena that are complex, researchers broke echolocation down into many small, simplified questions and special cases. Theories about how we hear the distance to an isolated wall or how we judge the size of a door opening are examples of special cases. At the current state of knowledge, the cognitive and perceptual sciences are more collections of disconnected theories and experiments than unified wholes. On the other

hand, when a blind person rides a bicycle in a city, that person is merging a great number of special cases into a holistic strategy. Navigating real spaces involves hearing walls, openings, passive acoustic objects, and extracting their relationship to the location and properties of sound sources. The whole is far larger than the sum of the parts. Space is experienced as an unconscious unity rather than as a collection of recognizably separable processes.

To appreciate the acoustic complexity of an urban street, consider that the environment is composed of multiple objects and numerous sound sources, some stationary, some mobile. Each traffic sign, parked automobile, or telephone pole has a surface that produces both sonic reflections when the sound source is in front of it and acoustic shadows when it is behind. A reflection may be heard as an echo if the sound is impulselike and the surface is more than 10 meters (33 feet) away, or as tonal coloration if the source is continuous and the surface is nearby. A sonic shadow may be diffuse and blurred for low frequencies, or sharp and clear for high frequencies. Sonic illumination is the visual equivalent of a space illuminated with multiple lights: some bright, some dim, some colored, some blinking, and some moving. In a real-life environment, the sound field is indeed complex.

Now consider that, because you have two ears separated by the width of your head, each ear senses sound at a slightly different location in space. By moving or rotating your head, you reposition your two ears at another location. The physical sound field actually varies in three dimensions: left-right, front-back, and up-down. Obviously, if we had more ears and if our heads were larger, the auditory cortex would acquire far more information about the spatial distribution of sound. But even with our limited abilities to sense a three-dimensional sound field, the sounds arriving at the two ears are often sufficient for the auditory cortex to build a perceptual model of the objects and geometries that could have produced those particular sounds. Perception is an unconscious inferential process that synthesizes a hypothetical collection of objects and geometries. This process is the result of having learned the subtle, ambiguous, and inexact relationship between auditory cues and spatial attributes. Those who have developed echolocation skills cannot describe how the spatial image suddenly appears in their consciousness.

Scientists are still probing for important clues and theories to explain echolocation. Since the phenomenon of echolocation was first recognized by Michael Supa, Milton Cotzin, and Karl M. Dallenbach (1944) at Cornell, explanations of its mysteries have been of periodic interest to small groups of researchers. The science of echolocation is far from the mainstream of auditory research, being supported mostly by those with an interest in rehabilitation of people with visual deficits.

Before reviewing what science has learned about echolocation, we need to explain the tentativeness of research conclusions. Scientists are wrestling with a confounding methodological problem: individual listeners are remarkably inconsistent in their abilities to hear space. Auditory spatial awareness ranges from raw sensation to unbelievably high levels, corresponding to an equally wide variability in sensitivity to acoustic cues and effective cognitive strategies. Is a scientist actually studying a general phenomenon, or the unique ability of specific individuals on specific tasks? In practice, scientists ignore this question when they use randomly selected subjects. Even within the sorted population of blind subjects, there is a wide range of abilities. Human echolocation is actually a collection of independent abilities to perform a variety of tasks, from hearing spectral changes produced by a nearby wall, to hearing the acoustic shadow produced by

a telephone pole, to hearing the reverberation arising from two coupled spaces. A given listener might be very good at one task but mediocre at another. Experiments are designed to focus on a single task under controlled conditions.

For example, blindfolded subjects might be asked to walk along a long hallway with a single continuous noise source located at the end. Because there is only one sound source and a very simple geometric, acoustic shadowing, diffraction, and reflections cannot exist. In this restricted case, the experimental paradigm is evaluating the degree to which a subject's cognitive strategy incorporates only auditory cues explicitly included in the experiment. Although good scientific studies produce modestly consistent results, it is unclear how or when such insights apply to real life.

Even with these limitations, scientific results explain certain aspects of echolocation. Daniel H. Ashmead and colleagues (1998) showed that blindfolded subjects walking through a hallway without colliding with the wall detected low-frequency tonal coloration near walls. The ear closer to the wall surface senses background coloration different from what the farther ear senses. In the center of the hall, the coloration is the same in both ears. Differential coloration corresponds to distance to the wall. The same mechanism allows subjects to detect when they are passing an open door, which is equivalent to a missing wall.

In addition to hearing an open door as the absence of a wall, the door's frame creates acoustic shadows of sounds originating from within the room. When presented with an open doorway of unknown width and height, subjects can estimate its dimensions relative to their own body size with remarkable accuracy (Gordon and Rosenblum, 2000). Walking past an open door into another room therefore involves at least two cues: the absence of coloration from the missing wall segment, and the sonic shadows produced by sounds emanating from the room. The relative contribution of each type of cue depends on the sonic illumination in each space. If one space has stronger sonic illumination and the other has none, only one of the two cues would be available for sensing the doorway. Moreover, if the door were partially open, the door surface would itself become a source of reflections, which would then become yet another set of cues. The door is an additional object, separate from, but related to, the open doorway and the doorframe. In this simplified example, a trained listener uses a cognitive strategy that melds three sets of cues into a single image of the space: a partially open door in a doorframe leading to another room.

Listeners can sense not only doors and walls, which are relatively large, but also small objects and small differences in larger objects when they are relatively nearby. Charles E. Rice (1967) showed that listeners can detect a difference of 1 centimeter (3/8 inch) in a 9-centimeter (3 1/2-inch) disk at a distance of 60 centimeters (2 feet). Winthrop N. Kellogg (1962) showed an even higher level of discrimination: listeners detected an area difference of 5 square centimeters (3/4 square inch) on a square of 60 square centimeters (9 square inches) at a distance of 2 meters (7 feet). One blind subject could reliably detect a 1-inch disk located at a distance of three feet (Rice, 1969, 1970). Even more remarkably, Steven Hausfeld and colleagues (1982) demonstrated that listeners could distinguish square, circular, and triangular objects. One blind subject was able to recognize a stop sign by its hexagonal shape. Kellogg (1962) found that on the most difficult discrimination tasks blind individuals performed significantly better than sighted subjects who were blindfolded.

Although only a few studies have been designed to explore why some individuals performed better than others, Connie Carlson-Smith and William R. Wiener (1996) showed that two specific aspects of auditory acuity were partial predictors of echolocation

ability. Those subjects who performed best at detecting spatial attributes were also better at sensing small changes in the amplitude and the frequency of continuous sounds. When a sound field is not uniform, moving through it converts spatial differences into time differences. As listeners move through the space, they hear spatial differences as temporal changes. Although the ability to detect soft or high-frequency sounds at threshold is not related to echolocation, the ability to hear and interpret small changes in sound is. Apart from genetic endowment, learning is the dominant component of acquiring echolocation skills. We are not, however, speaking of 20 hours of practice but of thousands of hours. Say you are a 20-year-old adult. You have already spent well over 100,000 hours listening to the physical world of spaces. If, during that time, you had also engaged in self-directed practice exercises, as would a blind person moving through life's spaces, you would likely have much improved your perceptual acuity to aural cues, and have become highly proficient both at inventing cognitive strategies and applying them to convert those cues into spatial perception. Like athletes who love sports, those who want to become more proficient in echolocation engage in com-of cue depends on the sonic illumination in each space. If one space has stronger sonic illumination and the other has none, only one of the two cues would be available for sensing the doorway. Moreover, if the door were partially open, the door surface would itself become a source of reflections, which would then become yet another set of cues.

The door is an additional object, separate from, but related to, the open doorway and the doorframe. In this simplified example, a trained listener uses a cognitive strategy that melds three sets of cues into a single image of the space: a partially open door in a doorframe leading to another room. Listeners can sense not only doors and walls, which are relatively large, but also small objects and small differences in larger objects when they are relatively nearby. Charles E. Rice (1967) showed that listeners can detect a difference of 1 centimeter (3inch) in a 9-centimeter (31 2-inch) disk at a distance of 60 centimeters (2 feet). Winthrop N. Kellogg (1962) showed an even higher level of discrimination: listeners detected an area difference of 5 square centimeters (3 4 square inch) on a square of 60 square centimeters (9 square inches) at a distance of 2 meters (7 feet). One blind subject could reliably detect a 1-inch disk located at a distance of three feet (Rice, 1969, 1970). Even more remarkably, Steven Hausfeld and colleagues (1982) demonstrated that listeners could distinguish square, circular, and triangular objects. One blind subject was able to recognize a stop sign by its hexagonal shape. Kellogg (1962) found that on the most difficult discrimination tasks blind individuals performed significantly better than sighted subjects who were blindfolded.

Although only a few studies have been designed to explore why some individuals performed better than others, Connie Carlson-Smith and William R. Wiener (1996) showed that two specific aspects of auditory acuity were partial predictors of echolocation ability. Those subjects who performed best at detecting spatial attributes were also better at sensing small changes in the amplitude and the frequency of continuous sounds. When a sound field is not uniform, moving through it converts spatial differences into time differences. As listeners move through the space, they hear spatial differences as temporal changes. Although the ability to detect soft or high-frequency sounds at threshold is not related to echolocation, the ability to hear and interpret small changes in sound is.

Apart from genetic endowment, learning is the dominant component of acquiring echolocation skills. We are not, however, speaking of 20 hours of practice but of thousands of hours. Say you are a 20-year-old adult. You have already spent well over

100,000 hours listening to the physical world of spaces. If, during that time, you had also engaged in self-directed practice exercises, as would a blind person moving through life's spaces, you would likely have much improved your perceptual acuity to aural cues, and have become highly proficient both at inventing cognitive strategies and applying them to convert those cues into spatial perception. Like athletes who love sports, those who want to become more proficient in echolocation engage in com-

of cue depends on the sonic illumination in each space. If one space has stronger sonic illumination and the other has none, only one of the two cues would be available for sensing the doorway. Moreover, if the door were partially open, the door surface would itself become a source of reflections, which would then become yet another set of cues. The door is an additional object, separate from, but related to, the open doorway and the doorframe. In this simplified example, a trained listener uses a cognitive strategy that melds three sets of cues into a single image of the space: a partially open door in a doorframe leading to another room.

plex sensory activities that simultaneously exercise a wide range of skills and methods. They invent methods to teach themselves how to become proficient—customized pedagogy. Formal training managed by a (usually sighted) teacher in a classroom is far more limited than a lifetime of training managed by the individual listeners themselves.

Sensory practice changes the brain. When examining blind subjects who had engaged in extensive practice, Brigitte A. Röder and colleagues (1999) found that their neurological responses to sounds in the peripheral field were significantly better than those of normal subjects. With enough practice, the improved ability of the blind subjects is observable in the neurological response of the relevant cortex. Similarly, Christo Pantev and colleagues (2001b) found that the brains of pianists who began their careers as children responded more intensely to piano notes than those who began later. Because immature brains have greater plasticity in their neurological wiring, practice produces larger brain changes during early developmental periods. Learning is far more specific to the task being practiced than you might expect, and acquired skills do not readily transfer from one task to another. Just as exercising one muscle group does not strengthen other muscles, exercising one sensory skill does not enhance other skills: each sensory skill involves specific brain substrates. An audio engineer who has acquired enhanced acuity to tonal coloration in reverberation is unlikely to transfer that skill to navigating a corridor without vision. Although the concept of task-specific learning is well understood, only a few isolated experiments confirm the phenomenon. A curious experiment on pitch discrimination dramatically illustrates the extreme specificity of auditory learning. Laurent Demany and Catherine Semal (2002) trained subjects over the course of 11,000 sessions to discriminate the pitch of a 3,000 Hz tone from tones at slightly different frequencies, a very specific task indeed. Subjects improved by a factor of 3, and would likely have improved further had training continued. Not only is it surprising that intensive practice produces improvement on such a basic psychophysical task; it is even more surprising that improvement at this one frequency did not transfer to other frequencies. Pitch discrimination at 8,000 Hz remained unchanged. Subjects were not learning generic pitch discrimination; they were learning pitch discrimination of 3,000 Hz tones. Although I believe that this result applies to a large number of other phenomena, scientific studies have not yet revealed the extent to which spatial cues can be learned with extensive practice.

These somewhat speculative conclusions have broad implications. First, extensive

practice produces dramatic changes in perceptual ability, and those changes are observable using neurological imaging techniques. Brains reflect how individual listeners live their lives. Second, a culture that motivates and rewards listeners to learn auditory spatial awareness is likely to have a population that can better appreciate aural architecture. And conversely, without such a population, aural architecture is likely to be irrelevant to the culture. Third, auditory spatial awareness is a collection of independent sensitivities. Some listeners may be acutely aware of reverberation and the enclosed volume of a space, whereas others may be aware of local objects and geometries in a navigational space. Finally, any discussion about aural architecture must include an understanding of various aural subcultures, each of which has its own idiosyncratic investment in the ability to detect and appreciate attributes of spaces.

Although our internal representation of space usually originates from an external reality, internal and external representations are not as tightly linked as you might expect. To use a misleading analogy, we often speak of an internal image as if it were a neurological “photograph” created by the brain. But internal images are not replicas of the external world. How does an external space become an internal space, and in what ways are these two spatial concepts related? The answer to this question involves cognition as well as perception and lifestyle as well as biology.

Although our knowledge of how the brain creates its internal representation of an external reality is, at best, rudimentary, a diverse collection of fragmentary insights reveals a consistent picture. Evidence shows that cognitive processing of spatial attributes is plastic, flexible, adaptive, and dependent on the way individual listeners conduct their lives. Evidence also shows that auditory spatial awareness merges with visual spatial awareness, together creating a holistic spatial awareness—a high-level cognitive process.

An internal spatial image is a cognitive map of space—a private construction that includes a mental response to sensory stimuli modified by personal experience. Roger Downs and David Stea (1973) provided a basic definition of cognitive mapping as a “process composed of a series of psychological transformations by which an individual acquires, stores, recalls, and decodes information about the relative locations and attributes of the phenomena of everyday life.” A cognitive map of a space is a combination of the rules of geometry as well as knowledge about the physical world. It is this extra environmental knowledge that allows us to perceive a ball as moving away from us rather than as simply shrinking. This knowledge associates reverberation with enclosed space, echoes with remote surfaces, and high frequencies with hard objects. These associations are learned. Because this knowledge is acquired in childhood and continually modified in our experience as adults, we are not conscious of its existence. When sensing a spatial environment, an individual builds a cognitive map of space using a combination of sensory information and experiences accumulated over a lifetime. The cognitive map of space in our consciousness is subjective, distorted, and personalized—an active and synthetic creation—rather than a passive reaction to stimuli visual object. But when sensory attributes are not aligned, you experience two objects, one with visual and the other with aural attributes. As neuroscience uncovers details about specific brain substrates, we find that some intellectual abstractions, such as cognitive maps, have an observable manifestation in the brain.

In attempting to solidify the vast collection of experimental data on sensory fusion, the neuroscientist Alvaro Pascual-Leone (2000) took the concept one step further. He argued for a metamodel of the brain where neural substrates act as “operators” to

implement a given functionality regardless of the sensory modality. In his conceptualization, there would be a spatial operator in a brain substrate that operated on aural and visual cues to create an internal representation of space. Similarly, there would be an emotional operator that created an affective response to that same space. As a rule, an operator appears to be dominated by a particular sense modality. Thus, for a sighted individual, a spatial operator might be dominated by visual inputs, and for a blind individual, by auditory cues. For a deaf individual, a speech operator might be dominated by visual or tactile cues. Dominance is far from universal or complete, and operators incorporate inputs from multiple senses without explicit awareness. For example, Beatrice de Gelder and colleagues (1999) showed that the emotional responses to hearing a voice and viewing a face influenced each other when the emotional content of the two modalities was not in agreement. We might expect that the emotional responses to hearing space and seeing space influence each other as well.

The separation of a cognitive map from its sensory inputs is illustrated by how individuals imagine an object when they have no visual input. Oliver Sacks (2003) observed that some blind individuals experience “deep blindness,” an inability to imagine the shape of an object without tracing it, whereas other individuals experience a “hallucinatory visual world,” rich and full with real and imagined objects. In one case, the visual cortex had atrophied, whereas in the other case it remained active using a combination of inputs from internal memory and the aural and tactile senses. Some part of the visual cortex may actually serve as a spatial operator. Sacks (2003) commented, “studies on the effects of blindness on the human cortex have shown that functional changes [in brain substrates] may start to occur in a few days, and can become profound as the days stretch into months and years.” Even after being blindfolded for only a few hours, sighted subjects begin to experience changes in spatial and object images. These changes reflect a rewiring of the spatial operators, thereby compensating for the lack of visual input. An internal representation (cognitive map) of space depends on the way you teach your brain to use all your senses. For Sacks, the visual cortex is only the “inner eye,” a concept that has nothing to do with sight itself.

Auditory and tactile information also contribute to the functioning of this inner eye. Because we use a visual vocabulary to describe spatial experiences rather than a sensory-neutral language, we assume that spatial experiences are visual both in origin and in representation. In common discourse, the word map itself means a visual picture of an environment. In fact, the “inner eye” is, not visual, but multisensory, “seeing” the present combined with the past.

Consider that perceived size and distance are not just a visual measure of a physical reality but also involve subjective and personalized concepts derived from multisensory data. The experience of large distances is also an indirect consequence of experiencing time, as exemplified by the time to walk from one place to another, or by the time for an echo to return from a distant surface. The vastness of an enclosed space is revealed by decaying reverberation. In contrast to distances that can be experienced as the passing of time, small distances can be measured in terms of the length of an arm. You experience the size of a doorway opening, not in terms of a ruler measurement, but in terms of its ability to accommodate the width of your body when walking through the opening. In an earlier discussion, we explored the concept of the acoustic horizon, which is also a measure of distance, using social spheres as the metric. The aural, visual, and tactile experiences of space contain different perceptual units for size, which are then fused into a single spatial map. Conversely, a single map can be converted into different units of sensory size: the object is at arm’s length, it takes ten strides to reach, or it returns an

echo in 100 milliseconds. We should think of spatial cognition as the process of fusing and reconciling overlapping contributions from all sensory modalities.

Having established that size and distance are multisensory abstractions that are fused into a single cognitive map of space, we now turn to the issue of spatial relationships among objects and the perceiver's relationship to those objects. This, too, is an abstraction that depends on a given reference point or viewpoint. A cognitive map of space implies a spatial framework. At the most basic level, saying that a boy is standing in front of the tree implies a specific location for the viewer, but saying that the boy is standing north of the tree implies an abstract spatial reference independent of the viewer. Where is the boy? In the first case, the relative location of the boy changes if the observer changes location. In the second, the boy's relative location changes only if the environment, including the observer, is rotated relative to the reference frame. In all spatial experiences, there are two perspectives: allocentric, from which objects are perceived relative to a fixed external framework; and egocentric, from which objects are perceived relative to the perceiver. Rotate a concert hall and, depending on which perspective you adopt, the relative location of the orchestra either changes or remains the same. Although mathematically equivalent, in that one reference frame can be converted into the other, each perspective is experienced differently. For example, musicians are at the front of the concert hall (allocentric), but the person with a large hat is sitting in front of you (egocentric). Cognitive maps of space contain aspects of both perspectives, but emphases vary from culture to culture. Because an allocentric framework situates you within a fixed external environment, philosophically, it implies that reality exists apart from your self. In contrast, an egocentric framework situates your self at the center of an experiential universe, where everything is interpreted relative to you. A cognitive map of space can be egocentric, allocentric, or some combination of both. The choice of framework modifies the experience of space.

There is evidence that the brain contains substrates for encoding space in a multiplicity of allocentric and egocentric perspectives (Behrmann, 2000). Although neural substrates exist to support both perspectives, cultural values and personality biases usually emphasize one over the other. One culture's language and religion may focus on egocentric representations of space; another's may focus on allocentric representations. It is easy, but presumptuous, to expect cognitive maps of space to be consistent across cultures, or even across individuals from the same culture.

Because a cognitive map is, by definition, entirely private, we have access to it only by observing behavioral differences among cultures, such as difference in the language of space, or in the ability to perceive spatial attributes. Benjamin Lee Whorf (1956) first advanced the thesis that language influences how we experience life, and vice versa. Although still controversial, his thesis remains a major component of cognitive theories (Lucy, 1997).

A manifestation of differences in cognitive maps of space can be observed by analyzing a culture's language, and by testing individuals on behavioral tasks. As Stephen C. Levinson (1999) notes, some languages do not employ the spatial notion of left-right-front-back but rely on north-south-east-west. These differences are more than merely linguistic. They are fundamentally different ways of viewing the world and placing oneself into the world. The type of cognitive map of space changes one's behavior on spatial tasks. For example, on various tests, Dutch subjects consistently performed better at encoding and referencing relative locations, which is characteristic of modern

cultures, whereas Tenejapan Mayan subjects performed better at encoding absolute locations, which is better for navigation and orientation in natural spaces. Similarly, Levinson (1999) observed that modern European languages favor using self-referencing body parts to identify building sections, such as the head, wings, back, or face of a structure. Other languages refer to component parts using absolute references, such as seaward or northerly.

The discussion on cognitive maps of space demonstrates that we cannot consider the navigational spatiality of aural architecture in isolation. And just as aural architecture is an inseparable component of sensory architecture, so aural spatial imaging is inseparable from spatial awareness, which is a high-level cognitive process separate from specific sensory modalities. The creation of a navigational space depends on the cognitive map of the aural architect, just as auditory spatial awareness depends on the cognitive map of the listener. Both designer and listener have acquired their maps through experiences.

Unfortunately, cognitive maps of space are difficult to observe, even though they are central to spatial experience. Although the ability to use auditory spatial awareness for navigating space is present in all human beings with adequate hearing, the degree to which that awareness contributes to cognitive maps of space is specific to individuals and their cultures.

Architecture is more than the design of a utilitarian space; architecture is also an expressive art form that communicates. Using the broadest definition of architecture, we also include decorations, ornaments, adornments, and embellishments as important elements of spatial design. These elements are aesthetic supplements to the utility of the spaces we occupy or live within. Although they are traditionally considered part of interior design, they are as relevant to the experience of a space as the structural framework that encloses a space. Every picture, statue, tapestry, archway, mirror, dome, textured surface, and ceiling molding, to name but a few, is an architectural embellishment.

There are embellishments that produce or admit light, such as candles, chandeliers, or frosted-glass panels, and there are embellishments that absorb light, such as dark tapestries or black walnut panels. There is no functionality in the aesthetic aspects of these adornments—flat white walls illuminated by industrial lamps are adequate for ordinary living—yet such embellishments enhance aesthetics by creating a pleasant or reflective mood. They may also convey symbolic meaning, such as wealth, political power, social status, or historic legitimacy.

Architecture includes aural embellishments in the same way that it includes visual embellishments. For example, a space we encounter might contain water spouting from a fountain, birds singing in a cage, or wind chimes ringing in a summer breeze—active sound sources functioning as active aural embellishments for that space. Producing aural rather than visual illumination, these are the aural analogues of decorative candles and lamps. In contrast, passive aural embellishments, such as interleaved reflecting and absorbing panels that produce spatial aural texture, curved surfaces that focus sounds, or resonant alcoves that emphasize some frequencies over others, create distinct and unusual acoustics by passively influencing incident sounds. Passive aural embellishments are the aural analogues of pictures, tapestries, mirrors, arches, and statues.

For both visual and aural embellishments, there are two independent oppositions: active versus passive and local versus global. A water fountain and a resonant alcove

are both aural embellishments, but the first serves as an active source of sounds whereas the second passively filters them. Similarly, a candle and a mirror are both visual embellishments, but the first actively generates light whereas the second passively reflects it. Affecting only an area of a larger space, fountain, alcove, candle, and mirror alike are local embellishments. We experience them only when we are relatively close. In contrast, affecting the entire larger space, both reverberation and diffuse lighting are global embellishments. We experience them throughout the space. Parallels between visual and aural embellishments are not generally recognized because visual objects are most often local, whereas acoustic objects are most often global.

Almost every visual embellishment has some acoustic influence. Thus a mirror, a statue, or a tapestry changes the acoustics of the space around it. If these changed acoustics are unintended, their role as aural embellishments may not be recognized or appreciated. Nevertheless, they are relevant to our experience of aural space. A large mirrored wall reflecting light also functions as a perfect reflector of sound. An elegant tapestry absorbs sound and a marble statue diffuses it. Conversely, a sonic diffraction grating designed as an aural embellishment might also be considered as a modern visual sculpture. Depending on the sensibilities of the designer or the perceiver, every embellishment can be either visual, aural, or both at the same time.

Aural embellishments give a space an aural personality. Without them, every space, be it bathroom, concert hall, military barracks, or other space, would sound like every other space of similar size and shape. In addition, without local aural embellishments, every area of a space would be aurally indistinguishable from every other area of that space. When you move into a new house, you add personal touches—visual embellishments—by your selection of art and furniture, thus making the space of the house visually unique. By analogy, and for quite the same reason, you also add aural embellishments, whether intentionally or not. The antique rug that contributes visual elegance also adds aural warmth. Customizing a space to give it a unique and personal feel, perhaps to make it a symbol of yourself (Cooper, 1974), operates both aurally and visually.

We are now ready to define aural embellishment. It is an acoustical object or geometry, whether local or global, that produces aesthetically recognizable acoustic attributes, adding aural richness and texture to the space. An alcove in a cathedral is a local embellishment, providing aural privacy. Extensive carpets and thick drapes, by removing high frequencies from reverberation, are global embellishments that create an aural sense of warmth. As a generalization, aural embellishments produce acoustic attributes that are not related to the functional aspects of an acoustic arena, spatial navigation, or musical aesthetics.

Because of the extensive interest and research in the architecture of musical spaces, many assumptions that apply to those spaces have been implicitly carried over to other applications of aural architecture. In the design of a concert hall, aural embellishments are considered to produce unwelcome acoustic effects and should be avoided whenever possible. According to our musical norms, the aural experience of a concert hall should ideally be uniform throughout the space. The acoustic shadows produced by a balcony, for example, are tolerated, but unwelcome. Similarly, specific global aural embellishments are unwelcome because the acoustics of a musical space, as extensions of the musical instruments, should match the musical repertoire. In contrast, aural embellishments are welcome in a social or religious space, providing aural variety, symbolic meaning, and spatial texture. Just as Japanese Noh drama and ancient Chinese opera convey little to an

inexperienced audience without extensive exposure and knowledge, so aural embellishments may convey little to inexperienced listeners. All three are art forms and serve as evolving vehicles for expressing our relationship to ourselves, the world, and the cosmos. Understanding their message requires experience with the cultural symbols they use to convey it.

Spatial Distortions in Aural Geometry

An aural architect can design a space such that the acoustics at selected areas magnify the aurally perceived size, mass, and intensity of a speaker. Unlike optical magnification, however, acoustic enlargement is inconspicuous, arising from the shape of the enclosing surfaces. Strong sonic reflections arriving shortly after the direct sound increase the apparent aural size of the sound source. Even when the total sound energy remains constant, shifting energy to the early sonic reflections enlarges the perceived size. In contrast, late sonic reflections are perceived as echoes or reverberation, degrading intelligibility. Concentrating sound in time and space is one means of creating local acoustics in aural architecture.

The same phenomenon is well recognized in musical spaces. When early sonic reflections from the sidewalls and ceiling reflectors are appropriately combined, musical instruments on the stage of a concert hall sound closer—aurally larger—than they would otherwise. The musicians playing on stage are, by their special location, like a judge sitting on a dais, a politician or lecturer standing at a podium, or a minister preaching from a pulpit. These individuals are deemed to have socially dominant status; their special locations should have acoustics consistent with their dominant status, their relative social prestige. Thus, to symbolize the social relationship, the acoustics of the podium area in a lecture hall should raise the aural status of the speaker, whereas those of the auditorium should lower the aural status of the listeners.

The same natural amplification that increases the apparent size of the speaker also increases the size of the acoustic arena. In addition to sounding larger, the voice of the dominant speaker covers a wider acoustic arena, and is heard by a larger audience. A socially dominant location thus has two acoustic attributes: larger aural size and a larger acoustic arena. From extensive research on concert hall design, the knowledge required to create such local acoustics is well known and readily transfers to social and religious spaces. Architectural design therefore includes, intentionally or incidentally, the aural symbolism of dominance.

Just as a visual architect specifies the shape of the physical space, an aural architect specifies the shape both of the acoustic arenas and of the areas where aural magnification occurs. Whereas physical boundaries clearly delineate a visual shape for the space, through a more complex process, those same boundaries determine the shape of acoustic arenas. Shaped acoustic arenas thus become tools of aural architects. While most of the following examples are, unfortunately, unrelated to the more conventional social spaces, they do illustrate the wide range of freedoms available. An analysis of these examples shows that they are the optical analogues of placing numerous lenses and curved mirrors about the space.

When a space has curved surfaces, its acoustics can readily change the aurally perceived geometry of that space. Like the side mirror of an automobile warning that (visual) objects are closer (larger) than they appear, curved surfaces also change the apparent location of

aural objects. Particular curved surfaces can focus sound such that the source appears aurally closer or farther, larger or smaller. We can think of these curved surfaces as distortions of a circular acoustic arena. Curved surfaces can also produce acoustic dead zones such that a source is inaudible, as if it were in an acoustically isolated arena. Aural privacy does not require walls. In contrast, some curved surfaces can give you the aural impression that a speaker is sitting on your right or left shoulder. Science museums often demonstrate how a parabolic sound reflector displaces a speaker 30 meters (100 feet) away to an aurally perceived distance of 3 cm (1 inch)—a thousandfold shift in location.

The concept of shaping an acoustic arena for aural effect is not new. In the early part of the last century, Wallace Clement Sabine (1922) described numerous examples of “whispering galleries,” large enclosed spaces where a listener could hear the whisper of a speaker at remote distances. The more famous ones in Sabine’s time included the Dome of Saint Paul’s Cathedral in London, Statuary Hall in the Capitol at Washington, D.C., Saint John Lateran in Rome, and the Ear of Dionysius at Syracuse in Sicily. Most, if not all, of these whispering galleries are architectural accidents resulting from curved surfaces presumably designed for their visual impact. The time delay for the sound to return from the ceiling, combined with its focused direction, gives the visitor standing in the center of such a gallery the “effect of an invisible and mocking presence.” This is not an echo. If you are the visitor, the sound of the distant speaker’s voice is focused directly at you, as if the speaker were right next to you. The experience is unforgettable.

Similar effects are found with elliptical enclosures, such as the Mormon Tabernacle in Salt Lake City. In such spaces, if you stand at one of the two foci, you can readily converse with someone at the other. Using our spatial language, we can describe the situation as two physically separate regions of space that are joined by a bilateral auditory channel into a single acoustic arena. Even widely separated or oddly shaped physical spaces can be acoustically joined. In an example described by Sabine, the Cathedral of Girgenti in Sicily, by an unlucky coincidence, one focus is located at the confessional; secrets of the most intimate nature are broadcast to a remote location in the church. There is a story, assuredly apocryphal, that Benjamin Franklin eavesdropped on visiting government dignitaries by placing them at one focus of the Capitol’s Statuary Hall and himself at the other. Besides creating an acoustically joined pair of foci, an elliptically curved surface can also create a sonic conduit, as sound bounces along its periphery hugging the wall. Geometrically complex spaces have complex acoustic arenas.

As these examples illustrate, visual and acoustic arenas of the same physical space can differ, sometimes surprisingly. But even though specific geometric designs can create acoustic arenas and aural distances that differ dramatically from their visual counterparts, the aural experience of distance between a speaker and a listener can change when they enter even a simple enclosure. In more complex spaces, visual proximity can correspond to aural remoteness, just the opposite experience to that in whispering galleries.

In summary, aural architecture determines the aurally perceived size and location and the acoustic arena of a speaker in each area of a physical space. Although these acoustic properties have a social meaning, however unintended, the most impressive historical examples of aural architecture are famous chiefly as spatial curiosities. Indeed, there is little evidence that the architects intentionally designed the acoustic arenas of these spaces based on the social, navigational, or aesthetic needs of those who were to use them.

Illusions of Expanded Spaces

Windows, mirrors, and pictures belong to a specific class of architectural embellishments: visual space manipulators. A window expands visual space by establishing a visual connection between the observer and an additional physical space; a mirror expands space by connecting the observer to a replica of the existing space; and a picture expands space by inserting the image of another environment. The size of the window, mirror, or picture determines the degree of coupling between two or more spaces.

When physical constraints force a traditional architect to work within a limited space, the art of visual illusion becomes important in making a space seem larger than it is. Mirrors, in particular, are visual space expanders: they create the visual illusion of added space. Small rooms with many mirrors give the impression of being far larger than their actual size. A mirror is a window into a virtual copy of the same room, located on the other side of the wall. The experience of the enclosing surfaces then disappears.

With mirrors on multiple surfaces, as in dance studios, replicated virtual spaces grow exponentially as if the visual space were infinite. Having drawn analogies between seeing and hearing, we can search for aural parallels to these visual space expanders. Although, from the physical perspective, light and sound waves closely parallel each other, from the experiential perspective, they diverge widely.

To understand this divergence between physics and experience, consider a crackling (noisy) candle, emitting both light and sound energy in a room with a mirror surface that reflects both forms of energy. The light and sound energy waves radiate spherically, following the same trajectories and producing the same reflections from the mirror. An observer sees the candle and its replicated image in the mirror. The image is equivalent to a virtual candle located in a virtual space. Similarly, a listener hears the direct sound from this noisy candle and hears the reflected sound from the mirror, which is physically equivalent to the sound that would have radiated from a virtual crackling candle in a virtual space. The optical and acoustical phenomena parallel each other closely. Indeed, parallels between light and sound in enclosed spaces are found in most elementary textbooks on spatial acoustics.

Because of physiological differences between hearing and seeing, however, the experience of reflected light and that of reflected sound diverge. Whereas multiple sonic reflections are generally perceived as a single fused sonic event even when sound arrives from different directions and at different times, multiple visual reflections always remain distinct. Under normal circumstances, aurally, we would perceive only a single noisy candle in our example, along with the reflecting wall; visually, we would perceive two candles, an actual and a virtual one.

A sonic reflection creates the illusion, not of a new virtual candle in a new virtual space, but rather of a louder (aurally larger) noisy candle—and it induces the aural perception of a solid wall. If the delay between the direct sound and its reflection is large enough to produce a distinct echo, and if we experience the echo as unbound from the direct sound, then, and only then, the sonic reflection creates the aural illusion of a separate virtual candle. But normally, we experience a distinct echo as bound to the original sound. A sound-reflecting surface is the aural equivalent of an opaque wall—a spatial boundary, a spatial reducer.

What kinds of acoustic objects and designs create the aural illusion of a larger space? What are the aural analogues to mirrors, pictures, and windows in creating this illusion? Unfortunately, such analogues remain in the hypothetical realm: they have not yet been realized in physical spaces.

To create the aural illusion of an expanded space, we must simulate the sound field at a virtual window, that is, we must replicate the sound field that would have been present if an additional space were actually present. Sound absorption is an aural space expander. Complete sound absorption would simulate a virtual window into an infinite, unbounded space, a space without the ability to respond to sonic illumination, and with no sound sources of its own. Thus a thick panel of dense, completely sound-absorbing materials, one that could absorb all sound waves that arrive, would aurally replicate a window into an absolutely open space. Sound arriving at the panel would completely disappear, as if it had actually encountered an open window into an absolutely open space—an infinite void.

But if the virtual space is to be the equivalent of an actual room, rather than an infinite void, the appended space must have its own sound-reflecting surfaces, sound absorption, and sound sources. The virtual space would reverberate sound entering from the real space through the virtual window. To experience the appended space as an actual environment, we would need to reproduce the appropriate sound field at the window. Hypothetically, we might create this illusion in a sequence of stages. First, to create the sound field, we might embed an array of small loudspeakers driven by a spatial synthesizer in a sound-absorbing panel. These loudspeakers would then duplicate at the surface of the panel the sound field of a space as it would appear at the virtual window. We might simulate the sounds of a bird sanctuary with chirping birds and babbling brooks together with its acoustics, including reverberation and sonic reflections. Our simulation would need to replicate the sound field only at the virtual window since listeners could not actually enter the virtual space. Walking near the panel with their eyes closed, they would have the impression of a window opening onto a bird sanctuary. The aural experience would be analogous to a visual picture of a bird sanctuary. In fact, if we had the panel also contain a visual display of the sanctuary, we would have a multisensory space expander.

Second, to refine our virtual window onto a virtual space, we would need the virtual space to respond to sound originating from the actual room. If the bird sanctuary were a real space, listeners could shout through the window into it and then hear the reverberation of their voices. Hypothetically, this is also possible. We might embed an array of microphones into the panel such that the sound waves arriving at that surface would feed a spatial synthesizer that created the virtual reverberation, which the loudspeakers would then reproduce.

Third, to make our virtual space simulate an extension of our actual room, we might expand the area of the sound-absorbing panel to cover an entire wall such that sound arriving from the actual room would be completely absorbed. This would effectively remove the aural perception of a wall. Sound would impinge on the absorbing surface and disappear. We might then have the spatial synthesizer create the sound field at the surface that would have been there had the actual room been larger. We might, for example, have the synthesizer add a delay of 10 milliseconds to the sound that arrived at the wall. The sound field would then be the same as that of an actual room 3 meters (10 feet) wider, with a wall 3 meters farther away. Or, using the same approach, we might simulate still larger and more complex spaces to create the illusions of larger and more complex actual rooms.

Our scenario is compelling if we assume the synthesized sound field could be made identical to its natural counterpart, paralleling an optical hologram, which re-creates the light field of an actual object at the surface of the holographic image. Primitive versions of an artificial acoustic wall have been demonstrated in the laboratory, but the technology has not yet sufficiently evolved to make the dream practical. I have no doubt that such hypothetical scenarios will eventually become reality if technology continues to advance at its current rate. Primitive versions are currently used to make musical spaces feel larger and more reverberant. Many concert halls now incorporate active acoustics with arrays of microphones and loudspeakers. Eventually, perhaps within a decade, aural architects will be able to use “acoustic holography” as an additional tool to create virtual space expanders.

Local Anomalies as Aural Texture

Nobody remembers a visual space that is without unique features. A rectangular room with blank walls and minimal furnishing acquires a unique visual personality only when embellishments such as pictures, wallpaper, colored surfaces, and mirrors are added. Likewise, a prosaic aural space acquires an aural personality only when aural embellishments are included. Openings such as windows and alcoves add aural personality; by absorbing sounds, thick drapes, large tapestries, and upholstered furniture create aural texture, as do statues, pillars, and complex geometries, which diffuse sounds. Such aural embellishments create local acoustic attributes, supplementing global ones such as reverberation. The aural personality of a space is especially apparent to blind persons, who experience embellishments chiefly by listening.

Although the concept of a local aural embellishment is not yet recognized as such, we can easily demonstrate its role in creating a personality for spaces. As children, many of us first experienced an aural embellishment when we placed a conch shell to our ear and listened to the sounds emanating from inside it. Because of the shell’s complex inner hollows and passageways, its interior creates resonances that filter background noise to produce a sound that resembles that of the ocean. The region of space near the opening of the shell creates an acoustic anomaly—a spatial filter that changes the spectrum of the background sound. The conch shell is a miniaturized version of a cave or alcove, which is also a hollow that can be experienced at its opening. There are examples of acoustic hollows other than caves and conch shells. Objects in the shape of a large vase with a narrow neck, called “Helmholtz resonators,” change the background sound at their openings. Depending on their construction, they can amplify or suppress particular frequencies. Archaeological and written evidence from ancient Greece and into the Middle Ages indicates that theaters and churches once had acoustic vases scattered about their auditoriums. Although scholars still argue about how effective such vases may have been in enhancing voices, those sitting or standing close enough would have likely heard them as some sort of aural embellishments.

The acoustic vase is the man-made equivalent of a conch shell, but with different resonant properties.

To appreciate the extent to which acoustic objects can create aural texture, first consider the visual analogy. Wallpaper produces visual texture because of nonuniformities in its visual pattern. At a distance, the details of the pattern may not be visible but they still create a texture that is quite different from a painted surface. When all surfaces of a space have the same hue, intensity, saturation, and reflectivity, the environment is visually sterile,

in contrast to the effect of elegantly decorated and richly textured wallpaper.

As the aural analogue of wallpaper, consider a wall that had a pattern of conch shells embedded in it, thus creating a pattern of resonances at different frequencies— like variations in aural color. Such a wall would have aural texture. By standing at the optimum distance, you would hear that texture. This example illustrates how small objects, each of which cannot be perceived individually, can be multiplied and extended to produce aural texture. We can take the idea further. An aural pattern might include small regions of absorbing mats, planar reflectors, dispersing wedges, and diffraction gratings. The art of aural wallpaper is as unlimited as that of its visual counterpart.

Besides creating a large acoustic surface from an array of small acoustic elements, we might also design larger acoustic objects that have a recognizable aural personality—the aural version of modern sculpture. After a search of the architecture literature, however, I failed to find any examples of acoustic objects characterized as aural embellishments. Yet many artistic and religious objects have acoustic properties that match our definition of aural embellishments, even though they were never intended to be aural art or acoustic sculptures. Although creative artists can design such objects for their explicit impact on listeners, there are also vast repositories of historical artifacts that have unusual acoustics. Combining mastery of both archaeology and acoustics, acoustic archaeologists have discovered ample physical evidence in ancient sites that older cultures valued objects and structures for their acoustics. Leaving extensive discussion of the cultural relevance and symbolic meaning of these objects and structures to chapter 3, let us briefly consider three examples of unintentional aural embellishments. Our first is from the Mayan culture. Acoustic consultant David Lubman (1998) discovered that, when illuminated by the sound of clapping hands at a particular location, the staircases at the Pyramid of Kukulka'n at Chiche'n Itza' produce chirplike echoes that bear an uncanny resemblance to the call of the Mayans' sacred bird, the resplendent Quetzal. This readily perceived resemblance most likely invested the staircases with special religious meaning.

Our second example, also a religious one, is the medieval shrine to Saint Werburgh in Chester, England. As described by Lubman (2004), the shrine's six recesses, where kneeling pilgrims would insert their heads while pleading their petitions, serve both as amplifiers and as filters, giving the petitioners' voices dramatic and emotional emphasis with only modest vocal effort. The shrine's recesses thus create uniquely private acoustic arenas that exclude external sounds without walls. (Their modern social counterparts might be alcoves designed for the aural intimacy of lovers.) Our third example is a sculpture by the respected twentieth-century Spanish minimalist artist Eusebio Sempere. Composed of a three-dimensional array of polished stainless-steel tubes, the sculpture rotates at its base, as shown in figure 2.4. The moving surfaces serve not only to dramatically reflect the sunlight but also to selectively filter transmission of particular frequencies of sound. Listeners on one side hear a tonal modification of sounds coming from the other side—the moving surfaces acting like the aural equivalent of colored glass prisms. Although scholars took several decades to recognize the sculpture's acoustic properties (Ma'rtinez-Sala et al., 1995; Sa'nchez-Pe'rez et al., 1998), there can be little doubt that, by changing the sounds that propagate through it, Sempere's work serves as an aural embellishment. If the artist had had a background in acoustics, we would assume he had intended to design a multisensory sculpture.

A search of the literature revealed that the phrase “aural sculpture” applies almost exclusively to experimental art based on active sound sources, often interacting with the

listeners and often prerecorded. Artists are sculpting the sound field by manipulating sources and their location. I did not find any reference, even using alternative search phrases, to any form of aural sculpture experienced by illuminating an object with the natural sounds of a living environment. Most likely, the aural effect is too subtle for a population more familiar with high-impact computer-generated sounds that do not occur in nature.

The previous examples also illustrate a plausible process by which aural embellishments come into existence. Without any formal knowledge of acoustics and aural perception, an artist creates an aural embellishment as an unintentional artifact of another design process. We, the listeners, are then left to sort and evaluate objects for their aural aesthetics. But with the appropriate knowledge, an artist can also explicitly create aural sculpture. And if that art, however created, is then included in a space, it becomes an embellishment of aural architecture. Nevertheless, the aesthetic value and symbolic meaning of these aural embellishments still depend on the attitudes of those who listen to them. Aural adornments can be overlooked, barely noticed, or even dismissed, appreciated or even revered.

Although aesthetic space, like social, navigational and musical space, is always a reflection of the prevailing culture, even when not recognized by auditory experts and professional architects, aesthetically pleasing aural spaces and their aural embellishment may still arise. They are there to be discovered. And they may be consciously experienced by those who have developed a refined sense of aural spatial awareness.

The Affect of Enveloping Reverberation

Aesthetically pleasing at an appropriate level in musical spaces and the label for millions of sonic reflections, reverberation can be mixed blessing in ordinary living or gathering spaces. Excessive reverberation degrades the intelligibility of spoken communication, raises the background noise level, and makes a living or gathering space aurally unpleasant, whereas inadequate reverberation makes a space seem aurally dead, unresponsive, and uninviting. Energy in the late-arriving sonic reflections reduces the size of the arena by creating corrosive noise, whereas energy in the early-arriving sonic reflections increases the size of the arena by amplifying and focusing a speaker's voice. In an unenclosed space with no sonic reflections, oral communication between speaker and listener is difficult unless they are close to and facing each other.

Each specific area of a space may have its own reverberation profile. An alcove with deep-pile rugs and a low ceiling has less reverberant energy than the large open space to which it is connected. The acoustic properties of a space are locally distinct when the profiles of early and late sonic reflections are not uniform throughout the space. Indeed, reverberation is uniform only when the space is large, open, and acoustically uniform. This is desirable for performance spaces, such as concert halls, but not necessarily appropriate for other kinds of spaces.

The physical properties of reverberation tell us little about their experiential meaning. From a social perspective, reverberation does not intrinsically produce a specific affect; rather, the affect is indirectly determined by the listeners' aural expectations. Spaces that match the listeners' aural expectations are pleasing to them; spaces that do not are not. Listeners have expectations about the way that reverberation should respond to sonic events (responsiveness), and about the way that reverberation should create acoustic arenas of particular sizes (social spheres). We cannot specify what the listeners' affective response

to reverberation will be—whether stress, anxiety, comfort, or well-being—without examining the social context; this aspect of reverberation in aural architecture is culturally relative.

Aside from its influence on acoustic arena size and listeners' spatial responsiveness, reverberation is unlike all other sounds. Because enveloping reverberation cannot be localized as a sound originating from a particular place, we refer to it as “enveloping aural ambience.” Just as we experience water visually, tactilely, and aurally as an enveloping environment when scuba diving, so we experience reverberation aurally as an enveloping environment when we find ourselves within it. The difference in affect between being underwater and on dry land parallels the difference between a cathedral and an anechoic chamber. For this reason, reverberation has an affective component apart from its associations with social expectations. How then is enveloping reverberation experienced, what properties should it have, and what role should it have in aural architecture?

The ability to determine the location and direction of a sound has undeniable survival value. When you hear a stampeding herd of animals, knowing which way to run can be a matter of life or death. If the natural acoustics of forest or savanna destroyed their ability to locate the direction of an approaching herd or predator, our ancestors would very likely never have survived. Fortunately, natural environments typically produce low-level sonic reflections, not enveloping reverberation. Over millions of years, our auditory cortex evolved the means to determine the location and direction of a sound source by using the direct sound, which is reliable, while disregarding sonic reflections from a multiplicity of surfaces. That process fails, however, when enveloping reverberation from enclosed spaces completely overwhelms a weak direct sound. Evolution could not adapt to the reverberation of enclosed spaces because they were the exception rather than the norm.

In a modern context, the ability to aurally localize the blaring siren of a fire truck in an acoustically complex metropolis is central to a city driver's safety. Moreover, the driver's feelings of anxiety upon hearing the siren are instantaneous and automatic. Unlocalized sounds are associated with potential danger; danger triggers either anxiety or a heightened state of arousal, which is a biological state of enhanced alertness. Reverberation gives rise to an interactive experience, with the space entering into an acoustic dialogue with its occupants. It is difficult to enter a reverberant space surreptitiously

because the sound of your footsteps produces an acoustic reaction for all to hear. Metaphorically, the reverberated sound of footsteps is the reactive voice of the space; the spatial acoustics of a reverberant space announce the presence of active life by responding with an audible hello, as either a whisper or a shout. The acoustics are like the voice of a receptionist, with aural architects determining how that voice should greet entering visitors. Aesthetically pleasing reverberation produces a dialogue that is neither unresponsive nor domineering—a pleasant voice. (As we will see in chapter 3, the idea that a space has a voice provides a plausible explanation of how prescientific cultures experienced spatial acoustics.)

Or, to use the metaphor of dining out, enter a space, and it responds to your footsteps with a serving of reverberation. But unlike dining in a restaurant, you cannot choose the taste of reverberation from a menu. But if you could select its taste, what would you choose? The clearest distinction among the choices involves the frequency content: at every frequency, reverberation fades away slower or faster—has a different decay time. Frequencies that

last longest dominate tonal color because the other frequencies have already decayed to inaudibility. Ideally, you would choose tonal color to match your mood and aesthetic taste. From our ordinary experience as Western listeners, we acquire associations to tonal color. We associate low frequencies with objects that are soft or malleable, and high frequencies with objects that are hard or brittle (Freed, 1990). In the language of experience, the two categories are often called “warm” and “cold,” respectively, even though they are unrelated to temperature. Although connections between physical objects and tonal color are, no doubt, learned, they are consistent across large populations for one simple reason. Objects that are soft and malleable, such as wood or fiber, produce weaker high frequencies when bent, hammered, or otherwise manipulated.

Hard materials, such as glass, steel, or porcelain, produce stronger high frequencies. The two categories of objects absorb sound in the same way that they create sound. A room with a deep-pile rug is heard as warm and soft; a barren room with hard plaster walls is heard as cold and hard. Interior decorating, which is part of aural architecture, determines the tonal color of reverberation. To the extent that enveloping reverberation is analogous to being underwater, tonal color can be thought of as the water temperature.

Although smaller spaces still produce reverberation, as a listening visitor, you experience it as changing the tonal color of the direct sound, not as enveloping you. The acoustic dialogue between you and the space changes, but it remains a dialogue nevertheless. The spatial acoustics of a shower stall may induce you to sing because a small space has numerous discrete resonances. When the pitch and overtones of your voice coincide with these resonances, its loudness is greatly enhanced; when they shift away from the resonances, the intensity of your voice decreases dramatically. Rather than remaining neutral, the space reacts to the presence of some frequencies and not to others. Spaces may thus be said to have tonal preferences. A singer is an aural detective exploring an environment the way a child explores a toy. Even though a space reacts to all sonic events with its own characteristic response, nobody from our modern cultures imagines that an enclosed space is actually alive

Using a similar concept, but without realizing that it still applies today, acoustic archaeologists speculate that ancient shamans heard cave acoustics as the voice of a cave's spirit. In ancient cultures, objects were animate, containing living spirits. Although, in modern terms, spatial acoustics have replaced animating spirits in describing the aural personality of a space, nevertheless, I prefer to believe that, however subliminally, some sense of spirits animating spaces resides within us even now.

Having explored some of the experiential attributes of auditory spatial awareness, we are now in a position to examine their relevance to aural architecture. Depending on which cognitive strategy they adopt, those who occupy or live within a space can experience it in any of four distinct modes: social, as an arena for community cohesion; navigational, as local objects and geometries that combine into a spatial image; aesthetic, as an enhanced aesthetic texture; and musical as an artistic extension of instruments. The four modes exist simultaneously for all listeners even if some listeners are aware of only one or two of them. Both the aural architect and the occupants or inhabitants of a space decide on the relevance of each mode, whether consciously or unconsciously. We experience a concert hall, for example, primarily as a musical space, but should the lights fail, we almost certainly would experience it as a navigational space as we tried to find an exit. When small tables and chairs replace the audience seats during Boston Pops performances, we experience a concert hall as a social space. And when attending to the local acoustics produced by statues and alcoves, we experience the hall as an aesthetic space.

Although the aural architect focuses on particular aspects of the aural design of any space, those who use the space control the nature of their aural experience. As a listener, you may be aware of the large spatial volume created by a high-domed ceiling at a given moment, but using those same cues at another moment, you may experience reverberation only as the blending of individual sounds. You aurally sense the location of nearby stairs, doors, walls, and low-hanging chandeliers; and when talking to your partner, you respond to an acoustic arena that is mismatched to the social sphere. Furthermore, those who use space also determine, consciously or unconsciously, its sonic illumination, which in turn influences their experience. A musical space requires music, a social space requires people having conversations, and a navigational space requires transient and continuous background noise. The inhabitants then are the final aural architects of a space.

When a space is being designed, the aural architect must balance how the range of physical properties specified by the acoustic engineer influences various aspects of experiential space: social, navigational, aesthetic, and musical. In many cases, spatial attributes produce conflicting experiences. Large, open spaces are weak on acoustic attributes that enhance navigational cues and local acoustic embellishments. Aural privacy in a multiplicity of small acoustic arenas conflicts with having a single public acoustic arena. A space with a socially dominant region that magnifies a speaker's aural size conflicts with an egalitarian space having uniform acoustics throughout. Conflicting requirements call for choices. For the aural architect, these choices depend on the values of the sponsors, as well as on the expected use of the space.

Of all attributes, throughout the history of architecture, the size of an enclosing space is, perhaps, the major source of conflict. Motivated by theology, economics, or politics, the need for large audiences dominates the architecture of public spaces. An intimate space for chamber music with an audience of 6,000 is impossible. For the same reason, the Protestant Reformation shifted to smaller churches, in part, as the means to elevate the importance of the spoken liturgy, which would have been unintelligible in the acoustics of a large cathedral.

Aural and visual architecture converge insofar as every object and every geometric shape has both visual and aural attributes. Because, however, we experience many architectural elements with more than one of our senses, not all of which can be best served at the same time, architects must make sensory trade-offs, which vary from culture to culture. For example, an open window couples one space to another by allowing the passage of light and air. But that same opening also provides a path for extraneous noise, and the opening functions as a perfect sound absorber with no reflected energy. Windows are thus multisensory acoustic structures. Similarly, statues are aesthetically pleasing to the eye as sculpture, but they also diffuse sound and may therefore affect the acoustics of a musical space. Panels suspended from the ceiling may produce welcome amplification through early sonic reflections, but may also produce an unwelcome visual sense of confinement. Where diffusion of sound is desirable, using an acoustic diffraction grating may simply be too visually unaesthetic to include in a space.

The aural and visual architecture of a space may diverge in other ways. Visual illumination is determined by the way that architects place lamps and windows; light sources are mostly static and built into the spatial design. In contrast, sonic illumination is mostly a consequence of some human activity. As a rule, then, aural architects have less influence than visual architects do over illuminating energy. As with any rule, however, there are

clear exceptions: visual architects sometimes give control of visual illumination to the users of a space and aural architects sometimes assume control of sonic illumination.

Aural architecture can influence, both directly and indirectly, the mood and emotions of those who occupy or live within a space. Such influence can be the direct consequence of how the space changes sounds: amplifying background noise to an uncomfortable level, creating enveloping reverberation, destroying aural localization cues, or pleasantly blending a sequence of musical notes. In these cases, listeners are responding to sounds modified by the aural architecture. And it can also be the indirect consequence of spatial acoustics: acoustic arenas that are too small to include the companion of a listener within the social sphere, a listener's personal associations to familiar aural embellishments, or a listener's comfort at navigating a space in the dark using strong aural cues. Listeners' responses to a space thus depend on the direct and indirect manifestations of spatial acoustics, as well as on culture and context and the listeners' individual biases, histories, and personalities.

In controlling the sonic illumination as part of the design process, an aural architect becomes a soundscape architect. This is seldom possible, however, because the dynamic and ephemeral activities of those who use a space are the dominant source of sound. Yet in certain art forms, the artist is also allowed to control sound. Japanese garden design, an ancient art form that stylizes and miniaturizes natural environments by creating the illusion of larger ones, includes the aural experience of space. Not only are objects and plants arranged for their visual pattern, but also for their ability to shadow and reflect sound from active sources. David A. Slawson (1987) mentions how muffling the sound of a waterfall makes it seem farther away, thereby enlarging the perceived size of the garden. By including the aural experience in its design, a Japanese garden becomes the artistic union of a landscape and a soundscape, and its designer a truly multisensory architect.